

WebSphere Application Server Enterprise Edition
Component Broker



Programming Guide

Version 3.0

WebSphere Application Server Enterprise Edition
Component Broker



Programming Guide

Version 3.0

Note

Before using this information and the product it supports, be sure to read the general information under "Notices" on page 531.

Sixth Edition (August, 1999)

This edition applies to:

WebSphere Application Server Version 3.0, Enterprise Edition Development Toolkit for AIX, program number 5765-E27

WebSphere Application Server Version 3.0, Enterprise Edition Development Toolkit for Windows NT, program number 5639-I07

WebSphere Application Server Version 3.0, Enterprise Edition for AIX, program number 5765-E28

WebSphere Application Server Version 3.0, Enterprise Edition for Windows NT and Gradient DCE, program number 5639-I08

WebSphere Application Server Version 3.0, Enterprise Edition for Windows NT and IBM DCE, program number 5639-I09

WebSphere Application Server Version 3.0, Enterprise Edition Encina Client for Windows NT/Windows 95, program number 5639-I10

WebSphere Application Server Version 3.0, Enterprise Edition MQSeries Application Adapter, program number 5639-I11

WebSphere Application Server Version 3.0, Enterprise Edition Oracle Application Adapter, program number 5639-I12

WebSphere Application Server Version 3.0, Enterprise Edition CICS/IMS Application Adapter, program number 5639-I13

and to all subsequent versions, releases, and modifications until otherwise indicated in new editions. Consult the latest edition of the applicable system bibliography for current information on these products.

Order publications through your IBM representative or through the IBM branch office serving your locality.

© **Copyright International Business Machines Corporation 1997, 1999. All rights reserved.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About this book.	ix	Object model	37
Who should read this book	ix	Applications	46
Conventions used in this book	ix	Key observations.	48
Notation	x		
Interface inheritance.	xi	Chapter 3. The managed object	
Implementation inheritance	xi	framework.	51
Associations	xi	Managed and non-managed objects	52
Aggregation	xii	Understanding MOFW objects	55
Qualified association	xiii	Component basics	55
State transition diagrams	xiv	Component state	57
Data flow diagrams	xv	Component attributes	58
How to send your comments.	xv	MOFW client programming model.	59
		Client view of Component Broker	
		applications	60
Chapter 1. Introduction	1	MOFW server programming model	62
Object-oriented application development		Application objects	62
paradigms	2	Using statics in business objects.	64
Forward engineering new functional		Circular references	65
requirements (top-down)	3	Releasing and deleting objects	65
Reverse engineering from the legacy		Navigating the name space using the	
(bottom-up)	4	Naming Service	70
Meet in the middle, incremental		A note on security	72
development	8	More copy helper topics	72
Combining the approaches	10	Using copy helpers	73
Programming roles and responsibilities	11	Copy helpers – sharing opportunities	75
Programming languages and conventions	11	Copy helpers and complex data types	76
Interface Definition Language	12	Moving object data in bulk	78
Java	13	Multiple interfaces to business objects.	80
C++	13	Transactions	81
Other languages (Windows only)	13	CosTransactions module	81
Naming conventions	14	Character set considerations	83
Coding conventions	14		
Application architecture	15	Chapter 4. MOFW C++ client	
Multi-tier architecture	15	programming model	85
Middle-tier architecture	17	Client programming model: basic tasks	86
Component overview	19	Initializing the client environment	86
Components	19	Finding a managed object.	87
Objects	23	Finding a managed object using the	
Another road map	29	Naming Service	88
		Finding a managed object using the	
		primary key	89
		Finding a managed object using another	
		managed object	91
		Using a managed object	91
		Creating a managed object	92
Chapter 2. Personal life insurance			
application example	31		
The object model (top-down).	32		
The application model (meet in the middle)	33		
Design model (bottom-up)	34		
Model details	37		

Creating a new object – create from key	92	Conventions and guidelines	151
Creating a new object – create from copy	94	Finding persistent objects	152
Using sets of objects.	95	Creating persistent objects.	152
Transient sets	98	The create_object() method	153
Specifying reference collection interfaces	98		
Remembering your favorite objects.	98		
Coding tips for proper CORBA memory		Chapter 7. MOFW - C++ server	
management	100	programming model – advanced	
Using object references.	100	concepts	157
Summary: The client programmer’s check		Extending a business object	157
list	101	Extending business object interfaces	158
		Essential state extensions	159
		Choosing an inheritance pattern.	160
		Implement the additional business logic	161
		Meet the MOFW IManageable	
		requirements	163
		MOFW requirements – IManagedServer	165
		More key classes	167
		More copy helper classes	167
		Extension summary	168
		Other variations to consider	168
		Object relationships	168
		Cardinality-1 relationships	169
		Cardinality-N relationships	184
		Creating specialized homes	192
		Extending the interface to IHome	192
		Implement the extended IHome interface	194
		Overriding specific methods on	
		specialized homes	199
		Thread safety	200
		Summary of home extension	201
		Creating UUID specialized homes	201
		Extending the interface.	201
		Implement the extended IHome interface	203
		CICS and IMS application adaptor	
		exception handling	205
		Chapter 8. MOFW – ActiveX client	
		programming model	211
		ActiveX client view of Component Broker	
		applications	211
		Developing the Component Broker ActiveX	
		client.	213
		Generating and registering DLLs	213
		Unregistering and moving DLLs	214
		Component Broker ActiveX client	
		application development information	215
		Exception handling	216
		Client programming model: basic tasks	216
		Initializing the Component Broker client	
		environment	216

Finding a managed object	217	Java exception handling	247
Finding a managed object using the Naming Service	218	Exiting Java applications when using the Abstract Window Toolkit	247
Finding a managed object using the primary key	219	Summary: The client programmer's check list	248
Finding a managed object using another managed object	221	Chapter 10. Java server programming	
Using a managed object	221	model	251
Creating a managed object	221	Developing a component	251
Creating a new object – create from key	221	Developing an interface to the business object	253
Creating a new object - create from copy	223	Selecting a pattern for handling essential state	258
Using sets of objects.	225	Implement business object methods	262
Transient sets	226	Implement the IManageable required methods.	268
Specifying reference collection interfaces	226	Implementing IManagedObject required methods.	270
Remembering your favorite objects.	227	Implementing the primary key class	273
Releasing and deleting objects	228	Implementing the optional copy helper class	275
When references explode	228	Loading C++ DLLs from Java business objects	277
Summary: The client programmer's check list	228	Additional information component creators should know	278
Chapter 9. MOFW - Java client		Where to next?	279
programming model	231	Chapter 11. MOFW - Java client	
Java client view of Component Broker applications	231	programming model - advanced concepts	281
Preparing to use VisualAge for Java for development of Component Broker Java clients	234	Transactions	281
Preparing managed objects for remote access	234	Transactions, exceptions and time-outs	284
Client programming model: basic tasks	235	Using transactions over reference collections	286
Initializing the Component Broker client environment	235	Session Service	286
Initialization requirements.	237	A simple example	287
Finding a managed object	237	Queries, iterations and specialized homes	289
Finding a managed object using the Naming Service	238	Using iterated homes-specific functions	290
Finding a managed object using the primary key	239	Using queryable homes-specific functions	295
Finding a managed object using another managed object	240	Using atomic transactions with query evaluator	297
Using a managed object	241	More on iterators.	297
Creating a new object	241	Using keyed reference collections	298
Creating a new object - create from key	241	Conventions and guidelines	301
Creating a new object - create from copy	242	Finding persistent objects	302
Using sets of objects.	243	Creating persistent objects.	302
Transient sets	244	Chapter 12. MOFW - Java server	
Specifying reference collection interfaces	244	programming model - advanced concepts	305
Remembering your favorite objects.	245	Extending a business object	305
Releasing and deleting objects	246	Extending business object interfaces	306
When references explode	246		

Essential state extensions	307	Managed objects and specialized homes	360
Choosing an inheritance pattern.	308	Data object implementation details	361
Implement the additional business logic	309	BOIM data object customization – static	
Meet the MOFW IManageable		SQL	364
requirements	311	BOIM data object customization – Cache	
MOFW requirements - IManagedServer	312	Service	371
More key classes	314	Transient data object customization –	
More copy helper classes	315	UUID key (production use)	379
Extension summary	315	Transient data object – any key	
Other variations to consider	316	(production use)	381
Object relationships	316	Summary of DataObject customization	386
Cardinality-1 relationships	317	Data object customization and	
Cardinality-N relationships	330	inheritance	389
Creating specialized homes	338	Data object customization for cardinality	
Extending the interface to IHome	339	relations.	395
Implement the extended IHome interface	340	Example scenario -	
Overriding specific methods on		createFromPrimaryKeyString	405
specialized homes	345	Packaging the application	406
Thread safety	346	Packaging for client and server (VA C++)	406
Summary of home extension	346	Packaging the DLL for the ActiveX	
Creating UUID specialized homes	346	Visual C++ Client	409
Extending the interface.	346	Packaging the Java client code	410
Details	347	Assembling and installing Java	
The implementation.	347	components	410
Meet MOFW IManageable requirements	349	Enabling additional clients	417
MOFW Requirements - IManagedObject		Configuring the CBConnector run time	418
interfaces	349	Introduction to container configuration	
Keys	349	parameters	418
Copy helper	349	Typical settings for container	
Leveraging server-provided essential		configuration parameters	425
state extensions	349	Summary of supported container	
Overriding the standard create interface	349	configurations.	430
		Configuring homes	433
Chapter 13. Assembling and installing		Expanding the client programming interface	434
components on Component		Application adapter quality of service	
Broker/Workstation	351	interfaces	435
Local only object implementation details	351	Using QOS interfaces for	
C++ local only	352	non-transactional support	437
Java local only	353	Behind the scene - an overview of	
Managed object implementation details	354	application adaptors	438
Business object methods in the managed		Behind the scene - the BOIM application	
object implementation	358	adaptor	438
OMG services methods in the managed		Adapting applications	439
object implementation	359	Managed object assembly	439
Application adaptor methods in the		Managed object	439
managed object implementation.	359	Business object	439
Special methods in the managed object		Data object	440
implementation	359	The life cycle of components	440
Handling component augmentation of		Qualities of service	440
OMG services methods.	360	Framework flows	442

Appendix A. Artifacts produced in building objects. 445

Appendix B. Interface Definition Language 447

IDL name scoping 447

Type and constant declarations 448

- Integral types 449
- Floating point types 449
- Character type 449
- Boolean type 449
- Octet type 449
- Any type 449
- Constructed types (struct, union, enum) 450
- Union type 450
- Template types (sequences and strings) 451
- Arrays 452
- Object types 452
- Constants 452

Interface declarations 453

- Constant, type, and exception declarations within an interface 454

Operation declarations 454

- "oneway" keyword 455
- Parameter list 455
- "raises" expression 456
- "context" expression 456

Attribute declarations 457

Exception declarations 457

IDL syntax 459

- Comments 460
- Include directives 461
- Pragma directives 461

Multiple IDL interfaces and modules 464

The idlc command 464

- Options for the idlc command 465
- Emitted file names 469
- IDLC_OPTIONS environment variables 470

The IDL-to-Java compiler 470

- Quick reference 470
- Emitting client and server bindings 471
- Specifying an alternate location for emitted files 472
- Specifying alternate locations for include files 472
- Inserting package prefixes 474
- Emitting makefiles and specifying the path separator character 475
- Defining symbols before compilation 476

- Preserving pre-existing bindings 476
- Viewing progress of compilation 476
- Displaying version information 476

The idl2com Command 477

- Quick reference 477
- Options for the idl2com command 478
- Generating interfaces using the idl2com command 478
- Emitted file names 479
- Data type restrictions 479
- idl2com generated makefile 480

Appendix C. CORBA programming 481

C++ bindings 481

- C++ bindings for constants 481
- CORBA types and business objects 482
- C++ bindings for data types 483
- Exceptions 503

Commonly used CORBA interfaces 506

- CORBA class interfaces 506
- CORBA::Object interfaces 507
- CORBA::ORB interfaces 507

Name scoping and modules in the C++ bindings 508

C++ bindings for interfaces 508

- Managing object references 510
- Widening object references 510
- Narrowing object references 511
- Narrowing to a C++ implementation 511

Storage management and _var types 512

- Argument passing considerations for C++ bindings 514
- C++ type mapping for argument passing 514
- Storage management responsibilities for arguments 516

C++ client bindings 520

C++ server bindings 521

C++ binding restrictions 523

Appendix D. Java ORB properties 525

- Property setting mechanisms 525
- Details of Java ORB properties 525

Notices 531

- Trademarks and service marks 533

Index 537

About this book

The *WebSphere Application Server Enterprise Edition Component Broker Programming Guide* describes the Component Broker programming model and presents sample code showing common tasks required to develop typical object-oriented applications.

This preface includes the following sections:

- “Who should read this book”
- “Conventions used in this book”
- “Notation” on page x

Who should read this book

The *WebSphere Application Server Enterprise Edition Component Broker Programming Guide* is intended for application developers who use the Component Broker environment to build robust, distributed object-oriented applications.

The examples are written in C++ and Java; therefore, programming experience in C++ or Java and a background in object-oriented programming is required.

This book is not a programming manual; it is for experienced programmers who are going to use this product.

Conventions used in this book

The following conventions distinguish different text elements:

plain Window titles, folder names, icon names, and method names.

monospace

Programming examples, user input at the command line prompt or into an entry field, user output, and directory paths.

bold Menu choices, push buttons, check boxes, radio buttons, group-box controls, drop-down list boxes, combo-boxes, notebook tabs, and entry fields.

italics Programming keywords, variables, and attributes, titles of information units, initial use of unique terms, and emphasis.

The following icons are used to indicate platform-specific sections.



Denotes a section that applies only to the Windows 95 or Windows NT platform. Do not interpret this symbol to denote that an equivalent section exists for any other platform.

Note: The Windows 95 platform only supports the Component Broker Java client.



Denotes a section that applies only to the AIX platform. Do not interpret this symbol to denote that an equivalent section exists for any other platform.



Denotes a section does NOT apply to OS/390 Component Broker. Do not interpret this symbol to denote that an equivalent exists in OS/390 Component Broker.

Notation

Even though there is a large degree of commonality between the content of various object methods, there is less agreement when it comes to the notations used to describe the content (that is, its form).

Throughout this document, the class diagrams use the Object Modeling Technique (OMT) notation developed by James Rumbaugh, with the exception that it is distinguished between interface and implementation inheritance.

Notations used in this document are:

- Interface inheritance
- Implementation inheritance
- Association
- Aggregation
- Qualified association

There is explanatory text in the examples to further clarify the graphical Rumbaugh notation.

If you are already familiar with these notations, skip ahead to “State transition diagrams” on page xiv or “Data flow diagrams” on page xv.

Interface inheritance

Figure 1 shows that a child class inherits just the interface, but not the implementation, of a parent class.




Figure 1. Class inheriting interface but not implementation

Implementation inheritance

The following notation shows that a child class inherits the implementation, not just the interface, of a parent class:




Figure 2. Class inheriting implementation and interface

This notation is usually reserved for OOD models, as OOA models normally focus on semantics, and therefore interfaces, rather than implementations.

Associations

The notation for associations, sometimes referred to as *by reference* or *uses-a* relationships, is the following:

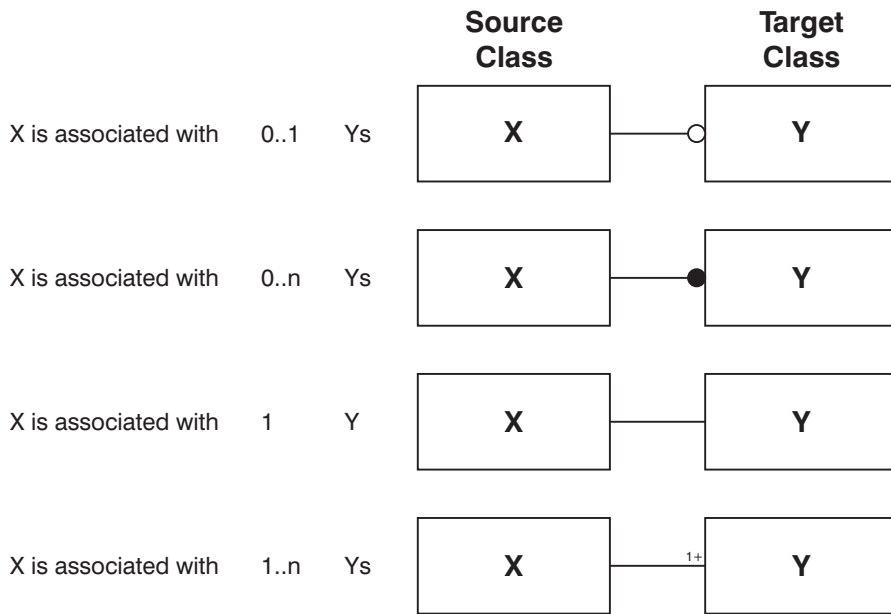


Figure 3. Associations notation

Aggregation

The notation for aggregations, which are sometimes referred to as containment or *has-a* relationships, is the following:

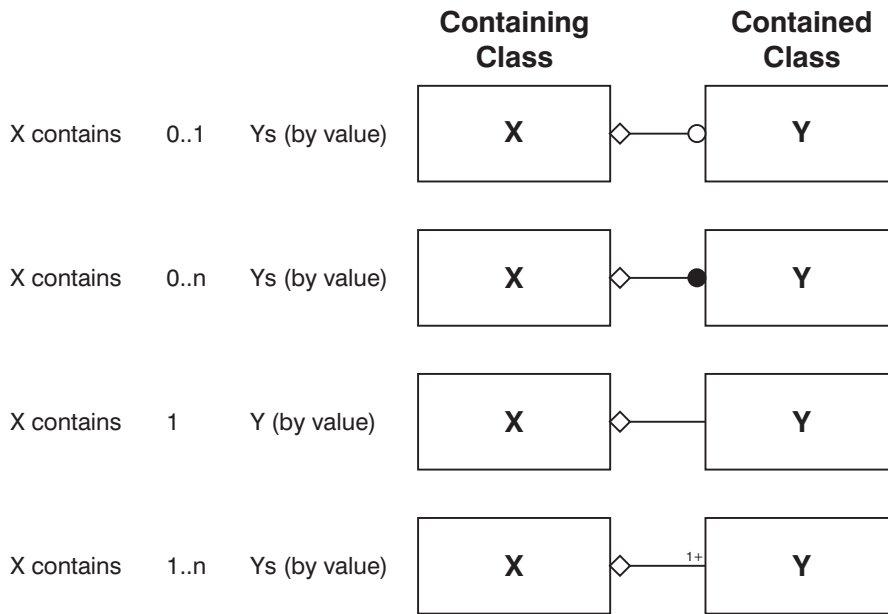


Figure 4. Aggregation notation

Qualified association

A qualified association is a way to constrain the targets in a relationship without requiring that a special subclass be created. The technique is to add an attribute to the relationship that can be considered to be a unique identifier with respect to the source class.

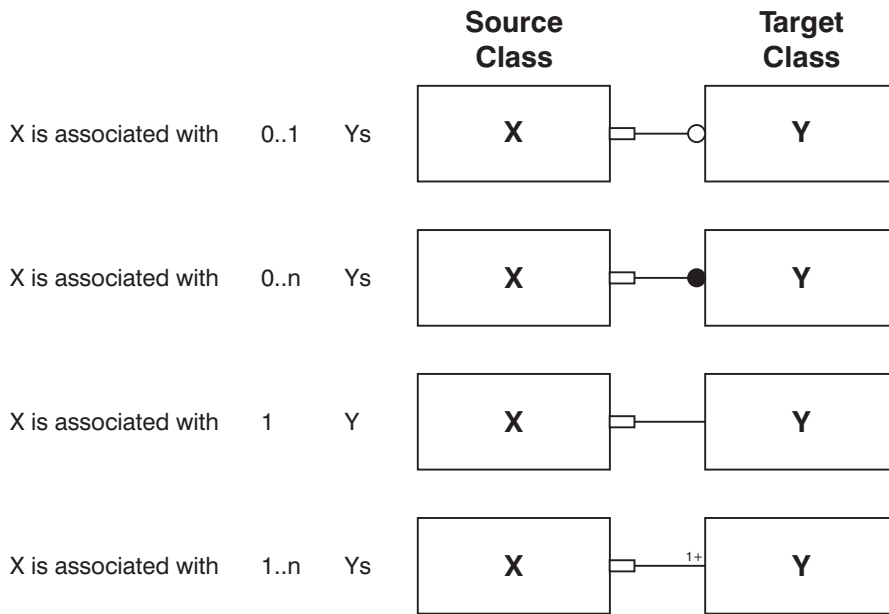


Figure 5. Qualified association notation

Sometimes the attribute used to constrain the relationship appears in the box on the relationship.

State transition diagrams

When an object exhibits a complex life cycle model, a simplified form of state-transition diagram (STD) is used. This notation makes no distinction between initial and final states (because that can be ascertained from the transitions); however, it separates the conditions from the actions, where the action is a method on the object in question either from the static or recursively, the dynamic model as follows:




Figure 6. State transition diagrams

When invoking the method on the object that can be assumed to be the *trigger*, no condition notation is required on the diagram although some programmers include one to show the objects that can invoke the method, and under what conditions, if any.

A loop back transition is no guarantee that a method is read-only as shown in the diagram. Of course, it does not change the overall state of the object as shown in the diagram. A recursive dynamic model might show a change of substate.

The method categorizations shown here are included as an introduction only. Their use within the Component Broker programming model to enable optimistic service implementations are described in “Chapter 13. Assembling and installing components on Component Broker/Workstation” on page 351.

Data flow diagrams

When the action in a transition, usually a method, is complex enough to warrant more than a simple text explanation or some basic pseudo-code, relatively standard Yourdon DeMarco data flow diagrams are used. The only difference is that objects appear as data source or *sink* objects, for example the box in the following diagram:

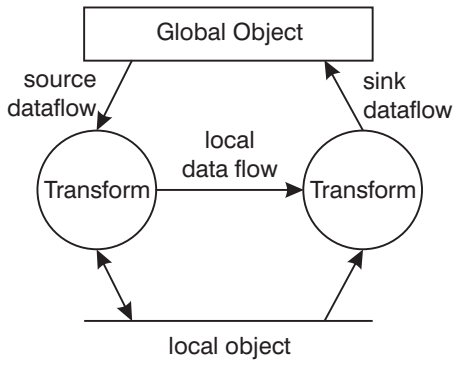


Figure 7. Data flow diagrams

As the diagram shows, a data flow can be labeled or unlabeled, depending on whether the entire object is used by the transform or not. Also, a data flow can be a read-and-write type data flow as required to reduce the amount of clutter.

How to send your comments

Your feedback is important in helping to provide the most accurate and highest quality information. If you have any comments about this book or any other WebSphere Application Server Enterprise Edition documentation, send your comments by e-mail to waseedoc@us.ibm.com. Be sure to include the name of the book, the document number, the version of WebSphere

Application Server Enterprise Edition, and, if applicable, the specific location of the information on which you are commenting (for example, a page number or a table number).

Chapter 1. Introduction

Component Broker is a client/server solution for object-oriented applications, in two respects:

- A deployed solution using Component Broker has a set of processes that serve objects to a set of client applications. This logical, client/server process model can be implemented in a one, two, or three tier client/server system, and a process may be both a client and a server.
- Object-oriented programming is client/server by nature. A server object provides an implementation of an interface. Other client objects implement new functions using the interfaces and data provided by server objects. An object may be a client of some objects and a server for others.

Due to the client/server nature of object-oriented applications, the Component Broker programming model is divided into two sub-models, the client programming model and the server programming model.

Component Broker provides frameworks and classes that implement common functions needed to build distributed object-oriented commercial applications. Some examples are transactions and the Naming Service. Much of the value in Component Broker derives from these supplied services, which enable developers to focus on their business logic and relieves them from the burden of implementing basic, cross-application functions.

In addition to its prepackaged services, Component Broker is extensible. End users, system integrators (SIs) and independent software vendors (ISVs) can extend Component Broker by providing additional sets of reusable services, or providing alternate pluggable replacements for the Component Broker services. Component Broker can be extended in many ways, but the most important extensions are:

- Add new *application adaptors*. An application adaptor provides interoperability between Component Broker objects and external environments such as DB2, CICS, and SAP. The application adaptor provides frameworks and services that:
 - Render or wrapper the external, non-object-oriented data and applications as Component Broker server objects.
 - Provide tailored implementations of managed object framework and Component Broker services that interoperate with the services provided by the external system. For example, a Component Broker application adaptor for DB2 provides an implementation of Object Transaction Services that interoperates with the databases transaction services.

- Component Broker comes with a set of application adaptors. They do not support all external environments. Therefore, customers, ISVs, and SIs can extend the capabilities of a Component Broker server by implementing new application adaptors.
- Add new Object Services and frameworks that can be used by business objects. Some examples could be frameworks or services for advanced transaction models or a business rules engine.

Component Broker is intended as an infrastructure or technical architecture which naturally supports object-oriented applications produced through object-oriented analysis and design. This document explains how Component Broker concepts are related to or derived from object-oriented analysis (OOA) and object-oriented design (OOD) concepts.

Object-oriented application development paradigms

Object-oriented application development can be approached in the following ways:

Top-down

Modeling and analysis define the classes and behaviors that must be implemented.

Bottom-up

Existing applications and databases provide functions and data that must be wrapped through encapsulating classes and instances.

Meet in the middle

A combination of top-down for new functions and bottom-up for reuse and incremental re-engineering of old applications.

The Component Broker programming model supports these scenarios.

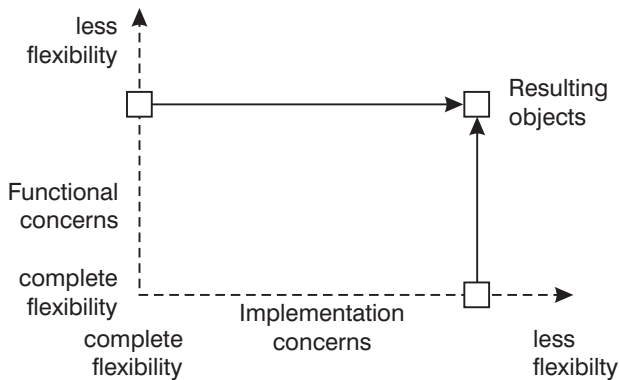


Figure 8. Degrees of freedom during application development

The models can be characterized by their balancing of the concerns of the functional domain versus those of the underlying implementation domain or design constraints. This balancing defines the degrees of freedom possible during object-oriented analysis and design. Usually the degree of freedom depends on the amount of legacy applications and databases involved along that dimension.

Typically, a given application of any complexity uses a mix of all three models.

Forward engineering new functional requirements (top-down)

Top-down development can be considered to be the ideal model for object-oriented applications because objects directly derived from functional requirements are the goal. At a high level, the main steps in the top-down process are:

1. Business process modeling (BPM) defines the application requirements and functions that must be performed. This includes requirements specification and use case analysis.
2. Object-oriented analysis and object-oriented design uses a methodology and design guidelines to produce an object model. The object model defines the relationships between classes, for example inheritance, aggregation, and usage and the methods and attributes of the classes. Detailed object-oriented analysis and design may also define state transitions and data flows. Finally, the model may contain semantic metadata information that describes more details of the model and the behaviors of classes. This metadata is used as the starting point for implementing the classes, and may be the basis for code generation.

- Object implementation involves defining the classes in an implementation language such as C++, Java, or Smalltalk. This step also involves implementing the objects as directly as possible in an underlying database and the use of other Object Services.

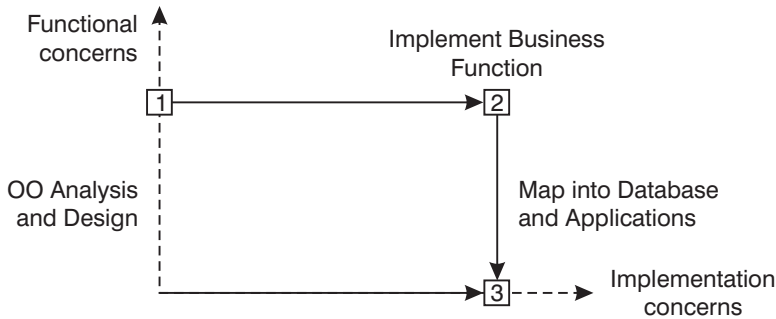


Figure 9. Top-down development

The Component Broker programming model and tools support top-down development by:

- Loading object models produced by popular front-end BPM and OOA/OOD tools into the Component Broker application development tools.
- Providing an Object Builder that eases the development of newly defined business objects that are derived from the managed object framework. The builder visually *tools* the Component Broker programming model and automates the implementation of some MOFW framework methods by generating code.
- Installing the new classes into an application adaptor to give the business objects access to Object Services, including persistence. This includes generation of a database scheme for persistent objects and emitting the implementation of methods and helper classes.

Reverse engineering from the legacy (bottom-up)

Bottom-up development involves the composition of new applications from assembly and reuse of existing functions and applications. One of the touted benefits of object-oriented programming is improved productivity through reuse of existing classes. In many cases, however, a non-object-oriented legacy application or database must be reused. When applied to object technology, this technique is often called *wrapping*. That is, the objects developed are usually just thin wrappers around legacy code or data that must be preserved in the new object-oriented application.

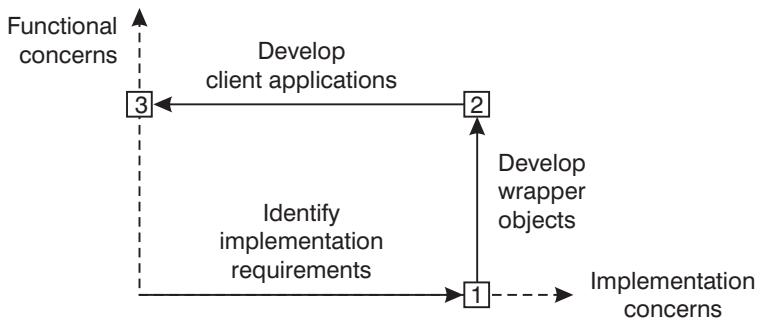


Figure 10. Steps involved in reverse engineering

When wrapping legacy data, the objects usually are little more than a set of public attributes, or get and set methods, corresponding to a logical grouping of data, such as a row in a database table or record in a file.

For legacy code, it is usually the case that multiple transaction programs share a given parameter that tacitly identifies an underlying object. For example, a bank might have the following transactions:

- Create and open a checking account.
- Close a checking account.
- Credit to checking account.
- Debit from checking account.
- Read and update taxpayer's Social Security number for the account.

In this simple example, the checking account number effectively identifies an object. The existing transactions are reused in the implementation of object methods. Transaction input and output parameters are mapped onto object state data, and the parameters of object methods.

The term operational reuse is used to describe the wrapping and reuse of existing databases and applications in the development of new object-oriented applications. For many readers of this document, the concept of reusing data is obvious. Why do customers want to reuse applications, instead of having the objects directly access the databases used by the old, legacy applications? There are two reasons for application reuse:

- The existing applications embody a substantial investment, and it is too expensive to rewrite the business logic. There are literally trillions of lines of existing application code which represents an investment of billions of dollars. The old applications need to be supported anyway, at least for a while, because not all client applications and terminals that use the legacy functions can be discovered and updated.
- In principle, new business logic could directly access the data. For example, the `CheckingAccount::debit()` method could directly update the database

record. However, most apparently simple applications, like the debit transaction, implement imbedded business rules and database integrity constraints over multiple databases, and have side effects. Objects directly accessing the underlying data must exactly implement these semantics. To do so effectively results in rewriting the existing application, which is too expensive and would require ensuring consistent behavior over two bodies of code.

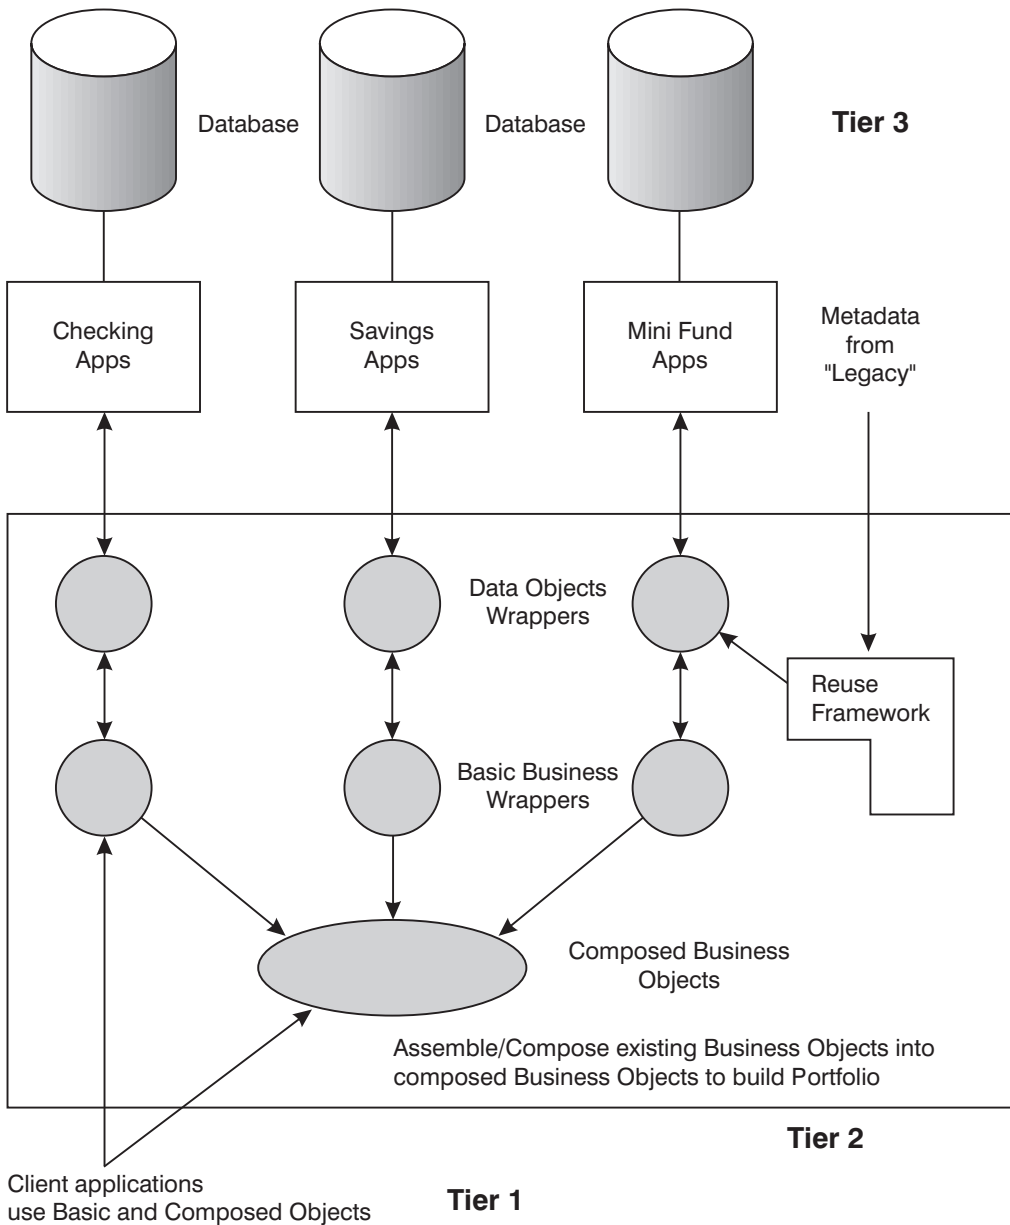


Figure 11. The operational reuse model

The bottom-up development process follows a well defined set of steps:

1. Wrapper existing applications and databases through encapsulating data objects. The Data Objects provide a direct object-oriented rendering of the existing applications and databases.

2. Implement business objects that provide a veneer over the data objects to convert them into views required by the object-oriented model.
3. Implement business objects that derive their state data and behavior from multiple or other business objects. The composed objects provide a more natural view of the existing, legacy defined model and provide a more natural object-oriented model for newly developed applications.
4. Develop new object-oriented applications.

The Component Broker programming model and supporting Application Development tools enable bottom-up object-oriented development by:

- Introducing the concept of application adaptors that provide interoperability between business objects and existing legacy environments. For example, Component Broker comes with a Relational Database application adaptor that enables object reuse of existing relational databases.
- Providing a set of application adaptor specific class libraries and frameworks that are supported by Object Builder extensions to facilitate the construction of the wrapper objects from existing metadata. For example, in the relational database application adaptor, this metadata is the database catalog that defines tables and rows.
- Providing a framework and visual tool in the Object Builder for construction by parts based definition and implementation of new business objects.
- Enablement of client applications that use server objects by generating client parts such as Java applets or ActiveX/COM components from the server class definitions.

Meet in the middle, incremental development

The first two paradigms represent the extremes of new object applications and direct rendering of existing applications. Most large development projects are a mix of both approaches. When top-down meets bottom-up, what happens? A developer would be lucky if the classes developed bottom-up exactly matched the interface and function requirements of the newly developed top-down model. In the entire history of object-oriented application development, no one has ever been that lucky. To meld top-down and bottom-up, it is necessary to create *mapping* objects that:

- Implement a facade that transforms the old object into one required by the new client objects.
- Are composed from existing bottom-up objects and reuse their implementation; can override existing functions and can add new methods.

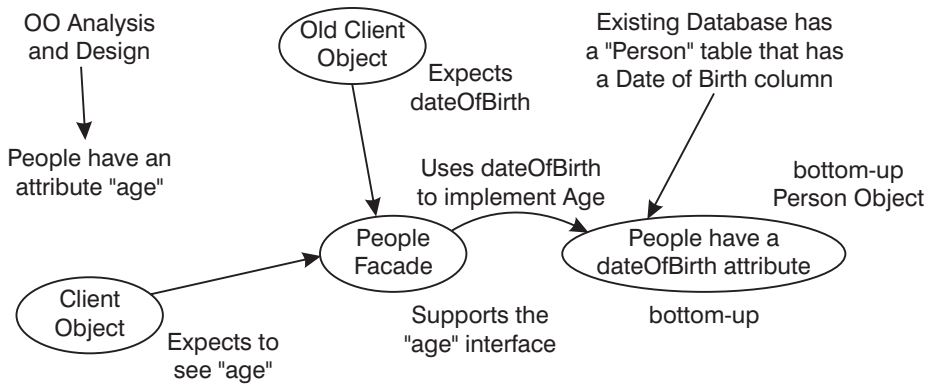


Figure 12. Meeting in the middle

This meeting in the middle problem is also present when new business processes and objects are added to an existing object model. The new client objects defined by OOA/OOD often require changes or extensions to the existing legacy server objects.

The developer could change the old objects, but the *existing old* client objects still require the interfaces and semantics. So, as subsequent development occurs either top-down or bottom-up, the functional view defined by the model begins to differ from the implementation view defined by what exists. An incremental approach is required, which involves defining a mapping object with the correct interface. The implementation of the mapping class uses composition to maximally reuse the functions of the old object.

Usually this approach makes heavy use of implementation inheritance and possibly delegation in an attempt to reuse the work done before. The idea is to inherit (and implement) the interface of the objects corresponding to the new functional requirements and extend the implementation of the existing legacy objects.

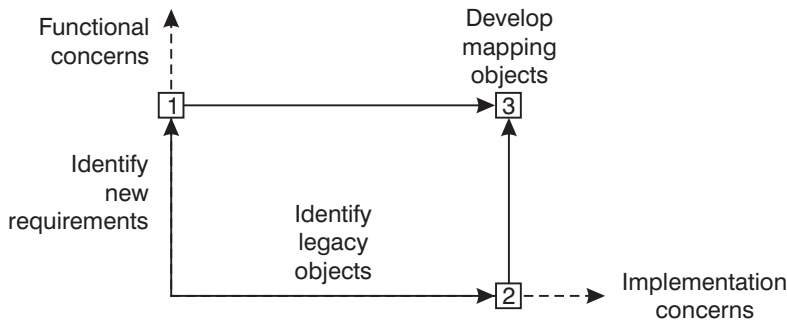


Figure 13. Meet in the middle

Component Broker supports incremental, meet in the middle development by:

- Supporting both top-down and bottom-up development.
- Composing new classes from existing classes.
- Implementing the new class as an aggregate composed of an instance of the new class and an imbedded instance of the existing class.
- Delegating behavior from the new classes onto the imbedded instance.

Combining the approaches

Component Broker is designed to fully exploit all three development paradigms, with the following considered to be a typical approach to migrating to and maintaining an object-oriented application:

1. Begin by bottom-up engineering the underlying legacy code and data into data objects and business objects that are contained in application adaptors.
2. Top-down engineer pure object-oriented applications to implement new business behavior.
3. Install the new objects in an application adaptor to support their business behavior with Object Services such as Concurrency Control, Transaction, Naming, Identity, and Persistence.
4. Use composition tools to build mapping and composed objects that join the bottom-up and top-down classes.
5. Publish the public server classes to client systems by generating client parts that can be loaded into client application development tools, for example, Java and Visual Basic. These client parts delegate their business and service implementations onto the server objects.

Component Broker defines frameworks for supporting the previous tasks and application development tools that facilitate using the frameworks.

Programming roles and responsibilities

The Component Broker programming model defines the following roles and set of tasks performed by each role:

- The Data Object Builder:
 - Implements the data objects that wrapper existing applications and data for bottom-up development.
 - Implements the data objects and scheme definitions that provide persistence and services for data objects that support state data introduced by newly-defined top-down business objects.
 - Implements the data objects in an application adaptor
- The Business Object Builder:
 - Implements business objects that map data objects onto the required object-oriented model.
 - Implements business objects that provide aggregate interfaces or views over other business objects. A simple example is implementing a Portfolio class that provides functions for sets of checking and savings objects.
- The Application Object Builder implements business processes or tasks on the server by scripting behavior over multiple business objects.
- A Service Implementor extends the basic Component Broker server by adding new frameworks and Object Services.
- An Application Adaptor Implementor extends Component Broker by enhancing the functions of an existing application adaptor, or extends a Component Broker server by adding a new application adaptor.
- The Client Part Provider generates and implements smart proxies for specific client environments.
- The Client Programmer builds client side applications by using Component Broker proxies and smart proxies.
- The Interface Builder develops an end-user interface using proxies and client applications.
- The Application Installer packages sets of client and server classes into application DLLs or sets of Java applets, and configures and installs the applications.

Programming languages and conventions

This section discusses the language options for Component Broker Programming, and conventions used in examples in this document.

Interface Definition Language

The primary method for defining Component Broker managed objects is the OMG Interface Definition Language (IDL). OMG IDL is a distributable, language-neutral form for defining interfaces, and can be mapped into almost any object-oriented language and many non-object-oriented languages.

The specific version used is CORBA 2.0. For more information about the syntax, see “Appendix B. Interface Definition Language” on page 447.

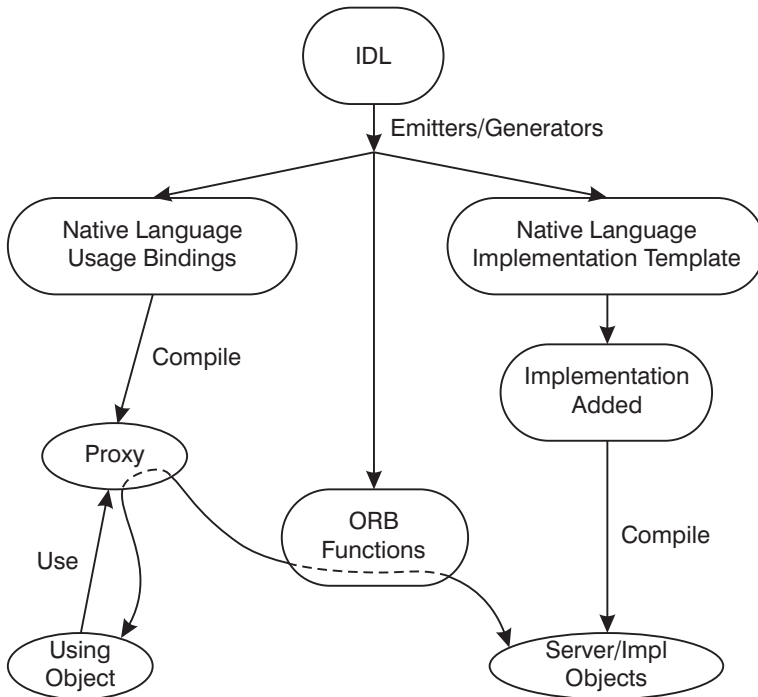


Figure 14. IDL, usage and implementation

Figure 14 is an overview of the relationship between IDL and application development languages. Object Providers use IDL to define the interfaces to their objects. The IDL may be directly defined by the Object Provider or may be produced under the covers in application development tools. Code emitters and generators produce the following:

- A *Usage Binding* that provides a native, client language rendering of the IDL. For example as a C++ class or Java Interface. The Usage Binding is also used to generate a proxy object that uses delegation to map the interface onto the server object providing the implementation.

- An *Implementation Template* that provides a native, server language class template into which method behavior can be inserted, for example, by editing the file and adding source code. Adding the behavior to the Implementation Template can be done in a tool.
- Implementation objects such as *skeletons* and *stubs* may also be emitted and compiled by the application development tools, if the client and server are in different processes, or in different languages. These implementation objects provide the functions necessary to make inter-language calls and remote method execution.

One often overlooked point about IDL is that all components in an IDL interface specification are considered to be just that - interfaces. Even attributes in IDL are a shorthand for get and set methods or just get methods for read-only attributes. There is no requirement that an underlying instance variable with the same name must exist.

This separation of interface from implementation is one key to understanding how a language which supports multiple inheritance, a construct not supported in all languages, can be considered to be language neutral.

Java

Java is considered by some to have combined the best features of Smalltalk-like languages (such as a well defined virtual machine interpreter with garbage collection) with those of C++ (the ability to drop out of pure objects for primitives and compiled native methods to get performance).

Component Broker provides support for using Java to develop client applications and server objects, and provides additional value over basic Java by supporting Java as a client application language that has access to Component Broker object functionality from a remote process, and supporting C++ to Java cross-language calls.

C++

Component Broker provides support for developing client and server applications in C++. Component Broker is also extensible, and supports the addition of application adaptors and Object Services, as well as tailoring of existing services and application adaptors. C++ is the language initial versions of Component Broker supported for extending Component Broker in these areas.

Other languages (Windows only)

Component Broker client applications can be developed in any language for which there are CORBA IDL Usage Bindings. The Component Broker application development tools also emit ActiveX/COM interfaces that wrap

the CORBA usage bindings, which enables Component Broker client development in other languages such as Visual Basic.

Naming conventions

Throughout this document and regardless of the language used, the class diagrams and code examples use the VisualAge C++ naming conventions with an additional convention that a class name is assumed to refer to an interface only (not an implementation class), unless it has a suffix of Impl. Therefore, Object as shown in the following figure refers to an interface, but ObjectImpl refers to an implementation (note the use of the interface inheritance symbol):




Figure 15. Naming conventions

The “I” prefacing the interfaces and implementation of Component Broker Objects stands for IBM. All of the interfaces provided as part of Component Broker begin with “I”.

Throughout this book and the rest of Component Broker, there are several conflicting method naming conventions. CORBA uses a convention that separates words with an underscore and uses lowercase while the convention used for Component Broker methods is mixed-case method and class names with no underscores between words. The only exception to this is when Component Broker extends CORBA interfaces with related methods. These are introduced using CORBA conventions.

Coding conventions

Coding conventions are beyond the scope of this book. However, material related to this topic is presented in “Appendix B. Interface Definition Language” on page 447 and in “Appendix C. CORBA programming” on page 481.

Application architecture

The three levels of architecture discussed here (the three-tier application architecture, the middle-tier layered architecture, and the component object architecture) result in an application that is reliable, extensible, and scalable. The development of such an application is made possible by the managed object framework, which provides the necessary code. The framework is supported by a suite of development tools, which allow you to make use of the framework to create components without going into details of inheritance and framework implementation.

Multi-tier architecture

In combination with existing resource managers, Component Broker functions as an object server by providing an application environment that lets clients access back-end-systems through object-oriented middleware. Middleware connects and adds value to the software capabilities in a system's client and resource tiers, whether the business logic is in a legacy system or a middle-tier application in a multi-tier system. This system provides a scalable infrastructure built on the platform-independent characteristics of the Internet, paired with a separation of business concerns from technological aspects, combined with a multi-tier software structuring and deployment topology that results in an infrastructure scalable enough to include everything from desktops to the largest cluster of mainframes. Platform-independent, you can design, develop, and deploy distributed object-oriented server applications for mission-critical solutions.

A Component Broker server process implements managed objects as an object that is managed by an application adaptor that is used by client applications to perform business functions. All Component Broker server objects are derived from managed object framework (MOFW). A managed object is an object that is managed by an application adaptor. Managed objects are used because one of the Component Broker server product's main contributions to object-oriented applications is the implementation of run-time management functions used by server objects. Some examples of these managed services are:

- Persistence, Transaction, and Security.
- Workload management and availability management over multiple servers.
- Object-oriented access to existing databases and applications.

The client application (tier-1), accesses business objects (tier-2), on the server. Business objects are instances of classes that implement the business logic application. Business objects support the interfaces of the MOFW which allows business objects to be installed onto, and managed by, a Component Broker server.

Component Broker enables operational reuse. Operational reuse focuses on re-engineering existing software so it can be used as building blocks for new applications. The Component Broker infrastructure enables the construction of these new building blocks. These building blocks achieve operational reuse and present an abstraction layer that can in turn be reused to build many new business applications.

This logical three-tier architecture forms the foundation for partitioning Component Broker applications. The first-tier is the client, the middle-tier is the application server, and the third-tier is the data store, such as databases, or other applications. In terms of these logical tiers, Component Broker provides components that are deployed in the middle tier, connecting the first-tier (client) with the third-tier (various kinds existing resources). The components are implemented as CORBA-based business objects or enterprise beans. They allow application logic to run on high-powered servers, and insulate client applications from the complexities of the various resource managers, such as CICS, IMS, and DB2. The client works with the components through a CORBA-compliant ORB, and the components work with the resource managers using whatever communication protocols are supported by the target resource manager (for example, TCP/IP).

The client, or first-tier, can consist of C++ programs, Java applets, or Visual Basic programs that a user interacts with directly, or they can consist of web servers or application servers that have their own clients. A typical client application consists of view objects, which provide the end-user interface interactions and a mapping to server components, and client objects, which implement business logic specific to the client, and provide integration with other desktop applications, such as spreadsheets and document processors.

From a Component Broker perspective, anything that can send an IIOP request is a potential client. The logical client tier can actually consist of any number of physical tiers.

The server, or middle tier, can also consist of many physical tiers, with components on different servers providing a unified front to the client. On the other extreme, the middle tier components could, in a simple case, be running on the same physical host that provides the data store, collapsing the middle and third tiers into a single physical tier.

The third-tier, representing the data store, could be running on many different physical hosts, and could be using application logic that itself provides additional physical tiers. On the other extreme, all the data could be stored on the same physical server that runs the components.

In summary, the Component Broker three-tier architecture can be implemented on any number of physical tiers. The logical partitioning allows

easy scaling of an application from a simple, one- or two-tier environment to an arbitrarily complex, multi-tier environment.

The three-tier architecture also makes an application more maintainable, by isolating the client (first-tier) from having to know about the data store or resource manager (third-tier). If you change the database you are using, only the server needs to be modified: the client is not affected. Because there are usually fewer copies of the server than the client, and because the servers are often in locations that are easier to update (for example, on central machines rather than on PCs running on users' desks), this simplifies the update procedure. This approach also provides additional security, because only servers need access to the data controlled by the resource manager.

Component Broker's main role, in this three-tier architecture, is in the definition of the middle-tier. Component Broker's support for the three-tiers is as follows:

First-tier

A programming model, which allows client applications to be implemented in C++, Java, or Visual Basic, and allows clients to access components on the server through a CORBA-compliant ORB.

Middle-tier

A programming model with full tools support, and a run-time environment, in which the components (CORBA-based business objects or enterprise beans) are deployed.

Third-tier

Application adaptors on the middle-tier allow components to access data in various resource managers, including DB2, Oracle, CICS, and IMS.

Component Broker server objects are used on the client through proxy objects. A proxy supports the interface of the server object, but implements the interface by using object remote call the server object. Component Broker uses OMG's CORBA specification as its distributed object infrastructure.

The Component Broker client programming model also supports access to components in a fashion that matches the client language and programming model. For example, Component Broker provides support for component proxies with ActiveX/COM objects to facilitate their use in Visual Basic applications.

Middle-tier architecture

The middle-tier of the three-tier architecture consists of components, implemented as CORBA-based business objects and enterprise beans. The

architecture for the middle-tier of an application has three layers, with three corresponding kinds of component: basic components, composite components, and application components.

Basic components are abstractions that have a reasonably direct mapping to existing data or procedural programs. This is the bottom layer of the middle-tier architecture. While these components may have dependencies on other basic components (one-to-one and one-to-many relationships), they are generally fine-grained enough to be highly reusable.

Composite components represent new abstractions that are easily usable and understandable by client programmers and by application component programmers. These compositions often represent an aggregation of basic components. Methods of the composite component typically delegate or bridge down to methods on the basic components which it aggregates. These mappings can be one-to-one (mapping directly to a method of an aggregated component), or one-to-many (executing methods on each of the aggregated components). While composite components are not as reusable as basic components, they are perhaps more valuable where they are reused, because they represent larger portions of the application.

Application components focus on business logic and usage of other components, in the way that some application programs do today. This is the top layer of the middle-tier architecture, this is the layer that will be accessed by client applications. Application components implement processes or tasks, as defined by object-oriented analysis and design. They implement any business logic that is not properly modeled as a method on other, individual components. Application components also provide the mechanism for moving application logic from the client to the server. Generally, application components represent the parts of the application architecture that are specific to an application, such as certain business processes or tasks. By separating out these elements, which are less suitable for reuse, you leave the rest of the application (composite components and basic components) as reusable as possible.

When you extend an existing Component Broker application, or create a new application that leverages the same data, you create new application components to provide application-specific business logic. However, you can reuse the basic components, and potentially the composite components. The basic components only need to change if the underlying data store changes. The composite components only need to change if the definition of the aggregation change. Even when change is necessary, the underlying component architecture allows the data store to change without affecting reuse of a component's behavior or interface. This is described in the following section on component architecture.

The architecture of a Component Broker application encapsulates access to the backend or data store, which sets the stage for reuse, and offers an adaptable structure that can evolve to meet new requirements.

Component overview

This section discusses components and objects used in Component Broker Programming.

Components

In Component Broker, a component consists of a distributed set of objects that client applications access as a single entity. To a client application, a component appears to be a single class, with methods and attributes and relationships like any other class. Behind this single interface, however, each component consists of multiple objects on both the client and the server. This separation provides flexibility and control in the way data is stored and accessed, and in the way that business processes are distributed. The objects can exist on any number of different servers and databases, but to the client they present a single interface, with a single set of attributes.

Typically, a component consists of the following objects in Object Builder:

- business object interface
- business object implementation
- data object interface
- data object implementation
- persistent object and schema
- key
- copy helper
- managed object

Component assembly

When you assemble a component, you can start from the business object, data object, or schema. Figure 16 on page 20 shows the relationships among component objects as you assemble them. When you configure the objects into a unified component, you select a subset of these objects to form a particular component on the server. The deployed component is accessed through its managed object, by client applications or other components that require access to the server data.

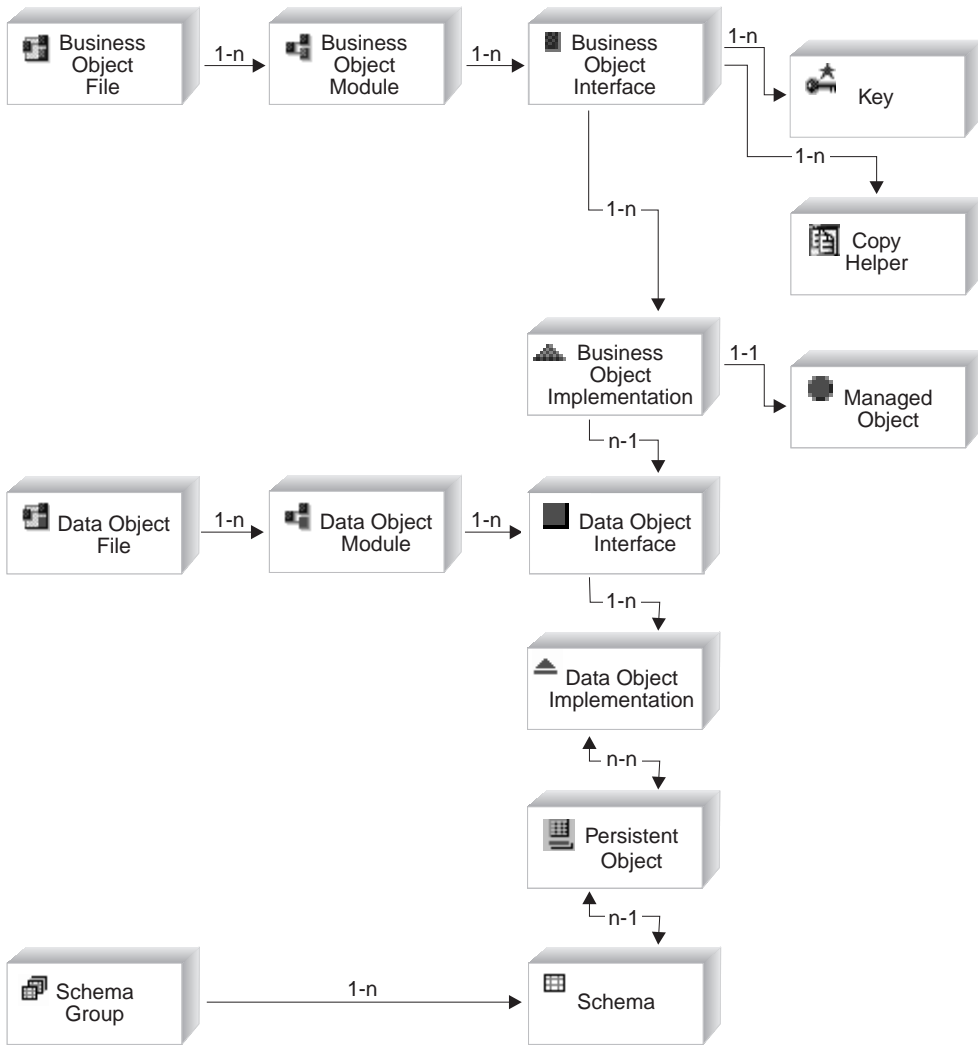


Figure 16. Component assembly

Component execution

At execution time, a call to the component on the server (for example, a call to a CarPolicy component to get the value of the attribute make) resolves in the following order:

1. The client calls the home to find or create the component assembly.
2. The managed object accepts the call (from the client or from another component).

3. The managed object calls its associated container for Object Services before passing the call on to the business object.
4. The business object accepts the call, and either returns the value of the attribute based on a cached copy of the data, or delegates the call to the data object.
5. The data object accepts the call, and either returns the value of the attribute based on a cached copy of the data, or delegates the call to the persistent object.
6. The persistent object accepts the call, and returns the value of the attribute based on a cached copy of the data if it exists.
7. The persistent object retrieves the value from a database using the schema, and returns it.
8. The value is returned up the component tree until it reaches the managed object, which calls the container again for Object Services before returning the value to the caller.

Execution

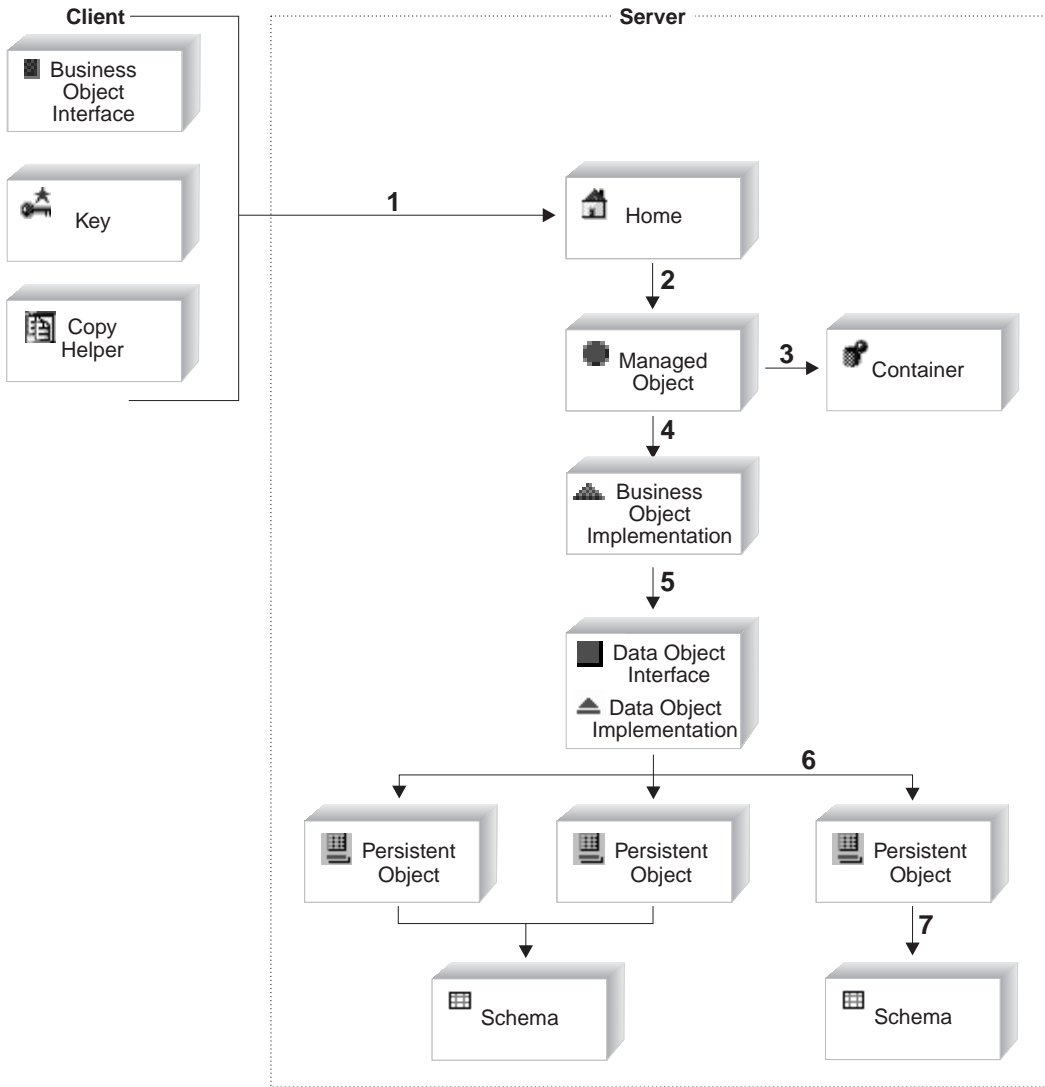


Figure 17. Component execution

State data

Every object has a state and a behavior, and presents an interface. An object's behavior is manifested in the implementation of the methods on the object's

private and public interfaces. An object's state is manifested in its public and private data members and can be divided into two categories: essential and nonessential. The essential state makes up the state data of an object and consists of data that is persistent and not calculated or derived from other data members. The nonessential state consists of transient data that can be recreated as required. The nonessential state is usually derived from other state data and complements the essential state.

Objects

This section discusses objects used in Component Broker Programming.

Application object

Application objects are business objects which are directing workflow and implementing some client initiated task. Methods often have an action verb associated with them which resembles the name of the use-case or end user transaction being performed. Application objects that are very process oriented are often stateless. They perform a particular task, calling other basic or composed business objects along the way, and then returning an indication of success or failure to the client application.

In Component Broker, an application object functions as part of an application component, which implements business logic and usage of other components, in the way that some application programs do today.

Business objects

A business object represents a business function. Business objects contain attributes that define the state of the object, and methods that define the behavior of the object. A business object also has relationships with other business objects. It can cooperate with other business objects to perform a specific task. Business objects are independent of any individual application. They can be used in any combination to perform a desired task. Typical examples of business objects are: Customer, Invoice, or Account.

In Component Broker, a business object functions as part of a component, which is a collection of related objects that work together to represent the logic and data relationships of the business function.

A business object's interface is defined in an IDL file. Its implementation can be in either C++ or Java.

Composed business object

A composed business object represents a facade or an abstraction that encapsulates basic business objects and other composed business objects.

Encapsulation means that the client of the composed business object is not generally aware of the various pieces that make up the composition. The data members of the composed business object implementation are usually references to other business objects. If these references can be calculated based on some key data, then the composition might not have any specific backing store information. If the composition is assembled at the object level and there is no way to calculate composition's components from the composition's key, then the object references which make up the composition must be stored persistently. Database tables for composite business objects are generally new tables and persistently capture these new relationships.

In Component Broker, a composed business object functions as part of a composite component, which is a collection of related business objects that work together to represent the logic and data relationships of the business function.

Compositions often have patterned method implementations. Each method exposed to the client might involve one or more calls to private methods, touching each of the basic business objects that are contributing to the composition.

Copy helper

A copy helper is an optional object that provides an efficient way for the client application to create new instances of the component on the server. The copy helper contains the same attributes as the business object, or a subset of them. Without a copy helper, the client might need to make many calls to the server for each new instance: one call to create the instance, and then an additional call to initialize each of the instance's attributes. With a copy helper, the client can create a local instance of the copy helper, set values for its attributes, and then create the server component and initialize its attributes in one call, by passing it the copy helper.

When you define a copy helper in Object Builder, its implementations are generated in both Java and C++.

Copy helper instances are created using the `_create` function. The client developer sets values locally, then creates a managed object from the copy helper using the `createFromCopy` function.

Data object

A data object is responsible for managing the persistence of a component's essential state information (state data). It provides an interface for getting and setting the state data. A data object isolates its business object from having to:

1. Know which of many datastores to use to make its state persistent.

2. Know how to access the datastore.
3. Manage access to the datastore.

A data object has two parts: the data object interface, which defines the state data of the component, and the data object implementation, which defines the form of persistence and access patterns for the data.

Handles

Component Broker supports a notion called object handles. A given object in the distributed system can be accessed by using an object handle using one of a variety of techniques. There are three main access patterns and there is handle class for each of them along with a base handle class present a consistent interface for accessing the object. It provides streaming support so handle strings can be stored persistently. Given a handle string user can locate/activate the object no matter which type of handle was used for the object.

Each of the access patterns have their merits and limitations. The Component Broker user can decide which pattern should be used for their circumstances.

SORHandle

This pattern encapsulates a stringified object reference (SOR). This pattern requires the most storage space, however it does provide the most efficient access. This pattern should be used only for objects which will never be relocated to other servers.

Note: In Component Broker for Windows NT and AIX, when SORHandle is used with objects that are not workload managed (WLM), the object is scoped to a specific server. When used with objects that are workload managed, the object can be moved to different servers in the same Server Group.

HomeKeyHandle

This pattern uses the home name and primary key of object. This pattern requires lesser amounts of storage but is slightly less efficient access due to additional path length required to find the home. However, this pattern provides a more flexible solution since homes, and therefore the objects that they manage, can be moved to other servers and hosts. This pattern should be used for all objects that may be required to be moved to other servers or hosts in the future.

Name Handle

This pattern is similar to the HomeKeyHandle pattern except it is explicitly registered in the name space and can therefore be assigned a well-known name to the user. This pattern should be used for well known objects only.

Key

A component's key object defines which attributes are to be used to find a particular instance of the component on the server. The key consists of one or more of the business object attributes, which must contain enough information to uniquely identify an instance.

The key is defined as a separate class (rather than simply flagging one or more of the attributes directly on the business object) for two reasons. First, CORBA does not permit passing a mix of different data types in one call; by defining the key in a separate class, you can mix multiple attributes and data types in whatever combination you require. Second, you gain the flexibility of changing the key, or having more than one key for different situations, without affecting the rest of the component.

When you define a key in Object Builder, its implementations are generated in both Java and C++.

Key assistant

A key assistant is a new key helper class. It is a concrete subclass that is associated with a component and has knowledge of the primary key that is configured for the component.

The interface of the key assistant supports creating keys from various existing objects. Currently, Object Builder supports creation of primary keys from a copy helper or a data object.

Object Builder optimizes performance by creating multiple copies of proxies of an object that every client can access. The way it does this is by generating a set of new `IManagedServer::IKeyAssistant` objects which introduce new sets of `idl`, `ih` and `.cpp` files.

Managed object

A class of objects that defines the set of methods that must be implemented by the business object to work with the appropriate application adaptor. Managed objects are enabled to work with two level storage containers and delegate to the data object the attributes of the object.

A managed object represents the component to the client application, and handles all calls from the client to the component on the server.

An application is defined by adding and configuring managed objects. By creating a managed object for a business object, you specify that it will be

installed on the server. The managed object handles communication with other classes, and initialization, de-initialization, activation, and passivation of the business object.

Persistent object

A persistent object is a C++ object that provides a mechanism for storing a component's state in a datastore. Every persistent object has an identifier or a key that is used for locating its corresponding record within the datastore.

There are two kinds of persistent objects: database (DB) persistent objects and procedural adaptor (PA) persistent objects.

Database persistent objects: This type of persistent object represents a record of a table or a view in a relational database. The component's state data that is stored in the relational database by means of a persistent object lasts longer than the execution time of the application that calls the component.

In a relational database such as an SQL database, all records are persistent because they are stored on disk in the form of database tables. A datastore, whether it is an object-oriented database management system (OODBMS) or a relational database management system (RDBMS), stores an object's persistent data.

There are two kinds of DB persistent object implementations:

- Persistent objects that use embedded static SQL to access and update the data they represent in a database. The generated files are of type .hpp and .sqx. You can edit the embedded SQL clauses that Object Builder provides for the methods of these objects, but Object Builder will not validate your entries.
- Persistent objects that use the caching capabilities of the server for accessing and updating data from the datastores they represent. The generated files are of type .hpp and .cpp.

In Object Builder, you can create these type of objects at the time of schema creation (top-down) or after a schema has been imported into Object Builder (bottom-up).

Procedural adaptor persistent objects: This type of persistent object encapsulates not only the data associated with an application but also the application's transaction logic. It is usually an embodiment of IMS and CICS transactions and data. A PA persistent object too is responsible for storing the data and transactions of a reusable application (which is bundled together as a component), much longer than the execution time of the application that calls the component.

Object Builder creates a PA persistent object for every Procedural Adaptor bean that is imported. The Procedural Adaptor bean exists as the PA schema in Object Builder and you can create additional PA persistent objects for a PA schema.

The generated files for a PA persistent object are of type .hpp and .cpp.

Schema group

A schema group is an organization of the different database schemas that you either create for a data object or import from the DDL file.

When the generate action is used on a schema group, an SQL file is created. It contains the subset of the definitions of the schemas that Object Builder requires to do table to persistent object mapping. Note the following points about schema groups in Object Builder:

- All schemas in a schema group must belong to the same database type. That is, within a schema group, you cannot have some schemas that are of the DB2 database type and others that are of the Oracle database type. However, you can arrange schemas that exist in the same database, into different schema groups.
- Schema groups cannot be nested.

Note the following points about schema group names:

- Group names must contain only alphanumeric characters, the blank space, and the underscore.
- They are case sensitive.

Schema

A schema is a description of the structure of a table or a view in a relational database. It is a structural and behavioral abstraction of the real physical data, and focuses on information relevant to users of the applications that use the database.

In CBCConnector, DB2 and Oracle are the relational datastores supported, and the schemas are described using the SQL language.

In Object Builder, there are two kinds of schemas: database (DB) schemas and procedural adaptor (PA) schemas.

DB schemas can either be created from data object implementations (the top-down approach) or imported from SQL DDL files (the bottom-up approach). For organizational purposes, DB schemas are found in schema

groups. All schemas within a group must be of a single database type. That is, they must all either be of the DB2 database type, or of the Oracle database type.

Follow these rules when you name a DB schema:

- The name must not exceed 18 characters for DB2; 30 characters for Oracle.
- All alphanumeric and DBCS characters are allowed, and there's no case sensitivity for names containing these characters.
- Non-alphanumeric names must be enclosed in double-quotes, and their case is maintained internally.

PA schemas are created when a procedural adaptor bean is imported into Object Builder from Enterprise Access Builder (EAB).

Another road map

The rest of this book provides you with the information necessary to write applications that use business objects, and to create, test, and install new business objects.

If you need to write C++ CORBA client applications that use a minimum set of the Component Broker-supplied client interfaces to business objects, read "Chapter 4. MOFW C++ client programming model" on page 85. Then see "Chapter 6. MOFW - C++ client programming model – advanced concepts" on page 133 for additional ways to work with managed business objects in Component Broker.

The basics of developing business objects can be found in "Chapter 5. MOFW - C++ server programming model" on page 103. This describes the basic abstractions that are necessary to build business logic and augment it with the necessary logic to be supported by the server. Advanced topics regarding business object development are discussed in "Chapter 7. MOFW - C++ server programming model – advanced concepts" on page 157.

If you need to write Java client applications that use a minimum set the Component Broker-supplied client interfaces to business objects, read "Chapter 9. MOFW - Java client programming model" on page 231. Then see "Chapter 11. MOFW - Java client programming model - advanced concepts" on page 281 for additional ways to work with managed business objects in Component Broker.

If the client application requirements specify the incorporation of ActiveX components into the client application, then refer to "Chapter 8. MOFW – ActiveX client programming model" on page 211.

If you want to develop business objects using Java, see “Chapter 10. Java server programming model” on page 251. Advanced topics about Java business object development are discussed in “Chapter 12. MOFW - Java server programming model - advanced concepts” on page 305.

The final step is to assemble and install the business objects and associated applications. A detailed description of this process can be found in “Chapter 13. Assembling and installing components on Component Broker/Workstation” on page 351.

Select and read the chapters pertaining to the roles and responsibilities that you are required to do. Application programmers will take one path, business object builders will take another. Each path includes options depending on how much complexity you desire (or require), and how many Component Broker systems capabilities are needed in your particular solution.

Chapter 2. Personal life insurance application example

One effective way to illustrate the programming model is through an actual example. For this guide, a personal life insurance application is used because it is a domain with which most readers are familiar and is complex enough to cover all of the analysis concepts and programming model concepts.

The scenario behind this example is as follows: an insurance company which has been selling business liability insurance decides to go into the personal life insurance market. The company needs new applications to support this new business opportunity, and has decided to develop new object-oriented applications. The new application must leverage a set of existing applications and databases.

The top-down model is presented first, starting with the object-oriented model for the new application and describing the new applications. Then, the constraints imposed by operational reuse of an existing set of applications (CICS Transactions) and databases (DB2 Tables) is described.

The remaining chapters in this document use this example to explain the programming model, as the basis for sample code, and to demonstrate the usage of tools.

The object model (top-down)

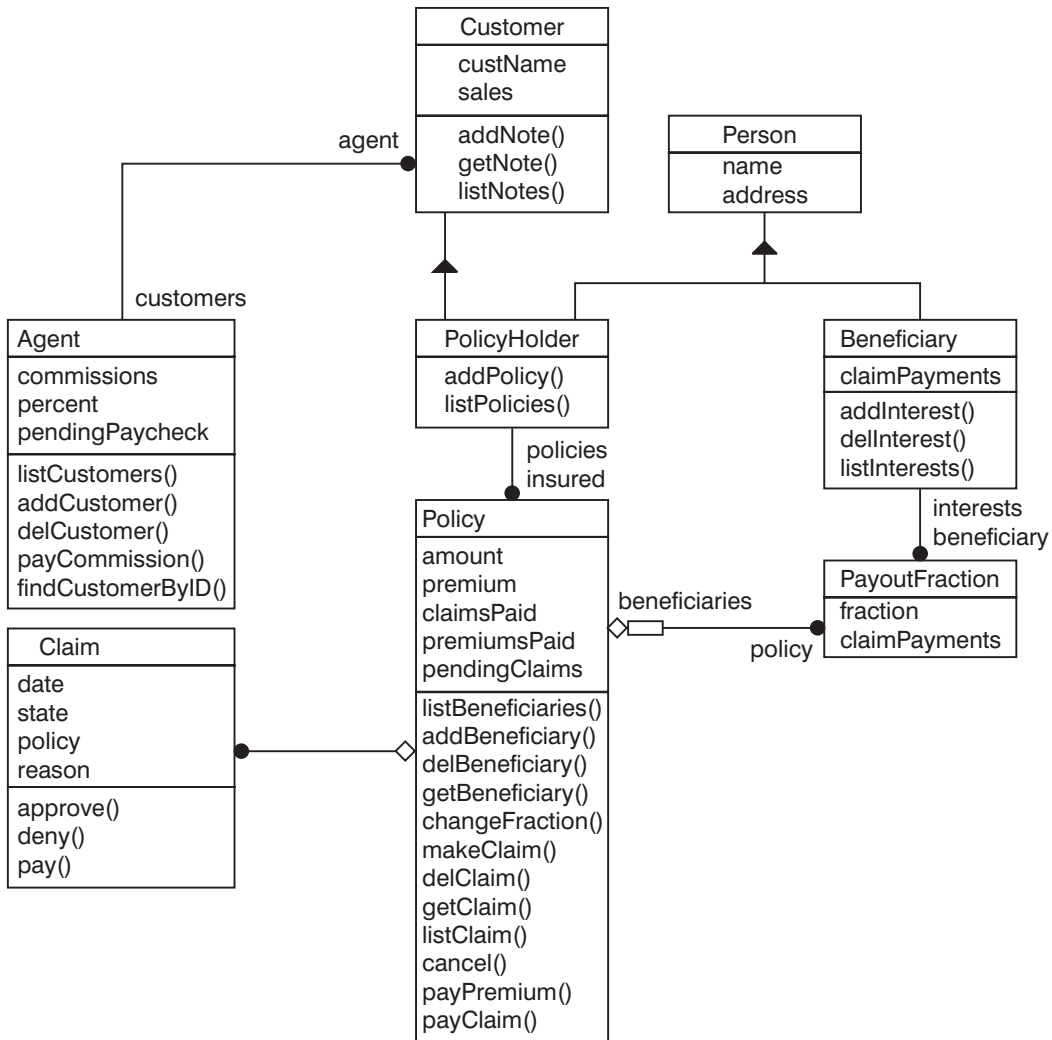


Figure 18. Top-down object model

Figure 18 presents the object model for the new insurance application. The model contains the following classes and relationships:

- The Person object contains information known about people, some of whom may be Customers or Beneficiaries. The insurance company has information about people who are neither Customers nor Beneficiaries. This information was obtained from purchased lists (for example, magazine subscription

lists) and other demographic services. The Life Insurance Application ensures that all Customers and Beneficiaries are in the Person database, and updates the database when necessary.

- The Customer object represents information and functions known for Customers of the insurance company. In the scenario, this information was obtained separately from the Person database, and contains separate information. The Customer information was built during years of business, while the Person database was recently acquired. A Customer is associated with an Agent.
- An Agent is associated with zero or more Customers. The Agent is a point of contact for the Customer and receives commission payments.
- A PolicyHolder is both a Customer and a Person, and can have from zero to many Policies. The new object-oriented application enforces the rule that both Customer and Person information must exist for a PolicyHolder. Therefore, when new Policies are created, or existing Policies are modified or have claims processed, any missing information is obtained.
- A Policy is owned by one PolicyHolder; it can have zero or more PayoutFractions, with one per Beneficiary.
- A PayoutFraction is associated with one Beneficiary and one Policy. The PayoutFraction defines the percentage of the Policy's value that is paid to the Beneficiary.
- A Beneficiary is a kind of Person that may hold interests in zero or more Policies (through a PayoutFraction).
- A Claim is associated with a Policy. It represents the request for payment for a Policy, and the state and resolution of this request.

In the terminology of "Multi-tier architecture" on page 15, the objects described in Figure 18 on page 32 are business objects.

The method and attribute descriptions for each of these objects begins in "Model details" on page 37.

The application model (meet in the middle)

In addition to the object-oriented model, the new line of business needs a set of Applications or Processes that implement tasks and functions for users of the Object Model. These Applications implement functions that transform the state of the Object Model by scripting task behavior over multiple business objects.

The Applications or Processes for the scenario are:

CreatePolicy

This application creates a new Policy, and ensures that:

- The Customer exists and is in the Person database.
- All Beneficiaries exist and are in the Person database.
- All PayoutFractions are in the range 0 to 100%, and that their sum is 100%.
- Commissions are paid to Agents.

ModifyPolicy

This application modifies a Policy by modifying:

- Attributes of the Policy, for example, *Value*.
- The PayoutFractions for a Policy, and ensuring that the fractions satisfy the constraints above.
- The Beneficiaries.

CreateCustomer

This application creates a new Customer object, and updates the Person database if necessary. This application associates an Agent with the new Customer.

ModifyCustomer

This application updates information about a Customer.

CreateBeneficiary

This application creates a new Beneficiary, and updates the Person database if necessary.

ModifyBeneficiary

Modifies information about a Beneficiary.

ProcessClaim

This application enters claims, moves a Claim through the states of Entered, Pending and Denied Approved and Paid. When a Claim is Approved, the Policy, Claim and PayoutFractions are marked as Paid.

In the terminology of “Multi-tier architecture” on page 15, the objects described previously are application objects.

The scripts for these applications are presented in “Applications” on page 46.

Clearly, a truly useful Life Insurance Application must have a richer set of functions, but this set is simple enough for illustrative purposes and covers the main Component Broker programming model concepts.

Design model (bottom-up)

Figure 19 on page 36 presents the “Design Model” or constraints on the new object-oriented application. These constraints and design points are:

- The Person class is defined by Operational Reuse of an existing relational database. The Person class is implemented in an object server that front-ends this class.
- There is an existing CICS mainframe based application Customer Management Application that is wrapped by the Customer and Agent classes. These classes are implemented in an object server.
- The Policy, PolicyHolder, Claim, PayoutFraction and Beneficiary classes are new objects whose state data is maintained in a newly defined relational database. In contrast to Person, the state data of the classes defines the database scheme instead of the database scheme defining the state of the classes. There are many possible choices for the persistence mechanism for new classes. Customers, ISVs and SIs are expected to expand the basic Component Broker solution by adding and extending application adaptors.
- The PolicyHolder class is also composed from an instance of Person and an instance of Customer (represented by the dotted line in Figure 19 on page 36).
- The Beneficiary class is composed from an instance of Person.
- The applications CreatePolicy, ModifyPolicy, CreateCustomer, ModifyCustomer, CreateBeneficiary, ModifyBeneficiary and ProcessClaim are implemented on a server.
- Agents and company employees use a GUI (Graphical User Interface) or a Web Browser to view business objects and implement applications. There may also be a set of thick clients that use the Applications and business objects.

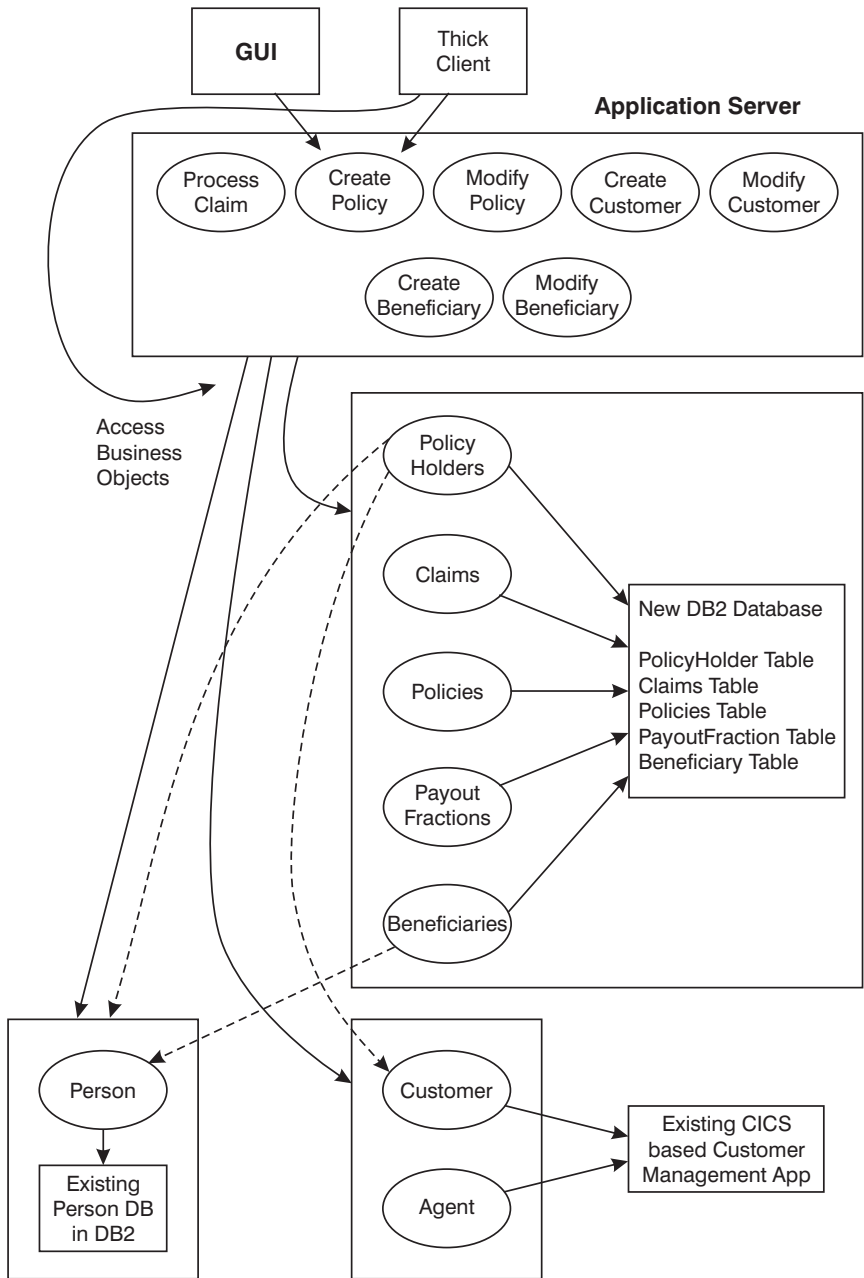


Figure 19. Design model

The existing applications place some constraints on the object model. For example, both PolicyHolder and Beneficiary subclass from Person. Because there are people who are neither PolicyHolders or Beneficiaries, and because

the classes are maintained on different servers, the PolicyHolder and Beneficiary classes support the Person interface, but implement this interface through composition and delegation onto instances of Person.

In “Design model (bottom-up)” on page 34, there are details on the CICS transactions to be reused, and the scheme for the Person database. The scheme for the new databases is defined by the new object model, and the database definition is explained as part of the programming model in “Chapter 4. MOFW C++ client programming model” on page 85 and “Chapter 5. MOFW - C++ server programming model” on page 103.

Model details

This section provides the details of the Object Model, Application Model and the existing CICS Customer Management Application and Person relational database.

Object model

Object Model includes the following topics:

- “Object Identity”
- “Agent” on page 38
- “Beneficiary” on page 39
- “Customer” on page 40
- “PayoutFraction” on page 41
- “Person” on page 41
- “Policy” on page 42
- “PolicyHolder” on page 44
- “Claim” on page 45

Object Identity

One of the most important concepts in an object-oriented model is Object *Identity*. The Identity of an object and its class uniquely identifies an instance of a class. Identity is a critical concept in the Component Broker programming model and for every class there must be a combination of attributes that uniquely identifies an instance of the class. For example, for the Person class, the unique ID is:

(ssNo, name)

In many cases, the ID must be explicitly defined as an attribute of the object, because no combination of attributes can be guaranteed to be unique. Where a class does not have an explicit ID field, the set of attributes that defines identity is explained.

The Component Broker model does not require that the Identity attributes be a public part of an interface and usable by clients. Identity is needed by Component Broker application adaptors, and to support the OMG Identity Service.

Component Broker introduces the notion of Keys in the programming model to assist in establishing identity and maintaining unique access paths to the business objects. See “Implementing the primary key class” on page 122 for further information.

Agent

The Agent object is reverse-engineered from the existing CICS based Customer Management Application. This object serves as a point of contact for a Customer, and therefore maintains a list of associated customers, as shown in the overview. The following IDL shows the next level of detail:

```
interface Agent
{
    readonly attribute float commissions;
        attribute string name;
    readonly attribute long id;
        attribute float pendingPaycheck;
        attribute float percent;

    void addCustomer (in Customer customer, in float bonus);
    void removeCustomer (in Customer customer);
    Customer findCustomerByID (in long customerNo);

    Iterator listCustomers();

    void payCommission (in float amount);
}
```


Attributes	<p>commissions The total commissions that have been paid to the Agent this year (does not include <i>pendingPaycheck</i>). This is a <i>read-only</i> attribute.</p> <p>name The name of the Agent.</p> <p>ID The ID of the Agent.</p> <p>pendingPaycheck The total earned commissions that have yet to be paid to the agent.</p> <p>percent The current percentage the Agent receives as a commission.</p>
Methods	<p>addCustomer Adds the Customer to the Agent's list of customers and gives a "signing" bonus by incrementing the <i>pendingPaycheck</i> attribute.</p> <p>removeCustomer Deletes the Customer from the Agent's list of customers.</p> <p>listCustomers Lists and iterates through the Customers associated with a given Agent.</p> <p>findCustomerByID When a Customer's ID is passed, returns a pointer to the Customer object, if the Agent is associated with this Customer. It returns NULL if the Agent does not manage the Customer.</p> <p>payCommission Pays the appropriate commission percentage to the Agent on the amount specified, incrementing the <i>pendingPaycheck</i> attribute.</p>

Beneficiary

The Beneficiary is a new object added for this Life Insurance Application. It inherits from Person and adds the association to the PayoutFractions shown in the overview. The IDL follows:

```
interface Beneficiary : Person
{
    attribute float claimPayments;
    readonly attribute long id;

    void addInterest (in PayoutFraction interest);
    void removeInterest (in PayoutFraction interest);
    Iterator listInterests ();
}
```

Attributes	<p>claimPayments The total amount of claims paid to the Beneficiary across all interests in all Policies.</p> <p>ID The ID.</p>
Methods	<p>addInterests Adds the "interest" in the form of a PayoutFraction to the Beneficiary's list of interests. That is, this method updates the Beneficiary to reference the PayoutFraction for a Policy.</p> <p>removeInterests Deletes the interest in a Policy.</p> <p>listInterests Checks on the interests a Beneficiary may have in various Policies.</p>

Customer

The Customer, like Person and Agent, is an object reverse engineered from an existing application. As with Agent, this class is backed by an existing CICS Customer Management Application.

```
interface Customer
{
    attribute Agent agent;
    attribute string custName;
    attribute float sales;
    readonly attribute long customerNo;
    attribute string ssNo;
}
```

Attributes	<p>agent The current Agent in charge of the Customer account.</p> <p>custName The name of the Customer.</p> <p>sales The total sales accrued by the Customer, that is, the total yearly premiums paid by this Customer.</p> <p>customerNo The unique customer ID.</p> <p>ssNo The Social Security number of the customer.</p>
------------	--

PayoutFraction

The PayoutFraction is a new class added especially for this application. The fact that it is attributes only (for a forward-engineered object) is an indicator that it is really more of a “link attribute” (in the manner of Rumbaugh) than a first class object. The IDL follows:

```
interface PayoutFraction
{
    readonly attribute Beneficiary beneficiary;
    readonly attribute Policy policy;
    attribute float claimPayments;
    attribute float fraction;
}
```

Identity	The Object Model allows only one PolicyFraction per pair of Beneficiary and Policy. Thus, the value of these two attributes uniquely identifies the PayoutFraction.
Attributes	<p>beneficiary The Beneficiary who is to be paid the specified fraction of the Policy amount.</p> <p>policy The Policy in which a Beneficiary has a fractional interest.</p> <p>claimPayments The total amount of claim payments made to the Beneficiary for this Policy.</p> <p>fraction The fraction of any claim that is to be paid to the specified Beneficiary.</p>

Person

Person represents another reverse-engineered object, this time from an existing relational database. In the example, it is attributes only. However, new business logic could be added to the class.

```
interface Person
{
    readonly attribute string name;
    attribute string street;
    attribute string town;
    readonly attribute string ssNo;
}
```

Identity	A Person is uniquely identified by: (ssNo, name)
----------	---

Attributes	<p>name The name of the Person.</p> <p>street The street number and name, for example: 30 Saw Mill River Road</p> <p>town The name of the town.</p> <p>ssNo Social Security Number.</p>
------------	---

Policy

The following IDL represents its “static” state (from the Static model), that holds regardless of the dynamic state:

```

interface Policy
{
    readonly attribute PolicyHolder    insured;
    readonly attribute long            policyNo;
    attribute float                    amount;
    attribute float                    premium;
    attribute float                    claimsPaid;
    attribute float                    premiumsPaid;
    attribute float                    pendingClaims;

    boolean addBeneficiary (in Beneficiary beneficiary, in float fraction);
    void removeBeneficiary (in Beneficiary beneficiary);
    Iterator listBeneficiaries ();
    Beneficiary getBeneficiary (in long id);

    boolean changeFraction(in Beneficiary beneficiary,in float fraction);

    void addClaim (in Claim Claim);
    void removeClaim (in Claim Claim);
    iterator listClaims ();
    Claim getClaim (in long claimNo);

    void cancelPolicy();
    void payPremium(in float amount);
    void payClaim(in float percent);
}

```

Attributes	<p>insured The PolicyHolder who is insured by this Policy. This attribute is <i>read-only</i>.</p> <p>policyNo The ID of the Policy.</p> <p>amount The total amount of the Policy, that is, the maximum amount for which the PolicyHolder is entitled to make claims.</p> <p>premium The amount that the Policy Holder is obligated to pay each year in order to keep the Policy In Force.</p> <p>claimsPaid The total amount of claims that have been paid so far against the Policy. It is used to determine when the Policy has been paid off, and is also useful in ROI calculations.</p> <p>premiumsPaid The total amount of premiums paid against the Policy over its lifetime. This figure is useful for ROI calculations.</p> <p>pendingClaims The amount of requested claims still left to be paid to the Beneficiaries if claims are approved.</p>
------------	---

Methods	<p>addBeneficiary Adds a Beneficiary to the Policy with the given fractional amount using a PayoutFraction object.</p> <p>removeBeneficiary Deletes the PayoutFraction from the Policy's list of beneficiaries and the Beneficiary's list of interests.</p> <p>listBeneficiaries Lists the PayoutFractions that show the Beneficiaries "interested" in this Policy.</p> <p>getBeneficiary Gets a Beneficiary by ID, if associated with this Policy.</p> <p>changeFraction Changes the PayoutFraction associated with a given Beneficiary.</p> <p>addClaim Associates a claim against the Policy by incrementing the <i>pendingClaim</i> attribute and creating a linkage between the Claim and Policy.</p> <p>removeClaim Removes the association between a Policy and Claim.</p> <p>listClaims Enables iteration through a set of Claims.</p> <p>getClaim Gets a claim by ID.</p> <p>payPremium Pays a specified amount towards premiums by incrementing <i>premiumsPaid</i>, either against the Policy In Arrears, or In Force. The Policy remains In Force or is placed In Arrears depending on whether the premium payments are up to date. In any event, premiums are not paid except while the Policy is In Force or In Arrears.</p> <p>payClaim Records that a Claim was paid against this Policy.</p> <p>cancelPolicy Moves the Policy into the cancelled state.</p>
---------	--

PolicyHolder

A PolicyHolder represents the Customer of a Personal Life Insurance application. Most of its data comes from the Customer and Person objects, and only the additional behaviors associated with a PolicyHolder are shown here.

```
interface PolicyHolder : Customer, Person
{
    void removePolicy (in Policy policy)
```

```

void addPolicy (in Policy policy);
Iterator listPolicies (in string query);
Policy getPolicy (in long policyNo);
}

```

Identity	Policy Holder derives its identity from Customer.
Methods	<p>removePolicy Removes a policy associated with a PolicyHolder.</p> <p>addPolicy Allows a new Policy to be associated with the target PolicyHolder. At present, the system allows for multiple Policies to be associated with a single PolicyHolder.</p> <p>listPolicies Allows multiple Policies associated with a given PolicyHolder to be sequentially accessed through an iterator.</p> <p>getPolicy Gets a Policy by its ID, if associated with this Policy Holder.</p>

Claim

A Claim tracks the state and processing of a request for payment on a Policy. Its interface is:

```

interface Claim
{
    readonly attribute Policy  thePolicy;
    readonly attribute long    claimNo;
    readonly attribute string  date;
    attribute enum    state;
    attribute string  explanation;

    void approve (in string explanation);
    void deny (in string explanation);
    void pay ();
}

```

Identify	The claim is identified by the claimNo.
----------	---

Attributes	<p>thePolicy The Policy for the Claim.</p> <p>claimNo The unique Claim number.</p> <p>date The date the Claim was created.</p> <p>state The current state of the Claim. The state transition diagram for the Claim is shown in Figure 20.</p> <p>explanation Additional information about the Claim.</p>
Methods	<p>approve Mark the Claim approved and set an explanation or comment.</p> <p>deny Mark the Claim denied and set an explanation.</p> <p>pay Mark the claim as Paid.</p>




Figure 20. Claim states

Applications

This section provides the pseudo-code and business logic for the applications.

Notes:

1. These applications are conversational, and involve many interactions between the end-user and the running application.
2. These applications are transactions and follow typical transaction guidelines for data integrity. On Commit, all databases are updated. On Abort, all changes are removed.

Create Policy

An Agent uses this application to create a new Policy. This application is defined by the following rules:

Pre-conditions	<ul style="list-style-type: none">• The PolicyHolder must exist
Post-conditions of Successful Completion	<ul style="list-style-type: none">• There will be a Person object for the Customer.• The Policy exists and has all attributes set to consistent values, that is, policyNo and PolicyHolder.
Invariants	The Customer is unchanged.
Application Logic	<ol style="list-style-type: none">1. The PolicyHolder is identified.2. A Policy is created and the attributes are set.

ModifyPolicy

Pre-conditions	The Policy must exist.
Post-conditions of Successful Completion	<ul style="list-style-type: none">• The Policy exists and has all attributes set to consistent values, that is, policyNo and PolicyHolder.• The PolicyHolder exists and is in a consistent state.
Application Logic	<ol style="list-style-type: none">1. The Policy is identified.2. Policy attributes are updated.

Create Customer

Pre-conditions	The Customer must not exist.
Post-conditions of Successful Completion	<ul style="list-style-type: none">• Customer exists and has attributes set.• Person exists and has attributes set.
Application Logic	<ol style="list-style-type: none">1. The Customer's name is obtained2. A Customer may or may not have an entry in the Person database. The Customer and Person classes are tied together by a note associated with the Customer through the Customer Management Application.3. If the Customer or Person is missing, it is created.

Modify Customer

This application enables modification of a Customer and the related Person information. It creates the Person if one does not exist and creates the binding between Customer and Person if one does not exist.

Create Beneficiary

Pre-conditions	The Beneficiary must not exist
----------------	--------------------------------

Post-conditions of successful completion	<ul style="list-style-type: none"> • The Beneficiary exists. • The Person exists and is linked with the Beneficiary.
Application Logic	<ol style="list-style-type: none"> 1. The Beneficiary's name is entered. 2. The Beneficiary is created if it does not exist. 3. The Person is created if it does not exist.

Modify Beneficiary

This application updates the Beneficiary and Person objects.

Process Claim

Pre-conditions	<ul style="list-style-type: none"> • The Policy must exist. • The PolicyHolder exists. • All Beneficiaries exist.
Post-conditions of Successful Completion	<ul style="list-style-type: none"> • If this is a new Claim: <ul style="list-style-type: none"> – A Claim is created. – The Claim is associated with a Policy. • If the Claim exists and is in the Created state: <ul style="list-style-type: none"> – The Claim may be Approved with an explanation. – The Claim may be Denied with an explanation.
Invariant	Existence of objects is unchanged.
Application Logic	<ol style="list-style-type: none"> 1. If this is a new Claim: <ol style="list-style-type: none"> a. The Policy is identified. b. A Claim is created. c. The Claim state is Entered. d. The Claim and Policy are associated. 2. If an existing Claim is being processed: <ol style="list-style-type: none"> a. The Claim is retrieved: <ul style="list-style-type: none"> • If the claimNo is known, the Claim is retrieved by number. • If the claimNo is not known, the Policy is retrieved and the Claim is retrieved by iterating through the associated claims to find the Claim with the desired properties. b. If the Claim is Entered, it may either be Approved or Denied. c. If the Claim is Approved, it may be Paid. d. If the Claim is Paid or Denied, it may only be viewed.

Key observations

The main points of this example are to demonstrate that:

- Both state data and attributes of objects can be derived from the in/out parameters of existing applications.

- Many object methods can be mapped onto existing applications.
- Reuse of applications may often involve complex mappings between transactions, menu states, and screens. Often, the mapping is simple and almost one-to-one between methods and transactions and attributes and the fields on display or update screens.
- Even if the data is not maintained in an RDBMS, the application usually provides support for limited query through selected attributes and iteration through the results. This is most often accomplished by simple query operators on keys.

Chapter 3. The managed object framework

This chapter introduces the managed object framework (MOFW) and its positioning within Component Broker. A systematic approach to explaining the details of the MOFW follows in succeeding sections.

The MOFW represents the set of interfaces, implementations, and conventions that must be followed in order to create and use business objects in Component Broker. The MOFW provides capabilities above and beyond those present in the basic CORBA ORB and Object Services defined by OMG. MOFW also provides simplified interfaces to some of the basic CORBA interfaces. The MOFW is not the only set of interfaces supported by Component Broker. Component Broker allows for additional frameworks which can be used by business objects and client programs.

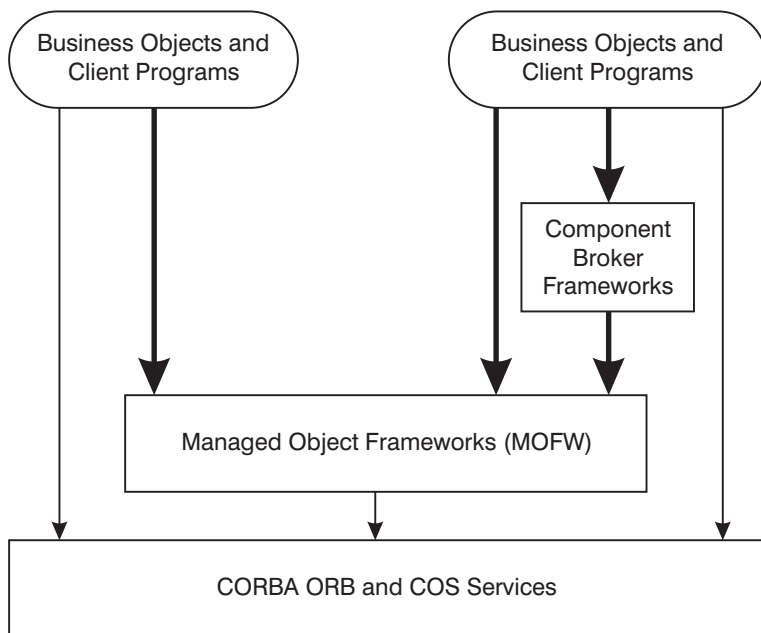


Figure 21. Component Broker and MOFW framework overview

As Figure 21 shows, business objects and client programs that use business objects can be written directly to the MOFW interfaces. The MOFW is not a complete layer over the CORBA services. It adds usability and function only in those places key to providing an integrated object server. Sometimes the existing CORBA ORB and Object Services provide what is needed to

implement the Component Broker server vision. In some cases, the Component Broker frameworks provide simpler access to the server. The object provider rather than the applications programmer chooses the interfaces. The Component Broker frameworks are not, however, a complete encapsulation of the MOFW interfaces. Be careful when mixing these sets of interfaces together. This document provides guidance on which combinations of Component Broker frameworks and the MOFWs make sense.

Examples of the abstractions found in the managed object framework (MOFW) include:

- `IManagedClient::IHome`
- `IManagedClient::IManageable`
- `IManagedLocal::ILocalOnly`
- `IManagedLocal::INonManageable`
- `IManagedLocal::IKey`, `IUniqueKey` and `IPrimaryKey`
- `IManagedLocal::IHandle`
- `IManagedServer::IManagedObject`
- `IManagedServer::IManagedObjectWithDataObject`
- `IManagedServer::IManagedObjectWithCachedDataObject`
- `IManagedServer::IDataObject`
- `IManagedCollections::IReferenceCollection`
- `IManagedCollections::IKeyedReferenceCollection`
- `IManagedCollections::Iterator` and `IManagedCollections::IMIterable`

Managed and non-managed objects

There are two kinds of objects in Component Broker, those that are managed by a Component Broker server and those that are not. All of the objects that client application programmers and business object builders deal with descend, directly or most often indirectly, from either `IManagedLocal::ILocalOnly` or from `IManagedClient::IManageable`. Component Broker has these two specific kinds of objects to ensure a minimum footprint client, separate server-only objects from those that may exist on clients and servers, and, most importantly, simplicity for the programmer. No extra methods need to be used or implemented based on this separation.

Those objects that are to be local only are descendants of `ILocalOnly` or `INonManageable`. Those objects that are to be accessed remotely and managed by a Component Broker server are subclasses of the `IManageable` interface. Figure 22 on page 53 shows the basic relationship between these MOFW interfaces and the CORBA Object Services interfaces.

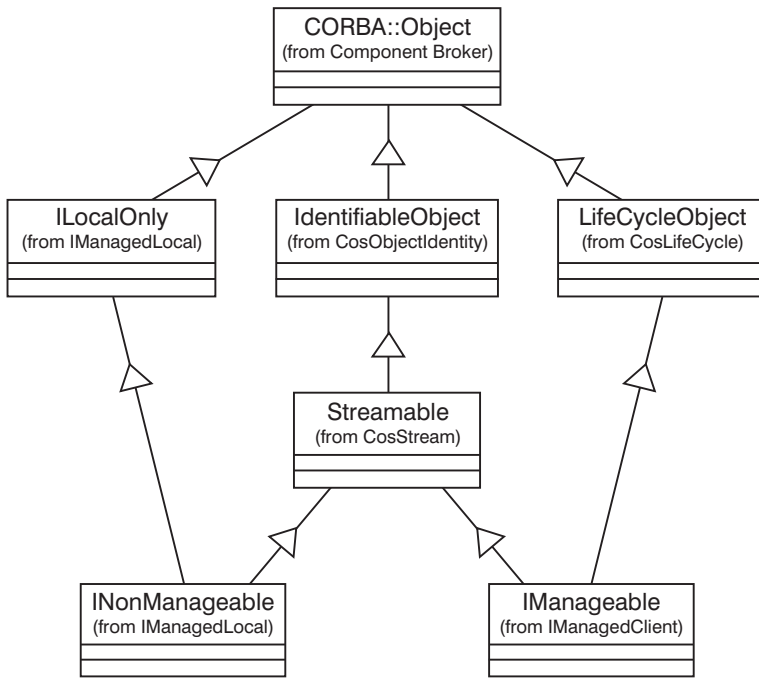


Figure 22. MOFW basic abstractions

The INonManageable interface comes from a module named IManagedLocal while the IManageable interface comes from the IManagedClient module. The Local Only Development Process contains other abstractions that are local only objects. While these are not to be accessed remotely, Component Broker made these descendants of CORBA::Object and treats them as much like CORBA objects as is practical for local only objects. INonManageable extends ILocalOnly by introducing methods that assist developers in making “stringified” versions of ILocalOnly objects.

IManagedClient is a set of abstractions that business object clients need to understand to some degree, and interact with when writing applications that use Component Broker business objects. Object providers subclass from, and often implement, abstractions in IManagedClient while client and applications programmers call some of the methods introduced by these abstractions.

Both of the key abstractions shown at the bottom of Figure 22 also get some additional interface from CORBA Object Services. Both INonManageable and IManageable are descendants of CosStream::Streamable, implying that they are also Identifiable objects. IManageable objects are also LifeCycle objects. This is described in “Implementing the IManageable required methods” on page 113.

The important fact is that there are two kinds of objects to be remembered, and that each has a slightly different ancestry and purpose in Component Broker.

All Component Broker objects inherit from CORBA::Object. This inheritance from here forward is not generally shown. INonManageable has a notable parent ILocalOnly. Component Broker programmers can always tell if a particular object is local only or accessible remotely. IManageable also implies a lot more. It means that descendants of this base class get not only remote accessibility, but also a full range of services from the Component Broker server. The Component Broker server takes care of the IManageable objects, allowing them to be persistent, accessed securely, participate in transactions, and take advantage of all of the additional Component Broker server features.

You need INonManageable and ILocalOnly because there are some internal Component Broker objects are ILocalOnly subclasses but do not need the additional interfaces implied by INonManageable.

Note: A Component Broker object cannot inherit from both IManageable and INonManageable. It can have only one of these parents.

All Component Broker objects are described in IDL for documentation and consistency. Each programming language that has a usage binding can be used to interact with Component Broker MOFW-based objects. For example, a C++ programmer deals with descendants of IManageable through a C++ usage binding to the IManageable's capabilities.

In addition to these two kinds of objects, any application using Component Broker also uses native language objects. A mixture of native language objects and usage bindings to Component Broker objects make up a particular application. Component Broker supports language usage bindings for Java and C++ and also provides special support for client programmers using ActiveX programming tools.

Using Component Broker MOFW-based objects is done with usage bindings as described previously. Building Component Broker MOFW-based objects is different. Component Broker features the construction of business objects using C++ or Java. The IBM ORB and emitters provide the basis for the implementation bindings that are used to construct implementations for the Component Broker MOFW-based business objects.

Understanding MOFW objects

All MOFW-based objects, both `IManageable` and `INonManageable`, inherit from `CORBA::Object`, even though this might not be apparent from the IDL.

`IManageable` objects:

- Are always created through homes, that are located using factory finders.
- Are accessible remotely through client usage bindings in Java and C++.
- Are implemented using C++ or Java.
- Are always managed by an application adaptor and reside on a Component Broker server.

`INonManageable` objects:

- Are always created using a static create method of the form `className::create()`, a method generated by the IDL compiler.
- Look and behave a lot like regular native language (Java or C++) objects.
- Are not managed by an application adaptor. They might be used on the server, transiently, but they are not persistent.
- Are written in C++ for use by the server and C++ client. They also must be written in Java if the Java-client feature of Component Broker is to be used.
- Are developed using a special local only development process. The emitters only generate the subset of the various bindings that are necessary to support local use. (For example, no dispatcher code, no server-side bindings, and so forth.) This is explained in more detail in “Local only object implementation details” on page 351.

Component basics

Each component in Component Broker consists of a number of pieces. The primary interface to the component and its implementation are shown in Figure 23 on page 56.

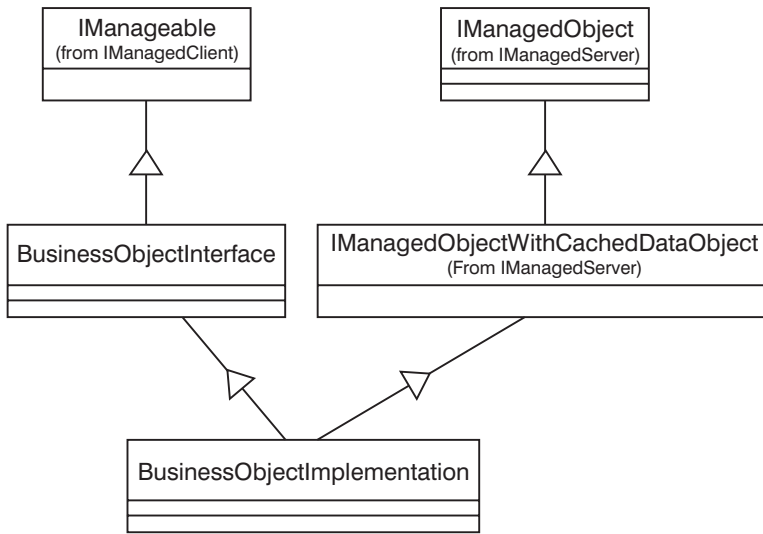


Figure 23. Business object basic structure

The BusinessObjectInterface abstraction in Figure 23 represents the interface as specified by the domain expert involved in the object-oriented analysis and object-oriented design (OOA/OOD) activities for the system under construction. This is represented by IDL and is an interface-only class from a run-time perspective. It is also the interface that you are most likely to interact with this business object. These types of interfaces always inherit from the IManageable interface in the IManagedClient module.

The BusinessObjectImplementation abstraction is an IDL file which is used to generate implementation bindings that are for the business logic that implements the BusinessObjectInterface described previously. The BusinessObjectImplementation inherits from one of the abstractions contained in the IManagedServer module. The example shows inheritance from the IManagedObjectWithCachedDataObject interface. This name indicates a number of things about how the BusinessObjectImplementation object will appear. The IManagedServer module interfaces prescribe additional responsibilities that must be met in the BusinessObjectImplementation that deal with Component Broker and the infrastructure.

The two key modules involved are IManagedClient and IManagedServer. IManagedClient represents those abstractions that are generally accessible to clients, and subclassed by object builders when they want to expose an interface to client programmers. As described in “Chapter 2. Personal life insurance application example” on page 31, this includes abstractions such as IHome, IHandle, and IManageable.

IManagedServer is a module that contains interfaces used only by object providers and customizing applications to run in particular environments. The abstractions in IManagedServer enable components to run in the Component Broker server and properly take advantage of the Component Broker server capabilities. Examples of this are the IDataObject abstraction and the abstractions that deal with enabling the basic component for a specific data object pattern. Understanding these two modules and the different purposes they serve is important.

Before a component can be tested, you need to complete the construction of its helper classes, namely primary keys and copy helpers. The additional steps involved in preparing a component to be installed in a Component Broker server for further testing, and eventual deployment, are discussed in “Chapter 13. Assembling and installing components on Component Broker/Workstation” on page 351. These are found in the module named IManagedLocal, and are the *local only* objects which assist in making MOFW work across clients and servers.

There are also other considerations with respect to factoring of the interfaces for various clients. Before proceeding with the steps to develop a simple component, it is important to establish some common terminology.

Component state

In object-oriented programming, an object has both behavior and state. An object’s behavior is manifested in the implementations of the methods on the public and private interfaces of the object. The state of an object, on the other hand, is predominantly manifested in the public and private data members of the object. The state can be divided into categories of *non-essential* or *essential*.

An object’s non-essential state consists of the subset of the state that can be calculated or derived from other state, and the subset of the state which is transient and does not need to be persistent because it can be recreated as necessary. A state that is not non-essential is considered to be essential state. To illustrate, assume a Policy component consists of the following:

- Policy number
- Amount (of coverage)
- Premium
- Insured (policy holder)
- Comment
- Risk calculator (helper object)

Furthermore, assume that the premium can be calculated by dividing the amount by 100. The risk calculator is a helper object which itself has no state, and a new one can be created if the Policy component is passivated and subsequently reactivated. The insured, on the other hand, is another

component of type PolicyHolder. If the Policy component is passivated and reactivated, the Policy component must find the same PolicyHolder as it had before passivation.

Given the previous assumptions, Figure 24 shows that the policy number, amount, insured, and comment are part of the essential state of the Policy component, whereas the premium and risk calculator are part of the non-essential state.

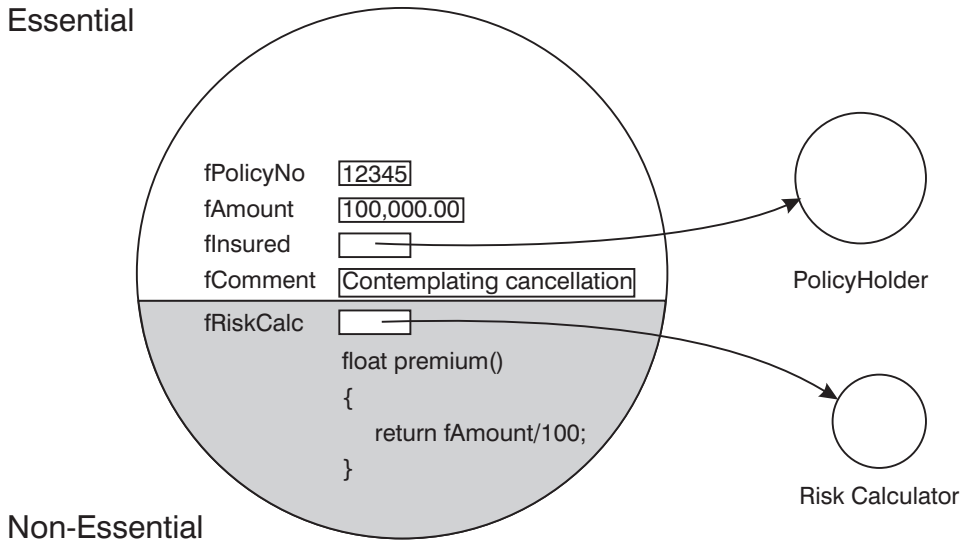


Figure 24. Essential and non-essential states

Component attributes

In C++, the term *attribute* is often used synonymously with the term *data member* to describe a piece of data which is part of an object's state. In CORBA or, more precisely, in IDL, attribute has a different meaning altogether. In CORBA, IDL describes the interface of an object. Because CORBA is all about distributed objects, and because there is no way to directly address data inside an object from across a network, IDL does not support the notion of having data as part of an object's interface. Thus, in IDL, attribute refers to a pair of methods for accessing and changing the value of a specific type. For instance,

```
attribute string comment;
```

defines the following methods in C++

```
virtual char* comment ();
virtual ::CORBA::Void comment (const char* comment);
```

The first of these two methods retrieves the value of the comment. The second allows the client to change it. A component often has a state which is necessary to support the implementation of the business logic, but which is not directly accessible by an attribute on the object's interface. In object-oriented design (OOD), this is referred to as *encapsulation*.

Now you should be able to go through the steps necessary to specify and implement a component based on the MOFW set of Component Broker interfaces. The goal of this chapter is to present enough material for you to create the minimum set of artifacts necessary to support the basic client programming model as described in "Client programming model: basic tasks" on page 86.

MOFW client programming model

This section describes the Component Broker managed object framework (MOFW) client programming model. The word *client* generally refers to a particular computer in a distributed environment and the relationship of that computer with other computers. In this context, the word *client* refers to any program and its relationship to a component. In other words, a client is any program executing in a process on a client or server computer.

Therefore, the client programming model explains how you can use components and how to develop objects that are clients of Component Broker components. Application programmers developing *tier-1* (client) or *tier-2* (server) applications use the client programming model when they use a component in the implementation of a new object. Client programmers use server objects to develop new applications by composing new application objects and scripting behavior over sets of server objects.

A Component Broker server process implements objects that are used by client applications to perform business functions. All Component Broker server objects are derived from the managed object framework. These objects effectively make up a component and are accessed through the part of the component called the managed object. The term *managed object* is used because one of the Component Broker server product's main contributions to object-oriented applications is the implementation of run-time management functions used in the implementation of server objects. Some examples of these management services are:

- Persistence, Transactions, and Security.
- Workload management and availability management over multiple servers.
- Object-oriented access to existing databases and applications.

Component Broker supports writing client applications in C++, Visual Basic, or Java. For information on writing client applications in C++, see "Chapter 4.

MOFW C++ client programming model” on page 85. For information on writing client applications in Visual Basic, see “Chapter 8. MOFW – ActiveX client programming model” on page 211. For information on writing client applications in Java, see “Chapter 9. MOFW - Java client programming model” on page 231.

Client view of Component Broker applications

Figure 25 presents a high-level overview of the client programming model. The way clients deal with components at the programming model level is consistent regardless of the underlying support mechanisms used to build those components.

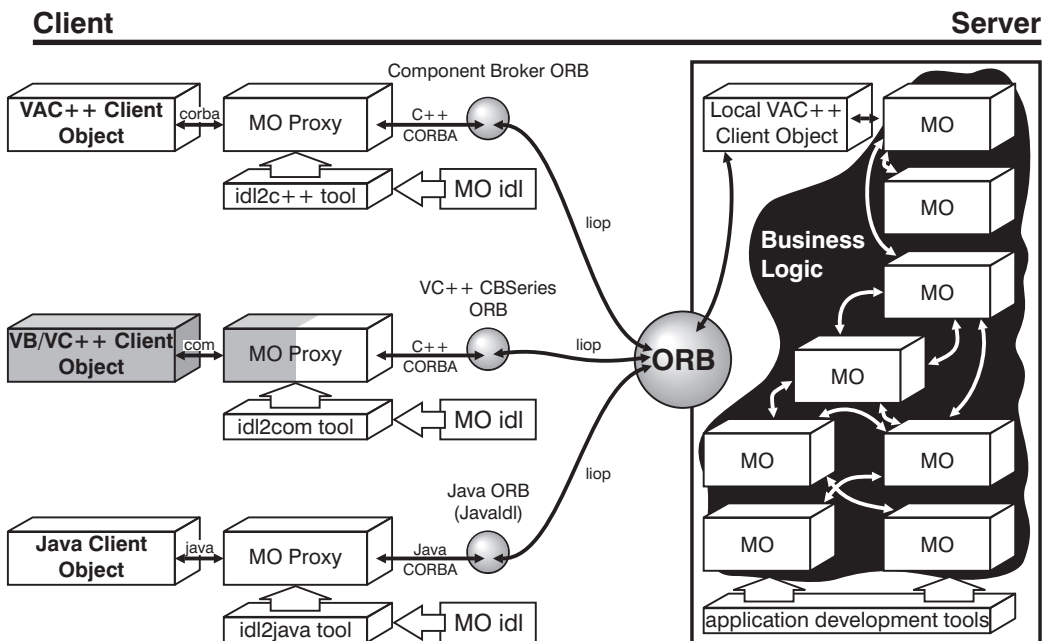


Figure 25. Client model overview

Component Broker provides functions that enable the development of ActiveX/COM *wrappers* for Component Broker managed objects. The wrappers implement an ActiveX/COM rendering of the component interface. Within the COM object is a Component Broker proxy, and this proxy delegates business logic calls from the COM object on the client to the component’s managed object on the server.

The *managed object implementor*, sometimes referred to as a *server application*, provides the client application with:

- A set of interface files that define the interface to a component and any *helper* class the client uses. Interface files are provided for each client programming language; for example, .hh files for C++ and .java files for Java. In addition, Interface Definition Language (IDL) is provided for the managed objects and helper classes.
- DLL and Java .class files that implement the classes in the interfaces and the helper classes. These include the proxy classes that enable remote/local object use.

To use an object-oriented application (a set of components) you need to understand its object model. Specifically, you need to understand the interfaces and their methods and attributes.

The basic client programming model assumes that you have compile time access to the server interface definitions (.idl files) and have documentation on the class relationships and the semantics of methods and attributes. The artifacts provided by the object provider define the interface syntax. What the interface does is provided in comments and in the documentation.

Component Broker supports a remote/local programming model for components. This is implemented by run-time libraries, proxy classes, and application development tools. The managed object implementor defines the interface to the component using CORBA IDL. The Component Broker application development tools emit the necessary language bindings and proxies to support access to the component from remote applications.

The Component Broker managed object framework allows object providers the opportunity to present a number of programming interfaces to clients. Determining which client interfaces to use is an advanced skill in client programming and not critical to getting started. The kinds of client programming interfaces provided and tips on how to select the correct one for a particular application are described in “Expanding the client programming interface” on page 434.

For this chapter, the container-type independent, business logic-only interface is used. This interface is the simplest with which to get started. The examples in this document use interface names like “claim” and “policy”. Using the client bindings for these interfaces provides access to all the methods that the object provider defined for client usage.

MOFW server programming model

The client application accesses *components* on the server. Components consist of instances of classes that implement the business logic of the application. Component objects support the interfaces of the managed object framework which allows components to be installed onto, and managed by, a Component Broker server. The part of the component that is installed on a Component Broker server is referred to as a *managed object* (MO).

If the managed object is in the same process and same language as the client application, the client directly accesses the object. If the object resides in another process or on another machine, the client uses a CORBA *proxy* object. This proxy object:

- Implements the interface of the component in the language and process of the client.
- Uses the ORB and CORBA object RPC to relay method calls to the component's managed object in the remote process.

Application objects

Application objects are objects which are directing workflow and implementing some client initiated task. Methods often have an action verb associated with them which resembles the name of the use-case or end user transaction being performed. Application objects that are very process oriented are often stateless. They perform a particular task, calling other basic or composed business objects along the way, and then returning an indication of success or failure to the client application. A stateless application object lives in a transient container and is a good candidate to be workload managed. Stateless application objects can be reused by clients. In fact, a popular architecture for stateless application objects is for a client to create one application object, keep the reference to it and use it each time it needs the services which are provided. In these cases, all relevant data need to accomplish the tasks is either passed into on the method signature or cached in a way that all instances of the application object can leverage it.

Application objects are often created by specialized homes that assign keys. These specialized homes either use the UUID Key support provided by Component Broker or some other mechanism to assign keys. There is little need for clients to create keys for these kinds of objects. The specialized home for the application object is also the place where pooling or reusing of instances (for stateless application objects) across clients can occur.

Application objects are sometimes called command objects. Command objects are a specific kind of application object that captures some state transiently and has a method that will undo what the main method has done in the

event that something unexpected occurs. This `undo()` method must compensate as necessary, calling the other business objects to reset or rollback to a previously known state.

Even more sophisticated application objects will enable and allow saving of a client session persistently. These are hybrid application objects that are a little like compositions and a little like applications.

Application objects usually initiate transactions and coordinate the work of one or more basic business objects or composed business objects. This means that the application objects themselves reside in No Object Services containers meaning that the application object itself does not participate in the transactions. Stateful application objects that are in No Object Services containers that are shared by multiple clients must be coded in such a way as to serialize access to any data. This only applies when there are stateful application objects that are used by multiple clients.

Application objects that are doing read-only tasks can be placed in containers that start new transactions for each method. This takes the responsibility for starting and committing transactions away from both the client programmer and the application object programmer. This configuration is only recommended for read-only situations and only when no special processing is required if the transaction were to be rolled back.

The following table summarizes the key attributes and differences between the primary kinds of business objects.

Table 1. Key attributes and differences between business objects

	Application Object	Business Object	Composed Business Object
Key	commonly UUID	based on state data	combination of one or more attributes in the composition which might be object references
State	commonly stateless, sometimes transient state, and rarely persistent	persistent state backed by third-tier resource manager, based on existing tables or definitions	state is largely references to basic business objects either calculated when composition activates or stored in new database tables

Table 1. Key attributes and differences between business objects (continued)

	Application Object	Business Object	Composed Business Object
Transactional Behavior	controls transactions or sessions and sometimes participates	always transactional or sessional; participates, doesn't initiate	participates, may need to initiate and control if composite parts have different transactional semantics (e.g., sessional versus transactional)
Clients	client applications (servlet, applet, application) and sometimes other application objects	composed business objects, related business objects, application objects and sometimes client applications;	application objects and sometimes client applications

Using statics in business objects

The usage of static variables is something supported by the C++ and Java programming languages. Static variables in general come with some usage rules and some restrictions. These apply when using static variables in the implementation of business objects. Some additional points need to be made as well.

First, static variables are going to apply to every instance of a business object that is active in a given server. This means that if there are multiple servers, for example, in a server group that support creation of the same business object type, there will be multiple copies of this static data. There will be one set of statics per server. For keeping class level data, this is seldom a problem.

Secondly, static variables last for the life of the server. The general pattern to use for static variables is to encapsulate their use in methods. These methods contain logic that checks for null and then sets them the first time. This keeps the business object programmer from having to worry about static initialization issues. This also keeps the usage of these statics encapsulated from the rest of the business object method implementations.

Static variables do not follow the CORBA memory model very cleanly. Static variables do not need to be released and have no notion of a reference count. The storage occupied by statics will not be freed up until the server is terminated.

Keeping object references cached for a long time introduces some important considerations. Specifically, if the server that has the object's implementation

is unavailable or that object implementation is placed in a different container in the server, then the object reference will not be usable and an exception can be expected. Placing object references into static variables does not change this. It does elongate the time period that cached object references may be expected to operate.

In general, static class variables should not be used unless there is a strong performance need. These are the object-oriented version of global variables and should be treated as such.

Circular references

Business object interfaces that have circular references should all be contained within a module. This means that if interface A includes interface B, and interface B refers to something from interface A, both of these interfaces are required to be in the same module. While not recommended, it is also possible to have circular references like this at the global scope.

Releasing and deleting objects

Eventually, you no longer need to use an object that you found or created. Component Broker supports two interpretations of “no longer needs”.

- The `remove()` method deletes the object and its persistent instance data.
- The `release()` method informs the object that the client application no longer plans to reference the object. The object still exists in the server, and other applications may be using it, but the client calling the `release()` method will no longer use the object.

In order to better understand what happens as a result of a `remove()` or `release()` request, and the difference between the two, refer to “Finding a managed object” on page 87 and “Creating a managed object” on page 92. To illustrate this, use an example which has a relational database table of information about people, the policy holders from the Life Insurance application. This table persistently stores the state data for objects of the `PolicyHolder` class.

As described in “Finding a managed object” on page 87, component objects may be found by invoking the `findByPrimaryKeyString()` method on a `home`. The result of this operation is two-fold:

- If an object matching the primary key is not currently active in the server, that is, in use by some other user or application, then one is activated and brought into memory on the server.

- An object reference, also referred to as a *proxy*, to the object on the server is returned to the client. In CORBA, a proxy is itself an object: it is an object on the client which identifies an object on the server, and has the same interface as the object on the server.

When creating a new component object, there is no object on the server, active or otherwise, with a matching primary key. If there is, the creation will fail. A new object is created in memory on the server, a new row is added to the database table, and a proxy is created on the client as shown in Figure 26.

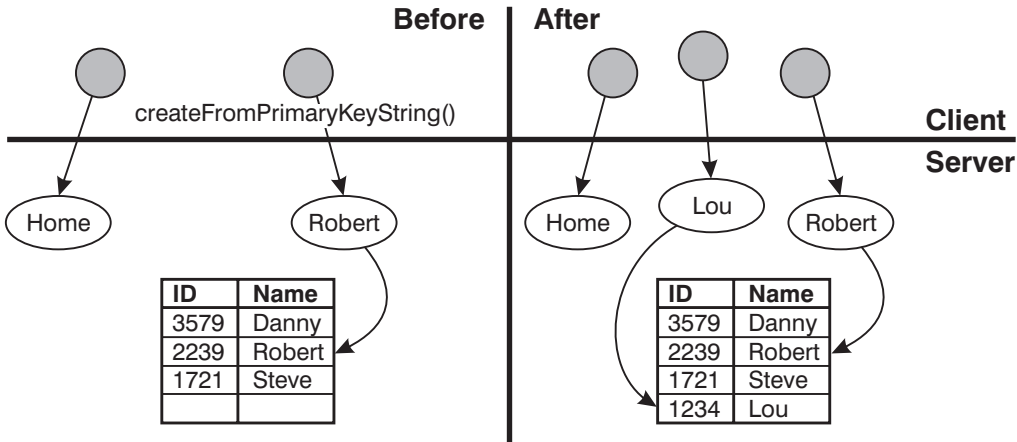


Figure 26. Creating components

Having found or created a component, the client application has a pointer to a proxy, not to the actual component. (A memory pointer to the actual object out in some server on the network has no meaning in the client's address space.) However, because the proxy has the same interface as the component on the server, the client application uses it just like the object itself. The Component Broker ORB and server take care of routing all method calls and parameters from the client proxy to the server object.

Figure 27 on page 67 shows how a client application changes the *name* attribute of an object on the server.

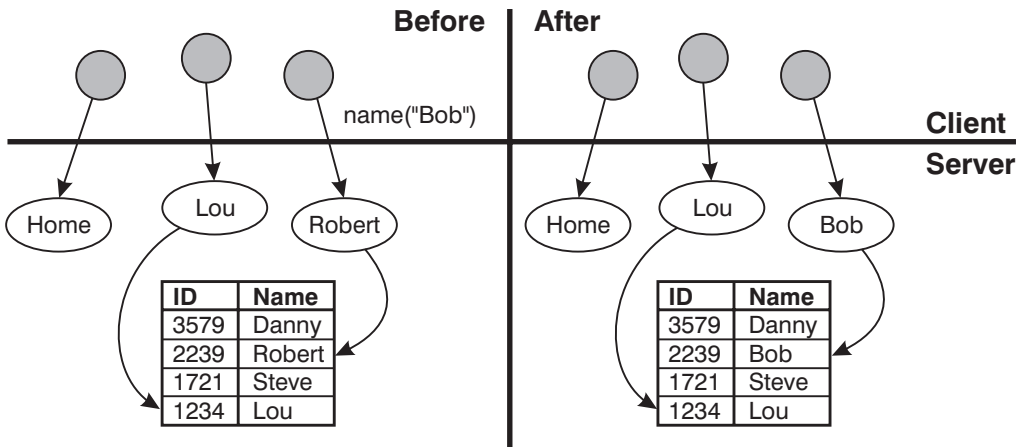


Figure 27. Using components

Now, assume that Robert or, as he prefers to be called, Bob calls up and says he would like to cancel his only insurance policy. At this point, the client application concludes that Bob is no longer a policy holder, and proceeds to delete the Bob component. The method for deleting an object in Component Broker is called `remove()`. Invoking the `remove()` method on a component's managed object has the following effect:

- The component's persistent state is removed from the database table.
- The component objects are removed from memory on the server.

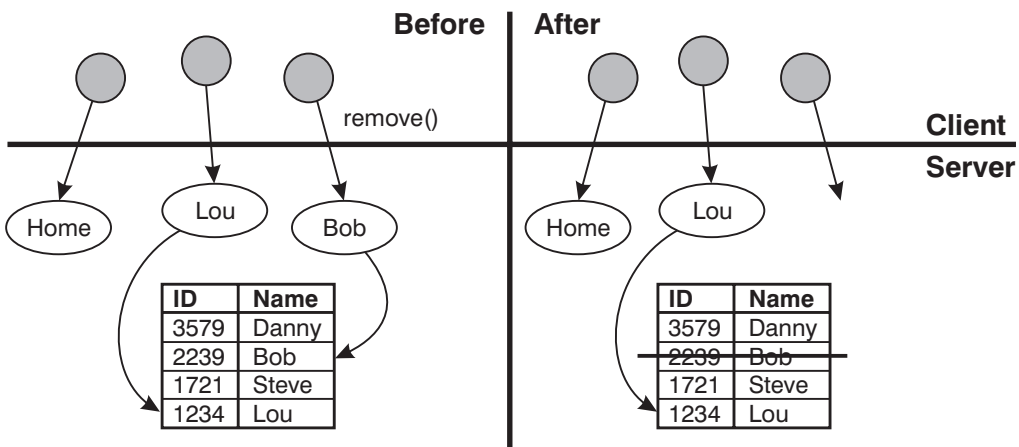


Figure 28. Removing components

There is one important thing to take note of here:

- The proxy (object) is not removed from memory on the client as a result of invoking `remove()` on the proxy. At this point, the proxy is not pointing to a valid object on the server. In order to delete the proxy (object) from memory, you must explicitly invoke `CORBA::release()`, passing it the proxy (object).

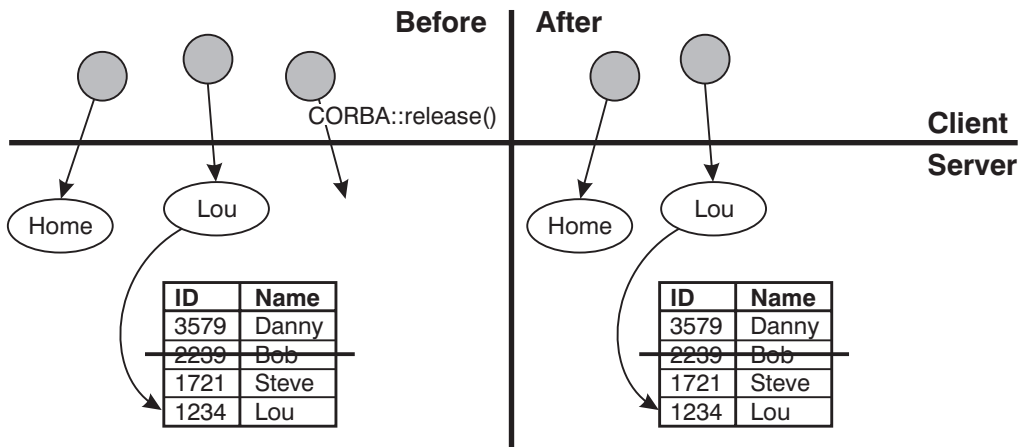


Figure 29. Releasing (removed) object references

Figure 29 shows the result of invoking `CORBA::release()` on the object reference which was removed. Of course, a client application will also want to release references to objects which it is not removing. This would be the case if customer Lou called up to increase his insurance coverage. In this scenario, the client application performs the following steps:

1. Invoke the `findByPrimaryKeyString()` method to find Lou's policy.
2. Invoke the `amount()` method on the policy to increase the coverage.
3. Release (not remove) the policy component object after calculating the new premium.

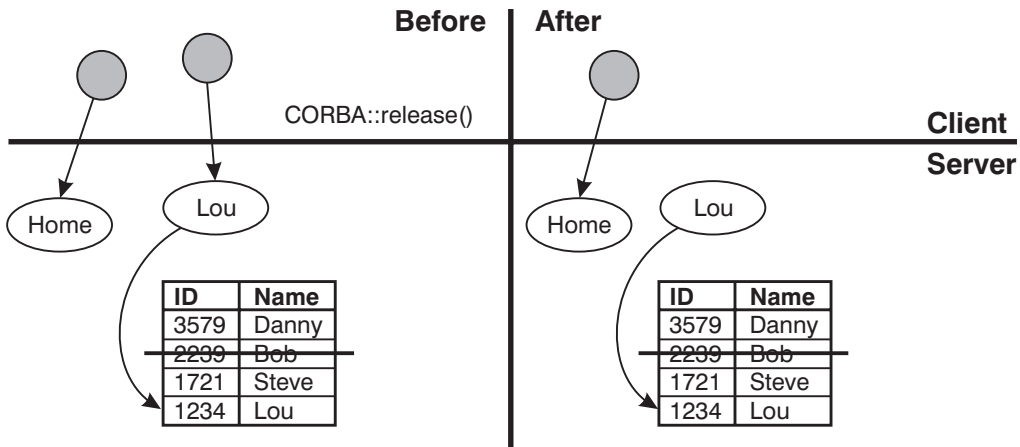


Figure 30. Releasing object references

Figure 30 shows that the proxy is deleted from the client’s memory space, but the component remains on the server. Eventually, if there are no other references to the component within the server, the server deletes the component objects from the server’s memory space but not from the database table. Note that the client application has no control over when the server decides to remove an active component (that is, one that has not been deleted by the `remove()` method). In fact, if the client application is waiting for user input, a server may decide to delete from memory a component to which the client application holds a reference. When the client application finally receives input from the user and attempts to use the component’s managed object, the server reactivates the component by putting it back in memory.

The `remove()` method is essential for supporting application logic. Client applications need a way to destroy and delete components and their instance data. In the previous example, Bob no longer exists as far as the applications are concerned. The `release()` method on the other hand, helps implement memory management. In a client/server system, applications span multiple servers, processes, and languages. Traditional memory management techniques such as garbage collection or relying on clean-up when a process terminates do not work. The client garbage collector does not see the server memory, and the server does not contain references to the objects because they are on the client. Therefore, the server needs to be told when clients no longer reference objects.

Because the server does not maintain a list of every object reference that every client application has requested, you might try to use an object reference to an object on which another client application has already invoked `remove()`. If this happens, an exception is triggered. The same thing would happen if a

client application tried to invoke a method on an object on which it had previously invoked the `release()` method.

Navigating the name space using the Naming Service

The Component Broker name space is hierarchical and similar in structure to a file system directory tree. As the nodes in a directory structure are files (either directories or leaf files), the nodes of the Component Broker name space are CORBA::objects (either NamingContexts objects or leaf objects). A NamingContext is an object that contains zero or more bindings of string name-object reference pairs. Each object, bound by name into a context, can be a leaf object or a subordinate NamingContext in the tree. Subordinate NamingContexts similarly can contain bindings of other NamingContexts and leaf objects.

Component Broker introduces the interface `IExtendedNaming::NamingContext` as an extension of the OMG defined `CosNaming::NamingContext` interface. This interface provides simplified methods for binding, unbinding, and resolving name-object pairs in the name space. Specifically, the `NamingStringSyntax::NameString` type allows programmers to specify the name space path in a string format similar to the way they use strings for specifying directory paths. For example:

```
"name1/name2/name3"
```

where:

name1 Identifies a NamingContext contained in the current NamingContext (and bound to "name1").

name2 Identifies a NamingContext within *name1*.

name3 Identifies a leaf object or another naming context bound in *name2*.

The following is an interface for `IExtendedNaming`. For more detailed list of operations, see the `IExtendedNaming` module in the *WebSphere Application Server Enterprise Edition Component Broker Programming Reference*.

```
module IExtendedNaming
{
    // This type resolves to a string that maps to a char*
    typedef NamingStringSyntax::NameString Name;

    interface NamingContext : CosNaming::NamingContext
    {
        void bind_with_string (in Name n, in CORBA::Object obj)
            raises (NotFound, CannotProceed, InvalidName, AlreadyBound);

        CORBA::Object resolve_with_string (in Name n)
            raises (Not Found, CannotProceed, InvalidName);
    }
}
```



```

void unbind_with_string (in Name n)
    raises (NotFound, CannotProceed, InvalidName);
    :
    :
};
};
};

```

The OMG Naming Service specification defines only the interface to the Naming Service and does not define any structure for the name space. Figure 31 provides an introduction to the Component Broker name space structure.

graphic to be completed

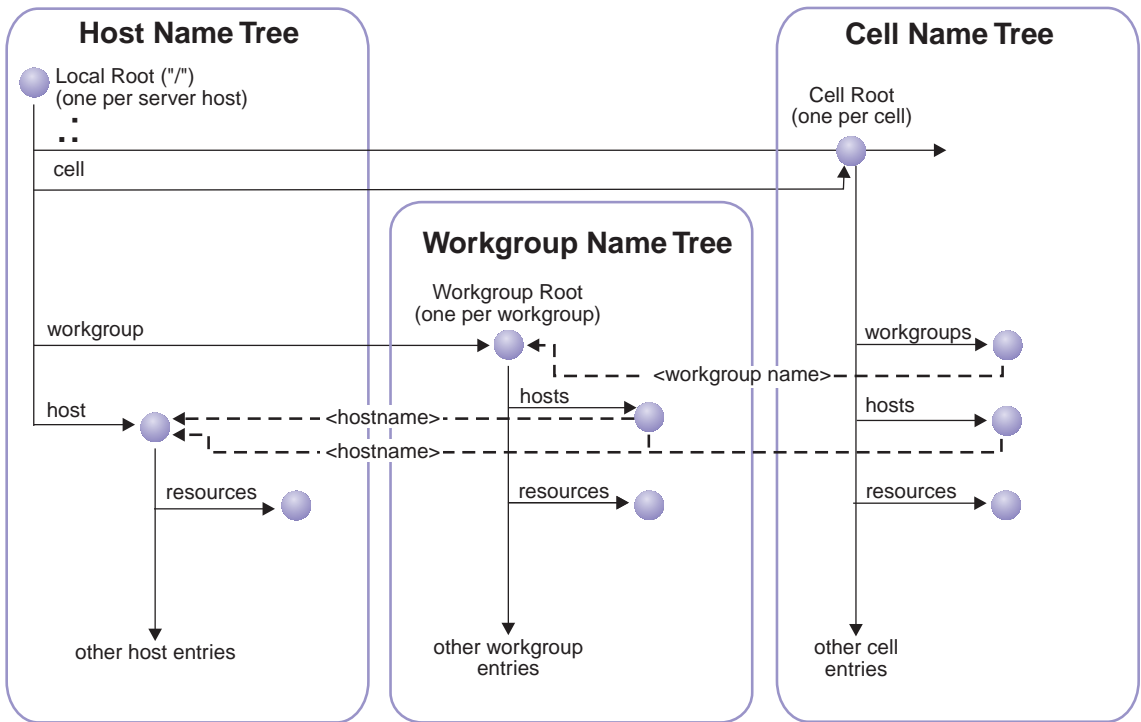


Figure 31. The name space structure

This structure defines the top level (root) naming contexts that exists in an enterprise (cell). Every host system participating in the network has a Local Root NamingContext. This is the NamingContext that is made available to applications through the CBSeriesGlobal::nameService() static method. This context is the anchor from which applications navigate the name space. Each host system belongs to one and only one work group and one cell.

A work group represents a subset of the systems participating in the network. The Workgroup Root NamingContext can be navigated from the Local Root NamingContext by resolving names prefixed with `workgroup/<rest of path>`. The NamingContexts of other systems within the work group can be navigated from the Local Root NamingContext by a path that traverses the work group NamingContext if the hostname of the desired system is known. For example, the name `workgroup/hosts/<hostname>/<rest of path>` provides access to the NamingContext of another host within the work group.

A cell represents the set of all systems and work groups participating in the network. The Cell Root NamingContext can be navigated from the Local Root NamingContext by resolving names prefixed with `cell/<rest of path>` or `./<rest of path>`. As can be seen from Figure 31 on page 71, the NamingContexts of any system or work group in the network can be navigated from the Local Root NamingContext by a path that traverses the Cell NamingContext.

The `resolve_with_string()` method returns a `CORBA::Object`. Before a client can use the returned reference as a NamingContext or as a specific class of leaf object, the client must narrow the object to the desired class.

A note on security

Setting the client up to work with a secure server is primarily a configuration and administration issue. For additional information, see Security under Concepts in Component Broker for Windows NT and AIX Online Documentation.

The client programming implications of security are fairly limited. Although there is much that you can do with security, there is little that you must do. The Concepts and Procedures sections under Advanced Topics in Component Broker for Windows NT and AIX Online Documentation describe many of the features available to someone needed to handle complex security issues. For the most part, however, you need to know that any request to a Component Broker object on the server can be denied if the client does not have the necessary security tokens.

More copy helper topics

Copy helpers can be used for more than just efficient creation mechanisms. This section provides more techniques that can be done using copy helpers.

Using copy helpers

A typical usage of a copy helper is used for efficient creation of objects as described in “Creating a new object – create from copy” on page 94. Copy helpers can also be used for more efficient reading and updating a large number of attributes of a business object after it has been created. More Copy helper classes will typically needed to be defined for this purpose.

For example, usage of the Copy helper for object creation that as shown in “Creating a new object – create from copy” on page 94 is shown below.

```
ClaimCopy claimCopy = ClaimCopyHelper._create();
// Now initialize the Claim's attributes. Note that these methods
// execute locally, within the same language.
claimCopy.claimNo(1234);
claimCopy.date("10/14/96");
claimCopy.state(entered);
claimCopy.reason(accident);
claimCopy.description("Side-swiped by teenager in a red convertible.");
```

This is the copy helper that is configured with the component and is used on the `createFromCopyHelper()` operation on the home of the business object.

The same copy helper cannot be used for updating attributes of this business object after the object has been created. The reason is the attributes that make up the primary key are *readonly* only and must not be changed after the business object has been created. If this copy helper was used for updating then these attributes could be changed. Therefore it is necessary to define other copy helpers classes for updating attributes after the business object has been created. In this example the *claimNo* attribute is the only attribute that makes up the primary key. Therefore a new copy helper class for updating can be defined that provides all attribute interfaces except *claimNo*.

The previous paragraph describes the first step, the second step is to define additional interfaces on the business object. For operations that read attributes from the business object the implementation is to create an instance of the correct copy helper class. Fill in the attributes, run a `toString()` operation on the copy helper and return the bytestring. For operations that update attributes to a business object the interface has an input parameter that accepts as input a bytestring that results from a `toString()` operation on the copy helper. The implementation of these operations is to create a new instance of a copy helper of the correct class, run the `fromString(byteString)` on that instance using the input data and then update the attributes supported by that copy helper.

For example, the client code to update attributes in the business object would can update all attributes but the ones that are *readonly*. These are *claimNo* and *date*.

```

// Use a Copy Helper Class that the Claim MO provider gave me for
// updating and reading.
ClaimCopyForUpdates claimCopyAlternate =
    ClaimCopyForUpdatesHelper._create();
// Now update the Claim's attributes. Note that these methods
// execute locally, within the same language.
claimCopyAlternate.state(entered);
claimCopyAlternate.reason(accident);
claimCopyAlternate.description(
    "Side-swiped by a little old lady in an Oldsmobile.");
byte[] claimString = claimCopyAlternate._toString();
theClaim.updateUsingCopyString(claimString );

```

the implementation of the business object method would be:

```

public void updateUsingCopyString(byte[] objString)
{
//Version identifier DCE:D3E893DD-6363-11d2-A40B-08005A49CF68:1
// Insert Method modifications here
ClaimCopyForUpdates copy = ClaimCopyForUpdatesHelper._create();
copy.fromString(objString);
IDataObject.state(copy.state());
IDataObject.reason(copy.reason());
IDataObject.description(copy.description());
// End Method modifications here
}

```

Similarly, the client interface for reading a number of attributes would provide reading of all attributes interesting to the caller. Note that the *claimNo* is used to find the business object. For efficiency, this information is not provided in the specific copy helper for reading attributes. Another alternative would be to just reuse the copy helper for reading that is also used for creation.

```

// Now use the Claim's attributes. Note that these methods
// execute locally, within the same language.
java.lang.String date = claimCopyAlternate.date();
State state = claimCopyAlternate.state();
long reason = claimCopyAlternate.reason();
java.lang.String description = claimCopyAlternate.description();

```

the implementation of the business object method would be:

```

public byte[] readUsingCopyHelper()
{
//Version identifier DCE:D3E893DB-6363-11d2-A40B-08005A49CF68:1
// Insert Method modifications here
ClaimCopyForReads copy = ClaimCopyForReadsHelper._create();
copy.date(iDataObject.date());
copy.state(iDataObject.state());
copy.reason(iDataObject.reason());
copy.description(iDataObject.description());
return copy._toString();
// End Method modifications here
}

```

For completeness, this example of a copy helper shows different copy helpers for both reading and updating because the business object contains *readonly* attributes that are not part of the primary key. If the business object did not contain these other *readonly* attributes then the same copy helper could be used for both reading and updating. See the samples that are provided with the product for an example of using copy helpers in this manner.

Copy helpers – sharing opportunities

Examination of the examples presented in this book might raise some questions about redundancy and duplication of interface and implementation. When the Policy example was originally described, there were separate classes for the Copy Helper Class (PolicyCopy) and for the actual interface that clients generally code to for business logic purposes (Policy). These classes look alike, yet have no inheritance or other meaningful relationship between them. They are essentially in separate class hierarchies partially because the CopyHelper classes are ILocalOnly types of classes while the real business objects are true CORBA objects.

While sharing implementations might not make sense, it is possible to share interfaces between the actual business object interface and the CopyHelper class. In the example, a new abstraction called PolicyCommon was introduced. This is an abstract class that introduces that which is common between the Policy Copy Helper and the Policy interface used for building the business object. Notice that Policy remains an IManageable, PolicyCopy remains an INonManageable, and PolicyCommon has no parent.

Although this example might not appear to simplify and reduce complexity, more realistic cases that include more methods and attributes might make this reasonable.

It is mostly attributes that go into the *xxxxBase* class, although methods could also be included. All attributes that are part of the real business object that are meaningful as part of a transient copy are candidates. If there are methods that apply to both the real business object and a transient copy, then those methods can go into this *xxxxBase* class. Methods that are used to validate and edit the attributes are the best candidates for inclusion into the *xxxxBase* class.

Figure 32 on page 76 shows an example of how this might be done.

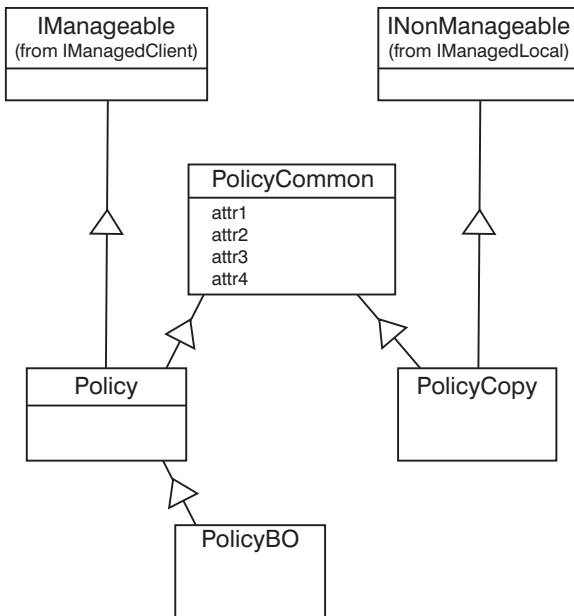


Figure 32. Example of interface sharing opportunities

Copy helpers and complex data types

The examples discussed so far have been done showing just simple data types such as *long* and *string* and are applicable to almost all the CORBA data types. There are several more complex data types that also need to be discussed when streaming those attributes for copy helpers. For example, if we defined an interface with the following data types.

```

interface ComplexPolicy : IManagedClient::IManageable
{
    enum Status {OPEN, PENDING, CLOSED};
    typedef sequence<float,5> OldPremiums;
    typedef Insurance::ComplexPolicy varArray[2][3];

    attribute float amount;
    readonly attribute long policyNo;
    attribute short policyID;
    attribute unsigned short lifetime;
    attribute unsigned long renewals;
    attribute double value;
    attribute char type;
    attribute wchar wtype;
    attribute string comments;
    attribute wstring wcomments;
    attribute boolean delinquent;
}
  
```

```

        attribute octet obyte;
        attribute Insurance::ComplexPolicy::Status currentStatus;
        attribute Insurance::ComplexPolicy::OldPremiums premiumHistory;
        attribute Insurance::ComplexPolicy::varArray z;
    };
    // end interface ComplexPolicy

```

The streaming methods of most of these data types is fairly straight forward using the Externalization Service interfaces that map directly to the data type. Other types must map to other Externalization interfaces.

```

public void externalize_to_stream(
    org.omg.CosStream.StreamIO targetStreamIO)
{
    // Insert Method modifications here
    targetStreamIO.write_float(iAmount);
    targetStreamIO.write_long(iPolicyNo);
    targetStreamIO.write_short(iPolicyID);
    targetStreamIO.write_unsigned_short(iLifetime);
    targetStreamIO.write_unsigned_long(iRenewals);
    targetStreamIO.write_double(iValue);
    targetStreamIO.write_char(iType);
    targetStreamIO.write_wchar(iWtype);
    targetStreamIO.write_string(iComments);
    targetStreamIO.write_wstring(iWcomments);
    targetStreamIO.write_boolean(iDelinquent);
    targetStreamIO.write_octet(iObyte);
    targetStreamIO.write_long(iCurrentStatus);
    com.ibm.IExtendedStream.StreamIO
    extendedStream=
        com.ibm.IExtendedStream.StreamIOHelper.narrow(targetStreamIO);
    {
        org.omg.CORBA.Any any;
        any <<= iPremiumHistory;
        extendedStream.write_any(any);
    }
    {
        org.omg.CORBA.Any any;
        any <<= ComplexPolicy.varArray_forany(iZ);
        extendedStream.write_any(any);
    }
}

public void internalize_from_stream(
    org.omg.CosStream.StreamIO sourceStreamIO,
    org.omg.CosLifeCycle.FactoryFinder there)
throws org.omg.CosLifeCycle.NoFactory,
org.omg.CosStream.ObjectCreationError,
org.omg.CosStream.StreamDataFormatError
{
    // Insert Method modifications here
    iAmount = sourceStreamIO.read_float();
    iPolicyNo = sourceStreamIO.read_long();
    iPolicyID = sourceStreamIO.read_short();
    iLifetime = sourceStreamIO.read_unsigned_short();
    iRenewals = sourceStreamIO.read_unsigned_long();
}

```

```

iValue = sourceStreamIO.read_double();
iType = sourceStreamIO.read_char();
iWtype = sourceStreamIO.read_wchar();
iComments = sourceStreamIO.read_string();
iWcomments = sourceStreamIO.read_wstring();
iDelinquent = sourceStreamIO.read_boolean();
iObyte = sourceStreamIO.read_octet();
iCurrentStatus = sourceStreamIO.read_long();
com.ibm.ExtendedStream.StreamIO
extendedStream=COM.IBM.IExtendedStream.StreamIOHelper.narrow(
    sourceStreamIO);
{
    org.omg.CORBA.Any any=extendedStream.read_any();
    ComplexPolicy.OldPremiums localCopy;
    any <<= localCopy;
    iPremiumHistory = localCopy;
}
{
    org.omg.CORBA.Any any=extendedStream.read_any();
    :ComplexPolicy.varArray_forany localCopy;
    any >>= localCopy;
    iZ = localCopy;
}
// End Method modifications here
}

```

The data types that are not mapped directly in this example use two different techniques. One simple technique is the mapping of *enum* data types to *long* which is feasible since this fits within the definition of this data types. Another technique involves using the *any* type support. This is a special data type that can store any data type. The receiver of this data from the stream knows what type of data was originally stored in that type. In this example we use *any* support for a *sequence* and *variable length array*.

Moving object data in bulk

One of the important design points in any distributed system is to “reduce trips over the wire”. This basic guideline will cause a number of low level design and implementations decisions to be made in object systems. There is always the risk of exposing too much data and violating encapsulation. However, performance often times takes precedence over design purity.

During the life-cycle of an object there are a number of points where the “reduce trips over the wire” guidelines should be considered. These will be explained in the following paragraphs.

Object creation is the first opportunity to reduce trips over the wire. The `createFromCopyString` method on all Component Broker homes is the simple answer to creating objects with multiple attributes from remote clients in an

efficient way. If a business object has five attributes, of which only one is the key, then it is obviously more efficient to use `createFromCopyString` as opposed to a `createFromPrimaryKeyString` followed by four setter methods on the newly-created business object. Specialized homes can also offer create methods that gather up as much from the client as is reasonable and take it over the wire on one method call to the specialized home.

Once created, the ability to move bulk data for objects becomes more complicated. Some of the possible patterns are described.

For situations where a subset of the state is needed to prime user interfaces a data array query (`evaluate_to_data_array` method on a query evaluator) should be considered. This returns tuples, not objects that represent the data for the objects. If usage of the actual objects is also expected to occur, then the object reference should be requested as a return value from the `evaluate_to_data_array` call. This reduces trips over the wire. However, there is another opportunity. Even when using a data array query, there is no way to update multiple object attributes in one call. The `evaluate_to_data_array` got lots of data, but there is no obvious way to use it to update the objects they represent.

In our example of a business object with five attributes, an `update(a1,a2,a3,a4,a5)` method on the business object could perform this 'update all attributes' function. The implementation of this method in the business object would probably validate that the key is the same and update all non-key attributes. Remember, keys cannot change on components once they have been created. A variation on update is to use an update of the form `update(copyHelperString)`. This does not change the trips over the wire, but offers a different style for client programmers to use.

As with any copies or caches, be very precise about when the copy is update to date and when it might be considered 'dirty'. Responsibility for keeping the business object synchronized with the copy falls to the business object provider. Additional logic may be appropriate in the `update()` method to ensure data integrity and to prevent unnecessary updates to the third-tier resource managers.

Treating object data as structures or buffers should be done as a performance optimization. It should not be a general practice when doing object-oriented programming due to the reduction in encapsulation that is introduced. This is a performance tuning and optimization technique.

Multiple interfaces to business objects

There are situations where a business object has different sets of clients. Good object-oriented analysis and design generally leads developers to interfaces that start with a base or minimal set of capabilities and then introduces additional interfaces that have an increasing set of methods. Be careful defining these interfaces and when handing them out to clients to avoid compromising the encapsulation that is necessary for a robust object design.

Figure 33 is an example of factoring the interface to be used by clients.




Figure 33. Example of interface factory

The previous example factors the Policy interface into PolicyForEveryone and PolicyForFriendsOnly. For clients that need access only to the PolicyForEveryone interface, the client side DLL (in the C++ case) should include only bindings for the PolicyForEveryone interface. Clients that want to access all of the methods in both the PolicyForEveryone interface and PolicyForFriendsOnly would need a client side DLL (in the C++ case) that contained C++ bindings for both of these interfaces.

Other factorings are possible.

Transactions

Transactions are like a contract that binds a client to one or more servers. They bracket a set of system behaviors or actions that have the following properties:

Atomicity

A group of actions behave in an all-or-none fashion as an indivisible unit of work.

Consistency

After a transaction is implemented, the system is left in a correct state, assuming that the system was in a correct state prior to the transaction.

Isolation

A given transaction's behavior is unaffected by other transactions and system activity occurring concurrently.

Durability

When a transaction commits, its effects are permanent or persistent. Because transactional models involve the concept of bracketing system behaviors, they provide commands for indicating transaction boundaries. All transactions have a start and an end that can involve committing to a change in system state or aborting the changes and rolling back to the pre-transaction state of the system.

The Component Broker Transaction Service provides standard primitives for transactional applications in a distributed object environment. When the Transaction Service is used in conjunction with a persistent storage medium and with some control over concurrent access to resources, you have the basics for creating robust business applications. Key elements of the Transaction Service are:

- Basic Transaction Service operations.
- Integration with the server run time.
- Client support.

Note: For additional information about the Transaction Service, see the *WebSphere Application Server Enterprise Edition Component Broker Advanced Programming Guide*.

CosTransactions module

The Component Broker CosTransactions module provides the transaction interfaces defined by OMG. Some of the operations supported by CosTransactions include:

- Controlling the scope and duration of a transaction.
- Allowing multiple objects to be involved in a single, atomic transaction.

- Coordinating the completion of transactions.
- Performing recovery of transaction states following restart of failed processes.

The primary interfaces in the CosTransactions module are:

Current Interface

The current interface defines methods that allow a client of the Transaction Service to explicitly manage the association between threads and transactions. It also defines methods that simplify the use of the Transaction Service for most applications. These methods can be used to begin and end transactions, and to obtain information about the current transaction (see “Transactions” on page 133).

TransactionFactory Interface

The TransactionFactory interface defines a single create() method that allows the transaction originator to create a new top-level transaction without it being associated with the current thread. It is primarily intended for creating transactions remotely.

Control Interface

The control interface defines methods to allow a program to explicitly manage or propagate a transaction. An object supporting the control interface is implicitly associated with one transaction only.

Terminator Interface

The terminator interface defines methods to complete a transaction, either by requesting commitment or demanding rollback. Typically, these methods are used by the transaction originator. An object that supports the terminator interface is implicitly associated with one transaction only.

Coordinator Interface

The coordinator interface provides common methods for top-level transactions and subtransactions. Participants in a transaction are typically either recoverable objects or agents of recoverable objects, such as subordinate coordinators. An object supporting the coordinator interface is implicitly associated with one transaction only.

RecoveryCoordinator Interface

The RecoveryCoordinator interface allows recoverable objects to drive the recovery process in certain situations. Each instance of this class is implicitly associated with a single resource registration, and can only be used by that resource in the particular transaction for which it is registered.

Resource Interface

The resource interface defines the operations invoked by the

Transaction Service, during transaction completion, on each resource registered using the `register_resource()` method of the coordinator interface.

SubtransactionAwareResource Interface

The `SubtransactionAwareResource` interface defines the operations invoked by the Transaction Service, during subtransaction completion, on each resource registered using the `register_subtran_aware()` method (or `register_resource()`) of the coordinator interface.

Synchronization Interface

The synchronization interface defines the operations invoked by the Transaction Service during transaction completion on each resource registered using the `register_synchronization()` method of the coordinator interface.

TransactionalObject Interface

The `TransactionalObject` interface is used by an object to indicate that it is transactional. By inheriting from the `TransactionalObject` interface, an object indicates that it wants the transaction context associated with the client thread to be propagated on requests to the object.

These interfaces show the breadth of function available within the `CosTransactions` module. The MOFW encapsulates the transactional interfaces and provides a programming interface to transactions. Most of what you need to know is shown in “Transactions” on page 133.

Character set considerations

To allow Component Broker to use multiple character sets, you must use only the Portable Character Set supported by the seven-bit ASCII code set when developing client and server code.

The Portable Character Set supported by Component Broker follows:

```
0 1 2 3 4 5 6 7 8 9
: ; < = > ? @ [ \ ] ^ _ ' ` { | } ! " # $ % & ( ) * + , - . / <space>
a b c d e f g h i j k l m n o p q r s t u v w x y z
A B C D E F G H I J K L M N O P Q R S T U V W X Y Z
```

Chapter 4. MOFW C++ client programming model

This chapter defines the Component Broker managed object framework (MOFW) client programming model. The word *client* generally refers to a particular computer in a distributed environment and the relationship of that computer with other computers. In this context, the word *client* refers to any program and its relationship to a component. In other words, a client is any program executing in a process on a client or server computer.

Therefore, the client programming model explains how you can use components and how to develop objects that are clients of Component Broker components. Application programmers developing *tier-1* (client) or *tier-2* (server) applications use the client programming model when they use a component in the implementation of a new object. Client programmers use server objects to develop new applications by composing new application objects and scripting behavior over sets of server objects.

A Component Broker server process implements objects that are used by client applications to perform business functions. All Component Broker server objects are derived from the managed object framework. These objects effectively make up a component and are accessed through the part of the component called the managed object. The term *managed object* is used because one of the Component Broker server product's main contributions to object-oriented applications is the implementation of run-time management functions used in the implementation of server objects. Some examples of these management services are:

- Persistence, Transactions, and Security.
- Workload management and availability management over multiple servers.
- Object-oriented access to existing databases and applications.

This chapter provides details on a client view of the structure of a Component Broker client/server application.

- “Client view of Component Broker applications” on page 60 contains details on how a client object uses a Component Broker component and explains what occurs if this component is implemented in a different language on a different system.
- “Client programming model: basic tasks” on page 86 contains a quick overview of the client programming model. This section explains what you need to know to get started developing a Component Broker client application. Specifically, it explains how to accomplish common programming tasks, and it contains sample code segments.

- “Summary: The client programmer’s check list” on page 101 contains a check list of information a client application programmer needs to start developing the client side of a Component Broker application.

Component Broker supports writing client applications in C++, Visual Basic, or Java. This chapter includes information on writing client applications in C++. For information on writing client applications in Visual Basic, see “Chapter 8. MOFW – ActiveX client programming model” on page 211. For information on writing client applications in Java, see “Chapter 9. MOFW - Java client programming model” on page 231 and “Chapter 11. MOFW - Java client programming model - advanced concepts” on page 281.

Client programming model: basic tasks

A client application can perform the following tasks:

- Find an object.
- Use an object.
- Create an object.
- Use a set of objects.
- Remember an object.
- Release or delete an object.

The following sections present a quick overview and samples of how a client application performs these tasks. However, before a client application can perform any of these tasks, it must initialize the Component Broker client environment.

Initializing the client environment

The Component Broker provides a convenient interface for initializing a Component Broker client. This interface is `CBSeriesGlobal`.

Initializing a Component Broker client requires only the following single line of code:

```
CBSeriesGlobal::Initialize();
```

After initialization, the client has access to the static methods `CBSeriesGlobal::orb()` and `CBSeriesGlobal::nameService()` that return references to objects of type `CORBA::ORB` and `IExtendedNaming::NamingContext`, respectively.

The `CBSeriesGlobal` interface is provided as a convenience to client programmers. A Component Broker client could define a set of CORBA calls that encapsulate the `Initialize()` method as follows:


```

interface CBSeriesGlobal
{
    void Initialize();

    CORBA::ORB orb();
    IExtendedNaming::NamingContext nameService();
};

```

This IDL is directly implemented in each client programming language (C++ and Java). The C++ interface is:

```

class CBSeriesGlobal
{
    public:
        static void Initialize();

        static CORBA::ORB_ptr orb ();
        IExtendedNaming::NamingContext_ptr nameService ();
};

```

Finding a managed object

When a client application starts, it has no knowledge of any objects. In a non-distributed environment, the application would create any objects it needs. These newly created objects would reside on the same system as the application. If the application needs these newly created objects to be available the next time the application runs, the application itself is responsible for making these objects persistent.

However, Component Broker is a distributed environment. In this environment, objects are not created on a client system. Objects are created from a client system on a server system. Object persistence is managed with the help of the Component Broker run time.

The client accesses the component through the components managed object. But first, it has to find the managed object.

There are two ways for a client application to find a managed object:

- Use the Naming Service.

If the object has a *Name*, the client can use the Naming Service to locate the object by its *Name*. In general, the Naming Service only contains a subset of the objects in a distributed system. This subset consists of well-known objects, such as collections of business objects or important objects in the object model.

- Use another object.

If the object does not have a *Name*, the client can find the object by using the Naming Service to find a well-known object, such as a *factory* or

collection, and then use this object to implement methods that return the object as a return value or output parameter.—

Both techniques involve the use of the Naming Service. In the Component Broker distributed environment, all work with objects in a client application must begin by using the Naming Service.

Finding a managed object using the Naming Service

Assume that the insurance company application example contains a few important Claim components in the Naming Service and that the client application is written with the knowledge that these Claim components can be located at a known path in the name space. The following code segment shows how to find the Claim component, belonging to a customer named "Lou".

```
{
  CORBA::Object_var temp;
  Claim_var louClaim;
  try
  {
    temp = CBSeriesGlobal::nameService()->resolve_with_string(
      "cell/Applications/LifeInsurance/Claim/LouClaim");
    louClaim = Claim::_narrow(temp);
    // use louClaim_ptr in various calls
  }
  catch(...)
  {
    // appropriate error recovery
  }
}
```

The path "cell/Applications/LifeInsurance/Claim/LouClaim" follows a path that traverses a NamingContext bound into the Cell Root NamingContext with a name of Applications and so forth to the desired LouClaim component.

The following code segment shows how the insurance application could have originally bound the LouClaim component.

```
Claim_var louClaim;

// Create Lou's Claim and initialize it prior to the following code segment
...
// Try to bind the 'louClaim' object into the name space.

try
{
  CBSeriesGlobal::nameService()->bind_with_string(
    "cell/Applications/LifeInsurance/Claim/LouClaim", louClaim);
}
```

```

catch(...)
{
    // appropriate error recovery
}

```

For additional information on the Component Broker Naming Service, see References in Component Broker for Windows NT and AIX Online Documentation.

Finding a managed object using the primary key

What if Lou's Claim is not in the Naming Service? How do you find a specific claim? This section explains how.

Homes are instances of the IHome class. You might decide to implement and provide a tailored subclass of IHome, or you might use an instance of the base class. The relationship between managed objects and collections is explained in "Using sets of objects" on page 95. For now, all that is important to know is that a home can be used to find objects that were previously created by that home.

As stated previously, a client application always starts with the Naming Service to find its first object. As such, it might seem reasonable to assume that a client application would find a home directly in the name space. However, the LifeCycle Service provides a facility for finding homes called FactoryFinder. Therefore, a client application uses the Naming Service to find a FactoryFinder, which is then used to find a home. FactoryFinders are discussed again in more detail in "Creating a new object – create from key" on page 92.

To continue the example, the following code segment finds the home for Claim objects.

```

CORBA::Object_var obj;
IExtendedLifeCycle::FactoryFinder_var ff;
IManagedClient::IHome_var claimHome;
try
{
    obj = CBSeriesGlobal::nameService()->resolve_with_string(
        "host/resources/factory-finders/host-scope");
    ff = IExtendedLifeCycle::FactoryFinder::_narrow(obj);
    obj = ff->find_factory_from_string("Claim.object interface");
    claimHome = IManagedClient::IHome::_narrow(obj);
    // use claimHome
}
catch(...) {
    // appropriate error recovery
}

```

Now you need to find Lou's Claim. Assume that Lou is on the phone and can tell you his claim number. In the example, the Claim's *claimNo* attribute uniquely identifies an instance of the Claim component, and is a primary key into the Claim's home. To facilitate the usage of a generic home, the Component Broker programming model introduces two types of helper class: primary key and copy helper. The copy helper is discussed further in "Creating a new object – create from copy" on page 94. Every component has a primary key that lets you find (and create) components of that type. An instance of a key is always local to the client's process and language. Keys, like all helper classes, are created with a static method on the class named `_create()`. This static method gets generated in the usage bindings of all interfaces that specify the *localonly* pragma in their IDL files. The same rule is in place for copy helpers and objects of other classes that are described in "Local only development process" on page 127.

When you create an instance of a primary key, the key must be set by one or more attributes on the primary key object. When all of the key attributes have been set, the primary key object is now usable. The Claim's home uses this primary key to find the previously created Claim component. Remember, the primary key is on the client system, but the Claim object and the Claim home are on the server system. If the client passes a primary key object as a parameter to the home, and the home is on a remote system, the remote system might get a proxy back to the client's primary key instance. This would turn the client into a server and unpredictable results could occur. Therefore, the Component Broker programming model uses strings as the method for passing keys to potentially remote objects.

Continuing the example, the following code segment would find Lou's Claim in the home (assuming Lou told you his number is 1234).

```
// Create an instance of the Key Helper Class
ClaimPrimaryKey_var claimPrimaryKey = ClaimPrimaryKey::_create();

// Set the claimNo attribute in the key
claimPrimaryKey->claimNo(1234);

// Must convert key object to a string to go over the wire to the server
ByteString_var claimString = claimPrimaryKey->toString();

IManagedClient::IManageable_var temp;

// Call find by key on the Home to find Lou's Claim
temp_var = claimHome->findByPrimaryKeyString(claimString);

// Narrow type to Claim
Claim_var louClaim = Claim::_narrow(temp);

// continue to use louClaim
```

The object provider of a public component always provides you with a set of helper classes for using the Homes that contain the component managed objects. There is always exactly one primary key helper class for each component. The object provider gives you:

- The interface definitions for the primary key class.
- An implementation of the primary key class.
- Documentation for its use.

The previous code segment creates a new instance of the key class `ClaimPrimaryKey` using the static `_create()` method, and sets its `claimNo` attribute to 1234. Then it creates a string version of the key and finds Lou's Claim managed object using the Claim primary key and the Claim home.

The Component Broker programming model mandates keys for components to make things easier for you. The keys are passed as strings. You need to use the available `setxxxx()` methods on the helper class to prepare the key information. No string manipulation to concatenate pieces of multi-valued keys is necessary. Using a key enables type errors to be detected at compilation time, and provides knowledge of the field ordering and algorithm for defining a key.

Finding a managed object using another managed object

Some components have interfaces that include methods or attributes that return other components. If this is the case, then when you have an object, you can use its methods to find related objects. Continuing the previous example, after finding Lou's Claim, you can find other objects that the Claim references. The following example returns a reference to Lou's Policy:

```
// Find Lou's Policy

Policy_var louPolicy; // declare a local variable
louPolicy = louClaim->policy();
```

Using a managed object

When you find a reference to a component's managed object, you can invoke methods on it. For example,

```
person->name("Lou Smith");
```

calls the `name()` method on the person component identified by the person object reference to set the value of the `name` attribute. The Component Broker server handles the use of remote objects in a way that is transparent to you.

Creating a managed object

Component Broker managed objects can be created in a number of ways. The following sections describe the default ways of creating managed objects.

Creating a new object – create from key

Every Component Broker managed object class has an instance of a factory associated with it. The factory provides a set of interfaces for creating instances of a managed object. The factory gets some of its interface from the base class `CosLifeCycle::GenericFactory`. The method used in `createFromPrimaryKeyString` is introduced in the `IManagedClient::IHome` interface supplied by Component Broker. This interface specializes the `CosLifeCycle::GenericFactory` interface and plays the role of factory for Component Broker managed objects. Object providers can implement and provide a tailored subclass of this interface, or they can use the implementation of `IHome` provided.

All you need to know is how to find the right `IHome` for creation. This is done using a factory finder. The input required for the factory finder is the name of the interface of the component that you want this factory to make instances of. The following code segment gets a reference to the Claim factory for the Life Insurance application.

```
CORBA::Object_var obj;
IManagedClient::IHome_var claimHome;
IExtendedLifeCycle::FactoryFinder_var myFinder;

obj = CBSeriesGlobal::nameService()->resolve_with_string(
    "host/resources/factory-finders/host-scope");
myFinder = IExtendedLifeCycle::FactoryFinder::_narrow(obj);

obj = myFinder->find_factory_from_string("Claim.object interface")
claimHome = IManagedClient::IHome::_narrow(obj);
```

Notice the usage of both the `IExtendedNaming::NamingContext` interface, the `resolve_with_string()` method and the `IExtendedLifeCycle::FactoryFinder` interface, the `find_factory_from_string()` method. These interfaces are extensions to the corresponding CORBA interfaces, which should be easier to use. The `IExtendedNaming::NamingContext` interface is described in “Navigating the name space using the Naming Service” on page 70 in the The managed object framework chapter. The `IExtendedLifeCycle::FactoryFinder` interface introduces the previous methods beyond the `CosLifeCycle::FactoryFinder` interface, providing simpler and more flexible ways to find a factory:

```
module IExtendedLifeCycle
{
    interface FactoryFinder : CosLifeCycle::FactoryFinder
    {
```

```

CosLifeCycle::Factory find_factory (in
    CosLifeCycle::Key factory_key)
    raises (CosLifeCycle::NoFactory);

CosLifeCycle::Factory find_factory_from_string (in
    FactoryKeyString factory_key)
    raises (CosLifeCycle::NoFactory,
        NamingStringSyntax::IllegalStringSyntax,
        NamingStringSyntax::UnMatchedQuote);

CosLifeCycle::Factories find_factories_from_string (in
    FactoryKeyString factory_key)
    raises (CosLifeCycle::NoFactory,
        NamingStringSyntax::IllegalStringSyntax,
        NamingStringSyntax::UnMatchedQuote);
};
};
};

```

Having found an IHome by a FactoryFinder, you must provide the IHome with information necessary to manufacture a new object instance. At a minimum, the primary key must be provided. An example of creating a new Claim with a *claimNo* of 1234 is shown in the following example:

```

// Create an instance of the Primary Key Class
ClaimPrimaryKey_var claimKey = ClaimPrimaryKey::_create();

// Set the claimNo attribute in the key
claimKey->claimNo(1234);

// Call createFromKey on the Factory to create Lou's Claim
IManagedClient::IManageable_var theMO;
Claim_var theClaim;

ByteString_var claimString = claimPrimaryKey->toString();
theMO = claimHome->createFromPrimaryKeyString(claimString);
theClaim = Claim::_narrow(theMO);

```

This example is almost identical to the example in “Finding a managed object using the primary key” on page 89. First, create a primary key object to define the identity of the component that will be made. Then, call `createFromPrimaryKeyString` on the home to pass the key as a string.

The `createFromPrimaryKeyString` method is defined by the IHome class, and all components can be created by this method. An object provider provides you with a subclass that introduces other, easier to use creation methods. Additional object creation techniques are described, with examples, in “Chapter 6. MOFW - C++ client programming model – advanced concepts” on page 133.

Creating a new object – create from copy

Setting and getting the attributes of a component (through its managed object) can be expensive. There are several reasons for this. First, if the managed object is implemented in another language, each get or set method is actually a cross-language call. Cross-language calls are more expensive than simple, same-language calls. The get and set overhead is even worse if the managed object is remote because each call is actually a remote procedure call and involves significant overhead. Consider the following code fragment:

```
// Let's assume that 'theClaim' is declared and created as in the
// previous code snippet.

// Creating 'theClaim' above required one RPC. Setting the rest of the
// object's attributes involves one RPC per attribute. The following
// lines of code show four such RPCs. This could, of course, be any
// number of RPCs, depending on the complexity of the object.

theClaim->date("10/14/96");
theClaim->state(entered);
theClaim->reason(accident);
theClaim->description("Side-swiped by teenager in a red convertible.");
```

This fragment could involve the following remote method calls:

- The client to the home of Claims to create the Claim.
- The client to the Claim managed object to set its *date* attribute.
- The client to the Claim managed object to set its *state* attribute.
- The client to the Claim managed object to set its *reason* attribute.
- The client to the Claim managed object to set its *description* attribute.

These calls could be reduced to a single remote method call by using the `createFromCopyString()` method on an `IHome` instead of `createFromPrimaryKeyString()`. In order for you to use the `createFromCopyString()` method, the object provider must provide you with a copy helper class. The following code segment rewrites the previous example using this design pattern.

```
// Create a new "local" Claim in my process and language.
// Use a Copy Helper Class that the Claim MO provider gave me.
ClaimCopy_var claimCopy = ClaimCopy::_create();

// Now initialize the Claim's attributes. Note that these methods
// execute locally, within the same language.
claimCopy->claimNo(1234);
claimCopy->date("1-14/96");
claimCopy->state(entered);
claimCopy->reason(accident);
claimCopy->description("Side-swiped by teenager in a red convertible.");

// Pass this local copy to the Home and have it return a new Claim MO
// whose attributes are initialized from the local copy's values. Because
// not all ORBs support Pass-By-Value, we first convert the local copy
```



```

// helper object to a string.
IManagedClient::IManageable_var theMO;
Claim_var theClaim;

ByteString_var claimString = claimCopy->toString();
theMO = claimHome->createFromCopyString(claimString );
theClaim = Claim::_narrow(theMO);

```

Like a key class, a copy helper class instance is always local to your process and implemented in the language you are using. The helper class interface and implementation is given to you by the object provider.

Copy helpers are especially useful if the client application needs to interact with an object during initialization, and then create a managed object from the attributes. A common scenario for this is entering data for the object from a GUI. The GUI updates the local copy helper, and then the `createFromCopyString()` method is called when the **Do** push button is pressed on the end user interface (EUI) to do the action.

Note: All of the attributes that make up the primary key must be set in the copy helper instance before using it for creation.

Using sets of objects

An `IHome` represents a set of component managed objects, all of the same type, whose relationship to one another is defined by the object provider, and maintained by the fact that they were all created in the same home. Sometimes an application needs to define (and maintain) the relationships between managed objects, based on the particular business task at hand. This might even include relationships between managed objects of different types (for example, `Policy Holder` and `Beneficiary`). This can be done using an `IManagedCollections::IReferenceCollection`, as shown in the following code segments.

First, create a reference collection. The following code shows how to do this from a client. Notice that because a reference collection is itself a managed object, it is created using a home. Component Broker provides a specialized home which makes it easy to create a reference collection. Specialized homes are described in greater detail in “Creating specialized homes” on page 192.

```

CORBA::Object_var obj;
IManagedCollections::ICollectionHome_var rcHome;
IManagedCollections::ICollection_var cc;
IManagedCollections::IReferenceCollection_var rc;

obj = myFinder->find_factory_from_string(
    "IManagedCollections::IReferenceCollection.object
    interface/PersistentReferenceCollectionFactory.object home");

```

```

rcHome = IManagedCollections::ICollectionHome::_narrow(obj);
cc = rcHome->createCollection();
rc = IManagedCollections::IReferenceCollection::_narrow(cc);

```

Note: The collection created in the previous example is capable of containing objects of any `IManagedClient::IManageable` subclass. If the collection is to contain elements of a specific type (for example, `PolicyHolder`), and you would like the collection to enforce this, you can use the `createCollectionFor` call instead of `createCollection`.

Having created the reference collection, it is now ready to be used. The following code segment shows how a client application adds a `PolicyHolder` object to the reference collection:

```

PolicyHolder_var policyHolder;
// ...
// Find or create the PolicyHolder object to add to the collection.
// ...
try
{
    rc->addElement(policyHolder);
}
catch (ICollectionsBase::IInvalidElement &ex)
{
    cout << "ERROR: Caught exception: " << ex.id() << endl;
    cout << "ERROR: Trying to add object to a reference collection";
}

```

Having possibly added many interesting components to the reference collection, it seems reasonable that at a later time it might be necessary to iterate over every managed object in the reference collection to perform some action. The following code shows one way of doing this:

```

IManagedCollections::IReferenceCollection_var theMixedCollection;
IManagedCollections::IIterator_var theIterator;
IManagedClient::IManageable_var theBO;

// Create an iterator on the reference collection that was found above.
// Note that when an iterator is created, it is automatically positioned
// preceding the first element.

theIterator = theMixedCollection->createIterator();
try
{
    // Loop through the collection. The "nextElement" method advances
    // the iterator to the next element (on the first invocation, this
    // will advance to the first element) and then return the element
    // pointed to by the iterator.

    while ( theIterator->more() )
    {
        theBO = theIterator->next();
        if ( theBO->is_a("PolicyHolder") )

```

```

        // Send him a bill
        if ( theB0->is_a("Beneficiary") )
            // Send him a check
    }
    catch (...)
    {
        cerr << "ERROR: Problem occurred using iterator." << endl;
    }

    // After iterating over the entire collection, the iterator is no
    // longer needed. As such, it must be removed.

    theIterator->remove();

```

The combination of the `IManagedCollections::IReferenceCollection` and the `IManagedCollections::IIterator` are used in the above code segment. An `IManagedCollections::IReferenceCollection` is a generalized collection of object references that can be iterated. `IManagedCollections::IIterator` supports advancement of the iterator and retrieval of elements by the `next()` method. `IManagedCollections::IReferenceCollection` supports adding and removing elements by `addElement()` and `removeElement()`. `IManagedCollections::IManagedReferenceCollection` is the most basic kind of collection supported in Component Broker. Combining this with the capabilities of `IHome` provides the basis for writing simple applications and the foundations for the more advanced query and collections capabilities provided by Component Broker. For more information on collections and query capabilities, see “Chapter 6. MOFW - C++ client programming model – advanced concepts” on page 133.

The `try/catch` block in the previous code segment is created after the iterator is created. The purpose of the `try/catch` block is to ensure, even if an exception is thrown during iteration, that the iterator that was created is removed when it is no longer needed. Without this `try/catch` block, if an exception is thrown during iteration, the `remove()` method is not invoked and the storage associated with the iterator on the server is not deallocated.

Note: The while loop in the example of iterating through the collection is shown as follows:

```

    while ( theIterator->more() )
    {
        theB0 = theIterator->next();
        ...
    }

```

Although this works, it is inefficient because it requires two calls to the server for each element that is retrieved. It would be better to use a more efficient approach when possible:

```

    while ( theB0 = theIterator->next() ) ...

```

Transient sets

Component Broker supports transient collections. Transient collections do not require transactions; they reside in transient containers and thus have no dependency or overhead on database connections. Database and transient collections provide identical programming interfaces. They differ only in their persistence characteristics. Users can find transient collections by passing an interface string to the factory finder as defined in the next section.

Specifying reference collection interfaces

A variety of interfaces for use with factory finders are provided for specifying the type of reference collections that you want to use.

For example, calling `find_factory_from_string` passing the argument `IManagedCollections::IReferenceCollection.object interface/PersistentReferenceCollectionFactory.object home`, as in the previous example, creates a DB2 backed reference collection. Using the generic string `IManagedCollections::IReferenceCollection.object interface` returns whichever collection is configured as the default.

Transient Collections

```
IManagedCollections::IReferenceCollection.object interface/  
TransientReferenceCollectionFactory.object home
```

Transient Collections Keyed

```
IManagedCollections::IKeyedReferenceCollection.object interface/  
TransientKeyedReferenceCollectionFactory.object home
```

Persistent Collections

```
IManagedCollections::IReferenceCollection.object interface/  
PersistentReferenceCollectionFactory.object home
```

Persistent Collections Keyed

```
IManagedCollections::IKeyedReferenceCollection.object interface/  
PersistentKeyedReferenceCollectionFactory.object home
```

Remembering your favorite objects

Assume you need to remember an interesting or important managed object instance. Ideally, you need to save the reference or pointer to the object. There are several complexities, however.

- The object might be implemented in another language or on a remote system.

- If the client application shuts down and restarts later, the object may have been removed from memory because of a lack of use. If the object is restarted, or reactivated, it probably will not be located in exactly the same place in memory.

Component Broker solves these problems, as does CORBA, by introducing the concept of an *object reference*. An object reference is opaque and you cannot set its internal structure. However, a reference always and uniquely refers to a managed object regardless of where it resides in the network. Continue the example from the previous section, and assume that you run the following code segment:

```
ofstream fout("SOMEFILE.DAT");

// Get a "string" version of my reference to Steve
// steve points to Steve or a proxy to Steve
char* steveStringifiedReference =
    CBSeriesGlobal::orb()->object_to_string(steve)

// Save the string to a file using a "pseudo" file routing
fout << steveStringifiedReference;
fout.close();

// I do not need Steve anymore
steve->release();
```

You saved a reference to Steve as a stringified object reference, and can use this string to re-access Steve at a later time. The following code segment is an example of re-accessing the Steve object.

```
char* infile = "SOMEFILE.DAT";
ifstream fin(infile);

char steveStringifiedReference[1024];
memset(steveStringifiedReference, 1024, '\0');

// Get back the string I saved
fin >> steveStringifiedReference;

// Make an object reference for Steve
CORBA::Object_var obj;
PolicyHolder_var steve;

obj = CBSeriesGlobal::orb()-> string_to_object(
    steveStringifiedReference);
steve = PolicyHolder::_narrow(obj);

// I can now work with Steve.
steve->name("Steven");
```

Coding tips for proper CORBA memory management

The rule for proper CORBA memory management is the following: The caller owns all storage.

The general model on the client is to use `_var` objects. This means that when an `_ptr` is returned, it should be placed into an `_var` by the client. The `_var` assumes responsibility for the storage pointed to the `_ptr` that is placed into the `_var`. The `_var` is a class and its destructor runs when the `_var` goes out of scope.

The other option for clients is to use the `duplicate()` and `release()` methods. The `duplicate()` method is available for making a copy of a proxy, while the `release()` method is used to free the local memory used by a pointer.

None of this should be confused with the LifeCycle Service `remove()` method. This has different semantics. It involves actually deleting the server object that is at the other end of a proxy.

More information on CORBA coding style and conventions, especially as they pertain to memory management, can be found in “Appendix C. CORBA programming” on page 481.

Using object references

Managing the storage of object references is one of the areas where proper memory management is required. You must use the `_duplicate` and `release` operations as described above or use `_var` variables.

There are also special considerations when passing object references as parameters. The caller is always responsible for allocating storage for object references. The caller is also responsible for releasing of all *inout* and returned object references.

For *inout* parameters the caller provides an initial value. If the callee wants to reassign the *inout* parameter, it must first all the `release()` operation on the initial input value. To continue to use an object reference passed as an *inout*, the caller must first duplicate the reference.

See the section titled “Storage Management and `_var` variables” in “Appendix C. CORBA programming” on page 481 for details on memory management.

Summary: The client programmer's check list

When using a set of components to implement an application, you need to know how to use the following things about components and the managed object framework which encapsulates and interoperates with the component objects. Table 2 is organized around the tasks outlined in "Client programming model: basic tasks" on page 86. Note that *xxx* is used to represent the name of the component. For example, Claim and Policy were used in this chapter as examples of domain-specific components. This also provides a good clue as to whether or not the classes were constructed by the object provider or whether they came as part of the Component Broker system.

Table 2. Summary of interfaces

Task	Class Types	Methods	Purpose	
General Use	StandardSyntaxModel	stringToName	Converts strings to CORBA name types.	
Find <i>xxx</i> component	NamingContext	resolve	Finds an object in the namespace	
	IHome	findByPrimaryKeyString	Finds an object in a home	
	<i>xxx</i> PrimaryKey	set <i>xxx</i> methods and toString	Sets the key value	
Create or delete an <i>xxx</i> component	FactoryFinder	findFactory	Finds the factory object.	
	IHome	createFromKeyString	Creates an object with a key only.	
	<i>xxx</i> PrimaryKey	set <i>xxx</i> and toString	Passes information on key to the server.	
	<i>xxx</i> CopyHelper	set <i>xxx</i> and toString	Passes copy information to the server.	
any <i>xxx</i> MO (managed object) subclass		remove and release	Deletes or releases object from memory.	
	Use sets of components	IReferenceCollection	createIterator	Creates an iterator object.
	IIterator	next	Iterates over a collection and gets the objects.	
	Remember interesting and important components	NamingContext	bind	Binds an object into the name space with a name.
resolve			Finds an object in the name space by its name.	
theORB		object_to_string	Creates a stringified form of an object reference.	
		string_to_object	Creates an object reference from a string.	

This is the basic set of methods needed to work with Component Broker components. More advanced methods for a C++ client are discussed in “Chapter 6. MOFW - C++ client programming model – advanced concepts” on page 133.

Chapter 5. MOFW - C++ server programming model

This chapter describes the interfaces and processes you follow to create a component based on the managed object framework (MOFW) interfaces. The first section introduces these abstractions and the relationship that they have to the domain-specific interfaces and implementations that make up the business logic and business data inherent in a component.

After this brief overview, the basic steps involved in developing a component are described.

Developing a component

The minimal set of required activities for developing a component is described in this section. The real components that you deploy will probably leverage some of the additional interfaces and capabilities of Component Broker described in subsequent sections. This section helps you get started and understand the basics involved in constructing a component. This section describes more details about the abstractions of the MOFW, how they are implemented, and what options exist when constructing component based on the MOFW.

The minimal set of steps is:

1. Develop an interface to the components.
2. Choose a pattern for handling essential state.
3. Implement the component methods, in a business object..
4. Implement the methods required by the MOFW interfaces.
5. Implement the necessary primary key class.
6. Implement the optional copy helper class. Implementation of a copy helper class is optional, but strongly recommended. Copy helper objects can greatly improve the performance of creating a component in a server from a client application.

At the end of these tasks, the component should be ready for testing.

Developing an interface to the component

In this and the following sections, the Policy component is used as an example. The Policy component has the following interface:

```
#include <IManagedClient.idl>
#include "Beneficiary.idl"

exception InvalidAmount {};
```

```

interface Policy : IManagedClient::IManageable
{
    void addBeneficiary(Beneficiary benRef);
    void delBeneficiary(Beneficiary benRef);

    readonly attribute long policyNo; // Primary key for 'Policy'

    void changeAmount(float newAmount) raises(InvalidAmount);
    float getAmount();

    attribute string comment;
    readonly attribute float premium;
}

```

This represents an interface that clients of the Policy component can look at, understand, and design to language bindings and is used for compiling client applications.

Note: The previous interface would be contained in a file named Policy.idl.

This is a standard looking IDL file. It declares some methods and attributes. What makes this a Component Broker MOFW-specific interface is the fact that it inherits from IManagedClient::IManageable, which is included from the IManagedClient.idl file.

This interface is of particular interest from a Component Broker MOFW perspective because of the *read-only* attribute called *policyNo*. “Implementing the primary key class” on page 122 describes how to develop a primary key class for the Policy component. However, even as early as this, the first step in developing a component, you should already be thinking about primary keys.

A primary key class encapsulates the data that makes one instance of a class of objects unique from other instances of the same class of objects. Occasionally, the data that uniquely identifies an object is not part of the object’s state. In such cases, the data which uniquely identifies an object would probably be something like a Universally Unique Identifier (UUID). However, typically, it is a subset of the state of an object which makes the object unique. In the Policy example, the policy number is used to distinguish among separate instances.

As described earlier, attributes on an object’s interface define methods for accessing the state of an object. For example, attribute long policyNo; defines the following methods in C++:

```

virtual ::CORBA::Long policyNo ();
virtual ::CORBA::Void policyNo (::CORBA::Long policyNo);

```

However, there needs to be a practical balance between module scoping and the number of interfaces defined in that module. All interfaces that are part of

the same module must also exist in the same actual file. Introducing a large number of interfaces in the same file can lead to performance problems with some of the tools, such as compilers. Therefore, it is recommended that you define no more than ten interfaces in the same module due to this practical limitation. The first method retrieves the value of the policy number. The second allows the client to set it. A component gets its identity set in its data object from its primary key as part of creation or reactivation. For more information, see “Implementing the primary key class” on page 122. Changing an object’s identity after creation by setting one of the attributes which make up the key to a new value is not allowed in Component Broker. Applying the *readonly* qualifier to the *policyNo* attribute prevents the generation of the set method shown in the previous example, which means a client application has no interface for changing a Policy object’s identity.

Changing an object’s identity after creation requires the following multi-step process:

1. Create a new object with the new identity.
2. Copy the state of the old object into the new object.
3. Delete the old object by invoking the `remove()` method.

This allows the application adaptor to properly keep track of, and manage, instances of components on a Component Broker server.

Module scoping

Although the examples used in this book do not show it, wherever possible, IDL interfaces and types should be enclosed inside a module scope. IDL declared outside of a module scope takes up name space in the global IDL name space and risks having name collisions with names declared by other IDL developers.

For example, if two groups in an organization write IDL interfaces without placing them in different modules, and happen to choose names for their interfaces that overlap with each other, the result is name collision when the IDL is loaded into a shared interface repository. However, if each group chooses a unique module name, perhaps based on a combination of company name and group name, each group’s IDL is able to co-exist with all other IDL interfaces.

This example is even more useful when Java and the Internet are considered. In the IDL-to-Java mapping specification adopted by the Object Management Group (OMG), IDL module names are mapped directly into Java package names. IDL declared outside a module is mapped into classes and interfaces that are not contained in any package. Because one of the purposes of packages in Java is to partition the Java name space for the entire Internet, it is particularly important that IDL that might be mapped into Java be

contained within a module. Similar reasoning applies to C++, but because C++ classes are typically not downloaded and shared among all users of the Internet, the risk is somewhat lessened.

For example:

```
module XYZCompany_Finance
{
    interface Receivable
    {
        attribute long customerID;
        attribute long amountCents;
    };
};
```

The Receivable interface has the fully-scoped name XYZCompany_Finance::Receivable.

Design tips for components

As discussed in “Component attributes” on page 58, in C++, the *term* attribute describes a piece of data which is part of an object’s state. Good C++ object-oriented design (OOD) principles discourage putting attributes on the public interface of an object. Public attributes allow the user of a C++ object to directly manipulate the state of the object, without any control by the object itself. C++ OOD uses the term *encapsulation* to mean putting an object’s data members in the protected or private sections.

Because CORBA and IDL do not allow specifying data as part of an object’s interface, public data is not a problem. In CORBA, attributes on an object’s interface (as specified using IDL) define methods for accessing the state of an object. Because IDL does not allow the specification of exceptions to be thrown as the result of an attribute, IDL attributes do not provide the same level of protection as encapsulation in C++. To see why, continue with the Policy example. For illustration purposes, assume that the insurance company does not want to issue a policy unless the amount is over \$10,000.00. As shown previously, if *amount* had been defined as an attribute, it would result in the following two methods being declared:

```
virtual ::CORBA::Float amount ();
virtual ::CORBA::Void amount (::CORBA::Float amount);
```

Look at the implementation of the method for setting the attributes value:

```
virtual ::CORBA::Void amount (::CORBA::Float amount)
{
    if ( amount > 10000.00 )
        fAmount = amount; // save value in private data member
    else
        // What shall we do here...?
};
```

IDL does not allow us to raise an exception on an attribute. The method signature, as emitted by the IDL compiler, has no return value or output parameters. All you can do is *not* save the new value, return to the caller, and hope that the caller notices that the value was not changed. A better way to truly encapsulate the amount is to not make it an attribute in IDL, but rather a pair of get and set methods. This lets the business object do constraint checking and communicate error conditions to the client of the business object.

On the other hand, this needs to be balanced against performance. The Query Service of Component Broker performs best when a query is specified in terms of attributes, and furthermore, that the attributes of the component's business object map one-to-one to the attributes of the component's data object. Doing so lets the Query Service of Component Broker push down the query to the underlying resource manager. Resource managers such as relational database managers perform queries much faster than Component Broker can without any other help.

Module and interface names must be different. Although identical names are valid IDL syntax, it gets mapped to nested classes by the C++ emitter, and the C++ compiler does not allow the name of a Container class to be the same as a Contained class.

Selecting a pattern for handling essential state

This step introduces another interface that looks like this:

```
#include "Policy.idl"
#include <IManagedServer.idl>

interface PolicyBO : Policy,
    IManagedServer::IManagedObjectWithDataObject
{
};
```

Note: This interface is contained in a file named PolicyBO.idl.

This interface is small and straightforward. By inheriting from `IManagedServer::IManagedObjectWithDataObject` a decision has been made about the pattern to be used for handling the essential state of this component. There are two interfaces in `IManagedServer` that correspond to the two patterns that the Component Broker MOFWs support for handling the essential state of a component.

IManagedObjectWithDataObject

This interface implies that the component has its state managed by a data object. Because the component has a data object, and no implied caching as in the next interface, this is called the *delegating pattern*. There are implications here for the implementor of the business logic

methods. These are described in “Implementing business object methods and attributes” on page 109.

IManagedObjectWithCachedDataObject

This interface implies the presence of a data object and support from Component Broker for fetching this data into a cached set of state data maintained by the component’s business object at the appropriate times throughout the lifetime of a component. This is called the *caching pattern*, because the business object is caching the state being maintained by the data object.

The data object

For the patterns described previously, a data object is necessary. A data object manages the essential state of a component. For the working example, assume that the data object for the Policy is as follows:

```
interface PolicyDO
{
    attribute long policyNo;
    attribute float amount;
    attribute string comment;
    attribute PolicyHolder insured;
}
```

The interface for the data object contains one attribute for each piece of the component’s essential state, as shown in Figure 24 on page 58, and that there are no attributes for the component’s non-essential state. Also notice that this interface is not a one-to-one mapping of the attributes on the business object. The *premium* attribute from the business object is not present here because it is not part of the essential state. Also, although *amount* and *insured* were not attributes on the business object, they are attributes on the data object because they are part of the component’s essential state.

This interface is enough to let you develop the component’s business logic, without having to actually implement the underlying data access methods. However, before the business logic can be tested, the data object interface must in fact be implemented. “Data object implementation details” on page 361 describes how to customize the data object to allow the component to be installed into a particular application adaptor.

Design tips for data objects

There are several considerations when defining and naming the data object and its attributes. First, be careful when naming the attributes for the data object. The Query Service of Component Broker operates more efficiently if the attribute names on the component’s business object are the same as the attribute names on the component’s data object. Second, while attributes on

the business object's interface may not be a good idea, depending on the amount of encapsulation desired, they are a convenient way of declaring get and set methods on the data object. Because the data object is used only by the business object, and by the application adaptor, encapsulation in the data object is not a consideration.

Implementing business object methods and attributes

In this step, the implementation bindings for PolicyBO are filled in with the business logic required to support the Policy interface specified in the Policy.idl file in C++ code. The implementation bindings generate a class called PolicyBO_Impl in this case. The emitted files, named PolicyBO.ih and PolicyBO_I.cpp, do *not* have any of the methods that are introduced by the parents of PolicyBO, but which need to be implemented by PolicyBO.

Note: This is a limitation of IDL as defined by CORBA: there is no way to specify in IDL whether an interface has been implemented. Without this knowledge, an IDL compiler must make an assumption. The IDL compiler provided with Component Broker assumes that every interface is implemented.

These methods must be introduced to PolicyBO.ih and PolicyBO_I.cpp after emitting.

The Object Builder is going to ensure that the .ih and _I.cpp files at the business object level get the method declarations and implementation stubs that are introduced at the level above. Additional C++ methods and data can be introduced as needed to meet the requirements of the implementation. However, anything additional that is added is not accessible remotely from clients because the interface used by clients is described in the client bindings. For C++, the client usage bindings are specified in .hh files according to the CORBA 2.0 C++ mappings.

A sample implementation of the addBeneficiary() method follows:

```
// somewhere at the top
#include "PolicyBO.ih"

// a business logic method implementation
::CORBA::Void PolicyBO_Impl::addBeneficiary(Beneficiary benRef)
{
    // Assume that beneficiaries was declared as a private C++ data
    // member of type IManagedCollections::IReferenceCollection_var.
    beneficiaries->addElement(benRef);
}
```

Each of the business logic methods needs to be implemented in a similar fashion. Handling of attributes from the interface is more involved, and is described in the following sections.

Using C++ 'this' references in business objects

Care must be taken when programming all methods in components that use references to themselves when communicating with other objects. Methods must use the programming model as described in this section when using these self references. The technique of using "this" is no longer supported in the programming model in these circumstances.

A local proxy class is created for each interface defining the managed object implementation, the managed object interface, the business object interface, and every other interface that they may support. Only instances of the local proxy of the managed object implementation are instantiated and these proxies must be used for self references. The `_this()` method can be used to access this proxy in the business object.

The home provides a copy reference to a local proxy for the `create()` and `findBy()` methods that return object references. The following example shows the set of rules to follow when an object passes itself as an argument or returns itself as a return argument:

```
I_ptr I::foo(Bar_ptr bar)
{
    // When using normal sequences and structs
    sequence[1] = _this();
    struct.i = _this();

    // We recommend using sequences that release object references
    // automatically. If you must use sequences that do not automatically
    // release their references, there must be another mechanism to
    // free them.
    I_var objref = _this();           // The var will release the objref upon
                                     // exiting from this scope. This is just
                                     // one of many possible ways to release
                                     // the object reference.
    specialSequence[1] = objref;     // Sequence instantiated with no release
                                     // indication.

    // return value
    return _this();
}
```

For backwards compatibility, the `_self()` method is retained in this release of Component Broker. The following example shows the set of rules to follow when using `_self()`:

```
I_ptr I::foo(Bar_ptr bar)
{
    // When using normal sequences and structs the reference must be
    // duplicated because these structures own the objects stored in them.
    sequence[1]= I::_duplicate(_self());
    struct.i = I::_duplicate(_self());
}
```



```

// For sequences which do not release the objects within it.
specialSequence[1] = _self(); // Sequence instantiated
                               // with no release indication

// The reference must be duplicated for returning the value
return (I::_duplicate(_self()));
}

```

Warning: We will be deprecating the `_self()` method in a future release of Component Broker in order to maintain CORBA compatibility using the `_this()` method. We encourage the user to migrate existing code, and to start using the `_this()` method in any new code.

Patterns for handling state (caching)

If `IManagedObjectWithCachedDataObject` was chosen as the data object pattern, then all essential state, and all non-derived non-essential state, is stored or cached in the component's business object. To store this state in the business object, protected or private data should be declared inside of the implementation binding header (.ih) file. For `Policy` and `PolicyBO` previously mentioned, add the following in the `PolicyBO.ih` file:

```

class PolicyBO_Impl : public virtual ::PolicyBO_Skeleton
{
    ...
    // methods being implemented go here
    ...
    protected:
        ::CORBA::Long fPolicyNo;
        ::CORBA::Float fAmount;
        PolicyHolder_var fInsured;
        ::CORBA::String_var fComment;
        RiskCalculator_var fRiskCalc; // C++ helper object
};

```

Attributes: The implementation bindings *getter* and *setter* methods for each of the attributes specified in an interface. In the `PolicyBO_I.cpp` file, the getter and the setter for the *comment* attribute are implemented as follows:

```

::CORBA::Void PolicyBO_Impl::comment (const char* comment)
{
    fComment = comment; // ::CORBA::String class will make copy
}

char* PolicyBO_Impl::comment ()
{
    return ::CORBA::string_dup(fComment);
}

```

The previous example is a simple implementation of getters and setters. More complex logic and exception handling may be required. In fact, getters and

setters might actually have logic in them to write to resource managers or get data from resource managers. However, as described previously under “Design tips for components” on page 106, if the logic required in the implementation of getters and setters includes any sort of bounds checking or error checking, then the attribute should be changed to a pair of methods in the IDL file.

Other methods: Other business logic methods which access state data do so similarly to the *comment* attribute mentioned previously: they use the data that is cached in the business object.

Patterns for handling state (delegating)

If `IManagedObjectWithDataObject` was chosen as the data object pattern, then all essential state is accessed by the data object: when the business logic needs to get or set its essential state, it does so by using the attributes on the data object. This is called the delegating pattern because the maintenance of the component’s state is delegated from its business object to its data object. All non-derived non-essential state is still cached in the business object. Using this pattern, the implementation binding header (.ih) file for `PolicyBO` (`PolicyBO.i`) would look something like this:

```
class PolicyBO_Impl : public virtual ::PolicyBO_Skeleton
{
    ...
    // methods being implemented go here
    ...
protected:
    PolicyDO_var fDataObject;
    RiskCalculator_var fRiskCalc; // C++ helper object
};
```

Attributes: The implementation bindings generate getter and setter methods for each of the attributes specified in an interface. In the `PolicyBO_I.cpp` file, the getter and setter for the *comment* attribute would be implemented as follows:

```
::CORBA::Void PolicyBO_Impl::comment (const char* comment)
{
    fDataObject->comment(comment);
}

char* PolicyBO_Impl::comment ()
{
    return ::CORBA::string_dup( fDataObject->comment() );
}
```

Other methods: Other business logic methods which access essential state data do so similarly to the *comment* attribute mentioned previously: they use the general `fDataObject->datamember` pattern.

An advantage of this programming model is that the business logic can be written independent of the exact implementation of the data object. At this point, you only know which data is to be persistently stored, not where it is stored or how it is accessed. This encapsulation of information inside of the data object makes business logic more stable.

Implementing the IManageable required methods

You need to implement the following methods, which are described in the IManagedClient::IManageable interface and its parents.

IManageable::getPrimaryKeyString

This method returns a ByteString that contains the contents (in the form of a string) of the IManagedLocal::IPrimaryKey subclass for the component type being implemented. This could be used to extract the identity of a component for the purpose of using it at a later time to locate the same component using a findByPrimaryKeyString() method on a home.

Note: A ByteString created from an object on one machine may not bitwise compare equal to a ByteString created from the same object on a different machine.

You have several options in implementing this method. The best way to implement this method is to create an instance of the IManagedLocal::IPrimaryKey subclass, set the values into this object, and return the ByteString value. Here's an example of how this would be done for the Policy example:

```
::ByteString* PolicyBO_Impl::getPrimaryString()
{
    //Insert Method modifications here
    PolicyKey_var policyKey = PolicyKey::_create();
    policyKey->policyNo(iDataObject->policyNo());
    return policyKey->toString();

    //End Method modifications here
}
```

The previous example assumes that the pattern for dealing with data is either caching, or that there is no data object. A delegating pattern would cause the `k->policyNo(fPolicyNo);` statement to change to a call to the data object. This method must be overridden. No default implementation is provided by the MOFW.

IManageable::getHandleString

CORBA provides the ability to invoke `object_to_string()` on any object. This might be useful for persistently storing a reference to an object. Component

Broker introduces the concept of a handle. Handles provide a way to keep track of an object using a mechanism different than that of a stringified object reference, if this is desired.

You are not required to implement handles. The Component Broker default implementation of this method uses a stringified object reference.

There are many ways to define a more stable handle. The following code is an example of an implementation of the `getHandleString` method for a handle that combines the primary key and the name of the home into a handle.

```
ByteString* Policy BO_Impl::getHandleString
{
    PKHomeHandle_var myHandle = PKHomeHandle::_create();
    myHandle->home(getHome()); // set value of home into handle
    myHandle->primaryKey(getPrimaryKeyString());
    // set value of primary key into handle
    return myHandle->toString();
}
```

CosStream::Streamable::externalize_to_stream

Component Broker components are always streamable. Streaming can be the basic mechanism used for copying and moving objects. This mechanism is used by the OS/390 Component Broker frameworks to copy objects for internal use. The other Component Broker platforms do not currently use this mechanism so this method is required to be implemented only for OS/390 Component Broker support. This mechanism is used by the OS/390 Component Broker frameworks to copy objects for internal use. The other Component Broker platforms do not currently use this mechanism so this method is required to be implemented only for OS/390 Component Broker support.

The `externalize_to_stream` method is the `CosStream::Streamable` method that writes the state data of an object into a stream. For components that are not using a data object or for those that use a cached data object, streaming is done as follows:

```
::CORBA::Void PolicyBO_Impl::externalize_to_stream
    (::CosStream::StreamIO_ptr targetStreamIO)
{
    targetStreamIO->write_long( fPolicyNo );
    targetStreamIO->write_float( fAmount );
    targetStreamIO->write_string( fComment );

    CORBA::String_var stringifiedInsured =
        CBSeriesGlobal::orb()->object_to_string( fInsured );
    targetStreamIO->write_string( stringifiedInsured );
}
```

For components that use a cached data object, use the data in the business object's cache instead of the data in the data object, because the data in the cache is likely to be more up-to-date.

Do not externalize the pointer to the data object that exists in the private data of the implementation interface in the cases where a data object is being used.

If the business object is delegating its state to a data object, it is the data in the data object that is extracted and placed into the stream. The implementation of this for the Policy example is as follows:

```
    ::CORBA::Void PolicyBO_Impl::externalize_to_stream
        (::CosStream::StreamIO_ptr targetStreamIO)
    {
        targetStreamIO->write_long( fDataObject->policyNo() );
        targetStreamIO->write_float( fDataObject->amount() );
        targetStreamIO->write_string( fDataObject->comment() );

        CORBA::String_var stringifiedInsured = CBSeriesGlobal::orb()->
            object_to_string( fDataObject->fInsured() );
        targetStreamIO->write_string( stringifiedInsured );
    }
```

If the data for an object involves references to other objects, there are two possible approaches: *shallow* and *deep*. In the shallow approach to streaming, references to contained objects are stringified, and then the string is written out to the stream. This is the approach shown in the previous two code segments. In the deep approach, contained objects are asked to stream themselves into the same stream as the rest of the object's attributes by invoking `externalize_to_stream()` on the contained objects.

You should use the shallow approach. If a contained object has many attributes, such as lots of state data, and if it contains references to other objects, it is likely that the deep approach is slower and results in a larger stream (object). In addition, the shallow approach makes internalization more straightforward, as discussed in the following section.

CosStream::Streamable::internalize_from_stream

The `internalize_from_stream` method on `CosStream::Streamable` needs to be written so that it can read the values that the `externalize_to_stream` method placed into the stream. The values must be read in the same order that they were written. This mechanism is used by the OS/390 Component Broker frameworks to copy objects for internal use. The other Component Broker platforms do not currently use this mechanism so this method is required to be implemented only for OS/390 Component Broker support. This mechanism is used by the OS/390 Component Broker frameworks to copy objects for internal use. The other Component Broker platforms do not currently use this

mechanism so this method is required to be implemented only for OS/390 Component Broker support. The example for the caching case and the case when no data object is present follows:

```

::CORBA::Void PolicyBO_Impl::internalize_from_stream(
    ::CosStream::StreamIO_ptr sourceStreamIO,
    ::CosLifeCycle ::FactoryFinder_ptr, there)
{
    ::CORBA::Long tempPolicyNo = sourceStreamIO->read_long();
    if ( fPolicyNo == tempPolicyNo )
    {
        fAmount = sourceStreamIO->read_float();
        fComment = sourceStreamIO->read_string();

        CORBA::String_var stringifiedInsured =
            sourceStreamIO->read_string();
        CORBA::Object_var obj = CBSeriesGlobal::orb()->
            string_to_object( stringifiedInsured );
        fInsured = PolicyHolder::_narrow(obj);
    }
    else
        // throw an exception which says that this is the
        // wrong object to load this stream into
}

```

Note: Do not change the key attributes of a Component Broker business object.

In the delegating to data object case, the code for the Policy example is as follows:

```

::CORBA::Void PolicyBO_Impl::internalize_from_stream(
    ::CosStream::StreamIO_ptr sourceStreamIO,
    ::CosLifeCycle ::FactoryFinder_ptr, there)
{
    ::CORBA::Long tempPolicyNo = sourceStreamIO->read_long();
    if ( fDataObject->policyNo()==tempPolicyNo)
    {
        fDataObject->amount(sourceStreamIO->read_float());
        fDataObject->comment(sourceStreamIO->read_string());

        CORBA::String_var stringifiedInsured =
            sourceStreamIO->read_string();
        CORBA::Object_var obj = CBSeriesGlobal::orb()->
            string_to_object( stringifiedInsured );
        fDataObject->insured(PolicyHolder::_narrow(obj));
    }
    else
        // throw an exception which says that this is the
        // wrong object to load this stream into
}

```

If the data in the stream is a stringified object reference, as it would be using the shallow approach described in the previous example, then run the

string_to_object method on the ORB with the string that is read in from the stream. This is shown in the previous two code segments.

If, however, the data in the stream represents the entire contents of a contained object, as would be the case in the deep approach mentioned previously, then the internalize_from_stream() method would have to do the following:

- Read all of the contained object's attributes from the stream.
- Create a primary key object for the contained object type.
- Initialize the primary key with data from the stream.
- Find a home for the contained object type.
- Find the contained object by invoking findByPrimaryKeyString() on the home.

The extra work involved in the deep approach could be minimized by writing the contained object's primary key string into the stream instead of its contents, but it would still not be as simple as the shallow approach, and would no longer really be a deep approach, but an alternate shallow approach.

The assumption for the internalize_from_stream method is that the object was created before reading in the stream. This also implies that the initForCreation() method was run on the business object.

Summary of IManageable methods

IManageable and its ancestors introduce a number of methods into the interface of a business object. Table 3 summarizes these methods and the implementation of each. Only the methods that you must implement or may choose to implement are enumerated here. Other methods exist on these interfaces which are not intended to be overridden by the object provider.

This table provides a summary:

Table 3. IManageable method implementations table

interface::MethodName	Production
IManageable::getPrimaryKeyString	object provider
IManageable::getHandleString	default provided by Component Broker
IManageable::getHome	default provided by Component Broker
Streamable::externalize_to_string	object provider
Streamable::internalize_from_stream	object provider

This might seem like a lot of methods to implement but they are simple methods to implement. Tools can assist in implementing all of the methods

described in Table 3 on page 117. The tool-generated implementation of these methods is sufficient to begin testing in most cases, and can be customized as needed for specific business reasons.

Implementing IManagedObject required methods

In addition to the business logic methods, there are MOFW-required methods that must be implemented regardless of which kind of IManagedObject class is chosen as the base class.

initForCreation() method

The managed object framework invokes this method on a component's managed object when the component is initially created. This method is the MOFW equivalent of a C++ constructor. Unlike a C++ constructor, by the time this method is called, the object's essential state has already been initialized. In Component Broker, an object's state is initialized by a primary key on a `createFromPrimaryKey()` method, or a copy helper object on a `createFromCopy()` method.

In addition, if you are inheriting from the IManagedObject interfaces that have a data object (IManagedObjectWithDataObject and the IManagedObjectWithCachedDataObject), then you must get the input parameter that is provided on this method. It is the first thing you should do in the method. The `::_narrow` is used so that the component has access to specific methods introduced for dealing with domain dependent state data.

```

::CORBA::Void PolicyBO_Impl::initForCreation(
::IManagedServer::IDataObject_ptr theDO);
{
    // keep the data object for later use
    fDataObject = PolicyDO::_narrow(theDO);
    // then put other initForCreation code here
    ...
}

```

If you chose a pattern for handling essential state that requires a data object, then you need to have a data member in the .ih file (for example, PolicyBO.ih) to hold the pointer to the data object. This is what is set in the previous code segment.

```
PolicyDO* fDataObject;
```

The data object can be used for any initialization tasks that need to be done. There is no guarantee as to the number of attributes in the data object that are validly set. That would be based on whether a copy helper or primary key was used to do the create. Caching business objects should take this opportunity to load the cache for the business object with any relevant values from the data object. At this point, the data object has values for any key or

copy helper data that was passed along during the create. In fact, the primary key information must be extracted from the data object and placed into the cache of the business object.

Given that the data object has a default constructor with reasonable initialization, as recommended and as generated by the tools, these lines of code should be in the `initForCreation()` method:

```
fPolicyNo = fDataObject->policyNo();  
fAmount = fDataObject->amount();  
fComment = fDataObject->comment();  
fInsured = fDataObject->insured();
```

uninitForDestruction() method

The managed object framework invokes this method on every managed object when the object is being destroyed. This method is the MOFW equivalent of a C++ destructor. If the managed object were managing its own resources, this is where it would remove them.

The data object can be used for anything that is needed. Information can be pulled out of the data object at this point. This is also the opportunity to enforce any referential integrity constraints. For example, this might mean removing an object to which it is pointed.

In addition to `initForCreation()` and `uninitForDestruction()`, the `IManagedObjectWithDataObject` and the `IManagedObjectWithCachedDataObject` interfaces introduce the `initForReactivation()` and `uninitForPassivation()` methods. These are described next.

initForReactivation() method

The managed object framework (MOFW) invokes this method on every managed object when the object is being recreated in storage. This is the MOFW equivalent of an operating system paging storage in from disk. Because the managed object is being put into storage, it needs to make sure that any resources that it was managing are ready to be used. For example, if the managed object were managing its own network connection, the implementation of this method would reestablish the connection.

In addition, if you are inheriting from the `IManagedObject` interfaces that have a data object (`IManagedObjectWithDataObject` and the `IManagedObjectWithCachedDataObject`), then you must get the input parameter that is provided on this method. This is the first thing you should do in the method. The `::_narrow` is used so that the business object has access to specific methods introduced for dealing with domain dependent state data.

```

::CORBA::Void PolicyBO_Impl::initForReactivation(
    ::IManagedServer::IDataObject theDO);
{
    // keep the data object for later use
    fDataObject = PolicyDO::_narrow(theDO);
    // then put other initForReactivation code here
    ...
}

```

While it is necessary to get the pointer to the data object in this method, the data object should not be used in any way during this method. This method has no access to the essential state stored in the data object.

uninitForPassivation() method

The managed object framework invokes this method on every managed object when the object is temporarily removed from storage. This is the MOFW equivalent of an operating system paging storage out to disk. Because the managed object is removed from storage, it clearly is not using any resources that it is holding. You have the choice of releasing any held resources at this time, or continuing to hold on to them. If the managed object is not managing any resources, or if you choose to hold onto those resources, then no implementation for this method is required.

The data object cannot be used or accessed during the execution of this method.

In addition to the methods already described, `IManagedObjectWithCachedDataObject` introduces two additional methods. These are the `syncToDataObject()` and `syncFromDataObject()` methods.

syncFromDataObject() method

The purpose of `syncFromDataObject()` is to load the business object with the data that is contained in the data object. In “`initForCreation()` method” on page 118 and “`initForReactivation()` method” on page 119, a pointer to the data object `fDataObject` is set. In this method, the data object must be accessed to load up the local state data into its cache, as the name of the base class for this business object implies (`IManagedObjectWithCachedDataObject`). The following code is placed in the `syncFromDataObject` for the Policy example:

```

fPolicyNo = fDataObject->policyNo();
fAmount = fDataObject->amount();
fComment = fDataObject->comment();
fInsured = fDataObject->insured();

```

Look in the implementation interface for the business object being constructed and ensure that all of the data is properly loaded from the data object. This method is called by the Component Broker server, but must be implemented by the object provider.

MOFW does not let you invoke other methods on the component's public interface from within the implementation of this method.

Observe that the following lines of code are the same in `initForCreation()` and `syncFromDataObject()`.

```
fAmount = fDataObject->amount();  
fComment = fDataObject->comment();  
fInsured = fDataObject->insured();
```

To avoid duplicating this code, introduce a method in your business object called `initializeState()` that is called from both `initForCreation()` and `syncFromDataObject()`.

syncToDataObject() method

The purpose of `syncToDataObject()` is to flush the data from the business object back to the data object. This puts the data object in a state in which it can deal with the underlying resource manager, and ensure that this data is properly stored persistently, using the right transaction interfaces to the underlying resource manager. The `syncToDataObject()` method for the Policy example would look like:

```
fDataObject->policyNo(fPolicyNo);  
fDataObject->amount(fAmount);  
fDataObject->comment(fComment);  
fDataObject->insured(fInsured);
```

This method is called by the Component Broker server, but must be implemented by the object provided.

The `syncFromDataObject()` and `syncToDataObject()` methods are called by the Component Broker server at the appropriate times. They should not be called by the business object in any of its methods. The server calls `syncToDataObject()` at all appropriate times before passivation to ensure the data is properly stored back to the underlying resource manager. It calls `syncFromDataObject()` at all appropriate times after reactivation to properly load the business object's cache. These methods should only contain logic that pertains to the values of the business object's cache.

Summary of IManagedObject methods

All of the methods defined in the `IManagedObject` subclass chosen must be implemented. They have been described in the previous sections. These

methods should be called only by the Component Broker server. See “Framework flows” on page 442 for further information.

Implementing the primary key class

Each component must have an associated primary key class to be used in Component Broker. There are several ways to develop a key class. The most prevalent way is to create a key class based on the MOFW base class `IManagedLocal::IPrimaryKey`.

The `IManagedLocal.idl` file provides a set of interfaces for keys. `IPrimaryKey` is the class that you want to subclass from in order to create a key class that works with your component. For example, in the insurance policy case, a key class like this:

```
interface PolicyKey : IManagedLocal::IPrimaryKey
{
    attribute long policyNo;
    #pragma meta PolicyKey localonly
}; // end interface PolicyKey
```

In this example, the primary key of the Policy component consists of a single attribute, *policyNo*. The primary key of an object can consist of one or more public attributes. Protected or private attributes cannot be used as part of a primary key. For example, the Primary Key of a Beneficiary component consists of both the name of the beneficiary and the policy number of the policy for which the person is a beneficiary.

To accommodate the key attributes, protected or private data should be declared inside of the implementation binding header file (.ih). For the `PolicyKey` in the previous example, put the following in the `PolicyKey.ih` file:

```
::CORBA::Long      fPolicyNo;
```

In the `PolicyKey_I.cpp` file, the attributes would be implemented as follows:

```
::CORBA::Void PolicyKey_Impl::policyNo(::CORBA::Long policyNo)
{
    fPolicyNo = policyNo;
}

::CORBA::Long PolicyKey_Impl::policyNo()
{
    return fPolicyNo;
}
```

CosStream::Streamable methods

The other main requirements for the implementation of this class are the streaming methods as shown in the following example:

```

::CORBA::Void PolicyKey_Impl::externalize_to_stream
    (::CosStream::StreamIO_ptr targetStreamIO)
{
    // Insert Method modifications here
    targetStreamIO->write_long(fPolicyNo);

    // End Method modifications here
}

::CORBA::Void PolicyKey_Impl::internalize_from_stream
    (::CosStream::StreamIO_ptr sourceStreamIO,
     ::CosLifeCycle::FactoryFinder_ptr there)
{
    // Insert Method modifications here
    fPolicyNo=sourceStreamIO->read_long();

    // End Method modifications here
}

```

Although the `CosStream::Streamable` methods and the `toString()` and `fromString()` methods introduced by Component Broker might appear to work independently, they actually work well together. Object providers generally implement the `internalize_from_stream()` and `externalize_to_stream()` methods on all objects that are constructed. For primary keys and copy helpers, call `toString()` and `fromString()` when an object must be flattened to bits or revived from bits. The `toString()` and `fromString()` are actually implemented by the MOFW by calling the `Streamable` methods. This ensures proper flattening of all CORBA types and proper handling of code pages when the bits are moved between machines. The `ByteString` data type used by `fromString()` and `toString()` is really a wrapper over the data format used by CORBA, and therefore Component Broker, to move data between processes.

Implementing `IKey::getName`

The implementation of this method returns the class name. For the `PolicyKey` class, the implementation would look something like this:

```

char* PolicyKey_Impl::getName()
{
    //Insert Method modifications here
    return CORBA::string_dup("PolicyKey");

    //End Method modifications here
}

```

Implementing `IKey::isEqualToKey`

Component Broker provides a default implementation of this method that compares the stringified values of the keys. It can be optionally implemented by the object provider. The implementation of this method should use the `get`

method on the keys and compare the values to see if the key is for the same component instance. The following would be an implementation of this for the PolicyKey class.

```
::CORBA::Boolean PolicyKey::isEqualToKey(IKey inKey)
{
    PolicyKey_var pk = PolicyKey::_narrow(inKey);
    if ( pk->policyNo() == fPolicyNo )
        return 1;
    else
        return 0;
}
```

Implementing IKey::isEqualToKeyString

Component Broker provides a default implementation of this method that compares the stringified values of the keys. It can be optionally implemented by the object provider. This method is similar to the previous one, but deals with the stringified version of the key. The implementation is shown in the following example:

```
::CORBA::Boolean PolicyKey::isEqualToKeyString(ByteString inString)
{
    PolicyKey_var tempKey = PolicyKey::_create();
    tempKey->fromString(inString)
    if (tempKey->policyNo() == fPolicyNo)
        return 1;
    else
        return 0;
}
```

Summary of key class construction

This primary key class enables you to use the createFromPrimaryKeyString() and findByPrimaryKeyString() methods that are part of the IManagedClient::IHome interface. To pull it all together now, the following shows the creation of a key, the creation of a component based on this key, and usage of this component.

```
// show a simple usage of createFromKeyString assuming you have a home

PolicyKey_var aPolicyKey = PolicyKey::_create();
aPolicyKey->policyNo(1234);

Policy_var aNewPolicy;
IManagedClient::IManageable_var aManageable;

aManageable = thePolicyHome->
    createFromPrimaryKeyString(aPolicyKey->toString());
aNewPolicy = Policy::_narrow(aManageable);

aNewPolicy->amount(123.45);
...
...
```

Information about how the value of a key gets mapped to a specific component is contained in Component Broker and in the data object implementation. Details on the data object implementation and this important mapping are in “Chapter 13. Assembling and installing components on Component Broker/Workstation” on page 351.

Other ways to get the key class

Some containers come with a predefined subclass of `IManagedLocal::IPrimaryKey` that works for all component types that are to be stored in a particular container. If you are planning to install the component into this kind of container, then you should consult the documentation for that container to determine the kind of key that needs to be used when dealing with component that live in that container.

See “Local only development process” on page 127 for information on building local only objects such as primary keys.

Implementing the optional copy helper class

In addition to the `createFromPrimaryKeyString()` method that is part of the `IManagedClient::IHome` interface, there is also a method called `createFromCopyString()`, which supports another creation method for a component. This method is available to make creating component more efficient. Rather than running `createFromPrimaryKeyString()` on the `IManagedClient::IHome` followed by a series of remote `setxxxx()` methods to load a component, `createFromCopyString()` enables the creation of a local transient object that is externalized and passed in its entirety to the server in one call.

To create a copy helper class, subclass from `IManagedLocal::INonManageable`. This is the base class for local only objects that interact with and act like CORBA objects, support streaming, but are not accessible remotely. Following is an example of the interface that gets declared in IDL for the `PolicyCopy` class.

```
interface PolicyCopy : IManagedLocal::INonManagable
{
    attribute long    policyNo;
    attribute float  amount;
    attribute float  comment;
}
```

In many ways, the development process for simple copy helper classes is the same as that required for implementing primary key classes. The only difference is the fact that typically more attributes are placed into a copy helper class. A pragmatic suggestion for building your first copy helper class

is to borrow pieces of the primary key subclass and modify it to become a copy helper. What is added is probably based on the pieces of the component's basic business object, for example Policy, that might be present in a copy.

This interface implies that the copy helper class allows the loading of all of the state of the Policy component interface.

This interface then has an implementation with setters and getters for each of the attributes. In addition, the `internalize_from_stream` and `externalize_to_stream` methods need to be implemented. These implementations are shown in the following code segments.

```
class PolicyCopy_Impl : public virtual ::PolicyCopy_Skeleton,
    public virtual IManagedLocal_INonManageable_Impl
{
public:
    // excerpts from the C++ interface used for implementation
    // which is probably named PolicyCopy.ih
    ...
    virtual ::CORBA::Void externalize_to_stream(
        CosStream::StreamIO_ptr targetStreamIO);

    virtual ::CORBA::Void internalize_from_stream(
        CosStream::StreamIO_ptr sourceStreamIO,
        CosLifeCycle::FactoryFinder_ptr there);

    // add getters/setters and isReady() to this list
    ...
    ...
protected:
    ::CORBA::Long      copyPolicyNo;
    ::CORBA::Float     copyAmount;
    ::CORBA::String_var copyComment;
};
```

Then, the implementation class for this copy object has this:

```
::CORBA::Void PolicyCopy_Impl::externalize_to_stream(
    CosStream::StreamIO_ptr targetStreamIO)
{
    targetStreamIO->write_long(copyPolicyNo);
    targetStreamIO->write_float(copyAmount);
    targetStreamIO->write_string(copyComment);
}

::CORBA::Void PolicyCopy_Impl::internalize_from_stream(
    CosStream::StreamIO_ptr sourceStreamIO,
    CosLifeCycle::FactoryFinder_ptr there)
{
    copyPolicyNo = sourceStreamIO->read_long();
    copyAmount = sourceStreamIO->read_float();
    copyComment = sourceStreamIO->read_string();
}
```



```
// add getters/setters and isReady() to this list
...
...
```

The copy helper is complete. It is a lightweight object that encapsulates the subset of the component's state data that is fed in on a `createFromCopyString()` call. In addition to this, the implementation of the `CosStream::Streamable` methods enables this class to use the Externalization Service to get the contents ready for sending to the server.

A copy helper class should have a default constructor that initializes values to appropriate default values. This should help in the case where only partial information is available when the copy is built. A copy helper must contain all the information necessary to create a primary key which allows a home to create a unique component from the copy data.

See "Local only development process" for more details on building these local only copy helper objects.

Local only development process

Both the copy helper classes and the primary key classes are developed using a variant of the development process used for building a component's implementations. It is actually a simpler process because these are local only objects which are not accessible remotely. All interfaces (specified in IDL), inheriting specifically from `IManagedLocal::ILocalOnly`, `IManagedLocal::INonManageable`, or any other subclass of `IManagedLocal::ILocalOnly` follows this development process. A more detailed description of this process is discussed in "Local only object implementation details" on page 351.

For special local only notes for copies and keys, every helper class must be written in IDL and generate a C++ binding (and implementation) because the Component Broker server is built using C++ and these classes get used on the server by the Component Broker run time. It is up to the object provider to determine if a Java version (binding and implementation) of the helper classes is also desired for use on a Java client.

Summary

This chapter explained the minimum important artifacts necessary to implement a component. At the end of these activities, you should have the following:

- IDL Files

- Interface for client (Policy.idl).
- Interface for implementation (PolicyBO.idl).
- Interface for essential state (PolicyDO.idl).
- Interface for primary key class (PolicyKey.idl).
- Interface for copy helper class (PolicyCopy.idl).
- Implementation files
 - Implementation of business logic and MOFW-required methods (for example, PolicyBO.ih and PolicyBO_I.cpp that contains the PolicyBO_Impl C++ class).
 - Implementation of the primary key class (for example, PolicyKey.ih and PolicyKey_I.cpp).
 - Implementation of the copy helper class (for example, PolicyCopy.ih and PolicyCopy_I.cpp).
- Client and server bindings (for example *xxxx_C.cpp* and *xxxx_S.cpp* files) for the interface classes and implementation classes. The IDL emitter generates these bindings. Because primary key and copy helper classes are local only, they only require *_C.cpp* files for bindings.
- Usage bindings
 - Policy.hh
 - PolicyBO.hh
 - PolicyKey.hh
 - PolicyCopy.hh

These are the artifacts necessary to begin preparing a real component, with a managed object to be installed into an application. Table 4 on page 129 summarizes what is specified and from where in the managed object framework the abstractions in the `IManagedServer` module are inherited.

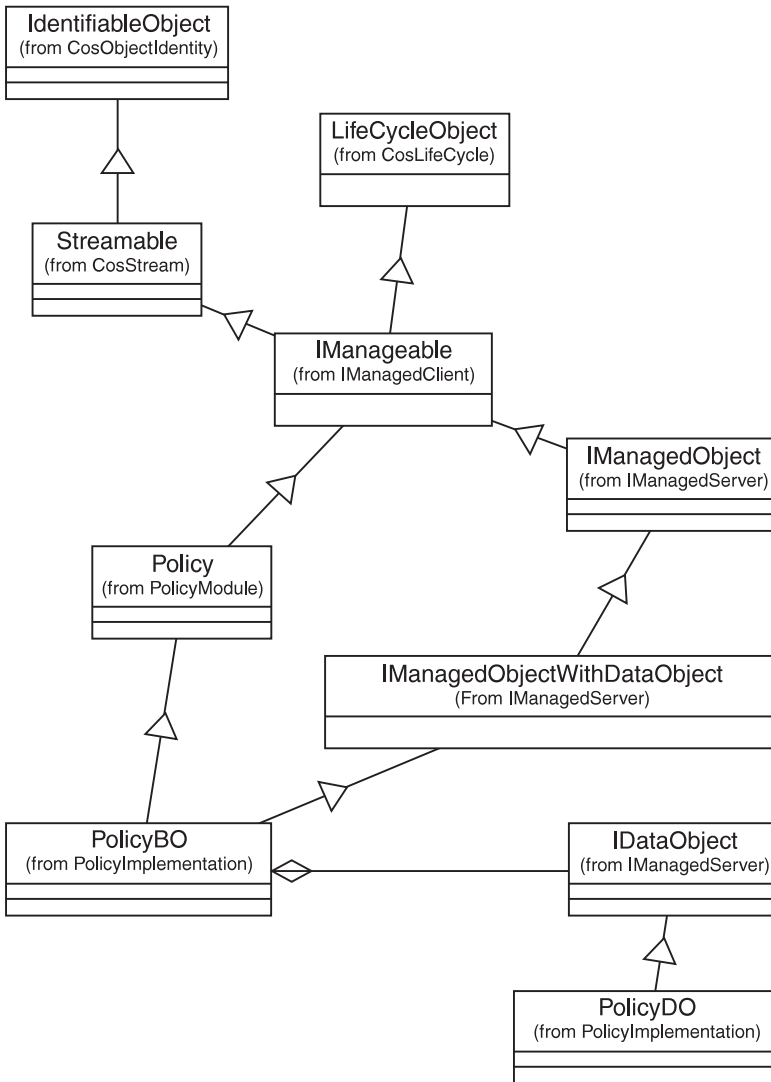


Figure 34. Basic artifacts in component development and their MOFW inheritances

Table 4. Artifacts table for business object

Interface Name	Implementor	Server DLL Name	Client DLL Name	Comments
Policy	Object provider	PolicyBO.dll should hold Policy_S.obj	PolicyClient.dll should hold Policy_C.obj	Interface only

Table 4. Artifacts table for business object (continued)

Interface Name	Implementor	Server DLL Name	Client DLL Name	Comments
PolicyBO	Object provider	PolicyBO.dll should hold PolicyBO_S.obj	See note 1	Real business logic and MOFW required methods
PolicyDO		n/a	See note 1	Documents essential state
Data object		n/a	See note 1	Interface only
Manageable	Component Broker provides some	Manageable.dll	ManageableClient.dll	Interface and some default implementation
	Object provider must do some in its PolicyBO	PolicyBO.dll		
Identifiable	Object Services	Identifiable.dll	IdentifiableClient.dll	
LifeCycle		n/a		
Streamable	Component Broker provides some	Manageable.dll	StreamableClient.dll	
	Object provider must do some in its PolicyBO	PolicyBO.dll		
IManagedObject			See note 1	Interface only
IManagedObjectWith DataObject	Object Provider (in its PolicyBO)	See PolicyBO.dll	See note 1	
PolicyCopy	Object Provider	PolicyCopy.dll	Same DLL on client and server	ILocalOnly
PolicyKey	Object Provider	PolicyKey.dll	Same DLL on client and server	ILocalOnly

Notes:

1. Not needed. You do not need a separate DLL or OBJ file for these as they are used only from inside the server and going through the _S.cpp is sufficient.

For the client DLL name you need to have the _C.cpp client side binding that corresponds to the IDL file when it is introduced regardless of where the actual implementation takes place. When it is server only, the _C.cpp included by the _S.cpp is sufficient.

Additional information component creators should know

Developing the business logic that implements business object interfaces (that is, component interfaces) is done using as much of the C++ language as is appropriate for optimally supporting the interfaces. Use native language class libraries to assist when needed. Choices on how to best implement the business logic are yours to make.

Do not create separate threads. The Component Broker server environment in which the components reside is complex. Component Broker must control the threading environment. Creation of additional threads from within business logic could cause unexpected results on the Component Broker server.

Where to next?

In order to have a component that can be tested and installed into a Component Broker server, some additional steps must be taken. If you want to provide more advanced features for clients to use and to leverage additional managed object framework features as part of the implementation of your component, then see “Chapter 6. MOFW - C++ client programming model – advanced concepts” on page 133. See “Chapter 7. MOFW - C++ server programming model – advanced concepts” on page 157 for information on how to add these advanced features to a business object.

Chapter 6. MOFW - C++ client programming model – advanced concepts

This chapter includes the following topics:

- “Transactions”
- “Session Service” on page 138
- “Queries, iterations and specialized homes” on page 142
- “Using keyed reference collections” on page 148
- “Conventions and guidelines” on page 151
- “The create_object() method” on page 153

For further information on quality of service interfaces, see “Expanding the client programming interface” on page 434.

Transactions

Code segments in this section show how a client could use the Transaction Service. The example uses policy objects.

The following code segment performs the initialization that is required to use transactions.

```
#include <CosTransactions.hh>

CBSeriesGlobal::Initialize();
CORBA::ORB_ptr orb = CBSeriesGlobal::orb();

CORBA::Object_var objPtr;
try
{
    objPtr = orb->resolve_initial_references ( "TransactionCurrent" );
}
catch (CORBA::ORB::InvalidName)
{
    cout << "Error while resolving to transaction current" << endl;
    exit();
}
TransactionCurrent_var currentTransaction = TransactionCurrent::_narrow (
    objPtr );
```

You can proceed to find, create, and use a factory as shown in earlier examples. Before attempting to create a policy managed object, the begin() method can be called on the *currentTransaction* pointer to indicate the start of a

transaction. When using transactions it is advisable to begin the transaction early in the processing cycle. All method calls that might result in a database access must be always called within the scope of a transaction. For example, methods such as `findByPrimaryKeyString()` and `createFromPrimaryKeyString()` may result in database accesses, it is safer to begin the transaction before making these calls. Beginning transactions early in the processing cycle will hide differences in implementation differences between application adaptors and other specific implementations in later releases.

```
PolicyKey_var keyVar = PolicyKey::_create();
keyVar->PolicyNo(55555);
theKeyString = keyVar->toString();

// Start the transaction now
currentTransaction->begin();

:

/* Assume that the policy home is found */

mVar = myPolicyHome->createFromPrimaryKeyString(*theKeyString);
myPolicy = Policy::_narrow(mVar);
myPolicy->amount(25000.00);
```

When the managed object has been created, methods can be called on it and its data can be manipulated as desired, all within the scope of the transaction that was started in the previous example. To indicate the termination of the transaction, the `commit()` method is called and the policy is released.

```
CORBA::Boolean inValue = 1; // Report heuristic exceptions

// End the transaction now
currentTransaction->commit(inValue);
```

The previous segments assume that the data manipulations resulted in desired updates to the data. The `commit()` method completes the current transaction by making these changes to the data permanent. But what if a condition occurred in which the client wanted to end the transaction without any changes?

The following code example shows how the `rollback()` method can be used to terminate a transaction without an update to the data. Two policies are created and have their data set. An if test determines whether the changes are committed or rolled back.

```
Policy_var aPolicy;
Policy_var aPolicy2;
IManagedClient::IManageable_var moVar;
IManagedClient::IManageable_var moVar2;
ByteString_var theKeyString;
IExtendedLifecycle::FactoryFinder_var myFinder;
CORBA::Object_var it;
```



```

ManagedClient::IHome_var policyHomeVar;
float inAmount;
CORBA::Boolean inValue;

CBSeriesGlobal::Initialize();
CORBA::ORB_ptr orb = CBSeriesGlobal::orb();

CORBA::Object_var objPtr;
try
{
    objPtr = orb->resolve_initial_references ( "TransactionCurrent" );
}
catch (CORBA::ORB::InvalidName)
{
    cout << "Error while resolving to transaction current" << endl;
    exit();
}
TransactionCurrent_var currentTransaction =
    TransactionCurrent::_narrow ( objPtr );
currentTransaction->set_timeout(600);

it = CBSeriesGlobal::nameService()->resolve_with_string(
    "host/resources/factory-finders/host-scope");
myFinder = IExtendedLifeCycle::FactoryFinder::_narrow(it);
it = myFinder->find_factory_from_string(
    "PolicyDefaultTransDB2.object interface");
policyHomeVar = IManagedClient::IHome::_narrow(it);

PolicyKey_var theKey = PolicyKey::_create();
theKey->policyNo(12345);

theKeyString = theKey->toString();

// Start the transaction now
currentTransaction->begin();

moVar = policyHomeVar->createFromPrimaryKeyString(theKeyString);
theKey->policyNo(99999);
theKeyString = theKey->toString();
moVar2 = policyHomeVar->createFromPrimaryKeyString(theKeyString);
aPolicy = Policy::_narrow(moVar);
aPolicy2 = Policy::_narrow(moVar2);

cout << "Enter amount for first policy" << endl;
cin >> inAmount;
aPolicy->amount(inAmount);
cout << "Enter amount for second policy" << endl;
cin >> inAmount;
aPolicy2->amount(inAmount);

if(aPolicy->amount() == aPolicy2->amount())
{
    // End transaction now with no changes to data
    try
    {

```

```

        currentTransaction->rollback();
    }
    catch(CosTransactions::NoTransaction)
    {
        // No Transaction to rollback. ...
    }
    catch(...)
    {
        // Rollback failed ...
    }
}
else
{
    // End transaction now and update data
    currentTransaction->commit(inValue);
}

```

Transactions, exceptions, and time-outs

“Transactions” on page 133 provides a basic usage view of the transactional interfaces. This section provides additional guidelines on application structuring in the area of time-outs and exception handling.

Consider the following code example to explain why the code is structured in this way:

```

// get current          See note 1
CBSeriesGlobal::Initialize();
CORBA::Object_var objPtr;
try
{
    objPtr = orb->resolve_initial_references ( "TransactionCurrent" );
}
catch (CORBA::ORB::InvalidName)
{
    cout << "Error while resolving to transaction current" << endl;
    exit();
}
TransactionCurrent_var current = TransactionCurrent::_narrow ( objPtr );
current->set_timeout(180);    // See note 3
// you could have a loop starting here...
try
{
    current->begin();        // See note 2
    .... do work
    if (businessLogicSaysCommit)
        current->commit(1);    // See note 2
    else
    try                    // See note 4
    {
        current->rollback();
    }
    catch(...)
    {

```

```

        // Rollback failed ...
    }
    catch (const ::CORBA::SystemException &se)
    {
        // See note 5
        cout << "System Exception" << se.id() << endl;
        try
        {
            current->rollback();
        }
        catch(...)
        {
            // Rollback failed ...
        }
    }
    catch (const ::CORBA::UserException &ue)
    { // See note 6
        cout << "System Exception" << ue.id() << endl;
        try
        {
            current->rollback();
        }
        catch(...)
        {
            // Rollback failed ...
        }
    }
    catch(...)
    {
        // See note 7
        try
        {
            current->rollback();
        }
        catch(...)
        {
            // Rollback failed ...
        }
    }
    // do whatever you want to do next..
}

```

What does all of this mean? In the previous example, the numbers in the comments correspond to the following list items:

1. Get the current. This is the same as always. Get it once for performance reasons.
2. Begin and Commit transactions. Bracket units of work in accordance with the requirements of your business logic and domain needs. In other words, if you want to be able to rollback a set of operations, put them in a begin or commit block. The issue is how much time is normally expected between the begin and commit. You must realize that there is some level of resource locking that goes on while there is a transaction in flight. Shorter transactions will increase throughput, and so on.

3. This is an important line of code. This code signals to the server that this is how long (in seconds) the server should wait before taking matters into its own hands. This means that the server will roll back without any user code having specified rollback. If a time-out is not specified using this method call, a platform-dependent value is used. This value is configurable using the System Management interfaces.
4. Rollback can be issued at any time based on the needs of the business logic. This undoes changes made since the last begin.
5. When a system exception occurs, you should code a rollback() in the catch block or some routine that the catch block calls.
6. When a CORBA user exception occurs, you should code a rollback() in the catch block or some routine that the catch block calls. This exception block would include exceptions such as IManagedClient::INoObjectWithKey.
7. The exceptions that make it to the final catch block, by definition, were thrown from within the client program. Anything that flowed over a wire would have been remapped to a CORBA exception and would have been caught in item number five. A rollback here will also work properly and release locks.

Using transactions over reference collections

You can encounter a transaction deadlock problem while running more than one instance of a program which in one transaction:

1. Adds a number of records to reference collections (or home collections)
2. Then retrieves the records either directly by creating an iterator on a collection or indirectly through queries on a collection.

This error could be reported either as CORBA::INTERNAL or CORBA::PERSIST_STORE.

To prevent transaction deadlock problems, split the operation between two transactions. Add records to collections in one transaction and commit the transaction. Then create a second transaction to retrieve the records from the collections.

Session Service



OS/390 Component Broker does NOT support Session Service.

The Session Service defines the notion of a session. It provides a mechanism for grouping a set of operations together as a logical unit. A session is conceptually similar to a transaction as defined and supported by the Transaction Service. Sessions differ from transactions in the following ways:

- A session is defined on an application scope rather than on an individual transaction scope. This provides a mechanism for checkpointing groups of persistent objects for which no externally coordinated transactional update is available.
- Sessions do not define an atomically recoverable commit scope as do transactions; sessions provide some services for checkpointing and clean conclusions of applications but do not provide the same level of application durability as transactions.
- A session can use an application profile. An application profile consists of attributes that define its properties, such as requirements for data concurrency, duration, visibility, update, execution priority, and resource dependencies.

Sessions and transactions can be used together. Transactions can be used within sessions to provide greater levels of durability within applications.

For additional information about the Session Service, see the *WebSphere Application Server Enterprise Edition Component Broker Advanced Programming Guide*.

A simple example

The following example demonstrates how a single-threaded client can begin a session, perform some work, and end the session, checkpointing any non-transactional work that occurs within the session.

```
// Get the current for this thread of execution
// from the ORB and narrow to the session current.
CORBA::Object_var objPtr;
try
{
    objPtr = orb->resolve_initial_references ( "ISessions::Current" );
}
catch (CORBA::ORB::InvalidName)
{
    cout << "Error while resolving to session current" << endl;
    exit();
}
ISessions::Current_var sessionCurrent =
    ISessions::Current::_narrow ( objPtr );
```

Set a time limit for all new sessions

The ISessions::Current interface has a setSessionTimeout() operation that enables the application to set a time limit for all sessions that are subsequently

started. The default time limit is zero. That is, sessions can run indefinitely. To set a time limit for all new sessions, invoke the `setSessionTimeout()` operation on the `ISessions::Current` object, passing the new time-out value.

```
// Set the session timeout.
sessionCurrent->setSessionTimeout(100);
```

Begin a session

You begin a session by invoking the `beginSession()` operation on the `ISessions::Current` interface. You should supply a text string representing the name of your application or the application profile under which you want the session to operate. If the specified profile cannot be found or if you specify an empty string, then a default application profile is used.

The application profile specifies certain expectations about how the session will behave. This information can be used in combination with the capabilities and policies of the running system to produce a set of execution decisions that optimize the total performance and throughput of the system. Once the session has been started you can perform any number of operations on business objects within the session. All operations invoked within the session are performed with the same session context.

```
// Begin a session context.
sessionCurrent->beginSession("LifeInsuranceApplication");

try
{
    // ... do the methods that will be executed under the
    // session ...
}
catch (ISessions::SessionResetForced)
{
    // The session was forced to reset mid-stream, probably the
    // session timeout tripped, or a session resource
    // encountered a significant error and had to force the session
    // reset.
};
```

End a session

You complete the session by invoking the `endSession()` operation on the `sessionCurrent`. Normally, you specify the *EndModeCheckpoint* end-mode with this operation. This drives all the sessionable resources used within the session to save their state changes persistently, through embedded operations on the underlying data system.

To reset the session, end it without saving any of the changes that occurred since your last checkpoint (or since the beginning of the session if you did not perform any checkpoints). Specify the *EndModeReset* end-mode with the `endSession()` request.

EndModeCheckpoint and *EndModeReset* have no bearing on any transactions issued within the session other than to ensure that the session is terminated before it ends. However, if you encounter severe errors in your processing, you can end the session with *EndModeResetForce*. This forces the session to be reset immediately, including rolling back any outstanding transactions.

```
// End the session context, including checkpointing any activity that
// occurred during the session.
```

```
try
{
    sessionCurrent->endSession(EndModeCheckpoint,1);
}
```

```
// Catch a variety of exceptions.
```

See the *WebSphere Application Server Enterprise Edition Component Broker Advanced Programming Guide* for details.

Other information

Further information on these and other topics regarding the Session Service can be found in the *WebSphere Application Server Enterprise Edition Component Broker Advanced Programming Guide*. For example, the following information is available:

Suspending and resuming sessions

There may be occasions when you want to switch the session under which you are operating. You can do this by suspending the current session and starting a new one. Later, you can resume the original session.

Explicit and implicit propagation of session context

This section describes different techniques of assigning context information for multi-threaded applications.

Checkpoint and reset a session context

Sessions can be checkpointed to save intermediate results to persistent data storage. Sessions can also be reset to revert the state of operations performed within the session.

Registering sessionable resources

An application can introduce resources and explicitly register the resources with a session.

Visibility rules

Sessions can be run concurrently. Visibility rules define how the data interactions between these concurrent sessions are defined.

Queries, iterations and specialized homes

“Chapter 4. MOFW C++ client programming model” on page 85 introduced generic homes as a facility for creating business objects and for finding business objects based on their primary key. This chapter discusses how a client of a business object could use other features that a home could support.

If the object provider has built a specialized home, the usage pattern for that home is different than that of a generic home, because a specialized home has additional methods that can be used by clients. Besides the additional methods available to the client, there is also a change in how this specialized home is found.

Using iterated homes-specific functions

Many times an application needs access to multiple objects. In “Using sets of objects” on page 95, an example of iterating through an arbitrary collection of business objects is presented. This iteration assumes that all the objects are inserted into an `IManagedReferenceCollection`.

If you wanted to examine a group of homogenous objects (for example, every `Claim`) and perform actions on those objects that met certain criteria, the Component Broker programming model extends the concept of iteration to `IHomes`. The following code segment shows the usage of an iterated home.

```
CORBA::Object_var obj;
IExtendedLifecycle::FactoryFinder_var myFinder;
IManagedAdvancedClient::IIterableHome_var policyHome;
IManagedCollections::IIterator_var theIterator;

obj = CBSeriesGlobal::nameService()->resolve_with_string(
    "host/resources/factory-finders/host-scope");
myFinder = IExtendedLifecycle::FactoryFinder::_narrow(obj);

obj = myFinder->find_factory_from_string(
    "PolicyDefaultTransDB2.object interface");
policyHome = IManagedAdvancedClient::IIterableHome::_narrow(obj);

if ( CORBA::is_nil(policyHome) )
{
    cerr << "ERROR: Application improperly configured. " <<
        "Home for Policies is not iterable." << endl;
}
else
```



```

{
    IManagedClient::IManageable_var aB0;
    // Repeat transactional setup.
    currentTransaction->begin();

    // Create a new iterator on the Claim home
    theIterator = policyHome-> createIterator();
    try
    {
        // Loop through the objects in the home
        while ( aB0 = theIterator->next() )
        {
            // Get the next element
            Policy_var aPolicy;
            aPolicy = Policy::_narrow(aB0);

            // Do something useful with it.
            cout << "Policy no = " << aPolicy->policyNo() << endl;
            cout << "Policy amount = " << aPolicy->amount() << endl;
        }
    }
    catch (...)
    {
        cerr << "ERROR: Problem occurred using iterator." << endl;
        try
        {
            current->rollback();
        }
        catch(CosTransactions::NoTransaction)
        {
            // No Transaction to rollback. ...
        }
        catch(...)
        {
            // Rollback failed ...
        }
    }
    // After iterating over the entire collection, the iterator is no
    // longer needed. Remove it.
    theIterator->remove();
    currentTransaction->commit(inValue);
}

```

In addition to the `next()` method, `IManagedCollections::IIterator` provides the following methods with similar functionality (Note that `next()` is much more concise.): `hasMoreElements()`, `more()`, `nextElement()`, `nextOne()`.

While the previous examples are functionally identical (with identical performance), there is another interface for iterating which is functionally similar, but with different performance characteristics. The `nextN()` method is similar to the `nextOne()` method, except that instead of returning only one element, `nextN()` returns as many elements as you request.

This interface might be useful if, for example, an application wants to display information about a fixed number of business objects at a time (limited by the size of a window on a screen). Here is an example of how to use this interface in the same application as the previous example:

```
char c;
ICollectionsBase::MemberList_var theB0list;

// Loop through the objects in the home.
while ( theIterator->nextN(5, theB0list) || theB0list->length() > 0 )
{
    for ( int n = 0; n < theB0list->length(); ++n)
    {
        // Narrow the next element obtained by nextN()
        aPolicy = Policy::_narrow(theB0list[n]);

        // Do something useful with it
        cout << "Policy # " << aPolicy->policyNo();
        cout << "amount is " << aPolicy->amount() << endl;
    }
    // Having displayed as many policies as would fit in the window,
    // the application now waits for user input before getting the
    // next N policies and displaying them.
    cout << "press any key (and hit enter) to continue." << endl;
    cin >> c;
}

// Because the nextN() call might not have been able to return each time.
```

The previous examples illustrate several ways to iterate through the objects in a home. To understand iteration better, a few general notes about iterators follow:

- When an iterator is created, it is positioned so as to precede the first element.
- The various types of next methods (`next()`, `nextElement()`, `nextOne()` and `nextN()`) all move the position of the iterator forward, then return the element at the current position.
- There is a `current()` method that returns the element at the current position without moving the position of the iterator forward. Because the initial position of an iterator precedes the first element, it is an error to invoke `current()` prior to invoking one of the next methods.
- There is a `reset()` method which moves the position back to its initial position.
- A home that is iterated is not guaranteed to return its elements in the same order on successive iterations, but it is guaranteed to return each element once only on a given iteration.

Members of Component Broker homes are always accessible through at least one key (the primary key), but may or may not be iterable. If the collection

is based on or wrappers a relational or object-oriented database, then both keys and iterator can be supported. The same is true for most data back-ends (such as files, IMS DL/1 or CICS File Control). If the collection wrappers a set of applications (such as the Customer Management Application in “Chapter 2. Personal life insurance application example” on page 31), iteration might not be supported by the application. Therefore, the wrapping collection cannot be iterated. Check with your System Administrator to see if your home supports iteration.

Because the LifeCycle object service factory finder interface is used to find homes and this interface has no way of telling an iterable home from a generic home, a client must be properly configured if its function depends on access to an iterable home. The client can protect itself against misconfiguration by verifying that the home returned by the factory finder supports the IManagedAdvancedClient::IIterableHome interface. A verification example is:

```

policyHome = IManagedAdvancedClient::IIterableHome::_narrow(obj);
if ( CORBA::is_nil(policyHome) )
{
    // Insert error message here to be displayed here.
}

```

Using queryable homes-specific functions

Iterating over the business objects in an IIterableHome is fine for some applications. Specifically, if the goal is to perform some operation on every element in the home, then iteration is a good approach. However, if the reason for iterating is to select a subset of all the business objects in the home (and then work with only that subset), then there is a more efficient way to accomplish the same thing: query.

Here the iteration example is used again; however, instead of iterating against all of the objects in a home, you select only the objects in which you are interested (by evaluating a query), and then you iterate over the subset in which you are interested. Because it is sometimes possible for the query to be performed in the database (without bringing into memory every object in the home), using query offers the possibility of significant performance increases over iterating against every object in the home.

```

CORBA::Object_var obj;
IExtendedLifeCycle::FactoryFinder_var myFinder;
IManagedAdvancedClient::IQueryableIterableHome_var policyHome;
IManagedCollections::IIterator_var theIterator;

obj = CBSeriesGlobal::nameService()->resolve_with_string(
    "host/resources/factory-finders/host-scope");
myFinder = IExtendedLifeCycle::FactoryFinder::_narrow(obj);

```

```

obj = myFinder->find_factory_from_string("PolicyDefaultTransDB2.object
interface");
policyHome = IManagedAdvancedClient::IQueryableIterableHome::_narrow(obj);

if ( CORBA::is_nil(policyHome) )
{
    cerr << "ERROR: Application improperly configured. "
        << "Home for Policies is not queryable."
        << endl;
}
else
{
    // Evaluate a query on the Policy home. The results of the query
    // are returned in the form of an iterator.
    int minPolicyNo = 3;
    char theQuery[80];
    sprintf(theQuery,"policyNo>%d",minPolicyNo);
    /* specifies the query itself */

    // Set up for transactions
    currentTransaction->begin();
    theIterator = policyHome->evaluate(theQuery);
    try
    {
        // Loop through the results of the query
        while ( aBO = theIterator->next() )
        {
            // Narrow the element obtained by nextOne()
            aPolicy = Policy::_narrow(aBO);

            // Do something useful with it.
            cout << "Policy no = "
                << aPolicy->policyNo()
                << endl;
            cout << "Policy amount = "
                << aPolicy->amount()
                << endl;
        }
    }
    catch (...)
    {
        cerr << "ERROR: Problem occurred using iterator."
            << endl;
        try
        {
            current->rollback();
        }
        catch(CosTransactions::NoTransaction)
        {
            // No Transaction to rollback. ...
        }
        catch(...)
        {
            // Rollback failed ...
        }
    }
}

```

```

    }
    // After iterating over the entire collection, the iterator is no
    // longer needed. Remove it.
    theIterator->remove();
    currentTransaction->commit(inValue);
}

```

The query language used to express the query is SQL with extensions and provisions for being able to query against objects (instead of data only in a database).

Using atomic transactions with query evaluator

The programming model for method duration transactions (called atomic or automatically start a new transaction) occurs when a method is invoked on the object; the method then will be wrapped within a transaction. That is, prior to invoking the method, a transaction will automatically start. When the method returns the transaction will be committed. The programming model has been extended so that creating or finding an object that is in a container (supporting method duration transactions then creation or finding) will also be wrapped in a transaction.

The programming model has not been expanded to include the Query Service.

Important: A transaction must have been started prior to using the Query Service.

For additional information about the Query Service, see the *WebSphere Application Server Enterprise Edition Component Broker Advanced Programming Guide* .

More on iterators

This information addresses the following:

- how iterators over homes and collections are used
- how query iterators work
- how atomic containers are configured
- side effect of end transaction

Managed object iterators over home collections, iterators over persistent reference collections, query iterators over persistent managed objects and query data array iterators over persistent objects become inappropriate at the end of the current transaction in which they were created.

When designing objects that will be configured into an atomic container, the object interface should not have attributes or methods which return to the client any of the above types of iterators.

If the object interface does have such iterators, they can be retrieved by a client only if the client explicitly starts a transaction and retrieves and uses the iterator in the same transaction scope.

Note the following example:

```
interface X
{
    IMangedCollections::Iterator methodX();
}
```

If this interface is implemented by an MO, configured into an atomic container and if my client does not start a transaction, the following statements will result in an exception because the iterator is being used outside the scope of a transaction.

```
IMangedCollections::Iterator_var i = x_ptr -> methodX();
IMangedClient::IMangeable_var mo = x_ptr->next();
```

However if the client does the following, it will be valid because the iterator is being retrieved and used in the same transaction scope.

```
currentTransaction->begin();
IMangedCollections::Iterator_var i = x_ptr -> methodX();
IMangedClient::IMangeable_var mo = x_ptr->next();
```

Using keyed reference collections

“Using sets of objects” on page 95 describes how to use a reference collection (that is, `IMangedCollections::IReferenceCollection`) to maintain references to a set of managed objects. If the elements in a reference collection require keyed access (that is, Hashtable-like access), a keyed reference collection can be used instead. Keyed reference collections are defined by the `IMangedCollections::IKeyedReferenceCollection` interface.

Keyed and non-keyed reference collections have many similarities. Both the `IReferenceCollection` and `IKeyedReferenceCollection` interfaces derive from the `ManagedCollections::ICollection` common base interface that defines a number of methods that apply to both keyed and non-keyed collections. The interfaces include many commonly used methods such as `containsElement()`, `isEmpty()`, `numberOfElements()`, and `createIterator()`. In addition, keyed reference collections are created using the same specialized home (that is, `IMangedCollections::ICollectionHome`) as non-keyed collections. The following code segment illustrates the creation of a keyed reference collection:

```
CORBA::Object_var obj;
IMangedCollections::ICollectionHome_var kcHome;
IMangedCollections::ICollection_var cc;
IMangedCollections::IKeyedReferenceCollection_var kc;
```

```

obj = myFinder->find_factory_from_string(
    "IManagedCollections::IKeyedReferenceCollection.object
    interface/TransientKeyedReferenceCollectionFactory.object home");
kcHome = IManagedCollections::ICollectionHome::_narrow(obj);
cc = kcHome->createCollection();
kc = IManagedCollections::IKeyedReferenceCollection::_narrow(cc);

```

This code is identical to that shown in “Using sets of objects” on page 95 to create non-keyed collections except that the `IReferenceCollection` interface name is replaced by `IKeyedReferenceCollection`.

The difference between keyed and non-keyed collections is the way objects are added and accessed. Instead of calling `IReferenceCollection::addElement()`, the `IKeyedReferenceCollection::addElementByString()` method is used to add elements to the collection. The `addElementByString()` method requires two arguments, the object to be added and a stringified key (that is, a `ByteString`) that is used for identifying the object in the collection. The key can be any appropriate subclass of `IManagedLocal::IKey`, either general purpose (for example, `StringKey`), application specific (for example, `PolicyHolderIdentifierKey`), or, for that matter, the primary key (`PolicyHolderPrimaryKey`).

For example, assume a general purpose `IKey` subclass has been defined:

```

interface StringKey : IManagedLocal::IKey
{
    attribute string value;
    #pragma meta StringKey localonly
}

```

Using this key, elements can be added to the keyed reference collection as follows:

```

// Get some policy holders
PolicyHolder_var policyHolderJohn;
PolicyHolder_var policyHolderKatherine;

// Create a key object
StringKey_var theKey = StringKey::_create();
ByteString_var keyString;

// Add policyHolderJohn to the collection with key "John"
theKey->value("John");
keyString = theKey->toString();
kc->addElementByString(policyHolderJohn, keyString);

// Add policyHolderKatherine to the collection with key "Katherine"
theKey->value("Katherine");
keyString = theKey->toString();
kc->addElementByString(policyHolderKatherine, keyString);

```

An element can be retrieved using the `getElementByString()` method. For example, the following code segment retrieves the policy holder "John":

```
theKey->value("John");
keyString = theKey->toString();
IManagedClient::IManageable_var mo = kc->getElementByString(keyString);
PolicyHolder_var thePolicyHolder = PolicyHolder::_narrow(mo);
```

To remove an element from the collection, the `removeElementByString()` method is used:

```
kc->removeElementByString(keyString);
```

Iterating through the elements in a keyed reference collection is done in exactly the same way as for non-keyed collections. In fact, if the collection is referenced using the `IManagedCollection::ICollection` base interface, as shown in the following code segment, the same code can be used to iterate over elements of either keyed or non-keyed collections. Note that if the variable *theCollection* were of type `ICollectionsBase::IMIterable_var`, the same code segment could be used to iterate over an iterable home.

```
// Get a keyed or non-keyed collection from somewhere
IManagedCollections::ICollection_var theCollection;

IManagedCollections::Iterator_var theIterator =
    theCollection->createIterator();

IManagedClient::IManageable_var theB0;
while ( theB0 = theIterator->next() )
{
    if ( theB0->is_a("PolicyHolder") )
        // Send him a bill
    if ( theB0->is_a("Beneficiary") )
        // Send him a check
}

try
{
    //Loop through the collection. The "nextElement" method advances
    //the iterator to the next element (on the first invocation, this
    //will advance to the first element) and then return the element
    //pointed to by the iterator.

    while ( theIterator->more() )
    {
        theB0 = theIterator->next();
        if ( theB0->is_a("PolicyHolder") )
            // Send him a bill
        if ( theB0->is_a("Beneficiary") )
            // Send him a check
    }
    catch (...)

    theIterator->remove();
```


Each element in the keyed or non-keyed collection is accessed only once and in no defined order.

A number of other methods are available on the `IKeyedReferenceCollection` interface:

`containsKeyString()`

Determines if an element with a given key exists.

`getElementKeyString()`

Determines the key of a given element.

`replaceElementWithKeyString()`

Replaces an element (that is associated with a specified key) with another object.

With these and the other keyed reference collection methods, an arbitrary set of keyed references to Component Broker managed objects can be maintained easily.

Conventions and guidelines

This section provides additional information on what is happening on the server when client programs are invoking methods on the business objects. This section does not describe new interfaces but provides additional technical details on what actually occurs. Some of these topics are guidelines or coding techniques that can be useful when interacting with the Component Broker server.

The Component Broker server plays a unique role among application servers: it serves objects. A database server, on the other hand, serves data and a file server serves files. However, the Component Broker server also interacts with database or file servers in cases where these are the chosen options for persistence. In a database server, interaction with the data is direct and the semantics of creating, deleting, and finding are straightforward because you use the methods to the resource manager directly. In the case of an object server, from a programming perspective you encapsulate interaction with a database to get persistence. However, to provide this encapsulation, the decisions about how the database is used are unknown to the clients. Sometimes expectations are not met.

The following sections discuss the interaction patterns that the Component Broker has with various resource managers and the coding patterns that you can use to meet various client requirements.

Finding persistent objects

When you want to find an object using the `findByPrimaryString()` method on the `IHome`, Component Broker follows a specific algorithm. This algorithm is:

1. Convert parameter into internal key format.
2. Look in the container cache of the active objects.
3. If not found, go to the database or database cache and look.
4. Return the object reference as soon as it is found.

The programming implications of this are:

- If the object exists, the object reference returned is valid and is ready to use.
- If the object was being cached by the container or the cache manager and had subsequently been deleted by a non-object program, an exception is thrown when the object reference is used.
- The probability of getting an exception when using the returned reference is directly proportional to the amount of deleting that is done by existing non-Component Broker applications that are running concurrently with Component Broker applications.

Creating persistent objects

When you want to create an object using the `createFromPrimaryString()` or other create methods, Component Broker follows a specific algorithm with respect to the creation.

1. Create the managed object and its parts.
2. Put it in the container's cache of active objects.
3. Return the object reference to the client.
4. Wait for transaction commit().
5. When the transaction commits, insert the row in the table.

The implications of this algorithm for you are as follows. If the row already exists in the underlying database, an exception is thrown when the transaction completes.

In many cases the exception is desired and should be planned for accordingly. To alter this behavior for create, you can either:

- Do a `findByPrimaryString()` before issuing the `createFromPrimaryString()`. Remember from the previous discussion that `findByPrimaryString()` looks in the database if necessary to find objects and in this instance would return `notFound`. If there is a slim chance that the object already exists, the performance penalty for the `findByPrimaryString()` might be too high. Select a technique based on the needs of the application.

- Bracket `createFromPrimaryKeyString()` in a transaction by itself. This ensures that the exception for `alreadyExists` is thrown before operations on the newly-created object are started.

These techniques do not reduce the amount of written code but are alternatives for structuring client code.

The `create_object()` method



The following `create_object()` method is platform-dependent and does NOT apply to OS/390 Component Broker.

You can use the `create_object()` method to replicate the function of the `createFromPrimaryKeyString()` and `createFromCopyString()` method. The reason for using the `create_object()` method is to be OMG compliant.

The `create_object()` method takes two parameters:

key Consists of:

kind The object interface for BOIM homes.

id The name of the business object interface, such as policy. It is the same value with which the home is configured for managed object class (interface name) for the particular managed object image in question.

criteria

A name-value pair. The name parameter tells the home what kind of creation is to take place.

BOIM homes support two names:

primary key string

A signal to use the contents of the value to perform a `createFromPrimaryKeyString()` using this data.

copy string

A signal to use the contents of the value to perform a `createFromCopyString()` using this data.

An example follows:

```
::CosLifeCycle::Key k;
::CosLifeCycle::Criteria crit;
```

```

k.length(1);// set the length of the key to 1

char* idString=new char[7];// allocate storage for the id string
strcpy(idString, "Policy");// set it to the object class interface name
char* kindString=new char[17];// allocate storage for the kind string
strcpy(kindString,"object interface");//set it to object interface
k[0].id=idString;// assign the id of the key
k[0].kind=kindString;// assign the kind of the key

crit.length(1);// set the size of the criteria to 1
char* critname = new char[19];// allocate space for criteria name
strcpy(critname,"primary key string");

// Do a createFromPrimaryKeyString using create_object
crit[0].name=critname;//assign the name

PolicyKey_var theKey = PolicyKey::_create();// build a key
theKey->policyNo(12345);// fill in needed information
ByteString* theKeyString = theKey->toString();//create a bytestring
crit[0].value<<=(*theKeyString);

// Assume that you already found the policy home in "myHome"
try
{
    ::CORBA::Object_var policy = myHome->create_object(k,crit);
    // call create_object
    // notice that a CORBA::Object is returned
}
catch(CosLifeCycle::NoFactory)
{
    // The key passed in does not match that of the home. Check
    // that the kind is "object interface" and the id is the same
    // as that of the managed object interface in the MO Image.

    // Insert recovery code goes here.
}
catch(CosLifeCycle::InvalidCriteria &ic)
{
    // The criteria length was 0.

    // Insert recovery code goes here.
}
catch(CosLifeCycle::CannotMeetCriteria &cmc)
{
    // Recovery:
    // 1) You may have sent in two name-value pairs that could both
    //    create objects, but the home can only create one object at a time.
    // 2) The ByteString may have failed to internalize into the Key
    //    for this type of object.
    // 3) A duplicate or invalid key error may have occurred.

    // In all cases, check the error log.
}

```

```

catch(CORBA::Exception)
{
    // The home may not have been configured.
    // An unknown error may have occurred.

    // Insert recovery code goes here.
}

```

To perform a `createFromCopyString`, the code is similar:

```

:

PolicyCopy_var theCopy = PolicyCopy::_create();
                        /* builds a key */
// Fill in the needed information.
theCopy->policyNo(12345);
// Fill in the other values.
theCopy->amount(10000.00);
theCopy->premium(250.00);
ByteString *theCopyString = theCopy->toString();
                        /* creates a bytestring */
crit[0].value<<=(*theCopyString);
CORBA::Object_var myObject = create_object(k,crit);

:

```

Chapter 7. MOFW - C++ server programming model – advanced concepts

The Component Broker server programming model is based on programming by framework completion. Component Broker introduces its APIs as sets of classes and frameworks. Developers implement their business functions by defining business objects that subclass from Component Broker frameworks and use the frameworks in their implementation.

This chapter provides business object builders with additional options for providing function in business objects. Some of the material in this chapter offers alternative ways of accomplishing some of the tasks described in “Chapter 5. MOFW - C++ server programming model” on page 103.

Some of the topics in this chapter assume the presence of specific application adaptors and other features within Component Broker. In other words, some of the techniques described might increase the cost of porting or targeting business objects and associated applications to other back-end databases and resource managers.

Extending a business object

“Chapter 5. MOFW - C++ server programming model” on page 103 presents a basic model for building business objects. Most of the discussion and examples in that chapter center around implementing a new business object interface. Careful examination of the examples shows that there are several meaningful layers of inheritance in the business objects that were presented. Often multiple levels of domain inheritance exist and need to be implemented.

This chapter addresses the addition of a subclass to existing business objects using data object inheritance as the implementation technique. This is the model supported by Component Broker through Object Builder. This is not as simple as adding a single class. This chapter revisits each of the steps used to develop a business object in the context of adding another subclass of domain functionality. These steps are:

1. Developing an interface to the business object.
2. Choosing an inheritance pattern.
3. Implementing the business object methods.
4. Implementing the methods required by the MOFW interfaces.
5. Implementing the key classes.

6. Implementing the copy helper classes.

When you create a child component (that is, a component that inherits behavior or data from another component in your application), the child component objects generally inherit from their equivalent parent objects:

- The child business object file must include the parent business object file.
- The child business object interface must inherit from the parent interface.
- The child key and copy helper can inherit from their equivalents in the parent component, or they can contain selected attributes of the parent interface, without inheriting from the parent key or copy helper.
- The child business object implementation must inherit from the parent implementation.
- The child data object interface must inherit from the parent data object interface.
- The child data object implementation must inherit from the parent data object implementation.
- The child managed object must inherit from the parent managed object.

For data inheritance to work, the type of persistence provided by the parent and child data object implementations must be the same.

Many variations involve extending a business object. The next sections give examples of an extreme case where it inherits as much interface and implementation as possible. There are variations that involve less implementation inheritance that can be extrapolated from the examples given.

Extending business object interfaces

The first step in building a business object is to construct the interface. In this case, the interface is extended. In the following example, a CarPolicy class extends the Policy.




Figure 35. Inheritance of interface for extended business object

The IDL for CarPolicy should look like the following example.

```
#include <Policy.idl>
interface CarPolicy : Policy
{
    attribute long year;
    attribute string make;
    attribute string model;
    attribute long serialNumber;
    attribute float collisionDeductible;
    attribute boolean glassCoverage;
    long riskQuotient( );
};
```

The interface looks like almost any other business object interface. However, because the Policy already inherits from IManagedClient::IManageable, the CarPolicy interface does not need to.

Essential state extensions

Extending the essential state is similar to extending the interface. The data object interface of Policy is extended as shown in the following figure.




Figure 36. Extending the interface for essential state

The IDL for this interface looks like the following example.

```
#include <PolicyDO.idl>
interface CarPolicyDO : PolicyDO
{
    attribute long year;
    attribute string make;
    attribute string model;
    attribute long serialNumber;
    attribute float collisionDeductible;
    attribute boolean glassCoverage;
    #pragma meta CarPolicyDO localonly ,abstract
};
```

“Data object customization and inheritance” on page 389 shows the decisions that must be made when customizing data objects from an implementation perspective.

Choosing an inheritance pattern

The choice of inheritance pattern is based on three concerns:

- Identity: whether parent and child have the same identity (that is, they share the same key)
- Performance tradeoffs: whether performance or space efficiency is more important.
- Form of persistence: whether the parent has data to be persisted, and where and how the parent’s and child’s data is persisted.

If the parent and child have different keys, you should probably use the *Attributes Duplication Pattern*. This means that the child’s datastore provides persistence for all of its data, including inherited data. The parent’s datastore only provides persistence for instances of the parent, never for instances of the child. If you do not use the overriding persistence pattern, the parent’s datastore will have two primary keys: the parent’s key for the parent’s data, and the child’s key for the child’s inherited data. It then becomes problematic to determine which data belongs to which object type.

If the parent and child have the same key, you can choose between the *Key Duplication Pattern* and the *Single Datastore Pattern*. The Key Duplication Pattern will generally be more efficient in its use of space (because the persistent objects for each component contain only the data required for that component), and the Single Datastore Pattern will generally provide faster look-up time (because both local and inherited data are mapped to the same persistent object and underlying datastore).

If the parent and child are both persisted in a database, you can compromise between the Key Duplication Pattern and the Single Datastore Pattern, by using the *Single Datastore with Views Pattern*. This pattern uses unique persistent objects for retrieval (the Key Duplication Pattern), and a shared persistent object for all other uses (the shared persistence pattern). This pattern is based on views of the underlying database table, and requires that there be some unique attribute of the child that can be used to select appropriate views of the database.

This example will use the Key Duplication Pattern for illustrations purposes. This is the default pattern supported by Object Builder.

See the “Inheritance” section in the *WebSphere Application Server Enterprise Edition Component Broker Application Development Tools Guide* for further information on the inheritance patterns.

Implement the additional business logic

Next, you must add the implementation of the business logic for the additional methods necessary in the subclass. Introducing another interface for the CarPolicyBO is shown in the following figure and in the following example.




Figure 37. Extending the business logic interface

```
#include <CarPolicy.idl>
#include <PolicyBO.idl>

interface CarPolicyBO : CarPolicy, PolicyBO
{
};
```

Consider using the same pattern (delegating versus caching) in the subclass as was used in the base class.

The actual implementation interface has the inheritance of the PolicyBO built right into it, as shown in the following example.

```
class CarPolicyBO_Impl : public virtual ::CarPolicyBO_Skeleton
,public virtual PolicyBO_Impl
{
public:

CarPolicyBO_Impl();

::CORBA::Long year();
::CORBA::Void year( ::CORBA::Long year);
char* make();
::CORBA::Void make(const char* make);
char* model();
::CORBA::Void model(const char* model);
```

```

::CORBA::Long serialNumber();
::CORBA::Void serialNumber( ::CORBA::Long serialNumber);
::CORBA::Float collisionDeductible();
::CORBA::Void collisionDeductible( ::CORBA::Float collisionDeductible);
::CORBA::Boolean glassCoverage();
::CORBA::Void glassCoverage( ::CORBA::Boolean glassCoverage);
virtual ::CORBA::Long riskQuotient();

virtual ::CORBA::Void initForCreation(
    ::IManagedServer::IDataObject_ptr theDO);
virtual ::CORBA::Void uninitForDestruction();
virtual ::CORBA::Void initForReactivation(
    ::IManagedServer::IDataObject_ptr theDO);
virtual ::CORBA::Void uninitForPassivation();
virtual ::CORBA::Void syncToDataObject();
virtual ::CORBA::Void syncFromDataObject();
virtual ::CORBA::Void externalize_to_stream(
    ::CosStream::StreamIO_ptr targetStreamIO);
virtual ::CORBA::Void internalize_from_stream(::CosStream::StreamIO_ptr
sourceStreamIO, ::CosLifeCycle::FactoryFinder_ptr there);
virtual ::ByteString* getPrimaryKeyString();

protected:

private:

CarPolicyDO* iDataObject;
::CORBA::Void initializeState();
};

```

This section is focused on business logic. The entire implementation interface is shown in the previous example. How the non-business logic or MOFW framework methods are to be handled is shown in upcoming sections.

The getters and the setters are implemented as they would be in any business object. The only difference in the other methods is that they can choose to access or use state data from the parent class. The following method implementation shows utilization of state data from both the Policy and CarPolicy class.

```

::CORBA::Long CarPolicyBO_Impl::riskQuotient ( )
{
    if ( iGlassCoverage() )
    {
        if ( iYear() > 1960)
            return 1;
        else if ( policyNo() < 1000)
            return 2;
        else
            return 5;
    }
    else
        if ( amount() > 1000 )

```

```
        return 10;
    else
        return 100;
}
```

In the previous code segment, the accessor or getter methods are used to access the state data that is needed from Policy. This is the most encapsulated way of doing this. However, if the members are declared as protected instead of private in the PolicyBO_Impl class, then direct access would be possible.

The previous example shows the caching data object case. In the delegating case, the same CarPolicyDO_var would be used to access all of the state data regardless of whether or not it resides in the CarPolicy or the Policy.

Meet the MOFW IManageable requirements

The following methods are required to be overridden in the simple case:

- getPrimaryKeyString
- getHandleString
- externalize_to_stream
- internalize_from_stream

Note: getHandleString() is not actually required to be overridden in the simple case, unless this object is going to be the target of a one-to-one relationship from another object.

The next section discusses considerations for the inheritance case.

getPrimaryKeyString

This method implementation in the inheritance case depends on the decision made about the key to be used for the new subclass. That is, what will the key be for the CarPolicy? If a new key class is introduced, then this method must be overridden. If the existing key class can be used, that is, use the PolicyKey for the CarPolicy class, then this method does not need to be overridden and can be inherited directly.

See “More key classes” on page 167 for help in determining if another key class is needed.

getHandleString

A default implementation of this method is provided by the managed object framework. Overriding this method in a subclass of a business object should be done based on the same criteria that are used to determine if an override is needed in the base class.

externalize_to_stream

This method needs to be implemented. The general direction is to call the parent class method and add those things which are necessary for the subclass.

```
::CORBA::Void CarPolicyBO_Impl::externalize_to_stream(
    ::CosStream::StreamIO_ptr targetStreamIO )
{
    // Insert Method modifications here
    PolicyBO_Impl::externalize_to_stream(targetStreamIO);

    targetStreamIO->write_long(iDataObject->year());
    targetStreamIO->write_string(iDataObject->make());
    targetStreamIO->write_string(iDataObject->model());
    targetStreamIO->write_long(iDataObject->serialNumber());
    targetStreamIO->write_float(iDataObject->collisionDeductible());
    targetStreamIO->write_boolean(iDataObject->glassCoverage());

    // End Method modifications here
}
```

While the previous example is for a caching business object, the pattern for a delegating data object is similar.

internalize_from_stream

This method needs to be implemented. The general direction is to call the parent class method and then add those things which are necessary for the subclass.

```
::CORBA::Void CarPolicyBO_Impl::internalize_from_stream(
    ::CosStream::StreamIO_ptr sourceStreamIO,
    ::CosLifecycle::FactoryFinder_ptr there)
{
    // Insert Method modifications here
    PolicyBO_Impl::internalize_from_stream(sourceStreamIO, there);

    iDataObject->year(sourceStreamIO->read_long());
    iDataObject->make(sourceStreamIO->read_string());
    iDataObject->model(sourceStreamIO->read_string());
    iDataObject->serialNumber(sourceStreamIO->read_long());
    iDataObject->collisionDeductible(sourceStreamIO->read_float());
    iDataObject->glassCoverage(sourceStreamIO->read_boolean());

    // End Method modifications here
}
```

While the previous example is for a caching business object, the pattern for a delegating data object is similar.

MOFW requirements – IManagedServer

The following methods must be overridden in the simple case:

- `initForCreation`
- `uninitForDestruction`
- `initForReactivation`
- `uninitForPassivation`
- `syncFromDataObject`
- `syncToDataObject`

The next section discusses considerations for the inheritance case.

initForCreation

Code this method following the same guidelines that were specified for building business objects independent of this inheritance case. If the subclass introduces additional state data, as the `CarPolicy` example does, then the data object must be set into a data member of the object.

If the subclass does not introduce additional state data, it does not need to save a pointer to the data object that is passed as a parameter. In the delegating case, the subclass might still want to hold a pointer to the data object rather than going through parent class get and set methods. However, regardless of whether the subclass introduces additional state or not, the parent class' `initForCreation()` method must be called at the beginning of the subclass' `initForCreation()` method.

```
    ::CORBA::Void CarPolicyBO_Impl::initForCreation(  
        ::IManagedServer::IDataObject_ptr theDO)  
    {  
        // Insert Method modifications here  
        PolicyBO_Impl::initForCreation(theDO);  
  
        iDataObject = CarPolicyDO::_narrow(theDO);  
  
        // End Method modifications here  
    }
```

uninitForDestruction

Implement this method following the guidelines described previously. It is also a good practice to call the parent class' `uninitForDestruction()` at the beginning of the subclass' `uninitForDestruction()` method if needed.

initForReactivation

Code this method following the same guidelines that were specified for building business objects independent of this inheritance case. If the subclass

introduces additional state data, as the CarPolicy example does, then the data object must be set into a data member of the object.

If the subclass does not introduce additional state data, it does not need to save a pointer to the data object which is passed as a parameter. In the delegating case, the subclass might still want to hold a pointer to the data object rather than going through parent class get and set methods. However, regardless of whether the subclass introduces additional state or not, the parent class' `initForReactivation()` method must be called at the beginning of the subclass' `initForReactivation()` method.

```
::CORBA::Void CarPolicyBO_Impl::initForReactivation(  
    ::IManagedServer::IDataObject_ptr theDO)  
{  
    // Insert Method modifications here  
    PolicyBO_Impl::initForReactivation(theDO);  
  
    iDataObject = CarPolicyDO::_narrow(theDO);  
  
    // End Method modifications here  
}
```

uninitForPassivation

This method should be implemented following the normal guidelines described earlier. It is also good practice to call the parent class' `uninitForPassivation()` at the beginning of the subclass' `uninitForPassivation()` method if needed.

syncFromDataObject

The pattern for implementing this method is similar to that used for the `internalize_from_stream()` method. The parent method must be called first, followed by the code necessary to prime the subclassing business object cache with values from the data object.

This method is also a good place for any initialization logic that is dependent on the presence of an active and usable data object. The act of loading up the business object cache ensures that the state data of the object is ready to be used.

syncToDataObject

The pattern for implementing this method is similar to that used for the `externalize_to_stream()` method. The parent method must be called first, followed by the code necessary to push the subclassing business object cache values back into the data object.

It is a good practice to always implement all of these `IManagedServer` methods even if they are not explicitly required based on the particular subclass being introduced. This practice results in consistently generated code which is less prone to error if future changes are made to the subclass that introduces new attributes.

More key classes

The MOFW requires that every managed object have a primary key class associated with it. When extending a business object, there are several possibilities you can use as a primary key class. The business object subclass can:

- Use its base class' primary key class.
- Extend its base class' primary key class.
- Introduce its own primary key class.

The simplest approach is to reuse the existing key. This approach is applicable if the attributes that uniquely identify objects of the subclass are the same ones that identify objects of the base class. This is most likely the case if the subclass introduces no additional state (that is, the subclass only re-implements the base class methods, or introduces new methods). Even if the subclass introduces an additional state beyond that of the base class, this additional state might not contribute anything to object identity.

If the subclass introduces additional state, some of which, combined with the key attributes of the base class, is used to uniquely identify objects of the subclass, then the best approach is to extend the base class' primary key class. When extending the primary key class of a base class, interface and implementation inheritance can be used. Implementation inheritance allows for the reuse of the base class key functionality (specifically, its getters and setters, as well as its streaming code).

Finally, if the attributes that uniquely identify objects of the subclass are neither the same ones as the base class, nor a superset of the base class' key attributes, then the subclass must introduce its own primary key class.

However, if the existence of new state data has altered the way in which the object is uniquely identified, then a new primary key class is necessary.

Note: If the subclass does not use the base class' primary key class, then the subclass' data object needs to be able to handle this extended or new key class.

More copy helper classes

If the subclassing business object introduces additional state data, then a new copy helper class might be useful. The data object needs to be able to handle

this new copy helper. The copy helper can inherit interface and implementation from the base class' copy helper, or it can be written from scratch.

Extension summary

In the managed object framework based programming model, there are a number of inheritance activities to follow. Interfaces should be inherited consistently. Implementations should be inherited when the base class' implementation can be reused to some degree. The simplest model is to inherit at all levels of the MOFW architecture and add in additional business logic as necessary. Additional requirements from the MOFW are also added in the new implementation subclass.

Managed object customization and data object customization are also different for business objects that inherit from other business objects. These topics are discussed in "Data object customization and inheritance" on page 389.

Other variations to consider

There are other variations to consider. Some are restrictions and others are tips for leveraging inheritance in the MOFW-based programming environment:

- Do not change data object patterns. It is possible to change data object patterns from caching to delegating at various levels of the data object hierarchy but this adds undue complexity in most cases.
- There might be cases when the subclasser does not know the data object pattern being used by the base class.

Object relationships

Component Broker applications often require persistent relationships between business objects. For instance, in the personal life insurance sample application, a Claim object has a relationship with a Policy object representing the insurance policy against which the claim has been filed. Also, a Policy object has a relationship with Claim objects for pending claims against the insurance policy.

Relationships between objects can be described in many ways. First, there is the cardinality of the relationship. If an object has a relationship to one other object at most, the relationship is considered to be cardinality-1. On the other hand, if an object has a relationship with more than one other object at a time, the relationship is considered to be cardinality-N. In a business object, relationships to other objects are implemented as object references (cardinality-1), or as collections of object references (cardinality-N).

A relationship can also be described as *optional* or *required*. If a relationship is optional, then an object is considered to be in a valid state even when it is not related to (linked to) another object. If the relationship is required, then the object must always be linked to another object. This distinction is often combined with cardinality to form the following combinations:

Class Relationship	Cardinality	Required
An instance of class X is related to 0..1 instances of class Y	-1	No
An instance of class X is related to 1 instance of class Y	-1	Yes
An instance of class X is related to 0..n instances of class Y	-N	No
An instance of class X is related to 1..n instances of class Y	-N	Yes

Finally, a relationship can also be described in terms of ownership. An object which is related to another object might or might not be considered to be the owner of the related object. If the first object is not considered to be the owner of the second object, then the relationship is often referred to as a *uses a* relationship, as in, for example, the first object *uses* the second object. This is sometimes also called an association. On the other hand, if the first object is considered to be the owner of the second object, then the relationship is often referred to as a *has a* relationship, as in, for example, the first object *has* (or contains) the second object. This is sometimes also called an aggregation.

The cardinality-1 or cardinality-N distinction results in much more fundamental differences in the business object code than the optional or required distinction and the uses or has distinction. Cardinality-1 relationships are discussed separately from cardinality-N relationships.

Cardinality-1 relationships

The following diagram shows an example of a cardinality-1 relationship, a simple link from a Claim object to a Policy object.



Figure 38. Cardinality-1 relationship

The relationship depicted in the previous figure is that of an optional cardinality-1 *uses a* relationship. In a business object interface, a cardinality-1 relationship is declared as a CORBA attribute whose type is a reference to the interface of the target object. For example, an attribute called *thePolicy* on a Claim business object could be used to link to the Policy object as follows:

```

interface Claim : ...
{
    attribute Policy thePolicy;
    ...
};

```

Clients can establish and traverse the link using the attribute set and get methods respectively:

```

Policy_var aPolicy = ... // Find or create a Policy object
Claim_var aClaim = ... // Find or create a Claim object

// Set the claim's policy to "aPolicy".
aClaim->thePolicy(aPolicy);
// Get the claim's policy as "somePolicy".
Policy_var somePolicy = aClaim->thePolicy();

```

Using the Component Broker delegating pattern, the business object implementation of the access methods for thePolicy passes the object pointer to or from the data object:

```

::CORBA::Void ClaimBO_Impl::thePolicy(Policy_ptr policy)
{
    fDataObject->thePolicy(policy);
}

Policy_ptr ClaimBO_Impl::thePolicy()
{
    return fDataObject->thePolicy();
}

```

Because the relationship is optional it is possible that the data object will return either a valid pointer to a policy object or a null pointer.

To speed up the performance of link access, caching the Policy pointer might be appropriate. The Claim business object implementation class could cache a pointer to the Policy object as shown in the following example:

```

class ClaimBO_Impl ...
{
    public:
        ...
    private:
        ...
        Policy_var fCachedPolicy;
        ...
}

```

The access methods for thePolicy would now use the cached pointer:

```

::CORBA::Void ClaimBO_Impl::thePolicy(Policy_ptr policy)
{
    fCachedPolicy = Policy::_duplicate( policy );
}

```

```

Policy_ptr ClaimBO_Impl::thePolicy()
{
    return Policy::_duplicate( fCachedPolicy );
}

```

In this case it is not necessary to invoke the `release()` method on the previous `Policy` object prior to saving the new one. This is because the `Policy` object is being saved in a `Policy_var` object. This `Policy_var` object ensures that the previous `Policy` object is released when it is no longer needed, including when it is being assigned a new `Policy` object.

The general Component Broker business object caching pattern uses the methods `syncFromDataObject` and `syncToDataObject` respectively to load and flush the cached values. See the description of “`syncFromDataObject()` method” on page 120 and “`syncToDataObject()` method” on page 121. The implementation of the `syncFromDataObject` method calls data object get methods to retrieve the `thePolicy` pointer as well as all other data contained in the `Claim` object. The implementation of the `syncToDataObject` method calls all the data object set methods. These methods appear as follows:

```

::CORBA::Void ClaimBO_Impl::syncToDataObject()
{
    ...
    fDataObject->thePolicy( fCachedPolicy);
    ...
}

::CORBA::Void ClaimBO_Impl::syncFromDataObject()
{
    ...
    fCachedPolicy = fDataObject->thePolicy();
    ...
}

```

It is not necessary to use the `_duplicate()` method in the implementations of the `syncToDataObject()` and `syncFromDataObject()` methods because the `Policy` object is neither an input parameter nor a return value of these methods.

Because links can be expensive to compute, and sometimes are not needed, a better pattern for caching object references is to leave them out of the `syncTo/FromDataObject` methods and instead compute and cache them in the `get` method on first use. This *lazy evaluation* approach can be implemented using the following pattern:

```

::CORBA::Void ClaimBO_Impl::thePolicy(Policy_ptr policy)
{
    fCachedPolicy = Policy::_duplicate( policy );

    // Synchronize the new Policy in the BO with (to) the DO

```

```

        fDataObject->thePolicy(fCachedPolicy);
    }

    Policy_ptr ClaimBO_Impl::thePolicy()
    {
        // If this is the first access of the Policy, get it from the DO.

        if ( CORBA::is_nil(fCachedPolicy) )
            fCachedPolicy = fDataObject->thePolicy();

        return Policy::_duplicate( fCachedPolicy );
    }

```

Assuming `fCachedPolicy` is initialized to `nil` in the constructor, this pattern results in the data object get method being called the first time the `Policy` is accessed, or not at all if the `Policy` is set in the same session as the get call. Subsequent accesses return the cached pointer value.

Optional or required cardinality-1 relationships

The previous section discusses how an optional relationship is implemented in a business object. An optional relationship is more flexible than a required relationship, and thus requires less code. However, if the relationship is required, the following diagram represents a required cardinality-1 *uses a* relationship.




Figure 39. Required cardinality-1 "uses a" relationship

If the relationship is truly required, then the link from a `Claim` to its `Policy` must exist throughout the life cycle of the `Claim`. Specifically, a `Policy` must be linked to a `Claim` as part of creating a `Claim`, and the link must be broken when a `Claim` is removed. The `Policy` that must be linked to a `Claim` as part of creation might exist prior to creating the `Claim`, or it might be created in the process of creating the `Claim`.

When reviewing a `Claim`, the link to its `Policy` must be broken, but the `Policy` is not necessarily removed. If the claim *uses a* (knows about a) `Policy`, then it is sufficient to release the `Policy` when a `Claim` is removed. However, if the `Claim` to `Policy` relationship is one of ownership (that is, a *has a* relationship), then when a `Claim` is removed, its associated `Policy` must also be removed.

When a `Claim` is removed, the link to its `Policy` must be broken. If the `Claim` is implemented using the caching pattern, then the link is automatically

broken if it is being cached in a `Policy_var`. Otherwise, it must be explicitly broken using `CORBA::release()`. Changes between an optional and a required relationship depend on one of four different scenarios (discussed over the next several pages). However, there is one change which is common to all four. Under any scenario, if the relationship is required instead of optional, the code for the set method for the *Policy* attribute must make sure that the link is not broken by setting the Policy reference to `nil`. The following example shows the changes necessary in the case of the delegating pattern. The changes for the caching pattern are analogous.

```

::CORBA::Void ClaimBO_Impl::thePolicy(Policy_ptr policy)
{
    if ( ! CORBA::is_nil(policy) )
    {
        fDataObject->thePolicy(policy);
    }
}

```

To understand how a link from a Claim to its Policy gets established as part of creating a Claim, it is necessary to recall that there are several ways to create a business object:

- Using a generic home configured for the appropriate type of business object. See “Creating a claim with an existing policy using a generic home” and “Creating a claim with a new policy using a generic home” on page 176.
- Using a specialized home developed by the business object provider. See “Creating a claim with an existing policy using a specialized home” on page 178 and “Creating a claim with a new policy using a specialized home” on page 180.

Creating a claim with an existing policy using a generic home: Using a generic home, a business object is created using the `createFromPrimaryKeyString()` method. If an object provider has provided a copy helper class, then a business object can also be created using the `createFromCopyString()` method. However, this has no effect on this discussion because a copy helper class’ attributes are defined as being a superset of the primary key class’ attributes. The only input to this method is a stringified version of the business object’s primary key helper object. This means that the primary key helper class for Claim must contain enough information in it such that the link to its (preexisting) Policy can be established. There are many ways in which this can be done, but the two most straightforward ways are the following:

- The primary key for Claim contains a reference to the Policy object. See “Primary key contains reference to related object” on page 174.
- The primary key for Claim contains all of the Policy object’s primary key attributes. See “Primary key contains key attributes of related object” on page 175.

Primary key contains reference to related object: In this case, the interface of the primary key for Claim would look like the following example:

```
interface ClaimKey : IManagedLocal::IPrimaryKey
{
    attribute long claimNo; // key attribute for Claim

    attribute Policy thePolicy; // associated Policy
}
```

The implementation binding header file (.ih) would include the following private or protected data members:

```
::CORBA::Long      fClaimNo;
Policy_var         fPolicy;
```

The attribute for the associated Policy object would be implemented as follows:

```
::CORBA::Void ClaimKey_Impl::thePolicy(Policy *thePolicy)
{
    fPolicy = Policy::_duplicate( thePolicy );
}

Policy ClaimKey_Impl::thePolicy()
{
    return Policy::_duplicate( fPolicy );
}
```

For a primary key to flow over the wire from the client to the server, the primary key must be able to externalize and internalize all of its attributes, including those which are necessary for establishing relationships to other objects. In the case of the primary key class for the Claim object, it must externalize and internalize a stringified object reference to the Policy object as follows:

```
::CORBA::Void ClaimKey_Impl::externalize_to_stream
 (::CosStream::StreamIO_ptr targetStreamIO)
{
    // Insert Method modifications here
    targetStreamIO->write_long(fClaimNo);
    CORBA::String_var policyRefString=
        CBSeriesGlobal::orb()-> object_to_string(fPolicy);
    targetStreamIO->write_string(policyRefString);

    // End Method modifications here
}

::CORBA::Void ClaimKey_Impl::internalize_from_stream
 (::CosStream::StreamIO_ptr sourceStreamIO,
  ::CosLifeCycle::FactoryFinder_ptr there)
{
    // Insert Method modifications here
    fClaimNo = sourceStreamIO->read_long();
    CORBA::String_var policyRefString = sourceStreamIO->read_string();
}
```



```

CORBA::Object_var policyRef =
    CBSeriesGlobal::orb()-> string_to_object(policyRefString);
fPolicy = Policy::_narrow(policyRef);
targetStreamIO->write_string(policyRefString);

// End Method modifications here
}

```

The streaming code in the previous example allows the primary key for Claim to flow from the client application to the Claim home on the server. When the primary key reaches the Claim home on the server, it eventually gets passed to the data object implementation for Claim. The link has been established by invoking the `string_to_object()` method on the ORB while internalizing the Claim primary key from its stream.

Primary key contains key attributes of related object: In the second case, the interface of the primary key for Claim would look like the following example:

```

interface ClaimKey : IManagedLocal::IPrimaryKey
{
    attribute long claimNo; // key attribute for Claim

    attribute long policyNo; // key attribute for Policy
}

```

The implementation binding header file (.ih) would include the following private or protected data members:

```

::CORBA::Long      fClaimNo;
::CORBA::Long      fPolicyNo;

```

The attribute for the associated Policy object would be implemented as follows:

```

::CORBA::Void ClaimKey_Impl::policyNo(::CORBA::Long policyNo)
{
    fPolicyNo = policyNo;
}

::CORBA::Long ClaimKey_Impl::policyNo()
{
    return fPolicyNo;
}

```

Again, in order for a primary key to go from the client to the server, the primary key must be able to externalize and internalize all of its attributes, including those which are necessary for establishing relationships to other objects. In this case, the primary key class for the Claim object must externalize and internalize the key attributes of its related Policy object as follows:

```

::CORBA::Void ClaimKey_Impl::externalize_to_stream
  (::CosStream::StreamIO_ptr targetStreamIO)
{
  // Insert Method modifications here
  targetStreamIO->write_long(fClaimNo);
  targetStreamIO->write_long(fPolicyNo);

  // End Method modifications here
}

::CORBA::Void ClaimKey_Impl::internalize_from_stream
  (::CosStream::StreamIO_ptr sourceStreamIO,
   ::CosLifeCycle::FactoryFinder_ptr there)
{
  // Insert Method modifications here
  fClaimNo = sourceStreamIO->read_long();
  fPolicyNo = sourceStreamIO->read_long();

  // End Method modifications here
}

```

As with the previous scenario, the streaming code in this example allows the primary key for Claim to flow from the client application to the Claim home on the server. However, that is the end of the similarity. The previous scenario shows that the cardinality-1 relationship is established as part of internalizing the Claim primary key inside the Claim home on the server.

In this scenario, because the Claim primary key class contains the key attributes of the related Policy object, the relationship must be established in the Claim data object implementation. The Claim data object implementation would do so by extracting the Policy key attributes from the Claim primary key, finding a home of Policy objects, and invoking the `findByPrimaryKeyString()` method on the home.

Creating a claim with a new policy using a generic home: In the previous scenario, the application domain specified a constraint that a Claim can only be created for an existing Policy. It is possible that the application domain, while requiring that a relationship between two objects exist, does not require one object to already be created when creating the second. While an insurance company which allows a Policy to be created at the same time as a Claim would probably not stay in business long. That scenario is used here for consistency.

Like the previous scenario, this scenario assumes that the Claim object provider has chosen not to develop a specialized home and instead expects clients of the Claim business object to use a generic home. Because the only input to the `createFromPrimaryKeyString()` method of a generic home is a stringified version of the business object's primary key helper object, this scenario could be implemented in a similar fashion to the previous scenario,

where the Claim primary key class contains the key attributes of a Policy. However, the one difference is what happens when the Claim primary key containing information about a Policy reaches the Claim home on the server.

In the previous scenario, where the Claim is being created with an existing Policy, there is no need for the Claim data object implementation to differentiate between when the business object is being created for the first time, and when it is being reactivated after having been previously passivated. In other words, in either case, the Claim data object implementation uses the Policy key attributes from its own key to find the existing Policy object.

In this scenario, however, the Policy is being created during the creation of a Claim. As such, the Claim data object implementation must distinguish between the creation and reactivation of a Claim as follows:

- If a Claim is being created, then a new Policy object must be created.
- If a Claim is being reactivated, then the previously created Policy object must be found.

A data object does not know if the business object is being created or reactivated. However, the business object itself does know if it is being created or reactivated. If the business object is being created the home invokes the `initForCreation()` method on it; if the business object is being reactivated the home invokes the `initForReactivation()` method.

In its implementation of the `initForCreation()` method, the Claim business object would do the following:

- Use a factory finder to find a Policy home
- Use the Policy key attributes from the Claim data object to create a new Policy by invoking the `createFromPrimaryKeyString()` method on the Policy home.

In its implementation of the `initForReactivation()` method, the Claim business object would do the following:

- Use a factory finder to find a Policy home.
- Use the Policy key attributes from the Claim data object to find its related Policy by invoking the `findByPrimaryKeyString()` method on the Policy home.

There is an alternate way to implement this scenario. Assuming that the client of a Claim is not required to provide the identity of its Policy, then the primary key class for Claim does not need to contain any information about the Policy. The primary key class for Claim would look like the following example:

```

interface ClaimKey : IManagedLocal::IPrimaryKey
{
    attribute long claimNo; // key attribute for Claim
}

```

In this case, it is the responsibility of the Claim business object, not the data object, to create the Policy during its own creation. A data object has no way to distinguish between object creation and object reactivation. However, a business object knows it is being created when its home invokes the `initForCreation()` method on it. Thus, the implementation of `ClaimBO_Impl::initForCreation()` would need to somehow create a Policy object. Because the Claim business object was not provided with any information on how to identify the Policy object, it would probably do one of the following:

- Create a primary key for Policy based on some combination (or function) of its own attributes.
- Create a primary key for Policy based on some random number generator or UUID (Universal Unique Identifier).
- Use a specialized home for Policy (if one was provided).

Because the relationship of Claim to Policy is required, if the Claim business object is unable to create a Policy for some reason, then it should cause its own creation to fail. The following example shows how this would be done:

```

::CORBA::Void ClaimBO_Impl::initForCreation(
    ::IManagedServer::IDataObject_ptr theDO)
{
    // Save the data object for later use
    fDataObject = ClaimDO::_narrow(theDO);

    try // to create the Claim's Policy somehow
    {
        fCachedPolicy = ...
    }
    catch(...)
    {
        throw IManagedServer::ICreationFailed();
    }

    // Other initialization...
}

```

Creating a claim with an existing policy using a specialized home: Using a specialized home, it is even easier to establish a link from a Claim to its Policy as part of creating the Claim. With a specialized home, there is no need to add any information about the Policy to the primary key for Claim. Adding information about one object to the primary key of another in the previous scenarios was done for the sole purpose of establishing the required

relationship. The primary key was being used not just to establish unique identity, but also as a vehicle for communicating information about the relationship from the client to the home. This is not what primary key classes were intended for, but was done out of necessity given that a specialized home was not provided.

Unfortunately, doing so introduced the following unwanted side-effect. Using the generic home configured for Claim objects at some later time to find a previously created Claim object required the client to have some information about the Policy too (at least some of the key attributes, if not a reference to the Policy object itself, depending on the scenario). This might not be an acceptable constraint in the application domain. If not, then the solution is to have the Claim object provider provide a specialized home as well.

Using a specialized home, the link between a Claim and its Policy is established in the specialized home itself, as opposed to in the Claim business object or data object. In order to support the required cardinality-1 relationship to a Policy, the specialized home for Claim would introduce new methods for creating a Claim. These new methods would have parameters which would allow the specialized home to establish the link between a Claim and its Policy. The following shows two examples of such create methods:

```

Claim_ptr ClaimHomeBO_Impl::createClaimWithPolicyRef(
    long claimNo,
    Policy_ptr policy)
{
    // Before we get too far, let's make sure that we have a valid
    // Policy reference for the required cardinality-1 relationship.
    if ( CORBA::is_nil(policy) )
        throw ClaimHome::MissingPolicy();

    ClaimKey_var claimKey = ClaimKey::_create();
    claimKey->claimNo(claimNo);
    ByteString_var claimKeyString = claimKey->toString();

    // For an inheriting specialized home, pass the claimKeyString
    // to the parent of the specialized home on creatFromPrimaryKey-
    // String(). For a delegating specialized home, pass the claim-
    // KeyString to the contained home on createFromPrimaryKeyString().
    Claim_var newClaim = ...

    // Now that the Claim has been created, establish the link to its
    // Policy using the Policy which was passed in on createClaim().
    newClaim->thePolicy(policy);

    return Claim::_duplicate(newClaim);
}

Claim_ptr ClaimHomeBO_Impl::createClaimWithPolicyNo(
    long claimNo,

```

```

        long policyNo)
    {
        // Assume that in the ClaimHomeBO_Impl::initForCreation()
        // the specialized home for Claim objects has used a factory
        // finder to find the home for Policy objects, and saved
        // its reference in fPolicyHome

        PolicyKey_var policyKey = PolicyKey::_create();
        policyKey->policyNo(policyNo);
        ByteString_var policyKeyString = policyKey->toString();

        Policy_var aPolicy; // Assume fPolicyHome is a specialized home...
        aPolicy = fPolicyHome->findByPrimaryKeyString(policyKeyString);

        ClaimKey_var claimKey = ClaimKey::_create();
        claimKey->claimNo(claimNo);
        ByteString_var claimKeyString = claimKey->toString();

        // For an inheriting specialized home, pass the claimKeyString
        // to the parent of the specialized home on createFromPrimaryKey-
        // String(). For a delegating specialized home, pass the claim-
        // KeyString to the contained home on createFromPrimaryKeyString().
        Claim_var newClaim = ...

        // Now that the Claim has been created, establish the link to its
        // Policy using the Policy which was passed in on createClaim().
        newClaim->thePolicy(policy);

        return Claim::_duplicate(newClaim);
    }

```

Creating a claim with a new policy using a specialized home: This scenario is similar to the previous scenario. However, because in this scenario the Policy is not created prior to creating a Claim, the createClaimWithPolicyRef() method from the previous scenario is not applicable and the implementation of the createClaimWithPolicyNo() method is different. Instead of invoking the findByPrimaryKeyString() method on the Policy home it must invoke the createFromPrimaryKeyString() method as illustrated in the following example:

```

Claim_ptr ClaimHomeBO_Impl::createClaimWithPolicyNo(
    long claimNo,
    long policyNo)
{
    // Assume that in the ClaimHomeBO_Impl::initForCreation()
    // the specialized home for Claim objects has used a factory
    // finder to find the home for Policy objects, and saved
    // its reference in 'fPolicyHome'.

    PolicyKey_var policyKey = PolicyKey::_create();
    policyKey->policyNo(policyNo);
    ByteString_var policyKeyString = policyKey->toString();

    Policy_var aPolicy; // Assume fPolicyHome is a specialized home...
    aPolicy = fPolicyHome->createFromPrimaryKeyString(policyKeyString);

```

```

ClaimKey_var claimKey = ClaimKey::_create();
claimKey->claimNo(claimNo);
ByteString_var claimKeyString = claimKey->toString();

// For an inheriting specialized home, pass the claimKeyString
// to the parent of the specialized home on creatFromPrimaryKey-
// String(). For a delegating specialized home, pass the claim-
// KeyString to the contained home on createFromPrimaryKeyString().
Claim_var newClaim      = ...

// Now that the Claim has been created, establish the link to its
// Policy using the Policy which was passed in on createClaim().
newClaim->thePolicy(policy);

return newClaim;
}

```

"Uses a" and "has a" cardinality-1 relationships

Now that the differences between optional and required cardinality-1 relationships have been described, especially as they pertain to developing a business object, the *uses a* and *has a* cardinality-1 relationships are described.

Previous sections discuss how *uses a* relationships are represented. The following figure illustrates an optional cardinality-1 *has a* relationship.




Figure 40. Optional cardinality-1 "has a" relationship

The following diagram, on the other hand, illustrates a required cardinality-1 *has a* relationship.




Figure 41. Required cardinality-1 "has a" relationship

A *uses a* relationship means that one object has a reference to another object. In the example, a Claim object has a reference to a Policy object. In fact, there might be more than one Claim object with a reference to the same Policy object and there might also be objects of other types with references to the same Policy object. In any case, any object with a reference to this Policy

object is capable of invoking any method on the Policy object's client interface. This includes the `remove()` method. Of course, things could get chaotic if any object with a *uses a* relationship to the same Policy object were allowed to invoke the `remove()` method on it. To avoid such confusion, Component Broker recommends that an object be removed only by its owner, and that an object be owned only by a single object. All other objects with references to that object should only release the reference.

A *has a* relationship is what is used to represent the concept of ownership. In other words, an object with a *has a* relationship to another object is said to own that object. In the previous two diagrams, a Claim object *has a* reference to a Policy object. Although the relationship would most likely be reversed in the insurance application domain, this scenario continues with the original example instead of introducing two new objects. Thus, a Claim is considered to be the owner of a Policy and as such is responsible for removing the Policy should this become necessary. One case in which it becomes necessary for a Claim to remove the Policy is when the Claim itself is being removed. A business object knows it is being removed when its home invokes the `uninitForDestruction()` method on it. The following example shows how the Claim business object would implement this:

```
::CORBA::Void ClaimBO_Impl::uninitForDestruction()
{
    // Remove the (owned) Policy object
    fCachedPolicy->remove();

    // If fCachedPolicy is declared as a Policy_ptr, then it is also
    // necessary to release the reference.
    fCachedPolicy->release();
    // If fCachedPolicy is declared as a Policy_var, then this happens
    // automatically when the BO is destructed.

    // Other un-initialization...
}
```

Another case in which it becomes necessary for a Claim to remove the Policy is if the Claim interface has a method that gives its clients the opportunity to request that the Policy be removed. This is similar to the `cancelPolicy()` method in the following example:

```
::CORBA::Void ClaimBO_Impl::cancelPolicy()
{
    // Remove the (owned) Policy object
    fCachedPolicy->remove();

    // If fCachedPolicy is declared as a Policy_ptr, then it is also
    // necessary to release the reference.
    fCachedPolicy->release();
    // If fCachedPolicy is declared as a Policy_var, then this happens
```



```

        // automatically when the BO is destructed.

        // Other clean-up associated with the Policy cancellation
    }

```

Of course, the `cancelPolicy()` method makes sense only in an optional *has a* relationship.

In any case, if one business object is responsible for the removal of another business object as the result of a *has a* relationship, then the reverse must not also be the case. In other words, in Component Broker, two objects might not have a *has a* relationship with one another. If such a bi-directional relationship is required, then one object must release the other, while the other object removes the first one as described in this section.

Making cardinality-1 relationships persistent

Because object references are really just memory addresses, they cannot be made persistent in that form. Therefore, object references must be converted to other forms that can be made persistent. Because persistence of all attributes, not just object relationships, is implemented in the data object, not the business object, this conversion is described further in “Data object customization for cardinality relations” on page 395.

As previously mentioned, in the Component Broker architecture, a business object’s persistent state is managed by its associated data object. Links (1-1 and 1-N object relationships) are no different. The persistent representation of a link is also typically managed by a data object. However, there might be cases where it makes more sense for the business object to manage the link itself. For instance, if the link can be computed based on some application domain business logic, then it makes sense not to burden the data object with managing the link.

For example, assume that a claim maintains a cardinality-1 link to the insurance policy with which it is associated. Assume as well that the design of the application is such that the primary key of a Claim (the claim number) includes the policy number of its associated Policy, for example, claim numbers are actually prefixed by their associated policy number. In this situation, implementing the Policy reference involves no additional persistent data. The Policy get method could be implemented directly in the business object as follows:

```

Policy_ptr ClaimBO_Impl::thePolicy()
{
    ByteString_var pkString;
    // get my primary key
    ClaimPrimaryKey_var myKey = ClaimPrimaryKey::_create();
    myKey->fromString(pkString = this->getPrimaryKeyString());
}

```

```

// extract the policy number from my primary key
long myPolicyNo = myKey->policyNo();

// get the policy home somehow (for example, using a factory finder)
IHome_var policyHome = ...

// create and initialize a policy primary key
PolicyPrimaryKey_var policyKey = PolicyPrimaryKey::_create();
policyKey->policyNo(myPolicyNo);

// get the policy from the policy home
IManagedClient::IManageable_var moPtr;
return Policy::_narrow(moPtr = policyHome->findByPrimaryKey(policyKey));
}

```

As shown, this algorithm does not involve an explicit representation of the converted pointer. It is extracted from the key and is based on the application-level knowledge that claim numbers are prefixed by their associated policy number. In this situation the persistent representation of the link is actually maintained as part of another one of the object's attributes. Also, when a link can be computed based on some business logic, the object reference is often considered to be a read-only attribute of the business object.

While it is possible for a business object to manage object references itself, more commonly the business object passes the pointer to the data object and it performs the required conversion. With this approach, every data object that is used with a particular business object is free to use whatever conversion algorithm (and persistent representation) it chooses. This approach allows the business object to remain de-coupled from the persistent storage mechanism and thus possibly be re-used in different scenarios with a wide range of back-end data stores.

Cardinality-N relationships

The previous section shows how to link an insurance Policy to a single Claim object. If, instead, the Policy object needs to reference multiple Claims, a cardinality-N relationship is required.




Figure 42. Cardinality-N relationship

There are two approaches for providing a type-safe implementation for a cardinality-N relationship in Component Broker. Both approaches rely on a

persistent collection object to manage the references, although one hides the collection in the implementation, while the other exposes it to clients as a first class object.

With the first approach, where the collection is hidden, the interface for adding and removing elements is provided on the Policy object itself. With this approach, the Policy business object interface would provide a set of methods for establishing, deleting, and accessing the links to Claim objects. The interface could be defined as follows:

```
interface Policy : ...
{
    void addClaim(in Claim claim);
    void removeClaim(in Claim claim);
    IMangedCollections::IIterator listClaims();
    ...
};
```

Clients would then use this interface to add, access, and remove claims, as follows:

```
Policy_var aPolicy = ...
Claim_var aClaim1 = ...
Claim_var aClaim2 = ...

//
// Add a couple of claims to the policy's set of claims
//
aPolicy->addClaim(aClaim1);
aPolicy->addClaim(aClaim2);

//
// Iterate through the policy's set of claims
//
// Get an iterator for the set of claims
IMangedCollections::IIterator_var iter = aPolicy->listClaims();
IMangedClient::IMangeable_var element;
while (element = iter->next())
{
    // iterate through the set of claims
    Claim_var aClaim = Claim::_narrow(element);
    // do something with (for example, process) "aClaim"
    ...
}

//
// Remove a claim from the policy's set of claims
//
aPolicy->removeClaim(aClaim1);
```

As shown, although the underlying implementation of the relationship methods might use a collection object to manage the references, the client

program is completely shielded from this fact. The Policy object encapsulates the collection behind the type-safe relationship interface methods, `addClaim()`, `removeClaim()`, and `listClaims()`.

The second approach for implementing the cardinality-N relationship exposes the collection in the client programming model. Instead of referencing multiple Claim objects directly, the Policy object references a single collection object using a cardinality-1 link. The collection, in turn, references the multiple claims.




Figure 43. Cardinality-1 link to a collection object that references claims.

This picture actually applies to the previous approach as well, although there the reference collection is hidden in the implementation. One difference is that when the collection is hidden, the Policy interface is the only client interface for adding elements to the collection and therefore it provides static type checking for the elements in the collection. If, on the other hand, the collection is visible to clients, the collection itself must prevent clients from adding objects other than Claims to the relationship.

When using the explicit collection approach for implementing a cardinality-N relationship, the Policy business object interface includes an attribute whose type is a reference to the collection.

```

interface Policy : ...
{
    readonly attribute IMangedCollection::IReferenceCollection claims;
    ...
};
  
```

The relationship of the collection object to the Policy object can be considered to be a standard cardinality-1 link as described in “Cardinality-1 relationships” on page 169, with one exception, the attribute is read-only. The *readonly* attribute allows clients to add and remove elements in the collection but prevents them from replacing the collection itself. In some situations this latter operation might be warranted, in which case, the *readonly* tag would be removed. In general, however, you should prevent clients from modifying the reference to the collection.

A client program uses the *claims* attribute access method to access the collection of claims. Methods on the collection can then be called to manipulate and access the actual Claim references.

```

Policy_var aPolicy = ...
Claim_var aClaim1 = ...
Claim_var aClaim2 = ...

//
// Get the set of claims;
ManagedCollections::IReferenceCollection_var claims =
    aPolicy->claims();

//
// Add a couple of claims to the policy's set of claims
//
claims->addElement(aClaim1);
claims->addElement(aClaim2);

//
// Iterate through the policy's set of claims
//
ManagedCollections::IIterator_var iter = claims->createIterator();
    // get a claims iterator
ManagedClient::IManageable_var element;
while (element = iter->next())
{
    // iterate through the set of claims
    Claim_var aClaim = Claim::_narrow(element);
    // do something with (for example, process) "aClaim"
    ...
}

//
// Remove a claim ("aClaim1") from the policy's set of claims
//
claims->removeElement(aClaim1);

```

Implementing the relationship interface

Depending on the interface requirements and possibly the back-end datastore being used, many different implementations of the Policy/Claim relationship are possible. Some common approaches include implementing the relationship using:

- A simple Non-keyed reference collection. See “Implementing a relationship with a simple reference Collection” on page 188.
- A Keyed reference collection. See “Implementing a relationship with a keyed reference collection” on page 189.
- A home or collection and some additional information that is used to identify a subset of the entries in the collection. Some examples include a non-unique secondary key supported by the home or collection and a query evaluation string if the home or collection is queryable. See “Subsetting a home or collection” on page 190.

The following sections describe each of these implementation approaches.

Implementing a relationship with a simple reference Collection: The relationship collection is exposed in the data object interface as a readonly attribute of type IReferenceCollection:

```
interface PolicyDO : ...
{
    readonly attribute IManagedCollections::IReferenceCollection claims;
    ...
};
```

The relationship interface methods, addClaim(), removeClaim(), and listClaims(), would be implemented as shown in the following example:

```
::CORBA::Void PolicyBO_Impl::addClaim(Claim_ptr claim)
{
    IManagedCollections::IReferenceCollection_var claims = claims();
    claims->addElement(claim);
}

::CORBA::Void PolicyBO_Impl::removeClaim(Claim_ptr claim)
{
    IManagedCollections::IReferenceCollection_var claims = claims();
    claims->removeElement(claim);
}

IManagedCollections::IIterator_ptr PolicyBO_Impl::listClaims()
{
    IManagedCollections::IReferenceCollection_var claims = claims();
    return claims->createIterator();
}
```

In each of the methods, the method claims is used to return a pointer to the reference collection containing the claims. Each method then delegates to its corresponding method on the reference collection.

The persistent object reference for the reference collection itself can be implemented using any of the design patterns described in “Cardinality-1 relationships” on page 169. Using the lazy evaluation caching pattern in the business object, the claims method would be implemented as follows:

```
IManagedCollections::IReferenceCollection_ptr PolicyBO_Impl::claims()
{
    // first time?
    if (fCachedClaims == IManagedCollections::IReferenceCollection::_nil())
        fCachedClaims = fDataObject->claims();
    return fCachedClaims;
}
```

Before the data object claims method can return a collection, a reference collection must actually be created. This is typically done in the data object claims method the first time it is called. As previously discussed, the reference collection used to implement a relationship may be required to guarantee that the type of elements added to it are of a specific type (for example, Claims). A

type-specific reference collection can be created by calling the `createCollectionFor` method on the `IManagedCollections::ICollectionHome` interface:

```
// get the collection home (for example, using a factory finder)
IManagedCollections::ICollectionHome_var cHome = ...

// create a reference collection that will only hold claims
IManagedCollections::IReferenceCollection_var rc =
    cHome->createCollectionFor(Claim::Claim_RID);
```

Alternatively, the simpler `createCollection` method can be used when no restriction on the type of collection element is required:

```
rc =    cHome->createCollection();
```

Implementing a relationship with a keyed reference collection: A Keyed reference collection can be used to implement relationships where the individual links are accessed using some kind of identifier. For example, if the claims associated with a given insurance policy are identified by, for example, a claim ID, the relationship interface might appear in the Policy business object as follows:

```
interface Policy : ...
{
    void addClaim(in Claim claim, in long id);
    void removeClaim(in long id);
    Claim getClaim(in long id);
    IManagedCollections::IIterator listClaims();
    ...
};
```

This relationship is most easily implemented using a keyed collection. For example, the data object interface might now include an attribute of type `IKeyedReferenceCollection`:

```
interface PolicyDO : ...
{
    readonly attribute
        IManagedCollections::IKeyedReferenceCollection claims;
    ...
};
```

The `addClaim` method could be implemented as follows:

```
::CORBA::Void PolicyBO_Impl::addClaim(Claim_ptr claim, ::CORBA::Long id)
{
    NumberKey_var key = NumberKey::_create();
    key.setValue(id);
    IManagedCollections::IKeyedReferenceCollection_var claims = claims();
    ::ByteString_var keyString = key.toString();
    claims->addElementByString(claim, keyString);
}
```

In this example, a key helper class (see “Chapter 4. MOFW C++ client programming model” on page 85), `NumberKey`, is used to add the element to the collection. The class `NumberKey` wrappers the ID in a key class that is capable of being stringified. When the key is created and initialized, the add operation is delegated to the reference collection.

The `removeClaim` and `getClaim` methods would be implemented similarly. The remaining method, `listClaims`, as well as the `claims` method that is used to access the collection object, would be implemented as shown in the non-keyed reference collection example.

Subsetting a home or collection: This approach uses a home augmented with information that identifies a subset of the objects in the home. This approach is particularly applicable in bottom-up development scenarios where related data already exists in legacy applications.

For example, an RDB-based insurance application might contain two related tables, one for policies and one for claims. The claims registered against a particular policy are identified by a `policy#` column in the claim table. This column, a foreign key in the claim table, is the primary key for the policy table.

POLICY TABLE

Policy#	Owner	...
12	"Joann"	
34	"John"	
56	"Katherine"	
78	"Katherine"	

CLAIM TABLE

Claim#	Policy#	...
87	100	
65	34	
43	101	
21	34	

Figure 44. RDB-based insurance application tables

In Component Broker object space, these tables represent a cardinality-N relationship between a policy object and its associated claims. For example, policy number 34 would have links to claim numbers 21 and 65.

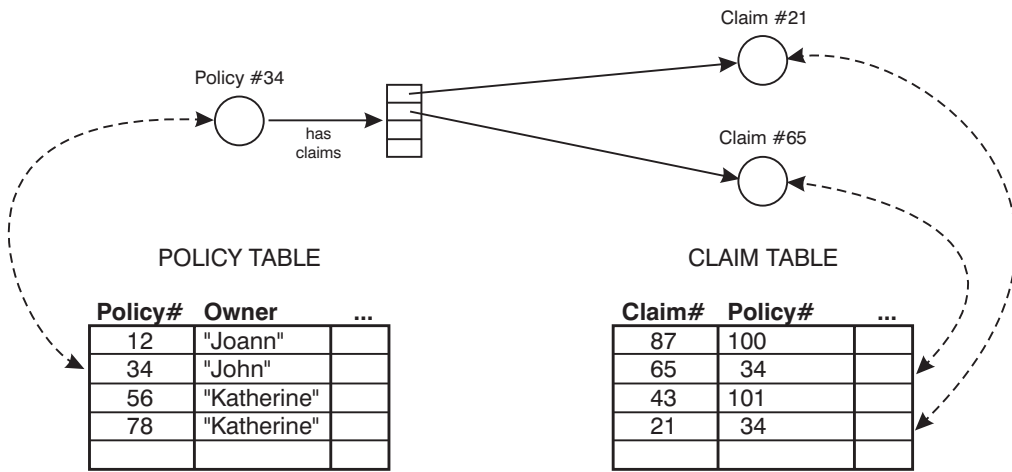


Figure 45. Links between policy and Claims

Normally these links are stored as converted pointers maintained by a persistent reference collection of some type. In this case, the relationship information stored in the claim table itself eliminates the need for a persistent reference collection. The set of claim links for a particular policy can be derived using a query on the claim table. In Component Broker object space, this can be implemented by evaluating an OOSQL query (for example, `policyNo == 34`) on a queryable claim home.

```

IManagedAdvancedClient::IQueryableIterableHome_var claimHome = ...

IManagedCollections::Iterator_var iter = claimHome->evaluate(
    "policyNo=34");

// iterate over the claims
...

```

In this example, the home object would be retrieved using the Naming Service. The home name could be hard coded in the application, available in an environment variable, or stored persistently in the policy table.

The particular pattern used for finding the home might or might not be dictated by legacy requirements.

Because the relationship is derived from other data, adding and removing claims often involve side effects. For example, adding a claim to the reference collection would require a change in state to the claim object (for example, `policy#` field must be updated). If the claim is already in a relationship with another policy, it would be implicitly removed from it as a result of the add operation here. To avoid this kind of change in semantics of the `addClaim`

method, the relationship interface could be changed to one that is more semantically consistent with the underlying implementation. For example:

```
interface Policy : ...
{
    Claim createClaim( in long id);
    void deleteClaim(in long id);
    Claim getClaim(in long id);
    IManagedCollections::IIterator listClaims();
    ...
};
```

By replacing the addClaim and removeClaim operations with createClaim and deleteClaim methods, the relationship semantics map well to the home-based implementation.

Creating specialized homes

Component Broker provides a default IManagedClient::IHome implementation. There might be specializations of this interface specific to an underlying application adaptor. These include IManagedAdvancedClient::IQueryableIterableHome and IManagedAdvancedClient::IIterableHome. This section explains how to extend a home with domain-specific methods for create and find.

There are a number of cases where the usage of IManagedClient::createFromPrimaryKeyString, IManagedClient::createFromCopyString and IManagedClient::findByPrimaryKeyString might not present the optimal interface for clients wishing to interact with the home to create and find business objects.

This section describes the process of extending the home and follows the pattern set out in “Extending a business object” on page 157.

Extending the interface to IHome

For the example, PolicyHome is an interface that specializes the IHome interface with some specific create and find methods. The goal is to provide methods specific to the insurance policy abstraction. These methods are shown in the IDL in the following example:

```
Policy create(in float premium, in float amount);
           // create a policy passing in the attribute values
Policy defaultCreate();
           // default create method - policyNumber is assigned
Policy createWithNumber(in long policyNo);
```

```

        // create a new policy with this number
    Policy findPolicyByNumber(in long policyNo);
        // find a policy by number

```

You must first decide which IHome interface should be specialized. This seems to have a simple solution: Inherit from the IHome in the managed object framework and get an interface that looks like this:

```

#include <IManagedClient.idl>
#include "Policy.idl"

interface PolicyHome : IManagedClient::IHome
{
    Policy create(in float premium, in float amount)
        raises(IManagedClient::IInvalidKey, IManagedClient::IDuplicateKey);
    Policy defaultCreate()
        raises(IManagedClient::IInvalidKey, IManagedClient::IDuplicateKey);
    Policy createWithNumber(in long policyNo)
        raises(IManagedClient::IInvalidKey, IManagedClient::IDuplicateKey);
    Policy findPolicyByNumber(in long policyNo)
        raises(IManagedClient::IInvalidKey, IManagedClient::INoObjectWKey);
};

```

Details

IManagedClient::IHome is being extended because it is the interface that is supported by all of the application adaptors provided by the server.

Exceptions have also been put on the methods that are introduced. For completeness, creation methods should raise the IManagedClient::IInvalidKey and IManagedClient::IDuplicateKey exceptions. Find related methods should raise the IManagedClient::IInvalidKey and IManagedClient::INoObjectWKey exceptions. Additional exceptions can also be introduced and raised as appropriate.

An alternative exception strategy is to have the specialized home implementations actually handle some of the exceptions. For example, catching the IManagedClient::IDuplicateKey exception on methods where the key is not passed in could be appropriate.

Alternatives to IManagedClient::IHome

Other alternatives for extending homes are IManagedAdvancedClient::IIterableHome and IManagedAdvancedClient::IQueryableIterableHome. This is true only if the additional interface supported by these extended homes is to be a proper superset of that which is available for the particular application adaptor configuration. Not all homes support these IManagedAdvancedClient interfaces.

Implement the extended IHome interface

The IHome itself is a managed object and the development of an extended home should be done like creating any subclass of a managed object. This is described in “Extending a business object” on page 157. This section highlights things that are specific to the IHome interface and extension.

Implementation interface

The PolicyHomeBO.idl file inherits according to the pattern described in “Extending a business object” on page 157. It is shown in the following example:

```
#include <PolicyHome.idl>
#include <IManagedAdvancedServer.idl>
interface PolicyHomeBO : PolicyHome,
    IManagedAdvancedServer::ISpecializedHome
{
};
```

The PolicyHomeBO.ih file contains the PolicyHomeBO_Impl class. This class inherits the implementation from the IManagedAdvancedServer::ISpecializedHome_Impl that it plans to extend. This might vary from what was described previously, but this example continues to use the extension to the default Home implementation provided by the managed object framework. The implementation interface is shown in the following example:

```
#include <IManagedAdvancedServer.ih>
#include "PolicyHomeBO.hh"
class PolicyHomeBO_Impl : public virtual ::PolicyHomeBO_Skeleton,
    public virtual
    IManagedAdvancedServer::ISpecializedHome_Impl
{
public:
    ::Policy_ptr create (::CORBA::Float premium, ::CORBA::Float amount);
    ::Policy_ptr defaultCreate ();
    ::Policy_ptr createWithNumber (::CORBA::Long policyNo);
    ::Policy_ptr findByPolicyNumber (::CORBA::Long policyNo);

    // Methods from IManagedServer
    virtual ::CORBA::Void initForCreation
        (::IManagedServer::IDataObject_ptr theDO);
    virtual ::CORBA::Void initForReactivation
        (::IManagedServer::IDataObject_ptr theDO);
    virtual ::CORBA::Void uninitForDestruction();
    virtual ::CORBA::Void uninitForPassivation();
    virtual ::CORBA::Void syncToDataObject();
    virtual ::CORBA::Void syncFromDataObject();
```

```

private:
    // superclass pointer of DO we are using
    IManagedAdvancedServer::ISpecializedHomeDataObject_ptr myDO;
};

```

To create a specialized home that has query and iterable capabilities, simply replace `IManagedAdvancedServer::ISpecializedHome` with `IManagedAdvancedServer::ISpecializedQueryableIterableHome` in the `PolicyHomeBO.idl` and replace `IManagedAdvancedServer::ISpecializedHome_Impl` with `IManagedAdvancedServer::ISpecializedQueryableIterableHome_Impl` in the `PolicyHomeBO.ih` file.

The implementation

The implementation of the new create methods involves careful usage of keys, copies, and the basic interface supported by home. All of the create methods end up using `createFromCopyString` or `createFromPrimaryKeyString`, passing a key or copy that has been loaded with the proper information to all proper creation of the object. The `findByNumber()` method follows a similar pattern.

Note: Care should be taken with CORBA types returned by these methods. Some types such as `CORBA::String` require special operations be used for return values. See the “Appendix C. CORBA programming” on page 481 for additional information.

create():

```

::Policy_ptr PolicyHomeBO_Impl::create
    (::CORBA::Float premium, ::CORBA::Float amount)
{
    CORBA::Long policyNo = getUnique();
    cout << "The pseudo-unique policyNo will be: " << policyNo << endl;
    // create a key from that number
    PolicyKey_var theKey = PolicyKey::_create();
    theKey->policyNo(policyNo);
    // do findBy with that number to ensure that it isn't a duplicate
    Policy_var tPolicy;
    try
    {
        tPolicy=Policy::_narrow(IManagedClient::IManageable_var moPtr =
            findByPrimaryKeyString(ByteString_var findKeyStr =
                theKey->toString()));
        if (tPolicy)
        {
            // The Object already exists..
            throw IManagedClient::IDuplicateKey();
        }
    }
    catch(IManagedClient::INoObjectWKey &nowk)
    {

```

```

    // The Object does not exist. Good we can create it
    // just eat this exception and continue
    // If object not found on server, go ahead and create one.
    // create a copy object
    PolicyCopy_var theCopy = PolicyCopy::_create();
    // load it up with the right stuff
    theCopy->policyNo(policyNo);
    theCopy->premium(premium);
    theCopy->amount(amount);
    // call createFromCopyString
    IManageable_var aManageable;
    aManageable=createFromCopyString(ByteString_var findKeyStr =
        theCopy->toString());
    return Policy::_narrow(aManageable);
}
}

```

defaultCreate():

```

::Policy_ptr PolicyHomeBO_Impl::defaultCreate ()
{
    ::CORBA::Long policyNo;

    // There is a 'userData' attribute on each home. User specific
    // information can be stored in this field. We have chosen to
    // store the key generation algorithm type that should be used
    // when creating a new key.

    // Go to my DO to get the userData attribute value out of the
    string_var phUserData = myDO->getConfigInfo();
    if (strcmp(phUserData,"KeyAlgorithm1") == 0 )
    {
        // use the first key algorithm
        policyNo = getUniqueKey();
    }
    else
    {
        // use the second key algorithm
        policyNo = 100 + getUniqueKey();
    }

    // create a key from that number
    PolicyKey_var theKey = PolicyKey::_create();
    theKey->policyNo(policyNo);
    // do findBy with that number to ensure that it isn't a duplicate
    Policy_var tPolicy;
    try
    {
        tPolicy=Policy::_narrow(IManagedClient::IManageable_var moPtr =
            findByPrimaryKeyString(ByteString_var findKeyStr =
                theKey->toString()));
        if (tPolicy)
        {
            throw IManagedClient::IDuplicateKey();
        }
    }
}

```

```

    }
    catch(IManagedClient::INoObjectWKey &nowk)
    {
        // object does not exist... we can create
        // call createFromPrimaryKeyString
        return Policy::_narrow(IManagedClient::IManegeable_var moPtr =
            createFromPrimaryKeyString(ByteString_var findKeyStr =
                theKey->toString()));
    }
}

```

createWithNumber():

```

::Policy_ptr PolicyHomeBO_Impl::createWithNumber (
    ::CORBA::Long policyNo)
{
    // create a key
    PolicyKey_var theKey = PolicyKey::_create();
    theKey->policyNo(policyNo);
    // call createFromPrimaryKeyString
    return Policy::_narrow(IManagedClient::IManegeable_var moPtr =
        createFromPrimaryKeyString(ByteString_var findKeyStr =
            theKey->toString()));
}

```

Note that this method does not perform exception checking. If something other than `IManagedClient::IDuplicate` should be thrown when the number provided as input is already in use, then a try/catch block and associated logic would be required.

findByNumber():

```

::Policy_ptr PolicyHomeBO_Impl::findByPolicyNumber (
    ::CORBA::Long policyNo)
{
    // create a key
    PolicyKey_var theKey = PolicyKey::_create();
    theKey->policyNo(policyNo);
    // call findByPrimaryKeyString
    return Policy::_narrow(IManagedClient::IManegeable_var moPtr =
        findByPrimaryKeyString(
            ByteString_var findKeyStr = theKey->toString()));
}

```

Meet MOFW IManageable requirements

Because this extension to the home introduces no new additional methods and no new key for the home itself, `getPrimaryKeyString`, `getHandleString`, `externalize_to_stream`, and `internalize_from_stream` do not need to be implemented.

MOFW requirements – IManagedObject interfaces

This section includes the following:

- “initForCreation()”
- “initForReactivation”
- “uninitForDestruction”
- “uninitForPassivation”
- “syncFromDataObject” on page 199
- “syncToDataObject” on page 199

initForCreation(): The initForCreation method is only required to call the parent, passing the dataObject that comes in as a parameter. Currently, homes are not actually created using this method. Homes are configured onto systems and brought into existence as part of server initialization. This code is not implemented; the example is included for completeness.

```
::CORBA::Void PolicyHomeBO_Impl::initForCreation
    (::IManagedServer::IDataObject_ptr theDO)
{
    // call my parent
    IManagedAdvancedServer::ISpecializedHome_Impl::initForCreation(theDO);
}
```

initForReactivation: Set the data object using the same pattern as initForCreation. Any other specialized home specific code that needs to be implemented when a home is activated goes in this method implementation.

```
::CORBA::Void PolicyHomeBO_Impl::initForReactivation
    (::IManagedServer::IDataObject_ptr theDO)
{
    //Call my parent
    ::IManagedAdvancedServer::ISpecializedHome_Impl::initForReactivation(
        theDO);
    // set my DO
    myDO = IManagedAdvancedServer::ISpecializedHomeDataObject::_narrow(
        theDO);
}
```

uninitForDestruction: uninitForDestruction also only calls the parent method.

```
::CORBA::Void PolicyHomeBO_Impl::uninitForDestruction()
{
    // call my parent
    ::IManagedSystemObject::IHome_Impl::uninitForDestruction(theDO);
}
```

uninitForPassivation: Call the parent following the same pattern as uninitForDestruction.

syncFromDataObject: Call the parent following the same pattern as `uninitForDestruction`.

syncToDataObject: Call the parent following the same pattern as `uninitForDestruction`.

Keys

The same class used for the MOFW home works here. Homes are found generally using factory finders and are not created. Keys are not used in the programming model. Keys are used in the internal server run time.

Copy helper

The same class used for the MOFW home works here. Homes are not created, and therefore no copy helper is needed or usable based on the life cycle of homes.

Leveraging server provided essential state extensions

A specially defined attribute is passed through from the Object Builder or DDL that represents this specialized home. Rather than having a separate data object for the specialized home, it is much more efficient to access this data when it is needed. This is done as follows:

```
char* x = myDO->getConfigInfo();
if ( /* some evaluation of x */ )
{
    // make up a number using srand or some program that helps with this
    policyNo = /*whatever;
}
else
    policyNo=/* whatever */
```

This data is read-only and cannot be changed.

Overriding specific methods on specialized homes

In addition to supplying specific methods that allow clients to create and find objects without using keys and copy helper, specialized homes also provide the mechanism for overriding other interfaces supported by the `IManagedClient::IHome` interface. The following specific methods that can be overridden are described below:

- `IManagedClient::IHome::createFromPrimaryKeyString`
- `IManagedClient::IHome::createFromCopyString`
- `IManagedClient::IHome::findByPrimaryKeyString`

The `IManagedClient::IHome::createFromPrimaryKeyString` and `createFromCopyString` can be overridden to prevent creation of objects of a

particular type. There are cases when using implementation inheritance and polymorphism where homes are configured onto a system even when the class is abstract. This is usually to facilitate polymorphic `findByPrimaryKeyString` operations. While `findByPrimaryKeyString` is desired as a polymorphic operation, the create capabilities may in fact be prohibited. Other reasons include special checking that needs to be done before the actual create is issued. While this can be done with specialized create methods, using the base programming model methods of `createFromPrimaryKeyString` and `createFromCopyString` may be desirable in some situations.

Overriding `IMangedClient::findByPrimaryKeyString` is done to facilitate special find logic. This is most common in cases where polymorphic relationships exist between business objects. For example, if an abstract class of "a" has subclasses of "b" and "c", it would be reasonable to try to find an "a" with a given key. While there are no "a" instance because it is abstract, it is reasonable to have an overridden `findByPrimaryKeyString()` method that would look at all concrete subclasses of "a" to properly implement the `find()` operation. By overriding `findByPrimayKeyString()` method instead of using a specialized find method, data objects that deal with polymorphic relationships can work unchanged. They depend on `findByPrimaryKeyString()` method as part of the attribute getter implementation. Another reason to override `IMangedClient::findByPrimaryKeyString` is to prevent standard access to objects by throwing an exception in this method. While this will work, and is allowed, a side effect is that the objects supported by this home cannot be resolved in object relationships using the standard Home/Key pattern. This is because the `findByPrimaryKeyString()` method will be called during the resolution of these objects. Since this method will now throw an exception the object reference will fail. Using an alternative object relationship pattern, such as SOR, will solve some of this problem, however ForeignKey relationships would still fail.

The contracts of these methods must be maintained even when they are overridden. The exceptions thrown must still be honored by the specialized home. For example, `IMangedClient::IDuplicateKey` must still be thrown by create and `IMangedClient::INoObjectWKey` should be returned. These are integral to much of the mainline programming practice in Component Broker client programs. This contract must be maintained even in the specialized home overridden versions of these methods. New exceptions cannot be introduced on these methods as there is no overriding or overloading of interface specifications allowed in IDL.

Thread safety

Special care needs to be taken when writing specialized homes that access and modify static data. This is because homes are more likely to be accessed from

many clients concurrently. Because of this they are required to be thread-safe in order to provide correct access and control of static data.

Summary of home extension

Extending a home is much like extending any business object but simpler. Refer to “Managed object implementation details” on page 354 for more details.

Creating UUID specialized homes

There are cases where a developer may want to use a transient object but does not have a unique value to use for a primary key. Component Broker provides a mechanism called UUID (Universally Unique Identifier) that can be used for this purpose. UUID support is most useful when there are short-lived components that do not need to be found after they are created. The developer can use the provided mechanisms to create a specialized home that can create these components. The Component Broker framework provides a mechanism to generate a UUID for the component. This value must be generated on the server in order to guarantee uniqueness of the value on all Component Broker platforms.

This section describes the process of creating this particular type of home and follows the pattern utilized in “Creating specialized homes” on page 192.

Extending the interface

For the example, AgentUUIDHome is an interface that specializes the IHome interface with some specific create methods that generate UUID objects. The goal, as with any specialized home, is to provide methods specific to the application. These methods are shown in the IDL in the following simple example:

```
interface AgentUUIDHome : IManagedClient::IHome
{
    AgentUUID createAgentWithKey(in float commPercent,
                                in float pendingPaycheck,
                                in string<100> agentName )
        raises (IManagedClient::IInvalidKey, IManagedClient::IDuplicateKey);

    AgentUUID createAgentWithCopy(in float commPercent,
                                  in float pendingPaycheck,
                                  in string<100> agentName )
        raises (IManagedClient::IInvalidKey, IManagedClient::IDuplicateKey);
};
```

As with any specialized home, you must first decide which IHome interface should be specialized. This seems to have a simple solution: Inherit from the IHome in the managed object framework. Note that there are other home interfaces that support iteration and query. This functionality is not typically required or useful for most UUID implementations.

Details

The details behind this specialized home are the same as Details except for the discussion on queryable and iterable homes. As described above, it is possible to use these classes for as the parent class for UUID specialized homes but it is not typical.

The AgentUUIDHomeBO.idl file is shown in the following example:

```
#include <AgentUUIDHome.idl>
#include <IManagedAdvancedServer.idl>

interface AgentUUIDHomeBO : AgentUUIDHome,
    IManagedAdvancedServer::ISpecializedHome
{
};
```

The AgentUUIDHomeBO.ih file contains the AgentUUIDHomeBO_Impl class. This class inherits the implementation from the IManagedAdvancedServer::ISpecializedHome_Impl that it plans to extend. This is the same as was previously described in Creating specialized homes. The implementation interface is shown in the following example:

```
#ifdef SOMCBNOLOCALINCLUDES
#include <IManagedAdvancedServer.ih>
#include <AgentUUIDHomeBO.hh>
#else
#include "IManagedAdvancedServer.ih"
#include "AgentUUIDHomeBO.hh"
#endif

class AgentUUIDHomeBO_Impl : public virtual AgentUUIDHomeBO_Skeleton,
    public virtual IManagedAdvancedServer_ISpecializedHome_Impl
{
public:
    AgentUUIDHomeBO_Impl();
    virtual Agent_ptr createAgentWithKey
        (::CORBA::Float commPercent, ::CORBA::Float
        pendingPaycheck, const char* agentName);
    virtual Agent_ptr createAgentWithCopy
        (::CORBA::Float commPercent, ::CORBA::Float
        pendingPaycheck, const char* agentName);
    virtual ::CORBA::Void initForCreation
        (::IManagedServer::IDataObject_ptr theDO);
    virtual ::CORBA::Void uninitForDestruction();
    virtual ::CORBA::Void initForReactivation
        (::IManagedServer::IDataObject_ptr theDO);
```

```

virtual ::CORBA::Void uninitForPassivation();
virtual ::CORBA::Void syncToDataObject();
virtual ::CORBA::Void syncFromDataObject();

protected:

private:
    IManagedAdvancedServer::ISpecializedHomeDataObject_ptr iDataObject;
};

```

Implement the extended IHome interface

The IHome itself is a managed object and the development of an extended home should be done like creating any subclass of a managed object. This is described in “Extending a business object” on page 157. This section highlights things that are specific to the IHome interface and extension.

The implementation

As with any specialized home, the implementation of the new create methods involves careful usage of keys, copies, and the basic interface supported by home. All of the create methods end up using `createFromCopyString` or `createFromPrimaryKeyString`, passing a key or copy that has been loaded with the proper information to all proper creation of the object. The unique aspect of creating a UUID specialized home is that the keys and copy helpers will subclass special UUID classes provided by the Component Broker framework that provide special UUID support.

```

Agent_ptr AgentUUIDHomeBO_Impl::createAgentWithKey(::CORBA::Float
commPercent,::CORBA::Float pendingPaycheck,const char* agentName)
{
    //Version identifier DCE:4441AB44-5FA5-11d2-A40B-08005A49CF68:1
    // Insert Method modifications here
    ByteString_var theKeyString;
    IManagedClient::IManageable_var myVar;
    Agent_var newAgent;
    ::CORBA::Boolean bduplicat = 0;
    // create a UUID key
    IManagedAdvancedServer::IUUIDPrimaryKey_var newKey =
    IManagedAdvancedServer::IUUIDPrimaryKey::_create();
    newKey->generateUuid();
    theKeyString = newKey->toString();
    try
    {
        myVar = createFromPrimaryKeyString(theKeyString);
    }
    catch (...)
    {
        throw;
    }
    newAgent = Agent::_narrow(myVar);
    if (CORBA::is_nil(newAgent))
    {

```

```

        throw IManagedClient::IInvalidKey();
    }
    else
    {
        newAgent->commPercent(commPercent);
        newAgent->pendingPaycheck(pendingPaycheck);
        newAgent->agentName(agentName);
    }
    return Agent::_duplicate(newAgent);
    // End Method modifications here
}

Agent_ptr AgentUUIDHomeB0_Impl::createAgentWithCopy(::CORBA::Float
commPercent,::CORBA::Float pendingPaycheck,const char* agentName)
{
//Version identifier DCE:4441AB54-5FA5-11d2-A40B-08005A49CF68:1
// Insert Method modifications here
ByteString_var theCopyString;
IManagedClient::IManageable_var myVar;
Agent_ptr newAgent;
// create a UUID key
AgentCopy_var newCopy = AgentCopy::_create();
newCopy->commPercent(commPercent);
newCopy->pendingPaycheck(pendingPaycheck);
newCopy->agentName(agentName);
theCopyString = newCopy->toString();
try
{
    myVar = createFromCopyString(theCopyString);
}
catch (...)
{
    throw;
}
newAgent = Agent::_narrow(myVar);
if (CORBA::is_nil(newAgent))
{
    throw IManagedClient::IInvalidKey();
}
return Agent::_duplicate(newAgent);    }

```

Meet MOFW IManageable requirements

Because this extension to the home introduces no new additional methods and no new key for the home itself, `getPrimaryKeyString`, `getHandleString`, `externalize_to_stream`, and `internalize_from_stream` do not need to be implemented.

MOFW requirements - IManagedObject interfaces

These requirements are the same as any specialized home, see “Creating specialized homes” on page 192.

Keys

Keys are not used to create homes, however a special primary key class is used in the home to create instances of components that this home creates. This class is `IManagedAdvancedServer::IUUIDPrimaryKey`.

Copy helper

Copy helpers are not used to create homes, however a special copy helper class is subclassed by the component that this home creates. This class is `IManagedAdvancedServer::IUUIDCopyHelperBase`.

Leveraging server provided essential state extensions

These requirements are the same as any specialized home, see “Creating specialized homes” on page 192. Overriding the standard create interfaces These requirements are the same as any specialized home, see “Creating specialized homes” on page 192.

CICS and IMS application adaptor exception handling

The CICS and IMS application adaptor allows the leveraging of business logic which is implemented in the procedural data store rather than in the business object. This logic is made available on the business object through the use of pushdown methods. When the Component Broker client calls such a method on the business object, control is passed from the method in the business object, through the data object, the persistent object, and the procedural adaptor object to the procedure stored in the CICS or IMS system.

Error conditions arising from the state of the data in the procedural data store or with the parameters passed on the method call may require exceptions to be thrown. This section describes how to catch these exceptions and raise them to the Component Broker client which invoked the method.

Assume that the following procedural adaptor object has been created using the procedure outlined in for the CICS and IMS application adaptor:

```
package paa.exception1;

public class Exception1 extends Exception {}
public class Exception2 extends Exception {}
public class Exception3 extends Exception {}

class MyPAO :extends com.ibm.ivj.eab.paa.EntityProceduralAdapterObject
{
    private int getKeyValue();

    public void del() throws com.ibm.ipaa.IDataKeyNotFoundException;
```

```

public void insert() throws
    com.ibm.ipaa.IDataKeyAlreadyExistsException;
public void retrieve() throws com.ibm.ipaa.IDataKeyNotFoundException;
public void update() throws com.ibm.ipaa.IDataKeyNotFoundException;

public void pushdown1(int p1, boolean p2) throws Exception1,
    Exception2;
public void pushdown2(boolean p1) throws Exception2, Exception3;
public void pushdown3(char p1);
};

```

The procedural adaptor object provides a method to get the key property called *keyValue*. It provides the four methods to create (or insert), retrieve, update and delete the object in the data store. It also provides three methods, *pushdown1*, *pushdown2*, and *pushdown3* which can be invoked by a client program. The first two of these may throw exceptions depending on the state of the object and the input parameters. These exceptions are defined above.

The following are the steps taken to define the error conditions on the business object:

1. Start Object Builder for a new model. From the pop-up menu for User-Defined PA Schemas, select **Import > Bean**.
2. Type the bean name `paa.exception1.MyPAO`, and click **Next**.
3. Select a **Connector Type**, click **Next**.
4. Select *fieldKeyValue* as the key attribute, and click **Finish**.
5. Expand User-Defined PA Schemas, and from the pop-up menu of **MyPAOPO**, select **Add Data Object**.
 - a. Change the following:
 - 1) **Interface File Name** to **AcctDO**
 - 2) **Interface Name** to **AcctDO**
 - 3) **Implementation File Name** to **AcctDOImpl**
 - 4) **Implementation Name** to **AcctDOImpl**

Note: You are using Acct here but MyPAO has nothing to do with the MenuCustomer procedural application object.
 - b. Click **Finish**.
6. Expand the User-Defined Data Objects folder, > **called AcctDO** file. From the pop-up menu of **AcctDO**, select **Properties**.
7. Refer to the Constructs page. From the pop-up menu for the Constructs folder, select **Add Exception**.
 - a. Name the exception: type `Exception1`.
 - b. Repeat by adding two more exceptions, naming them `Exception2` and `Exception3`, respectively.
 - c. Click **Finish**.

8. From the pop-up menu for the data object interface `AcctDO`, select **Properties** again.
9. Refer to the `Methods` page and expand method `pushdown1`, and from the pop-up menu of the `Exceptions` folder, select **Add**.
 - a. Select `AcctDO AcctDO Exception1` as the name of the exception.
 - b. Repeat to add the exception `AcctDO AcctDO Exception2` to `pushdown1`.
 - c. Using these steps, also add exceptions `AcctDO AcctDO Exception2` and `AcctDO AcctDO Exception3` to method `pushdown2`.
 - d. Click **Finish**.
10. Select `AcctDOImpl` under the data object interface `AcctDO`.
11. In the `Methods` window, locate the User-Defined Method `void pushdown1(in long arg1, in boolean arg2)`. From the pop-up menu, select **Properties**.
12. Select the **Use the implementation defined in the Source pane** radio button. Click **Finish**.
13. In the `Source` window, type the following:

```

::CORBA::Void AcctDOImpl_Impl::pushdown1
  (::CORBA::Long arg1, ::CORBA::Boolean arg2)
{
    try {
        iMyPAOPO.pushdown1(arg1, arg2);
    } catch (MyPAOPOIF::Exception1 &exc) {
        AcctDO::Exception1 doExc;
        throw doExc;
    } catch (MyPAOPOIF::Exception2 &exc) {
        AcctDO::Exception2 doExc;
        throw doExc;
    } catch (CORBA::UserException &cue) {
        throw CORBA::UNKNOWN(0, CORBA::COMPLETED_NO);
    }
}

```

Note: The line `iMyPAOPO.pushdown1(arg1, arg2);` invokes the `pushdown1()` method on the procedural adaptor object shown above. It is not possible to catch the exceptions thrown by the procedural adapter object's `pushdown()` method since they are defined in Java, and the code above is written in C++. Instead, a class called `MyPAOPOIF` is defined to facilitate this. This class defines C++ versions of the exceptions, and converts the Java exceptions to their corresponding C++ exceptions.

14. In the `Methods` window, locate the User-Defined Method `void pushdown2(in boolean arg1)`. From the pop-up menu, select **Properties**.
15. Select the **Use the implementation defined in the Source pane** radio button. Click **Finish**.
16. In the `Source` window, type the following:

```

::CORBA::Void AcctDOImpl_Impl::pushdown2(::CORBA::Boolean arg1)
{
    try {
        iMyPAOPO.pushdown2(arg1);
    } catch (MyPAOPOIF::Exception2 &exc) {
        AcctDO::Exception2 doExc;
        throw doExc;
    } catch (MyPAOPOIF::Exception3 &exc) {
        AcctDO::Exception3 doExc;
        throw doExc;
    } catch (CORBA::UserException &cue) {
        throw CORBA::UNKNOWN(0, CORBA::COMPLETED_NO);
    }
}

```

17. In the **AcctDO** file, > User-Defined Data Objects folder.
 - a. Select **AcctDO** (the data object).
 - b. From the pop-up menu, select **Add Business Object**.
18. Refer to the Name and Data Access Pattern page, and select **Delegating for the pattern to Handle State Data**.
19. Refer to the Implementation Language page, and select **Java**.
20. Refer to the Attributes page..
 - a. In the **Add Attribute** section, select the **key** and the **copy helper** check boxes.
 - b. Click **Finish**.
21. Expand the User-Defined Business Objects folder > **Acct** file. From the pop-up menu of the business object interface **Acct**, select **Properties**.
22. Refer to the Constructs page > Constructs folder.
 - a. From the pop-up menu, select **Add Exception**.
 - b. Name the exception: type Exception1.
 - c. Repeat this process to add exceptions called Exception2 and Exception3.
 - d. Click **Finish**.
23. From the pop-up menu for the business object interface **Acct** , and select **Properties** again.
24. Refer to the Methods page and expand the method called pushdown1.
 - a. Select **AcctDO AcctDO Exception1**.
 - 1) Change the name of this exception by selecting **Acct Acct Exception1** in the **Exception Name** combination box.
 - b. Select **AcctDO AcctDO Exception2**.
 - 1) Change the name of this exception by selecting **Acct Acct Exception2** in the **Exception Name** combination box.
25. Expand the pushdown2 method.
 - a. Select **AcctDO AcctDO Exception2**.

- 1) Change the name of this exception by selecting **Acct Acct Exception2** in the **Exception Name** combination box.
 - b. Select the exception called **AcctDO AcctDO Exception3**.
 - 1) Change the name of this exception by selecting **Acct Acct Exception3** in the **Exception Name** combination box.
 - c. Click **Finish**.
26. Select the business object implementation called **AcctBO**.
27. In the Methods window, locate the User-Defined Method **void pushdown1(in long arg1, in boolean arg2)**.
- a. From the pop-up menu, select **Properties**.
 - b. Select the **Use the implementation defined in the Source pane** radio button. Click **Finish**.
 - c. In the Source window, type the following:


```
public void pushdown1(int arg1, boolean arg2) throws
    AcctPackage.Exception1, AcctPackage.Exception2
{
    try {
        iDataObject.pushdown1(arg1, arg2);
    } catch (AcctDOPackage.Exception1 doexc) {
        throw new AcctPackage.Exception1();
    } catch (AcctDOPackage.Exception2 doexc) {
        throw new AcctPackage.Exception2();
    }
}
```
28. In the Methods window, locate the User-Defined Method **void pushdown2(in boolean arg1)**.
- a. From the pop-up menu, select **Properties**.
 - b. Select the **Use the implementation defined in the Source pane** radio button.
 - c. Click **Finish**.
29. In the Source window, type the following:


```
public void pushdown2(boolean arg1) throws
    AcctPackage.Exception2, AcctPackage.Exception3
{
    try {
        iDataObject.pushdown2(arg1);
    } catch (AcctDOPackage.Exception2 doexc) {
        throw new AcctPackage.Exception2();
    } catch (AcctDOPackage.Exception3 doexc) {
        throw new AcctPackage.Exception3();
    }
}
```
30. From the pop-up menu for the data object implementation called **AcctDOImpl**, select **Properties**.
31. Refer to the Key and Copy helper page.

- a. Select **AcctKey** for the key.
 - b. Select **AcctCopy** for the copy helper.
 - c. Click **Finish**.
32. Add the managed object and generate the code. Create the makefiles, and then compile.

The exceptions thrown by the procedural adaptor object can now be caught by the Component Broker client as `Exception1`, `Exception2`, or `Exception3` defined on the business object interface `Acct`.

Inheritance of exceptions is not supported; all exceptions should derive directly from `java.lang.Exception`. Also, `Java.lang.Exception` is reserved for IBM use and should not be thrown by code written by the user; user written code should only throw exceptions defined by the user.

Chapter 8. MOFW – ActiveX client programming model

This chapter discusses the specifics of:

- Using an ActiveX client to access managed objects on a Component Broker server.
- Developing code using managed object proxies in an ActiveX environment

ActiveX client view of Component Broker applications

The intent is to keep the interface to Component Broker objects as familiar to the ActiveX developer as possible.

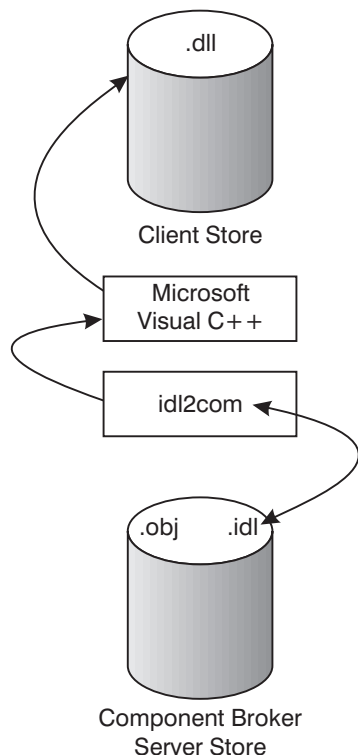


Figure 46. ActiveX client view

The development view of the managed object (MO) proxies that are used on the client platform starts on the server. When developers are programming managed objects for Component Broker, they are obliged to create IDL files

which represent those managed objects. Typically, the Component Broker developers use the Object Builder to help them coordinate these files, but that is not required.

The key, for an ActiveX client developer, is that the IDL for the MOs must be run through the `idl2com` compiler. The `idl2com` compiler takes the interface definitions in the IDL, and produces the managed object proxies that are used from the client. The proxies are produced in C++. Supporting files that need to be compiled are placed on the client for client applications to use. See *Developing the Component Broker ActiveX client* for further information.

As in Figure 46 on page 211, the route to producing the C++ files for the managed object proxies is by calling `idl2com`.

Once the C++ files for the managed object proxy are available, the proxy class can be treated as any other C++ class, with the exception of the specific requirements added to the class based on its involvement with the CORBA based Component Broker server. These are detailed throughout this chapter. Look at the run-time elements of the solution.

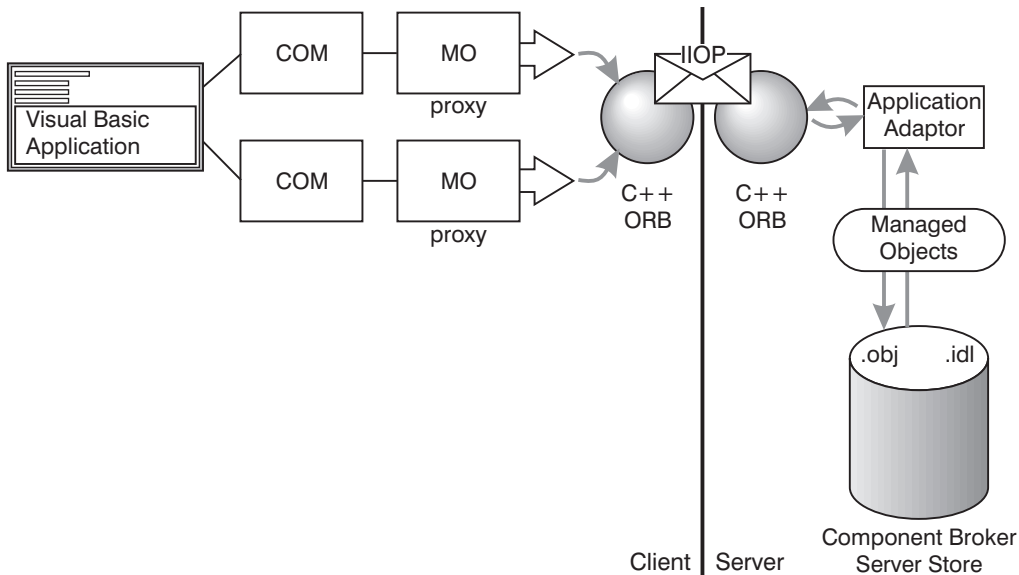


Figure 47. ActiveX client run-time scenario

A run-time scenario for an ActiveX client begins when the managed object proxy is accessed. All managed objects accessed from the ActiveX client are *wrapped* at run time by a COM object. The first order of business is to connect to the ORB and establish a reference to the root of the system name space. From there, the issues of finding and using the managed objects through the

proxy are issues related to how the Component Broker programming model is set up. Only the syntax changes for the ActiveX environment.

Developing the Component Broker ActiveX client

To produce a DLL containing an ActiveX accessible interface to a CORBA object, perform the following basic steps:

1. Create your IDL.
2. Process IDL using `idlc` to produce the client side bindings.
3. Process IDL to generate the `.ih` and `_I.cpp` files if needed.
4. Generate a GUID using `guidgen.exe`.
5. Process IDL using `idl2com` to produce the ActiveX accessible interfaces.
6. Compile and link the bindings.
7. Register the interfaces using `regsvr32`.

At this point the interfaces are now available for application usage.

Generating and registering DLLs

More detailed steps on generating and registering DLLs follow:

1. Create your IDL either manually or by using Object Builder. See “Object Builder” in *WebSphere Application Server Enterprise Edition Component Broker Application Development Tools Guide* and *WebSphere Application Server Enterprise Edition Component Broker for Windows NT and AIX Online Documentation* for details. If you choose to use Object Builder to create your IDL, it is important to know that you cannot use the client-side usage bindings generated by Object Builder. You can only use the `*.idl`, `*_I.cpp` and `*.ih` files. The `*_I.cpp` and `*.ih` files are only needed when you have defined `localonly` objects.
2. Emit the client-side usage bindings from the IDL. Use the `idlc` command to emit the client-side usage bindings from the IDL, specifying the `-mcpponly` and `-suc:hh` options on the command. For example:

```
idlc -mcpponly -suc:hh Policy.idl
```

Important: If you did not emit the client-side usage bindings into your ActiveX build directory, you must copy or move them there now.

See The `idlc` command for details on `idlc`.

3. If you have defined `localonly` objects (that is your IDL contains `#pragma localonly`), you need to create the implementation headers and templates from the IDL. This step is **only** necessary if you have defined `localonly` objects. You have two options:
 - If you used Object Builder to create the IDL, use the `*.ih` and `*_I.cpp` files that Object Builder automatically generated.

- Otherwise, manually generate the *.ih and *_I.cpp files using the idlc command.

If using the idlc command to produce the *.ih and *_I.cpp files from the IDL, specify the -mcpponly and -sih:ic options on the command. For example:

```
idlc -mcpponly -sih:ic Policy.idl
```

You must now copy or move the *_I.cpp and *.ih files into your ActiveX build directory.

4. Generate a GUID using the guidgen.exe utility included with Microsoft Visual C++.
5. Use the idl2com command to produce the ActiveX accessible interfaces from the IDL. For example:

```
idl2com -g AE3E2131-C6DE-11d0-92AF-08005ACE818D Policy.idl
```

See The idl2com Command for details on idl2com.

6. Use nmake -f [filename] to compile and link the ActiveX accessible interface DLL for the CORBA object. Use the IDL name, appended with the .mak extension for [filename]. For example:

```
nmake -f Policy.mak
```

When producing the DLL for the CORBA object, ensure that all required IDL for the object (including IDL referenced by the object's IDL) has been processed by the idl2com command into the same directory. This ensures that the correct header and library files are available when the .mak file is processed.

7. Use regsvr32 [filename] to register the DLL in the Windows system registry. Use the IDL name, appended with the .dll extension for [filename]. For example:

```
regsvr32 Policy.dll
```

Unregistering and moving DLLs

If you find in the future that you wish to move or no longer need the DLL(s) you have registered, perform the following steps:

1. Use regsvr32 /u [filename] to unregister the DLL. Use the IDL name, appended with the .dll extension for [filename]. If you no longer need the DLL, you are finished.
2. If the DLL produced in step 6 above is to be moved to a different directory, do the following steps:
 - a. Move the DLL file and its corresponding TLB file to the new directory. The TLB file was generated during the makefile processing (step 4 above).

- b. Reregister the DLL.

Component Broker ActiveX client application development information

Using COM or OLE objects

Conformance to the OMG's COM-CORBA Interworking Part A specification is not complete. CORBA objects can be accessed through COM and OLE automation-produced interfaces, but CORBA objects cannot access COM or OLE objects.

Using OLE automation interfaces

While the produced OLE Automation interfaces are intended to be generic OLE automation interfaces available to any OLE controller, only Visual Basic 5.0 has been used to test these at this time.

Using remote CORBA objects

During installation, several OMG COM-CORBA Interworking Specification interfaces are installed and registered. Of key importance are the CORBA Factory interfaces: `GetObject` and `CreateObject`. Use these interfaces to get started using remote CORBA objects. The samples provide guidance on using these interfaces.

Using IBM-supplied COM wrappers

During installation, several pre-built COM/OLE Automation interface DLLs for some of the Object Services are provided and registered for you. These DLLs are contained within the installed bin directory, and supporting TLB, LIB, and header files are installed as well. The specific list of DLLs which are shipped is contained in the `RegActX.bat` file found in the installed bin directory. This file registers the DLLs.

Using VBScript and Internet Explorer

A sample is provided that shows the use of a CORBA object from VBScript and Internet Explorer. This sample is available in the samples subdirectory (`InstallVerification/ProgrammingModel/Applications/ActiveX`) under the directory where Component Broker was installed.

Only the following data types may be passed to a CORBA object from a VBScript script:

- primitives (ie. long, short, char, boolean, etc)
- strings
- objects
- structures containing only primitives, strings or objects
- unions containing only primitives, strings or objects

Do not use arrays or sequences in a VBScript for this release.

Exception handling

In Visual Basic, there are two means of working with exceptions:

1. You can pass an “exception object” (a VARIANT type) as the last parameter to a method. The last parameter is always optional. If an exception occurs, this object will be filled in with the necessary exception information (a CORBAException object.).
2. You can pass a null for the last parameter which tells the bindings that there is no exception object to be filled in. The method instead needs to pass back an HRESULT which is then mapped by Visual Basic. The Visual Basic application has to put in place the ON ERROR code (similar to a C++ or Java try/catch block) to catch the exception.

Client programming model: basic tasks

In “Chapter 4. MOFW C++ client programming model” on page 85, the following tasks that a client is likely to want to do are discussed:

- Find an object.
- Use an object.
- Create an object.
- Use a set of object.
- Remember an object.
- Release or delete an object.

This chapter explores the same topics (with the same examples) but presents them from the perspective of ActiveX development. The examples are primarily in Visual Basic. An ActiveX client application developer can use the COM class output of the idl2com compiler directly from C++, and following COM rules, as well.

Initializing the Component Broker client environment

Initializing the Component Broker client environment is discussed in detail in “Chapter 4. MOFW C++ client programming model” on page 85. Initialization provides access to the ORB and Naming Service.

The following Visual Basic code accomplishes this task:

```
Dim corbaFactory as Object
Dim orb as Object
Dim NameService as Object
Dim genericObject as Object

Set corbaFactory = CreateObject("CORBA.Factory")
Set orb = CreateObject("CORBA.ORB")
```

```
Set NameService = CreateObject("IDL:IEExtendedNaming.NamingContext")

Set genericObject = orb.ResolveInitialReference("NameService")
NameService.narrow genericObject
```

After initialization, the client has access to the *orb* and *nameService* variables. You can see how the client might use these in subsequent sections.

Remember that the *IEExtendedNaming* is used to make things go a little easier.

The previous example shows how you can initialize the environment, and specifically, how access to the *NameService* can be achieved. However, see “Finding a managed object using the Naming Service” on page 218 for information on how access to objects can be obtained without these steps, if the object is well-known in the name space. The ActiveX client run time performs the previous sequence of operations automatically on *CreateObject/GetObject* methods calls to the *Corba.Factory COM* object.

Finding a managed object

When a client application starts, it has no knowledge of any objects. In a non-distributed environment, the application would create any objects it needs. These newly created objects would reside on the same system as the application. If the application needs these newly created objects to be available the next time the application runs, the application itself is responsible for making these objects persistent.

However, Component Broker is a distributed environment. In this environment, objects are not created on a client system. Objects are created from a client system on a server system. Object persistence is managed with the help of the Component Broker run time.

The client accesses the component through the components managed object. But first, it has to find the managed object.

There are two ways for a client application to find a managed object:

- Use the Naming Service.

If the object has a *Name*, the client can use the Naming Service to locate the object by its *Name*. In general, the Naming Service only contains a subset of the objects in a distributed system. This subset consists of well-known objects, such as collections of business objects or important objects in the object model.

- Use another object.

If the object does not have a *Name*, the client can find the object by using the Naming Service to find a well-known object, such as a *factory* or collection, and then use this object to implement methods that return the object as a return value or output parameter.—

Both techniques involve the use of the Naming Service. In the Component Broker distributed environment, all work with objects in a client application must begin by using the Naming Service.

Finding a managed object using the Naming Service

Assume the insurance company in the example placed several important Claim objects in the Naming Service. The following Visual Basic code example shows how to find such a Claim, belonging to a customer named Lou.

```
' nameService initialized as above

Dim tempObj as Object
Dim louClaim as Object
Set louClaim = CreateObject("IDL:Claim")
Set tempObj = nameService.resolve_with_string
                ("cell/Applications/LifeInsurance/Claim/LouClaim")
louClaim.narrow tempObj
Set tempObj = Nothing
...
' No longer need louClaim
Set louClaim = Nothing
```

The Component Broker server run time initializes the global object instance nameService to refer to the root of the installation's Naming Service.

For a refresher on determining the naming context, and details on how the name space is specified, see "Chapter 4. MOFW C++ client programming model" on page 85.

Recall that in addition to the resolve_with_string() method, Naming Contexts also support the bind_with_string() method which associates a name with an object instance. The following Visual Basic code example could have been used to name the Lou Claim object.

```
' Declare and create louClaim prior to the following code segment
...
' Add Lou to the Name Space
nameService.bind_with_string
                ("cell/Applications/LifeInsurance/Claim/LouClaim", louClaim);
```

For additional information on the Component Broker Naming Service, see References in Component Broker for Windows NT and AIX Online Documentation.

Finding a managed object using the primary key

The preceding example was simplified in that LouClaim was in the Naming Service. In the event that you do not have a well-known name for the claim, you can use homes to find a specific claim.

Homes are instances of the IHome class. The managed object provider might decide to implement and provide a tailored subclass of IHome, or might use an instance of the base class. The relationship between managed objects and collections is explained in “Using sets of objects” on page 225.

For the overview, you can find the home for Claim objects by the following Visual Basic code segment

```
Dim claimHome as Object
Dim myFinder as Object
Dim tempObj as Object

Set claimHome = CreateObject("IDL:Claim")
Set myFinder = CreateObject("IDL:IExtendedLifecycle.FactoryFinder")
Set tempObj = nameService.resolve_with_string
               ("cell/Applications/LifeInsurance/Homes/Claim")
myFinder.narrow tempObj
Set tempObj = Nothing
Set tempObj = myFinder.find_factory_from_string("Claim.object interface")
claimHome.narrow tempObj
Set tempObj = Nothing
...
' claimHome is now usable
...
' No longer need myFinder
Set myFinder = Nothing
```

Now you need to find Lou’s Claim. If you know the Claim number, all you need is the primary key helper class for the Claim. Every managed object class has a set of local helper classes that let you use its keys. An instance of a key helper class is always local to the client’s process.

The provider of the managed object that you are working with has given you access to source for, or actual .DLL files containing the code for the key helper class. Regardless, it is up to you to ensure that you have access to them and can use them in your applications.

Key helpers, like all helper classes are created with a static method on the class named _create(). This static method gets generated by the bindings that accompany all subclasses of ILocalOnly. As an ActiveX client programmer, you can create a helper object by instantiating the ActiveX wrapper for the helper class.

Having created an instance of a primary key, the key must be set by one or more attributes on the primary key object. When all of the *key* attributes have been set, the primary key object is now usable. The Claim home uses this primary key to find the previously created Claim object. Remember, the primary key is on the client system, but the Claim object and the Claim home are on the server system. If the client passes a primary key object as a parameter to the home, and the home is on a remote system, the remote system might get a proxy back to the client's primary key instance. This would turn the client into a server and unpredictable results could occur. Therefore, the Component Broker programming model uses strings as the method for passing keys to potentially remote objects.

Continuing the example, the following Visual Basic code segment would find Lou's Claim in the home (assuming his number is 1234).

```
' Create an instance of the Key Helper Class
Dim ClaimPrimaryKey as Object
Set ClaimPrimaryKey = CreateObject("IDL:ClaimPrimaryKey")

' Set the claimNo attribute in the key
claimPrimaryKey.ClaimNo = 1234

' Get the data out of the Stream to go onto the wire to the server
Dim claimStringvar as Variant.
claimStringvar = claimPrimaryKey.toString()

' Turn the variant returned by toString() method into a safearray of shorts
Dim claimString() as Integer
ReDim claimString(UBound(claimStringVar))
For counter = 0 to UBound(claimStringVar)
    claimString(counter) = claimStringvar(counter)
next counter

' Call find by key on the Home to find Lou's Claim
Dim tempObj as Object
Dim louClaim as Object = CreateObject("IDL:Claim")
Set tempObj = claimHome.findByPrimaryKeyString(claimString)
louClaim.narrow tempObj
Set tempObj = Nothing
...
' No longer need louClaim
Set louClaim = Nothing
```

The object provider of a public managed object always provides you with a set of helper classes for using the homes that contain his managed objects. There is always exactly one primary key helper class. The object provider gives a client developer the:

- The interface definitions for the key classes. In addition to the primary key, there might be a secondary key helper classes. A secondary key might also uniquely identify an object, or multiple instances may have the same value.
- An implementation of the key classes.

- Documentation for their use.

Finding a managed object using another managed object

Once you have an object, you can use its methods to find related objects. Continuing the previous example, after finding Lou's Claim, you can find other objects that the Claim references. The following Visual Basic example gets you a reference to Lou's Policy:

```
' Find Lou's Policy
Dim louPolicy as Object
Set louPolicy = louClaim.policy
```

Using a managed object

When you find a reference to a managed object, you can invoke methods on it. For example,

```
person.set_Name("Lou Smith");
```

Calls the set_Name method on the person object identified by the person. The Component Broker internal implementation handles the use of remote objects.

Creating a managed object

Component Broker managed objects can be created in a number of ways. The following sections describe the default way in which you can easily create managed objects.

Creating a new object – create from key

Every Component Broker managed object class has an instance of a Factory associated with it. The Factory provides a set of interfaces for creating instances of a managed object. The Factory gets some of its interface from the base class CosLifeCycle::GenericFactory. The createFromPrimaryKeyString method is introduced in the IManagedClient::IHome interface supplied by Component Broker. This interface specializes the CosLifeCycle::GenericFactory interface and plays the role of factory for Component Broker managed objects. Object providers might implement and provide a tailored subclass of this interface, or might use the implementation of IHome provided. You need to know how to find the right IHome for creation. Homes are at well-known locations in the Naming Service. The input required for the factory finder is the name of the interface of the class that you want this factory to make instances of. The following Visual Basic code fragment gets a reference to the Claim Factory for the Life Insurance Application.

```

Dim myFinder as Object = CreateObject(
    "IDL:ExtendedLifeCycle.FactoryFinder")
Dim tempObj as Object
Set tempObj = nameService.resolve_with_string
    ("cell/Applications/LifeInsurance/FactoryFinders")
myFinder.narrow tempObj
Set tempObj = Nothing

Dim claimHome as Object = CreateObject("IDL:ClaimHome")
Set tempObj = myFinder.find_factory_from_string("Claim.object interface")
claimHome.narrow tempObj
Set tempObj = Nothing
    ...
' No longer need claimHome
Set claimHome = Nothing

```

You need to provide the IHome with information necessary to manufacture a new object instance. At a minimum, the primary key must be provided. A Visual Basic example of creating a new Claim with a *claimNo* of 1234 is:

```

' Create an instance of the Primary Key Helper Class
Dim ClaimPrimaryKey as Object
Set ClaimPrimaryKey = CreateObject("IDL:ClaimPrimaryKey")

' Set the claimNo attribute in the key
claimPrimaryKey.ClaimNo = 1234

' Get the data out of the Stream to go onto the wire to the server
Dim claimStringvar as Variant.
claimStringvar = claimPrimaryKey.toString()

' Turn the variant returned by toString() method into
' safearray of shorts
Dim claimString() as Integer
ReDim claimString(UBound(claimStringVar))
For counter = 0 to UBound(claimStringVar)
    claimString(counter) = claimStringvar(counter)
next counter

' Call createFromPrimaryKeyString on the Factory to create Lou's Claim
Dim tempObj as Object
Dim theClaim as Object = CreateObject("IDL:Claim")
Set tempObj = claimHome.createFromPrimaryKeyString(claimString)
theClaim.narrow tempObj
Set tempObj = Nothing
    ...
' No longer need theClaim
Set theClaim = Nothing

```

The previous two examples are almost identical. Create a primary key object to define the identity of the object that is made. Then, the `createFromPrimaryKeyString` call is made on the home and the key is passed as a string.

The `createFromPrimaryKeyString` method is defined by the `IHome` class, and all business objects can be created by this method.

An object provider might provide you with a subclass that introduces other, easier to use creation methods. Additional create methods are described, with examples, in “Chapter 6. MOFW - C++ client programming model – advanced concepts” on page 133.

Creating a new object - create from copy

Setting and getting the attributes of a managed object can be expensive. There are two main reasons for this. First, if the managed object is implemented in another language, each get or set method is actually a cross-language call. Cross language calls are more expensive than simple, same language calls. The get and set overhead is even more expensive if the managed object is remote because each call is actually a remote procedure call and involves significant overhead. Consider the following code segment:

```
' Assume that 'theClaim' is declared and created as in the
' previous code segment.

' Creating 'theClaim' above required one RPC. Setting the rest of the
' objects attributes involves one RPC per attribute. The following
' lines of code show four such RPC's. This could, of course, be any
' number of RPC's, depending on the complexity of the object.

' Now initialize the Claim's attributes
theClaim.date = "10/14/96"
theClaim.state= entered
theClaim.reason = accident
theClaim.description = "Side-swiped by teenager in a red convertible."
```

This segment could involve the following remote method calls:

- The client to home of Claims to create the Claim.
- The client to Claim managed object to set its *date* attribute.
- The client to Claim managed object to set its *state* attribute.
- The client to Claim managed object to set its *reason* attribute.
- The client to Claim managed object to set its *description* attribute.

The previous calls could be reduced to a single remote method call by using the `createFromCopyString()` method on an `IHome` instead of the `createFromPrimaryKeyString()` method. To use the `createFromCopyString()` method, the object provider must provide you with a copy helper class. The following Visual Basic code segment presents the code from the previous example rewritten using this design pattern.

```
' Create a new "local" Claim in my process and language.
' Use a Copy Helper Class that the Claim MO provider gave me.
Dim ClaimCopy as Object
Set ClaimCopy = CreateObject("IDL:ClaimCopy")
```

```

' Now initialize the Claim's attributes. Note that these methods
' execute locally, within the same language.
claimCopy.date = "10/14/96"
claimCopy.state = entered
claimCopy.reason= accident
claimCopy.description= "Side-swiped by teen-ager in a red convertible."

' Pass this local copy to the Home and have him return a new Claim
' MO whose attributes are initialized from the local copy's values.
' Since not all ORBs support Pass-By-Value, we first convert the
' local copy helper object to a string.
Dim tempObj as Object
Dim theClaim as Object = CreateObject("IDL:Claim")
Dim claimStringvar as Variant.

claimStringvar = claimCopy.toString()

' Turn the variant returned by toString() method into
' safearray of shorts
Dim claimString() as Integer
ReDim claimString(UBound(claimStringVar))
For counter = 0 to UBound(claimStringVar)
    claimString(counter) = claimStringvar(counter)
next counter

Set tempObj = claimHome.createFromCopyString(claimString)
theClaim.narrow tempObj

Set tempObj = Nothing
Set classClaimCopy = Nothing
Set claimCopy = Nothing
...
' No longer need theClaim
Set theClaim = Nothing

```

Like a key helper class, a copy helper class instance is always local to your process and implemented in the language you are using. The object provider gives you the interface and implementation of the helper class.

Copy helper classes are especially useful if the client application needs to interact with an object during initialization, and then create a managed object from the attributes. A common scenario for this is entering data for the object from a GUI. The GUI updates the local Copy Helper Object, and then the `createFromCopyString()` method is called when you click **Do** on the end user interface (EUI).

Using sets of objects

An IHome represents a set of managed objects, all of the same type, whose relationship to one another is defined by the object provider, and maintained by the fact that they were all created in the same home. Sometimes an application needs to define (and manage) the relationships between managed objects, based on the particular business task at hand. This might even include relationships between managed objects of different types (for example, PolicyHolder and Beneficiary). This can be done using an IManagedReference Collection, as shown in the following Visual Basic code segment:

```
Dim mixedCollection as Object
Set mixedCollection = CreateObject
                        ("IDL:IManagedCollections.IReferenceCollection")

' Find a collection in the name space which contains PolicyHolders and
' Beneficiaries

' Create an iterator on the reference collection that was found above.
'When an iterator is created, it is automatically positioned preceding
'the first element.
Dim anIterator as Object
Set anIterator = mixedCollection.createIterator()

Dim element as Object
Set element = CreateObject("IDL:IManagedClient.IManageable")

' Loop through the collection. The "next" method advances the iterator
' to the next element (on the first invocation, this will advance to the
' first element). If the iterator is now past the end of the collection,
' the "next" method returns FALSE; otherwise, returns TRUE and sets the
' output parameter to the element at which it is now positioned.
Do While anIterator.next(element)
    If element.is_A(PolicyHolder) Then
        ' Send him a bill
    End If
    If element.is_A(Beneficiary) Then
        ' Send him a check
    End If
Loop
```

The combination of the IManagedCollections::IReferenceCollection and the IManagedCollections::IIterator were used in the previous code segment. An IManagedCollections::IReferenceCollection is a generalized collection of object references that is iterable. IManagedCollections::Iterator supports advancement of the iterator and retrieval of elements by the next() method. IManagedCollections::IReferenceCollection supports adding and removing elements using the addElement() and the removeElement() methods. IManagedCollections::IManagedReferenceCollection is the most basic kind of collection supported in Component Broker. Combining this with the capabilities of IHome provides the basis for writing simple applications and

the foundations for the more advanced query and collections capabilities provided by Component Broker. For more information on collections and query, see “Chapter 6. MOFW - C++ client programming model – advanced concepts” on page 133.

Transient sets

Component Broker supports transient collections. Transient collections do not require transactions; they reside in transient containers and thus have no dependency or overhead on database connections. Database and transient collections provide identical programming interfaces. They differ only in their persistence characteristics. Users can find transient collections by passing an interface string to the factory finder as defined in the next section.

Specifying reference collection interfaces

A variety of interfaces for use with factory finders are provided for specifying the type of reference collections that you want to use.

For example, calling `find_factory_from_string` passing the argument `IManagedCollections::IReferenceCollection.object interface/PersistentReferenceCollectionFactory.object home`, as in the previous example, creates a DB2 backed reference collection. Using the generic string `IManagedCollections::IReferenceCollection.object interface` returns whichever collection is configured as the default.

Transient Collections

```
IManagedCollections::IReferenceCollection.object interface/  
TransientReferenceCollectionFactory.object home
```

Transient Collections Keyed

```
IManagedCollections::IKeyedReferenceCollection.object interface/  
TransientKeyedReferenceCollectionFactory.object home
```

Persistent Collections

```
IManagedCollections::IReferenceCollection.object interface/  
PersistentReferenceCollectionFactory.object home
```

Persistent Collections Keyed

```
IManagedCollections::IKeyedReferenceCollection.object interface/  
PersistentKeyedReferenceCollectionFactory.object home
```

Remembering your favorite objects

Component Broker allows you to remember a managed object instance, by introducing the concept of an object reference. An object reference is opaque and you cannot set its internal structure. However, a reference always and uniquely refers to a managed object regardless of where it resides in the network.

Continuing the example, assume that you perform the following Visual Basic code segment.

```
' This example just creates the string containing the object reference
' This string could be written to a file using normal VB conventions if
' desired

' Get a "string" version of my reference to Robert
' robert points to Robert or a proxy to Robert

Dim corbaFactory as Object
Set corbaFactory = CreateObject("CORBA.Factory")

Dim orb as Object
Set orb = corbaFactory.GetObject("CORBA.ORB.2") ' Get access to the orb

Dim robertStringifiedReference as String
robertStringifiedReference = orb.ObjectToString(robert)

' Save the string to a file using normal VB file IO

' I do not need Robert anymore
set robert = nothing
```

If you save a reference to Robert as a Stringified Object Reference, then you can use this string to re-access Robert at a later time. The following Visual Basic code segment presents an example of re-accessing the Robert object.

```
' Get back the string from the file I saved earlier into
' robertStringifiedReference

' Make an object reference for Robert
Dim tempObj as Object
Dim robert as Object = CreateObject("IDL:Robert")
Set tempObj = orb.StringToObject(robertStringifiedReference)
robert.narrow tempObj

' I can now work with Robert.
robert.Name()
```

Releasing and deleting objects

Eventually, you no longer need to use an object that was created or found. Component Broker supports two interpretations on “no longer needs”.

- The `remove()` method deletes the object and its persistent instance data.
- The `release()` method informs the object that the client application no longer plans to reference the object. The object still exists in the server, and other applications may be using it, but the client calling the `release()` method is done with the object.

For a detailed description of the differences between releasing and removing objects, see “Releasing and deleting objects” on page 65.

When references explode

Another application may call the `remove()` method on objects referenced by a client application. Component Broker supports a Concurrency Control Service to regulate sharing of objects. If a client does not lock an object, it can be deleted. Language environments like ActiveX prevent this from happening in simple, single process applications, but achieving the same level of function is impossible in a cross-system, multi-language environment. If you are operationally reusing data and applications, some legacy code not under the object-oriented application’s control can delete the instance data for objects. An object to which you have a reference can vanish unless you use Concurrency Control.

Therefore, if you do not use Concurrency Control, you need to be prepared for exceptions. If an inappropriate object reference is used, an exception is thrown. While not perfect, this is a substantial improvement over the semantics of languages like C++ that can result in unpredictable behavior if inappropriate references are used.

Summary: The client programmer’s check list

When using a set of components to implement an application, you need to know how to use the following things about components and the managed object framework which encapsulates and interoperates with the component objects. Table 5 on page 229 is organized around the tasks outlined in “Client programming model: basic tasks” on page 216. Note that `xxx` is used to represent the name of the component. For example, Claim and Policy were used in this chapter as examples of domain-specific components. This also

provides a good clue as to whether or not the classes were constructed by the object provider or whether they came as part of the Component Broker system.

Table 5. Summary of interfaces

Task	Class Types	Methods	Purpose
General Use	StandardSyntaxModel	stringToName	Converts strings to CORBA name types.
Find <i>xxx</i> component	NamingContext	resolve	Finds an object in the namespace
	IHome	findByPrimaryKeyString	Finds an object in a home
	<i>xxx</i> PrimaryKey	set <i>xxx</i> methods and toString	Sets the key value
Create or delete an <i>xxx</i> component	FactoryFinder	findFactory	Finds the factory object.
	IHome	createFromKeyString	Creates an object with a key only.
	<i>xxx</i> PrimaryKey	set <i>xxx</i> and toString	Passes information on key to the server.
	<i>xxx</i> CopyHelper	set <i>xxx</i> and toString	Passes copy information to the server.
	any <i>xxx</i> MO (managed object) subclass	remove and release	Deletes or releases object from memory.
Use sets of components	IReferenceCollection	createIterator	Creates an iterator object.
	IIterator	next	Iterates over a collection and gets the objects.
Remember interesting and important components	NamingContext	bind	Binds an object into the name space with a name.
		resolve	Finds an object in the name space by its name.
	theORB	object_to_string	Creates a stringified form of an object reference.
		string_to_object	Creates an object reference from a string.

Chapter 9. MOFW - Java client programming model

This chapter deals with the specifics of using a Java client to access managed objects on a Component Broker server.

Throughout this chapter, there are references to one of the components of the solution as the `idl.toJava` (`com.ibm.idl.toJava.Compile` is the full name) compiler. IBM's compiler is based on Java.

Java client view of Component Broker applications

The development view of the managed object (MO) proxies that are used on the client platform starts on the server. When developers are programming managed objects for Component Broker, they are obliged to create IDL files that represent those managed objects. Typically, the Component Broker developers use the Object Builder to help them coordinate these files, but that is not required.

The key, for a Java client developer, is to run the IDL for the managed objects through the `idl.toJava` compiler. The `idl.toJava` compiler takes the interface definitions in the `idl`, and produces the managed object proxies that are used from the client. The proxies are produced in the form of `.java` files, that need to be compiled and placed either in a location where the Web server has access for applets, or on the client for applications.

As in Figure 48 on page 232, there are two routes to producing the `.java` files for the managed object proxies. Whether you process the IDL files generated by the Object Builder tool, or process IDL files that you wrote, the results should be the same.




Figure 48. Java client view

When the .class files for the managed object proxy are available, the proxy class can be treated as any other Java .class, with the exception of the specific requirements added to the class, based on its involvement with the CORBA-based Component Broker server. These are explained in this chapter. Look at the run-time elements of the solution.

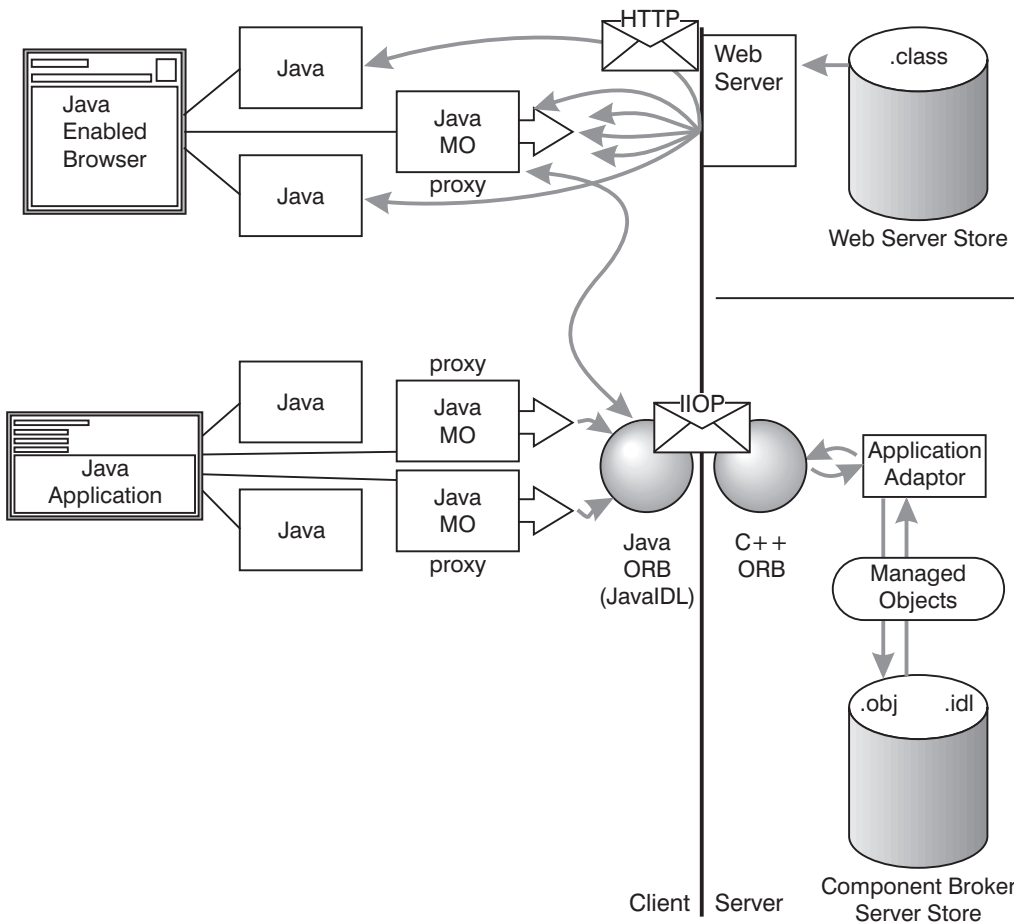


Figure 49. Java client run-time scenario

A run-time scenario for a Java client begins when either the applet, or the application containing the managed object proxy is accessed. In either case, the first order of business is to connect to the ORB, and establish a naming context. From there, the issues of finding and using the managed objects through the proxy are really issues related to how the Component Broker programming model is set up. Only the syntax changes for the Java environment.

Since only one Object Request Broker (ORB) is started per Java Virtual Machine (JVM), the ORB will be shared among all applets running in the JVM. Applets intended to run with other applets in the same Java Virtual Machine should be designed to efficiently and safely use the shared ORB.

Preparing to use VisualAge for Java for development of Component Broker Java clients

Before using VisualAge for Java for development of Component Broker Java clients, it is necessary to import the IBM Java ORB into a VisualAge for Java project. Do this using the following steps:

1. Start VisualAge for Java.
2. Insert your Component Broker Client CD into your CD drive.
3. From the VisualAge for Java Workbench, select **File > Import**.
4. Select "Jar file" as the input source.
5. Browse to your CD drive and navigate to the \jclient directory.
6. Select the file somojor.zip.
7. Select .class and resource as the file types to import.
8. Enter the name of a new project or browse to an existing one.
9. Click **Finish**.

Since somojor.zip is a large file, it may take a few minutes for the importation process to complete.

Preparing managed objects for remote access

The client programming model is based on the idea of having remote proxies available for managed objects on the server which are represented in IDL files. Depending on how managed objects are created in your organization, you may not be the one who performs this step. Even if that is the case, it is helpful to know what is going on.

The first thing that you need to do is to make sure that the proxy is available to the client. Given that there is an IDL definition for the managed object, there are two ways to generate the Java proxy for the client. The first way to generate a proxy is by instructing Object Builder that you would like a Java client proxy for the managed object you are dealing with. The second way to generate the Java client proxies is to use the The IDL-to-Java compiler directly.

In most cases, the invocation of this compiler is straightforward. Regardless of which method you use, the tools should generate several Java proxy files. The rest of this chapter describes how to use these files.

Depending on whether you intend to use the proxies in an application, or in an applet served up by a Web server, you have to decide where to put, or where to tell the tools to put, the proxies. When you have made these decisions, and generated the proxies, you can start developing the code that uses the proxies for the managed objects on the server.

Client programming model: basic tasks

“Chapter 4. MOFW C++ client programming model” on page 85 describes the following tasks that a client is likely to want to do:

- Find an object.
- Use an object.
- Create or delete an object.
- Use a set of objects.
- Remember an object.
- Release and delete objects.

This chapter explores the same kinds of tasks as in “Chapter 4. MOFW C++ client programming model” on page 85, with the same examples, but presents them from the unique perspective of a Java client developer. There is, however, one additional step that is discussed prior to getting into the programming model for the client.

Initializing the Component Broker client environment

The `CBSeriesGlobal` interface is provided as a convenience. A Component Broker client could make the same set of calls as is encapsulated in the `Initialize()` method independently, and will work in some cases, however it is highly recommended that the `CBSeriesGlobal` interface be used in the Component Broker client. To use the `CBSeriesGlobal` interface, you must include the following line:

```
import com.ibm.CBCUtil.CBSeriesGlobal;
```

The `CBSeriesGlobal` interface provides a number of different `Initialize` methods depending on your needs as a Java *applet* or application. All of these `Initialize` methods encapsulate the calls to initialize the ORB and get an initial reference to the Naming Service. Therefore, initializing a Component Broker Java *application* client requires a single line of code, one of the following:

```
CBSeriesGlobal.Initialize(host);  
CBSeriesGlobal.Initialize(host, port);  
CBSeriesGlobal.Initialize(arguments);  
CBSeriesGlobal.Initialize(arguments, properties);
```

To initialize a Component Broker Java *applet* client also requires a single line of code:

```
CBSeriesGlobal.Initialize(applet);  
CBSeriesGlobal.Initialize(applet, properties);
```

Only after the client has called the Initialize method, can they use the static methods *orb()* and *nameService()*. Subsequent sections discuss how the client might use these methods.

The CBSeriesGlobal interface encapsulates the call to initialize the ORB, which appears in CBSeriesGlobal as:

```
java.util.Properties props = new java.util.Properties();
props.put("org.omg.CORBA.ORBClass", "com.ibm.CORBA.iiop.ORB");
orb = (com.ibm.CORBA.iiop.ORB) ORB.init (naught, props);
```

Where *naught* is a null array of strings, which could be the command line arguments passed in to initialize the ORB. *props* is a properties object filled in with the host name and port of the Bootstrap server to which the client should connect. The Bootstrap server is installed as part of the Component Broker server. If you are unsure of the host and port for this Bootstrap server, contact your Component Broker system administrator. The property names for the Bootstrap server's host and port are com.ibm.CORBA.BootstrapHost and com.ibm.CORBA.BootstrapPort. These properties would be added to the properties object in the following manner:

```
props.put("com.ibm.CORBA.BootstrapHost", host);
props.put("com.ibm.CORBA.BootstrapPort", port);
```

Where *host* and *port* are Java String objects containing the values for the Bootstrap server's host and port. In addition to *host* and *port*, to avoid conflicts with other Java ORBs, a property is added to the property list to make a direct reference to the Component Broker Java ORB. This prevents problems with inadvertently using other ORBs. This problem can occur most often when using Java *applets* in Web browsers. To avoid the problem, modify the code for your Java *applet* to include an explicit cast to the Component Broker Java ORB, com.ibm.CORBA.iiop.ORB.

Within CBSeriesGlobal, the Naming Service is resolved in the following manner:

```
obj = f0rb.resolve_initial_references("NameService");
fNameService = NamingContextHelper.narrow(obj);
```

As you can tell, there are many ways to use these interfaces. The Initialize() methods allow the user to use CBSeriesGlobal to eliminate concern over such intricacies.

IExtendedNaming is used to make things go a little easier. If you need a refresher on it, refer to References in Component Broker for Windows NT and AIX Online Documentation.

Initialization requirements

CBSeriesGlobal is a convenience interface that is required for most client programs. If the client program uses either a copy helper or primary key class with an attribute that is an object reference, then initializing CBSeriesGlobal is a requirement. This is because the implementation of the copy helper and primary key depend on CBSeriesGlobal:orb() when using the ORB object_to_string() operation. It is also a requirement to use CBSeriesGlobal if security or object handles will be used in the Component Broker client. In general it is a good idea to use CBSeriesGlobal in all Component Broker client programs.

Finding a managed object

When a client application starts, it has no knowledge of any objects. In a non-distributed environment, the application would create any objects it needs. These newly created objects would reside on the same system as the application. If the application needs these newly created objects to be available the next time the application runs, the application itself is responsible for making these objects persistent.

However, Component Broker is a distributed environment. In this environment, objects are not created on a client system. Objects are created from a client system on a server system. Object persistence is managed with the help of the Component Broker run time.

The client accesses the component through the components managed object. But first, it has to find the managed object.

There are two ways for a client application to find a managed object:

- Use the Naming Service.

If the object has a *Name*, the client can use the Naming Service to locate the object by its *Name*. In general, the Naming Service only contains a subset of the objects in a distributed system. This subset consists of well-known objects, such as collections of business objects or important objects in the object model.

- Use another object.

If the object does not have a *Name*, the client can find the object by using the Naming Service to find a well-known object, such as a *factory* or collection, and then use this object to implement methods that return the object as a return value or output parameter.—

Both techniques involve the use of the Naming Service. In the Component Broker distributed environment, all work with objects in a client application must begin by using the Naming Service.

Finding a managed object using the Naming Service

Assume that the insurance company in the example placed several Claim objects in the Naming Service. The following code segment shows how to find such a Claim, belonging to a customer named Lou.

```
Claim louClaim = Claim.narrow(
    CBSeriesGlobal.nameService().resolve_with_string(
        "cell/Applications/LifeInsurance/Claim/LouClaim" ) );
```

The Component Broker client and server run times initialize the global object instance nameService to refer to the root of the installation's Naming Service.

For a refresher on determining the naming context, and details on how the name space is specified, refer to "Chapter 4. MOFW C++ client programming model" on page 85.

Recall that in addition to the resolve_with_string() method, Naming Contexts also support the bind_with_string() method which associates a name with an object instance. The following code segment could have been used to name the Lou Claim object.

```
// Assume all of the appropriate imports have been done.
// Do a bunch of stuff with a reference to Lou's Claim

Claim louClaim; // declared somewhere prior to the next set of code
louClaim.
...
// Get a reference to the Life Insurance Application's
// Claim Naming Context using IBM's Extended Naming
// (com.ibm.IExtendedNaming is syntactically optional)

com.ibm.IExtendedNaming.NamingContext nc;

// getting the naming context for Claim Objects

nc = com.ibm.IExtendedNaming.NamingContextHelper.narrow(
    CBSeriesGlobal.nameService().resolve_with_string(
        "cell/Applications/LifeInsurance/Claim" ) );

// Add Lou to the Name Space

nc.bind_with_string("LouClaim",louClaim);
```


Finding a managed object using the primary key

When you have an object, you can use its methods to find related objects. Continuing the previous example, after finding Lou's Claim, you can find other objects that the Claim references. The following example gets you a reference to Lou's Policy:

```
// Find Lou's Policy

Policy louPolicy; // declare a local variable
louPolicy = louClaim.policy();
```

The example is simplified in that LouClaim is in the Naming Service. In the event that you do not have a well-known name for the claim, how do you find a specific claim? Homes are instances of the IHome class. The managed object provider may decide to implement and provide a tailored subclass of IHome, or might use an instance of the base class. The relationship between managed objects and collections is explained in "Using sets of objects" on page 243.

You can find the home for claim objects by using the following code segment:

```
import com.ibm.IManagedClient.* ;

IHome claimHome;
claimHome = IHomeHelper.narrow(
    CBSeriesGlobal.nameService().resolve_with_string(
        "cell/Applications/LifeInsurance/Homes/Claim" ) );
```

Now you need to find Lou's Claim. If you know the Claim number, all you need is the Primary Key Helper class for the Claim. Every managed object class has a set of local helper classes that let you use its keys. An instance of a Key Helper Class is always local to the client's process and language.

The provider of the managed object that you are working with has given you access to source for, or actual, .class files for the helper class. Regardless, it is up to you to ensure that you have access to them and can use them in your applications.

Key Helpers, like all helper classes are created with a static method on the class named `_create()`. This static method gets generated by the bindings that accompany all subclasses of `ILocalOnly`. The same rule is in place for copy helpers and objects of other classes.

Having created an instance of a Primary Key, the key must be set by one or more attributes on the Primary Key object. When all of the key attributes have been set, the Primary Key object is now usable. The Claim Home uses this Primary Key to find the previously created Claim object. Remember, the Primary Key is on the client system, but the Claim object and the Claim Home

or on the server system. If the client passes a Primary Key object as a parameter to the home, and the home is on a remote system, the remote system might get a proxy back to the client's PrimaryKey instance. This would turn the client into a server and unpredictable results could occur. Therefore, the Component Broker programming model uses strings as the method for passing keys to potentially remote objects.

Continuing the example, the following code segment finds Lou's Claim in the home (assuming his number is 1234).

```
import <package>.ClaimPrimaryKeyHelper;
// Package containing ClaimPrimaryKey and Helper;
// Create an instance of the Key helper Class
ClaimPrimaryKey claimPrimaryKey = ClaimPrimaryKeyHelper._create();

// Set the claimNo attribute in the key
claimPrimaryKey.setClaimNo(1234);

// Must get data out of the Stream to go onto the wire to the server
String claimString = claimPrimaryKey._toString();

// Call find by key on the Home to find Lou's Claim
Claim louClaim = ClaimHelper.narrow
    (claimHome.findByPrimaryKeyString(claimString) );
```

The object provider of a public managed object always provides you with a set of helper classes for using the homes that contain his managed objects. There is always exactly one Primary Key helper class. The object provider gives a client developer:

- Interface definitions for the Key Classes. In addition to the Primary Key, there may be Secondary Key Helper Classes. A Secondary Key may also uniquely identify an object, or multiple instances may have the same value.
- Implementation of the Key Classes.
- Documentation for their use.

The Component Broker programming model mandates Key Helper Classes for managed objects. The helper classes make things easier for you. The keys are passed as strings. You only need to use the available setxxxx() methods on the Helper Class to prepare the key information. String manipulation to concatenate pieces of multi-valued keys is not necessary. Using a Key Helper Class enables type errors to be detected at compilation time, and alleviates the need for you to know the field ordering and algorithm for defining a key.

Finding a managed object using another managed object

Some components have interfaces that include methods or attributes that return other components. If this is the case, then when you have an object, you can use its methods to find related objects. Continuing the previous

example, after finding Lou's Claim, you can find other objects that the Claim references. The following example returns a reference to Lou's Policy:

```
// Find Lou's Policy

Policy louPolicy; // declare a local variable
louPolicy = louClaim.policy();
```

Using a managed object

When you have found a reference to a managed object, you can use it by invoking methods on it. For example:

```
person.setName("Lou Smith");
```

This example calls the setName() method on the person object identified by *person*. The Component Broker internal implementation handles the use of remote objects.

Creating a new object

Component Broker managed objects can be created in a number of ways. The following sections describe the default ways to easily create managed objects.

Creating a new object - create from key

Every Component Broker managed object class has an instance of a Factory associated with it. The Factory provides a set of interfaces for creating instances of a managed object. The Factory gets some of its interface from the base class `CosLifeCycle.GenericFactory`. The `createFromPrimaryKeyString()` method is introduced in the `IManagedClient.IHome` interface supplied by Component Broker. This interface specializes the `COSLifeCycle.GenericFactory` interface and plays the role of factory for Component Broker managed objects. Object providers may implement and provide a tailored subclass of this interface, or they might use the implementation of `IHome` provided.

You need to know how to find the right `IHome` for creation. Homes are at well-known locations in the Naming Service. The input required for the factory finder is the name of the implementation class that this factory makes instances of. The following code fragment gets a reference to the Claim Factory for the Life Insurance Application.

```
import com.ibm.IExtendedLifeCycle.*;
import org.omg.CosLifeCycle.;

// (com.ibm.IExtendedLifeCycle is syntactically optional)
com.ibm.IExtendedLifeCycle.FactoryFinder myFinder;

myFinder = (com.ibm.IExtendedLifeCycle.FactoryFinderHelper.narrow(
```

```

CBSeriesGlobal.nameService().resolve_with_string(
    "cell/Applications/LifeInsurance/FactoryFinders"));
IHome claimHome = IHomeHelper.narrow(myFinder.find_factory_from_string(
    "Claim.object interface"));

```

You need to provide the IHome with information necessary to manufacture a new object instance. At a minimum, the Primary Key must be provided. An example of creating a new Claim with a claimNo of 1234 is:

```

// Create an instance of the Primary Key Class
ClaimPrimaryKey claimKey = ClaimPrimaryKeyHelper._create();

// Set the claimNo attribute in the key
claimKey.setClaimNo(1234);

// Call createFromKey on the Factory to create Lou's Claim
Claim theClaim;
theClaim = ClaimHelper.narrow
    ( claimHome.createFromPrimaryKeyString(claimKey._toString() ) );

```

This example is almost identical to the previous example. First, create a Primary Key Object to define the identity of the object to be made. Then, call `createFromPrimaryKeyString` on the home to pass the Key.

The `createFromPrimaryKeyString` method is defined by the IHome class, and all business objects can be created by this method. An object provider might provide you with a subclass that introduces other, easier to use creation methods. Additional create methods are described, with examples in “Chapter 6. MOFW - C++ client programming model – advanced concepts” on page 133.

Creating a new object - create from copy

Setting and getting the attributes of a managed object can be expensive. There are two main reasons for this. First, if the managed object is implemented in another language, each get or set method is actually a cross-language call. Cross-language calls are more expensive than simple, same language calls. The get and set overhead is even worse if the managed object is remote because each call is actually a remote procedure call and involves significant overhead. Consider the following code segment:

```

// Assume that 'theClaim' and 'theClaimHome' are declared in the previous
// code segment, and that 'theClaimHome' is located as in the previous
// code segment. Create a new Claim with a Home generated claimNo
theClaim = theClaimHome.create();

// Now initialize the Claim's attributes
theClaim.setDate("10/14/96");
theClaim.setState(entered);

```

This fragment could involve the following remote method calls:

- The client to Claim Home (specialized IHome) to create the Claim.
- The client to Claim managed object to set its *date* attribute.
- The client to Claim managed object to set its *state* attribute.

An object provider must provide you with a Copy Helper Class for use with the `createFromCopy()` method on an IHome. The following code segment rewrites the previous example using this design pattern.

```
import <packagename>.ClaimCopyHelper;
// Package with ClaimCopy and Helper;
// Create a new "local" Claim in my process and language.
// Use a Copy Helper Class that the Claim MO provider gave me.
ClaimCopy claimCopy = ClaimCopyHelper._create();

// Now initialize the Claim's attributes.
claimCopy.setDate("10/14/96");
claimCopy.setState(entered);

// Pass this local copy to the Home and have him return a new Claim
// MO whose attributes are initialized from the local copy's values.
// Since not all ORBs support Pass-By-Value, we first convert to a string.
Claim theClaim = ClaimHelper.narrow
    (claimHome.createFromCopyString(claimCopy._toString()) );
```

Like a Key Helper Class, a Copy Helper Class instance is always local to your process and implemented in the language you are using. The object provider gives you the helper class interface and implementation.

Copy Helper Classes are especially useful if the client application needs to interact with an object during initialization, and then create a managed object from the attributes. A common scenario for this is entering data for the object from a GUI. The GUI updates the local Copy Helper Object, and then `createFromCopy()` is called when the **Do** button is pushed on the end user interface (EUI).

Using sets of objects

An IHome represents a set of managed objects, all of the same type, whose relationship to one another is defined by the object provider, and maintained by the fact that they were all created in the same home. Sometimes an application needs to define and manage the relationships between managed objects, based on the particular business task at hand. This might even include relationships between managed objects of different types (that is, PolicyHolder and Beneficiary). This can be done using an IManagedReference Collection, as shown in the following code segment:

```
import com.ibm.IManagedCollections.*
com.ibm.IManagedCollections.IReferenceCollection mixedCollection;

// Find a collection in the name space which contains PolicyHolders and
```

```

// Beneficiaries
...
// Create an iterator on the reference collection that was found above.
// When an iterator is created, it is automatically positioned preceding
// the first element.

IManagedCollections.IIterator anIterator = mixedCollection.createIterator();

IManagedClient.IManageable element;
// Loop through the collection. The "next" method advances the iterator
// to the next element (on the first invocation, this advances to the
// first element). If the iterator is now past the end of the collection,
// the "next" method returns NULL.; otherwise, the "next" method returns
// the element the iterator at which it is now positioned.

while ( anIterator.next(element) )
{
    if ( element.is_A(PolicyHolder) ) // Send a bill
    if ( element.is_A(Beneficiary) ) // Send a check
}

```

The combination of the `IManagedCollections.IReferenceCollection` and the `IManagedCollections.IIterator` are used in the previous code segment. An `IManagedCollections.IReferenceCollection` is a generalized collection of object references that is iterable. `IManagedCollections.Iterator` supports advancement of the iterator and retrieval of elements by the `next()` method. `IManagedCollections.IReferenceCollection` supports adding and removing elements by `addElement()` and `removeElement()`. `IManagedCollections.IManagedReferenceCollection` is the most basic kind of collection supported in Component Broker. Combining this with the capabilities of `IHome` provides the basis for writing simple applications and the foundations for the more advanced query and collections capabilities provided by Component Broker. For more information on collections and query, see “Chapter 6. MOFW - C++ client programming model – advanced concepts” on page 133.

Transient sets

Component Broker supports transient collections. Transient collections do not require transactions; they reside in transient containers and thus have no dependency or overhead on database connections. Database and transient collections provide identical programming interfaces. They differ only in their persistence characteristics. Users can find transient collections by passing an interface string to the factory finder as defined in the next section.

Specifying reference collection interfaces

A variety of interfaces for use with factory finders are provided for specifying the type of reference collections that you want to use.

For example, calling `find_factory_from_string` passing the argument `IManagedCollections::IReferenceCollection.object interface/PersistentReferenceCollectionFactory.object home`, as in the previous example, creates a DB2 backed reference collection. Using the generic string `IManagedCollections::IReferenceCollection.object interface` returns whichever collection is configured as the default.

Transient Collections

```
IManagedCollections::IReferenceCollection.object interface/  
TransientReferenceCollectionFactory.object home
```

Transient Collections Keyed

```
IManagedCollections::IKeyedReferenceCollection.object interface/  
TransientKeyedReferenceCollectionFactory.object home
```

Persistent Collections

```
IManagedCollections::IReferenceCollection.object interface/  
PersistentReferenceCollectionFactory.object home
```

Persistent Collections Keyed

```
IManagedCollections::IKeyedReferenceCollection.object interface/  
PersistentKeyedReferenceCollectionFactory.object home
```

Remembering your favorite objects

The Component Broker allows you to remember an interesting or important managed object instance by introducing the concept of an object reference. An object reference is opaque and you cannot set its internal structure. However, a reference always and uniquely refers to a managed object regardless of where it resides in the network.

Object references are available to the Java programmer without regard to what kind of Java programming you are doing. Applet restrictions on the use of the file system, however, render this feature unusable for the applet programmer. The following example assumes that the work is being done from an application.

Continuing the example, assume that you run the following code segment.

```
File output_file = new File("SOMEFILE.DAT");  
FileOutputStream output = new FileOutputStream(output_file);  
  
// Get a "string" version of my reference to Robert  
// robert points to Robert or a proxy to Robert  
String robertStringifiedReference = CBSeriesGlobal.orb.object_to_string(  
    robert)  
  
// Save the string to a file using a "pseudo" file routing  
output.write(robertStringifiedReference);
```

```
output.close();

// I do not need Robert anymore
robert.release();
```

You saved a reference to Robert as a stringified object reference, and can use this string to re-access Robert at a later time. The following code segment presents an example of re-accessing the Robert object.

```
String infile = "SOMEFILE.DAT";
File input_file = new File(infile);
FileInputStream input = new FileInputStream(input_file);

// Get back the string I saved

String robertStringifiedReference;
input.read(robertStringifiedReference); // returns the bytes read, or -1,
// if you'd like to check that.

// Make an object reference for Robert
robert = CBSeriesGlobal.orb.string_to_object(robertStringifiedReference);

// I can now work with Robert.
robert.setName();
```

Releasing and deleting objects

Eventually, you no longer need to use an object that you created or found. Component Broker supports two interpretations on “no longer needs”.

- The `remove()` method deletes the object and its instance data.
- The `release()` method informs the object that the client application no longer plans to reference the object. The object still exists in the server, and other applications may be using it, but the client calling the `release()` method is done with the object

For a detailed description of the differences between releasing and removing objects, see “Releasing and deleting objects” on page 65.

When references explode

Concurrency control is not presently extended to the client proxies. You need to be prepared for exceptions. If an inappropriate object reference is used, an exception is thrown.

Java exception handling

The Java exception handling model has some slight differences from the C++ exception handling model. Within C++ a method can throw two types of exceptions, either exceptions that are explicitly declared on the “raises” clause of the method or any of the CORBA standard exceptions.

In the Java Programming Language, Java exceptions are primarily checked exceptions, meaning the compiler checks your method only throws exceptions that have been declared as throwable in the “raises” clause. These classes extend the `java.lang.Exception` class. Other run time exceptions and errors extend the classes `RuntimeException` and `Error`, which are unchecked exceptions, meaning (much like the CORBA Standard exceptions) they can be thrown from any method.

Exceptions such as `IDataKeyAlreadyExists` and `IDataKeyNotFound` are checked exceptions while other exceptions, such as `IDataObjectFailed` extends `java.lang.RuntimeException` and are unchecked exceptions and can be thrown from any method.

When doing exception handling while using Java, be aware that exceptions other than those listed on the “raises” clause of methods can be thrown by any method at any time. These unchecked exceptions typically occur as a result of some drastic run time error.

Exiting Java applications when using the Abstract Window Toolkit

If you use Abstract Window Toolkit (AWT) components in your application, note that the AWT library spawns non-daemon threads to handle screen updates, event queues, etcetera. These threads are not terminated when the AWT object is no longer in use (no more references to this object) or when the AWT object has been disposed. By convention, the Java virtual machine (JVM) will not terminate if there are non-daemon threads which have not terminated. It may appear as though your application has hung and will not complete. At this time, the only known method of terminating these threads is to use `System.exit()` at the end of the “main” method of the application class.

An example of where this situation exists is when using Secure Sockets Layer (SSL) based authentication between Java clients and application servers. The client principal will be presented with a login panel prompting for a userid and password. The login panel uses AWT, therefore, you would need to use `System.exit()` at the end of the “main” method of your application class to allow it to terminate the AWT threads and exit when complete.

Note that you would not want to use `System.exit()` in an applet as this would attempt to terminate the Java virtual machine of the browser.

Summary: The client programmer's check list

When using a set of components to implement an application, you need to know how to use the following things about components and the managed object framework which encapsulates and interoperates with the component objects. Table 6 is organized around the tasks outlined in "Client programming model: basic tasks" on page 235. Note that *xxx* is used to represent the name of the component. For example, Claim and Policy were used in this chapter as examples of domain-specific components. This also provides a good clue as to whether or not the classes were constructed by the object provider or whether they came as part of the Component Broker system.

Table 6. Summary of interfaces

Task	Class Types	Methods	Purpose
General Use	StandardSyntaxModel	stringToName	Converts strings to CORBA name types.
Find <i>xxx</i> component	NamingContext	resolve	Finds an object in the namespace
	IHome	findByPrimaryKeyString	Finds an object in a home
	<i>xxx</i> PrimaryKey	set <i>xxx</i> methods and toString	Sets the key value
Create or delete an <i>xxx</i> component	FactoryFinder	findFactory	Finds the factory object.
	IHome	createFromKeyString	Creates an object with a key only.
	<i>xxx</i> PrimaryKey	set <i>xxx</i> and toString	Passes information on key to the server.
	<i>xxx</i> CopyHelper	set <i>xxx</i> and toString	Passes copy information to the server.
	any <i>xxx</i> MO (managed object) subclass	remove and release	Deletes or releases object from memory.
Use sets of components	IReferenceCollection	createIterator	Creates an iterator object.
	IIterator	next	Iterates over a collection and gets the objects.

Table 6. Summary of interfaces (continued)

Task	Class Types	Methods	Purpose
Remember interesting and important components	NamingContext	bind	Binds an object into the name space with a name.
		resolve	Finds an object in the name space by its name.
	theORB	object_to_string	Creates a stringified form of an object reference.
		string_to_object	Creates an object reference from a string.

Chapter 10. Java server programming model

This chapter describes how to use Java to implement managed objects that run in a Component Broker server.

Some aspects of server Java programming were introduced in “Chapter 9. MOFW - Java client programming model” on page 231. Managed objects access the CBSeriesGlobal object, Object Services, and other managed objects just like client programs. This could be called the client aspect of server programming, because the managed object acts as a client to the service or to another object. In Component Broker, the client aspect of server Java programming is almost identical to programming in a Java client. However, when generating Java classes from IDL, the IDLC compiler is used for server Java but the idl.toJava compiler is used for the client.

Server Java also introduces some new constructs and procedures that have no counterparts in client programming. Implementing a business object, adapting it to the server execution environment, and managing its persistent data are all tasks that are specific to the server environment. This chapter deals with those server-specific issues.

Developing a component

Component Broker lets you use Java to implement managed objects by performing the following steps. The steps involved are the same as those listed in “Chapter 5. MOFW - C++ server programming model” on page 103; differences are discussed in this chapter:

1. Develop an interface to the business object.
2. Choose a pattern for handling essential state.
3. Implement the business object methods.
4. Implement the managed object framework methods
5. Implement the necessary primary key class.
6. Implement the optional copy helper class.

You will then be ready to package the Java business object for running on a server (for information about packaging the Java business object, see “Assembling and installing Java components” on page 410). This procedure changes only slightly from the non-Java procedures described in “Chapter 13. Assembling and installing components on Component Broker/Workstation” on page 351

on page 351, and mainly involves using the Component Broker tools to generate and compile C++ code to connect the business object into the server run-time environment.

Figure 50 shows how the tools are used to construct the pieces of a server-installable Java business object.

graphic to be completed

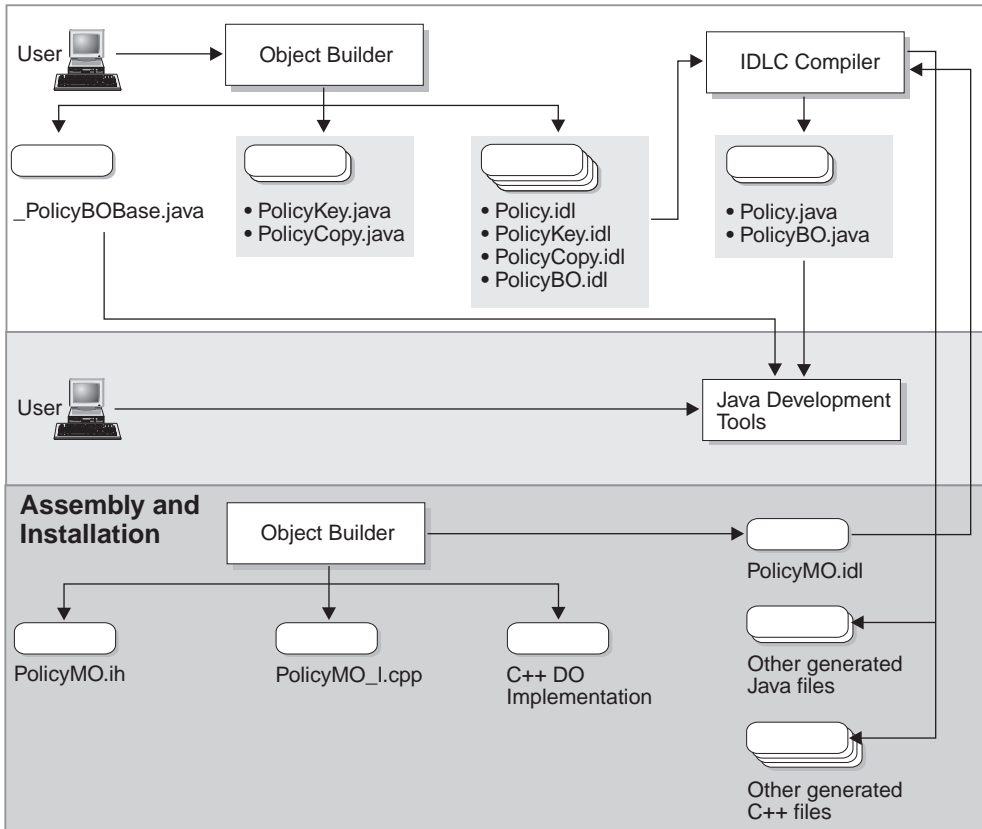


Figure 50. Using tools to construct a server-installable Java business object

The top part of the figure represents the activities described in this chapter. The object builder tool and the IDLC compiler are used to create the Java source code for the business object (`_PolicyBOBase.java` in the figure), plus a collection of other related Java source files. Included here are the Java Key and Copy helper classes required by Java clients of the business object.

Finally, you return to the object builder tool to generate other C++ and Java code that adapts the business object to the server environment and connect it

to your chosen persistent stores, if any. For more information, see “Assembling and installing Java components” on page 410.

Developing an interface to the business object

The procedures and tips in “Chapter 5. MOFW - C++ server programming model” on page 103 apply equally well to the development of Java business objects. The interface for a business object is defined in a way that is independent of the language used to implement it and yields an IDL interface definition. As before, the interface should inherit from the `IManagedClient::IManageable` interface. If you follow the module scoping advice from “Chapter 5. MOFW - C++ server programming model” on page 103, the Policy interface might look similar to the following segment:

```
#include <IManagedClient.idl>
#include "Beneficiary.idl"

module XYZCompanyInsurance
{
    exception InvalidAmount {};

    interface Policy : IManagedClient::IManageable
    {
        void addBeneficiary(in Beneficiary benRef);
        void delBeneficiary(in Beneficiary benRef);

        readonly attribute long policyNo;          // Primary key

        void changeAmount(in float newAmount) raises (InvalidAmount);
        float getAmount();

        attribute string comment;
        readonly attribute float premium;
    };
};
```

This IDL is mapped into Java according to the CORBA IDL-Java mapping specification when you run the IDLC compiler with the `-suj` option, as follows:

```
idlc -suj Policy.idl
```

The `uj` in the example stands for user Java, because it tells IDLC to generate the Java classes and interfaces needed for a client to use the declared IDL interfaces and types. Even though they are called the user (or client) bindings, these Java artifacts are also needed on the server. For the sample `Policy.idl`, the generated files, placed in directory `XYZCompanyInsurance`, are:

- `InvalidAmount.java`
- `InvalidAmountHelper.java`
- `InvalidAmountHolder.java`

- Policy.java
- PolicyHelper.java
- PolicyHolder.java
- PolicyOperations.java
- _PolicyStub.java

See the IOM (Interlanguage Object Model) online documentation for more information about these files. See “Assembling and installing Java components” on page 410 for information about the `-sbj` option that generates additional business-object Java used only on the server.

The IDL exception becomes a Java final class extending a standard CORBA class for user-defined exceptions:

```
package XYZCompanyInsurance;

public final class InvalidAmount extends org.omg.CORBA.UserException
{
    public InvalidAmount() {};
}
```

The `UserException` class, in turn, extends `java.lang.Exception`, and so its subclasses can all be used in a Java `throw` statement, similar to:

```
throw new XYZCompanyInsurance.InvalidAmount();
```

This example also illustrates the following rules for Java package names:

- CORBA-defined classes and interfaces, like `UserException` above, are in the `org.omg.CORBA` package.
- IBM-defined extensions all appear in packages whose names begin `com.ibm`.
- User-defined types and interfaces appear in a Java package that is the name of the enclosing module.

These rules are illustrated again in the Java interface created from the IDL `Policy` interface:

```
package XYZCompanyInsurance;

public interface Policy extends com.ibm.IManagedClient.IManageable
{

    void addBeneficiary(Beneficiary benRef);

    void delBeneficiary(Beneficiary benRef);

    /**
     * Getter method for attribute "policyNo"
     */
    int policyNo();
}
```



```

void changeAmount(float newAmount) throws XYZCompany.InvalidAmount;

float getAmount();

/**
 * Getter method for attribute "comment"
 */
java.lang.String comment();

/**
 * Setter method for attribute "comment"
 */
void comment(java.lang.String comment);

/**
 * Getter method for attribute "premium"
 */
float premium();
}

```

The mapping of names and types is according to the CORBA specification of which the following table is a brief summary:

- Names that conflict with Java keywords have a single underscore prepended to them. The affected names are:

abstract	double	int	super
boolean	else	interface	switch
break	extends	long	synchronized
byte	final	native	this
case	finally	new	throw
catch	float	package	throws
char	for	private	transient
class	goto	protected	try
const	if	public	void
continue	implements	return	volatile
default	import	short	while
do	instanceof	static	

- IDL names ending with the suffixes `Helper`, `Holder`, or `Package`, when used to define an IDL type or interface, also have an underscore prepended in Java. Therefore, the following IDL segment:

```

module XYZCompanyInsurance {
    interface PolicyHelper {

```

Yields the following Java:

```

package XYZCompanyInsurance;
public interface
    _PolicyHelper extends com.ibm.IManagedClient.IManageable

```

3. IDL operation names that are the same as methods of the class `java.lang.Object` also have underscores prepended. These names are:

<code>clone</code>	<code>getClass</code>	<code>notifyAll</code>
<code>equals</code>	<code>hashCode</code>	<code>toString</code>
<code>finalize</code>	<code>notify</code>	<code>wait</code>

4. Types are mapped according to the following table:

IDL Type	Java Type	IDL Type	Java Type
boolean	boolean	long unsigned long	int int
char	char	float	float
wchar	char	double	double
octet	byte	array	array
string	<code>java.lang.String</code>	sequence	array
wstring	<code>java.lang.String</code>	void	void
short	short	any	<code>org.omg.CORBA.Any</code>
unsigned short	short	Object	<code>org.omg.CORBA.Object</code>

Note: `org.omg.CORBA.Any` is an abstract class and cannot be instantiated. Each Java-CORBA ORB provides an implementation of Any. To create an instance in a Component Broker environment, use:

```
org.omg.CORBA.Any any =
    com.ibm.CBCUtil.CBSeriesGlobal.orb().create_any();
```

Module scoping

Although the examples used in this book do not show it, wherever possible, IDL interfaces and types should be enclosed inside a module scope. IDL declared outside of a module scope takes up name space in the global IDL name space and risks having name collisions with names declared by other IDL developers.

For example, if two groups in an organization write IDL interfaces without placing them in different modules, and happen to choose names for their interfaces that overlap with each other, the result is name collision when the IDL is loaded into a shared Interface Repository. However, if each group chooses a unique module name, perhaps based on a combination of company name and group name, each group's IDL is able to coexist with all other IDL interfaces.

This example is even more useful when Java and the Internet are considered. In the IDL-to-Java mapping specification adopted by the Object Management

Group (OMG), IDL module names are mapped directly into Java package names. IDL declared outside a module is mapped into classes and interfaces that are not contained in any package. Because one of the purposes of packages in Java is to partition the Java name space for the entire Internet, it is particularly important that IDL that might be mapped into Java be contained within a module.

However, there needs to be a practical balance between module scoping and the number of interfaces defined in that module. All interfaces that are part of the same module must also exist in the same actual file. Introducing a large number of interfaces in the same file can lead to performance problems with some of the tools, such as compilers. Therefore, it is recommended that you define no more than ten interfaces in the same module due to this practical limitation.

For example:

```
module XYZCompany_Finance
{
    interface Receivable
    {
        attribute long customerID;
        attribute long amountCents;
    };
};
```

The Receivable interface has the fully-scoped name XYZCompany_Finance::Receivable.

Design tips for components

The term attribute describes a piece of data which is part of an object's state. Good object-oriented design (OOD) principles discourage putting attributes on the public interface of an object. Public attributes allow the user of an object to directly manipulate the state of the object, without any control by the object itself. OOD uses the term *encapsulation* to mean putting an object's data members in the protected or private sections.

Because CORBA and IDL do not allow specifying data as part of an object's interface, public data is not a problem. In CORBA, attributes on an object's interface (as specified using IDL) define methods for accessing the state of an object. Because IDL does not allow the specification of exceptions to be thrown as the result of an attribute, IDL attributes do not provide the same level of protection as encapsulation in Java. To see why, continue with the Policy example. For illustration purposes, assume that the insurance company does not want to issue a policy unless the amount is over \$10,000.00. As shown previously, if *amount* had been defined as an attribute, it would result in the following two methods being declared:

```
virtual float amount ();
virtual void amount (float amount);
```

Look at the implementation of the method for setting the attributes value:

```
virtual void amount (float amount)
{
    if ( amount > 10000.00 )
        fAmount = amount; // save value in private data member
    else
        // What shall we do here...?
};
```

IDL does not allow us to raise an exception on an attribute. The method signature, as emitted by the IDL compiler, has no return value or output parameters. All you can do is not save the new value, return to the caller, and hope that the caller notices that the value was not changed. A better way to truly encapsulate the amount is to not make it an attribute in IDL, but rather a pair of get and set methods. This lets the business object do constraint checking and communicate error conditions to the client of the business object.

On the other hand, this needs to be balanced against performance. The Query Service of Component Broker performs best when a query is specified in terms of attributes, and furthermore, that the attributes of the component's business object map one-to-one to the attributes of the component's data object. Doing so lets the Query Service of Component Broker push down the query to the underlying resource manager. Resource managers such as relational database managers perform queries much faster than Component Broker can without any other help. Module and interface names must be different. Although identical names are valid IDL syntax, it gets mapped to nested classes by the emitter, and the compiler does not allow the name of a container class to be the same as a contained class.

Selecting a pattern for handling essential state

This step introduces another interface that looks like this:

```
import Policy;
import com.ibm.IManagedServer.*;

interface PolicyB0 : Policy,
    IManagedServer::IManagedObjectWithDataObject,
    IManagedServer::IWrappable /* only for Java */
{
};
```

This step proceeds as described in "Chapter 5. MOFW - C++ server programming model" on page 103, with one significant difference: When the Java business object is installed in the server, there is a companion C++ object

created to help tie it in to the server's C++ infrastructure. In order to correctly interconnect the Java object and its C++ companion, the Java object needs to implement the `IManagedServer::IWrappable` interface.

Note: This interface is contained in a file named `PolicyBO.idl`.

This interface is small and straightforward. By inheriting from `IManagedServer::IManagedObjectWithDataObject` a decision has been made about the pattern to be used for handling the essential state of this component. There are two interfaces in `IManagedServer` that correspond to the two patterns that the Component Broker MOFW support for handling the essential state of a component.

IManagedObjectWithDataObject

This interface implies that the component has its state managed by a data object. Because the component has a data object, and no implied caching as in the next interface, this is called the *delegating pattern*.

There are implications here for the implementor of the business logic methods. These are described in "Implementing business object methods and attributes" on page 109.

IManagedObjectWithCachedDataObject

This interface implies the presence of a data object and support from Component Broker for fetching this data into a cached set of state data maintained by the component's business object at the appropriate times throughout the lifetime of a component. This is called the *caching pattern*, because the business object is caching the state being maintained by the data object.

Patterns for handling state (caching)

If `IManagedObjectWithCachedDataObject` was chosen as the data object pattern, then all essential state, and all non-derived nonessential state, is stored or cached in the component's business object.

```
class _PolicyBOBase implements PolicyBO
    extends com.ibm.IManagedClient._IManageableBase
{
    ...
    // methods being implemented go here
    ...
    protected:
        int fPolicyNo;
        float fAmount;
        PolicyHolder fInsured;
        java.lang.String fComment;
};
```

Attributes: The implementation bindings *getter* and *setter* methods for each of the attributes specified in an interface. The *getter* and the *setter* for the *comment* attribute are implemented as follows:

```
void comment (java.lang.String comment)
{
    fComment = comment;
}

java.lang.String comment ()
{
    return fComment;
}
```

The previous example is a simple implementation of getters and setters. More complex logic and exception handling may be required. In fact, getters and setters might actually have logic in them to write to resource managers or get data from resource managers. However, as described previously under “Design tips for components” on page 257, if the logic required in the implementation of getters and setters includes any sort of bounds checking or error checking, then the attribute should be changed to a pair of methods in the IDL file.

Other Methods: Other business logic methods which access state data do so similarly to the *comment* attribute mentioned previously; they use the data that is cached in the business object.

Patterns for handling state (delegating)

If *IManagedObjectWithDataObject* was chosen as the data object pattern, then all essential state is accessed by the data object: when the business logic needs to get or set its essential state, it does so by using the attributes on the data object. This is called the delegating pattern because the maintenance of the component’s state is delegated from its business object to its data object. All non-derived nonessential state is still cached in the business object. Using this pattern, the implementation binding header (.ih) file for *PolicyBO* (*PolicyBO.ih*) would look something like this:

```
class _PolicyBOBase implements PolicyBO
extends com.ibm.IManagedClient._IManageableBase
{
    ...

    // methods being implemented go here

    ...
protected:
    PolicyDO fDataObject;
};
```

Attributes: The implementation bindings generate getter and setter methods for each of the attributes specified in an interface. In the `PolicyBO_I.cpp` file, the getter and setter for the *comment* attribute would be implemented as follows:

```
void comment (java.lang.String comment)
{
    fDataObject.comment(comment);
}

java.lang.String comment ()
{
    return fDataObject.comment();
}
```

Other Methods: Other business logic methods which access essential state data do so similarly to the comment attribute mentioned previously; they use the general `fDataObject.datamember` pattern.

An advantage of this programming model is that the business logic can be written independent of the exact implementation of the data object. At this point, you only know which data is to be persistently stored, not where it is stored or how it is accessed. This encapsulation of information inside of the data object makes business logic more stable.

The data object

For the patterns described previously, a data object is necessary. A data object manages the essential state of a component. For the working example, assume that the data object for the Policy is as follows:

```
interface PolicyD0
{
    attribute long policyNo;
    attribute float amount;
    attribute string comment;
    attribute PolicyHolder insured;
}
```

The interface for the data object contains one attribute for each piece of the component's essential state, as shown in Figure 35 on page 158, and that there are no attributes for the component's nonessential state. Also notice that this interface is not a one-to-one mapping of the attributes on the business object. The *premium* attribute from the business object is not present here because it is not part of the essential state. Also, although *amount* and *insured* were not attributes on the business object, they are attributes on the data object because they are part of the component's essential state.

This interface is enough to let you develop the component's business logic, without having to actually implement the underlying data access methods.

However, before the business logic can be tested, the data object interface must in fact be implemented. Data object customization describes how to customize the data object to allow the component to be installed into a particular application adaptor.

Implement business object methods

In this step an implementation of the business object interface defined in “Selecting a pattern for handling essential state” on page 258 is created. In “Chapter 5. MOFW - C++ server programming model” on page 103 this was done using C++, but for this implementation it must be done in Java.

Because the business object client interface was called XYZCompanyInsurance::Policy (in IDL), the derived interface would be named XYZCompanyInsuranceBO::PolicyBO. Component Broker server Java requires that the implementation of this interface be provided by a class named XYZCompanyInsuranceBO._PolicyBOBase. The Base suffix is a Component Broker convention, and the underscore ensures that this class does not collide with any IDL types whose names end with Base.

The implementation class is also required to implement the PolicyBO interface and to extend the base class IManagedClient._IManageableBase. If you use the object builder to create your Java code it is defined with the correct inheritance.

The implementation class also needs to contain implementations for the business methods of the object, as well as some of the framework methods. Other framework methods are implemented in _IManageableBase, and do not need to be overridden. The framework methods that do require implementation are listed, but if you use the object builder to create your implementation class it automatically generates the correct subset of framework methods.

The implementation of the PolicyBO interface could look like the following segment:

```
package XYZCompanyInsuranceBO;

public class _PolicyBOBase
    implements PolicyBO                // See note 1
    extends com.ibm.IManagedClient._IManageableBase
{
    public void addBeneficiary(Beneficiary ben)
    {
        try
        {
            beneficiaries.addElement(ben);        // See note 2
        }
    }
}
```



```

        catch(com.ibm.ICollectionsBase.IInvalidElement e)
        {
        }
    }

    public void delBeneficiary(Beneficiary ben)
    {
        try {
            beneficiaries.removeElement(ben);
        }
        catch(com.ibm.ICollectionsBase.IElementNotFound e)
        {
        }
    }

    public int policyNo() {return iPolicyNo;}    // See note 3

    public void changeAmount(float newAmount)
        throws InvalidAmount
    {
        if(newAmount < 0.0)
            throw new InvalidAmount();
        iAmount = newAmount;                    // See note 4
    }

    public float getAmount()
    {
        return iAmount;
    }

    public String comment()
    {
        return iComment;
    }

    public void comment(String comment)
    {
        iComment = comment;
    }

    protected
        com.ibm.IManagedCollections.IReferenceCollection beneficiaries;
    protected int iPolicyNo;                    // See note 5
    protected float iAmount;
    protected string iComment;
    protected float iPremium;

    /* ... framework methods go here, we will see them later */

};

```

This segment shows one possible implementation of the business methods of the Policy interface. For clarity, none of the framework methods are included in the example.

Notes about the implementation of the PolicyBO interface:

1. This class implements the PolicyBO interface and extends the standard implementation class mentioned earlier. `implements XYZCompanyInsurance.Policy` could have been written out in full, but because this class is part of the same Java package as the interface, the qualifying package name can be omitted.
2. The PolicyBO interface was specified to use the `IManagedObjectWithDataObject` pattern, so the business object maintains copies of all state data including the `IReferenceCollection` used to implement the beneficiaries relationship. The use of collections to represent relationships is described in “Chapter 7. MOFW - C++ server programming model – advanced concepts” on page 157 and is not repeated here.

The use of a try block in the `addBeneficiary` routine should be explained. In the IDL for the Policy interface, there was no “raises” clause specified for this routine, so that when the IDLC compiler generated the Java interface for Policy it did not include a “throws” clause on the corresponding Java method. According to the rules of the Java language, implementations of `addBeneficiary` are then not permitted to throw any user exceptions.

However, the `addElement()` method of the `IReferenceCollection` interface can throw the `InvalidElement` exception. Either those exceptions must be caught and dealt with inside the `addBeneficiary()` method as is done in this example, or the Policy IDL must be modified to say that the `addBeneficiary()` method can raise the `InvalidElement` exception.

3. Because the Policy IDL declared `policyNo` as a read-only attribute, only a get-accessor method is provided. Non read-only attributes also get a set-accessor. Because IDL attributes cannot take “raises” specifications, accessors can throw only subclasses of `java.lang.RuntimeException`. CORBA system exceptions do map to subclasses of `RuntimeException`, but user-defined IDL exceptions do not.
4. The correct pattern to use when setting or getting attributes depends on the mapped Java data type. If the attribute is not an IDL interface type but maps to a modifiable Java object, copies should be created in both the getter and setter methods. This rule is also appropriate for methods like `changeAmount()` that do not correspond to IDL attributes but do set or return business object instance data.

In the case of the `changeAmount()` method, the data type involved is “float” which is not an object type, and so simple assignment is appropriate. The same is true for IDL types like `char`, `short`, and `long`. It is also true for IDL string and enum types because, although they map to Java objects, the values of the objects cannot be modified.

Contrast this with an IDL struct, which maps into a Java class with public fields. If a setter method saved a reference to the passed-in object instead

of copying it, the caller could later modify the shared object. To avoid this the business object can save a copy of the passed-in value object and make another copy when returning a value. This is usually appropriate for objects corresponding to the IDL struct, union, array, sequence, exception, and any complex data types. Also, the copy should be what is usually called a deep copy so that, for example, copying an array of struct objects yields an array of copies of all the original objects.

Any convenient deep copying technique is acceptable. Here is a standard technique that relies on the support for marshalling these objects, using the write and read methods of the corresponding Helper class:

```
org.omg.CORBA.portable.OutputStream os =
    com.ibm.CBCUtil.CBSeriesGlobal.orb().create_output_stream();
TypeHelper.write(os, object);
Typecopy = TypeHelper.read(os);
```

5. These are the cached attributes as described in “Chapter 7. MOFW - C++ server programming model – advanced concepts” on page 157. They can be specified as private or as protected if it is likely that a derived business object needs direct access to them. They should not be public or package access (that is, not keyword access).

Managing memory

If you have already read “Chapter 5. MOFW - C++ server programming model” on page 103 and “Chapter 7. MOFW - C++ server programming model – advanced concepts” on page 157, you have seen several mentions of methods called `_duplicate()` and `release()` and the distinction between `_var` and `_ptr` references and may be wondering where are the corresponding complications in Java.

The answer is that, except for the hint in the previous section about copying values, there are none. Those complications are related to management of memory in C++, which is a manual or at best a semi-automated process.

In Java, those worries are handled automatically by the Java garbage collector. You do not need to do anything special.

Tip: Sometimes the garbage collector can be fooled into believing that an object is not garbage, even though you do not need it anymore, because your program still has a variable somewhere that refers to it. This could cause a performance degradation if a large amount of storage is involved. You can minimize the effect by explicitly assigning null to references to large objects when you do not need them.

```
int[] iarray = new int[1000000];
/* work with iarray */
iarray = null;      /* nullify the pointer */
```

Using 'this' references in business objects

Care must be taken when programming all methods in business objects that use references to themselves when communicating with other objects. Methods must use the programming model as described in this section when using these self references. The technique of using "this" is no longer supported in the programming model in these circumstances.

A local proxy class is created for each interface defining the managed object implementation, the managed object interface, the business object interface, and every other interface that they may support. Only instances of the local proxy of the managed object implementation are instantiated and these proxies must be used for self references. The `Helper._self(this)` method can be used to access this proxy in the business object.

The home provides a copy reference to a local proxy for the `create()` and `findBy()` methods that return object references. The following example shows the set of rules to follow when an object passes itself as an argument or returns itself as a return argument:

```
public class I
{
    I foo()
    {
        // Use IHelper._self(this)
        sequence[i] = IHelper._self(this);
        struct.i = IHelper._self(this);
        return IHelper._self(this);
    }
}
```

Reference scoping

Local references to Java business objects are not valid outside the remote method call which created the reference. For example, it is not legal to set the value of a Java business object into a static variable and then retrieve that object from the static on a subsequent method call and use it. The passivation scheme for Java BO's and the mixin technology of the server require that interposing being done by the C++ MO before running outside the context of the initiating remote method is not allowed.

Java exceptions

All exception classes in Java are subclasses of `java.lang.Exception`. If such a class is thrown within a method, it must either be caught in that method or declared in the `throws` clause of that method. Only subclasses of `java.lang.RuntimeException` are free of this constraint. The intent was that subclasses of `java.lang.RuntimeException` are thrown under conditions not anticipated by the business logic (such as null pointer and number format

exceptions) while all other exceptions classes are thrown as a result of conditions detected by the business logic (such as InvalidAmount).

Figure 51 demonstrates how CORBA exceptions fit into this scheme. Exceptions thrown from CORBA business logic should directly or indirectly inherit from `org.omg.CORBA.UserException`. Exceptions that subclass `org.omg.CORBA.SystemException` do not have to be caught or declared in a throws clause since these inherit from `java.lang.RuntimeException`. Subclasses of `SystemException` representing specific error conditions are listed in the *WebSphere Application Server Enterprise Edition Component Broker Programming Reference*.

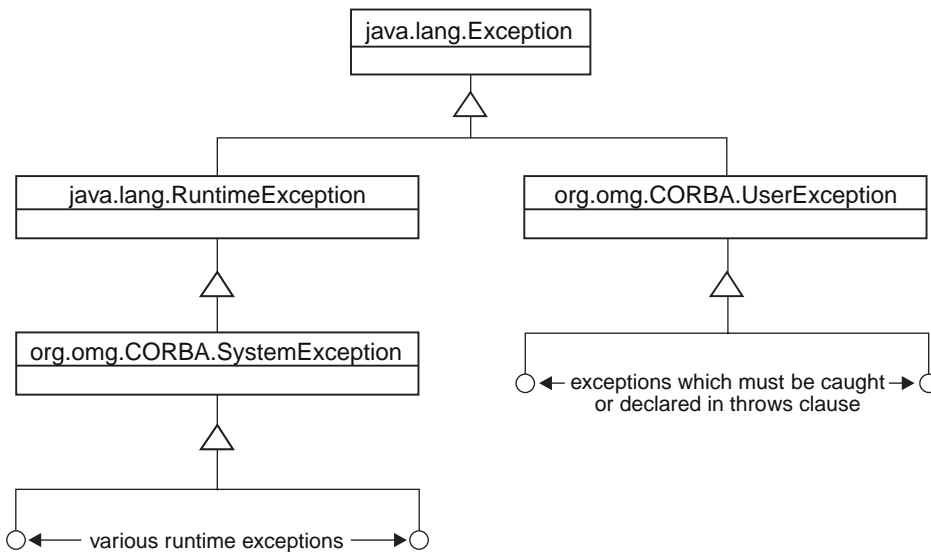


Figure 51. CORBA exception classes in Java

`SystemException` classes are usually thrown from Component Broker framework code rather than business logic. However, one can add or remove `SystemException` throws in business logic and get a clean compile without having to modify both IDL and Java method signatures. Aside from special cases like debugging or throwing a `SystemException` for its intended use, business logic should generally throw exceptions derived from `UserException`. All exception classes declared in IDL inherit from `UserException`.

`UserException` classes are emitted to Java with two constructors. One of these is the default constructor. The other one takes as parameters all the exception's data members ordered by their appearance in the exception's IDL. Unfortunately, in order to maintain language neutrality, `UserException` classes which come across a CORBA connection do not have their `getMessage()` method overridden as some Java client applications might expect.

Implement the IManageable required methods

The framework methods that need to be implemented for a Java business object are the same as those described in “Chapter 5. MOFW - C++ server programming model” on page 103, but as you would expect the Java code looks slightly different. This section includes the Java versions of those methods for the Policy business object.

IManageable::getPrimaryKeyString

The IDL return type of this method is `ByteString`, which is an IDL typedef. Because there is no typedef capability in Java, types are mapped according to the underlying resolved type. `ByteString`'s underlying type is *sequence of octet*, which maps to `byte[]` in Java.

```
public byte[] getPrimaryKeyString()
{
    PolicyKey key = PolicyKeyHelper._create();
    key.policyNo(iPolicyNo);
    return key._toString();
}
```

The `PolicyKey` is a locally-implemented interface that is discussed in more detail later in this chapter. The `PolicyKey` interface is defined in IDL and can be implemented in either C++ or Java. You can use the `_create()` method of the Helper class to create an instance because that method is available whatever the implementation language.

The underscore on the `_toString()` method illustrates a rule from list item 3 on page 256, where the IDL method name `toString()` would conflict with a method introduced by `java.lang.Object`.

IManagedClient::IManageable::getHandleString

The Java version of this method is quite straightforward. In the following example, the stringified object reference (SOR) standard version of the `Handle` is used:

```
public byte[] getHandleString()
{
    com.ibm.IHandlesImpl.ISORHandle iSORHandle =
        com.ibm.IHandlesImpl.ISORHandleHelper._create(this);
    return iSORHandle._toString();
}
```

Note: This is the default implementation provided by the Framework. If this is sufficient (and it should be) there is no need to override this in the business object unless this business object will participate in relationships that are to be managed using other handle patterns.

CosStream::Streamable::externalize_to_stream

Component Broker components are always streamable. Streaming can be the basic mechanism used for copying and moving objects. This mechanism is used by the OS/390 Component Broker frameworks to copy objects for internal use. The other Component Broker platforms do not currently use this mechanism so this method is required to be implemented only for OS/390 Component Broker support. This mechanism is used by the OS/390 Component Broker frameworks to copy objects for internal use. The other Component Broker platforms do not currently use this mechanism so this method is required to be implemented only for OS/390 Component Broker support.

Here is one way to externalize the Policy implementation. The following segment stores a stringified reference to the beneficiaries object:

```
public void externalize_to_stream(org.omg.CosStream.StreamIO s)
{
    s.write_long(iPolicyNo);
    s.write_float(iAmount);
    s.write_float(iPremium);
    s.write_string(iComment);
    String benString = com.ibm.CBCUtil
                      .CBSeriesGlobal.orb()
                      .object_to_string(beneficiaries);
    s.write_string(benString);
}
```

CosStream::Streamable::internalize_from_stream

The `internalize_from_stream` method on `CosStream::Streamable` needs to be written so that it can read the values that the `externalize_to_stream` method placed into the stream. The values must be read in the same order that they were written. This mechanism is used by the OS/390 Component Broker frameworks to copy objects for internal use. The other Component Broker platforms do not currently use this mechanism so this method is required to be implemented only for OS/390 Component Broker support. This mechanism is used by the OS/390 Component Broker frameworks to copy objects for internal use. The other Component Broker platforms do not currently use this mechanism so this method is required to be implemented only for OS/390 Component Broker support. This method corresponds to the `externalize_to_stream()` method:

```
public void internalize_from_stream(org.omg.CosStream.StreamIO s,
                                   org.omg.CosLifeCycle.FactoryFinder ff)
    throws org.omg.CosLifeCycle.NoFactory,
           org.omg.CosStream.ObjectCreationError,
           org.omg.CosStream.StreamDataFormatError
{
    if(s.read_long() == iPolicyNo)
```

```

    {
        iAmount = s.read_float();
        iPremium = s.read_float();
        iComment = s.read_string();
        String benString = s.read_string();
        beneficiaries = BeneficiariesHelper.narrow(com.ibm
            .CBCUtil.CBSeriesGlobal.orb()
            .string_to_object(benString));
    }
    else
    {
        throw new org.omg.CosStream.StreamDataFormatError();
    }
}

```

As noted in “Chapter 5. MOFW - C++ server programming model” on page 103, this method should not modify the key attribute *iPolicyNo* of the business object. Instead, it throws an exception if the key does not match. Note that an arbitrary user exception cannot be thrown here; it must be one that was declared in the “raises” clause in the IDL that introduced this method, which was the OMG standard interface `CosStream::Streamable`.

Note also the use of `BeneficiariesHelper.narrow()` instead of a Java cast to convert from `org.omg.CORBA.Object` (the return type of the `string_to_object()` method) to the `Beneficiaries` interface. The Java cast is not reliable if the object is located in a remote server, or is implemented in a language other than Java. In those situations the cast may fail when the `narrow` function would succeed. It is good practice to use `narrow` when dealing with any IDL-defined interfaces.

The streamable method implementations shown here are examples. There is no required implementation at this time because there is no dependency on these methods from within Component Broker.

Summary of `IManagedObject` methods

All of the methods defined in the `IManagedObject` subclass chosen must be implemented. They have been described in the previous sections. These methods should be called only by the Component Broker server. See “Framework flows” on page 442 for further information.

Implementing `IManagedObject` required methods

In addition to the business logic methods, there are MOFW-required methods that must be implemented regardless of which kind of `IManagedObject` class is chosen as the base class.

initForCreation() method

If the pattern chosen for handling essential state requires a data object, this method needs to save a reference to the supplied data object. Also, if the pattern is `WithCacheDataObject`, as it is in the `PolicyBO` example, the cached values in the business object need to be initialized. Because that function is needed elsewhere, you could split it into a separate private function as illustrated in the following segment:

```
private PolicyDO theDO;

private void initializeState()
{
    iPolicyNo = theDO.policyNo();
    iAmount = theDO.amount();
    iPremium = theDO.premium();
    iComment = theDO.comment();
    beneficiaries = theDO.beneficiaries();
}

public void initForCreation(com.ibm.IManagedServer.IDataObject do)
    throws com.ibm.IManagedServer.ICreationFailed
{
    theDO = PolicyDOHelper.narrow(do);

    initializeState();
}
```

uninitForDestruction() method

Nothing needs to be done for this method:

```
public void uninitForDestruction()
    throws com.ibm.IManagedServer.IDestructionFailed
{
}
```

initForReactivation() method

In the example this method saves the data object reference:

```
public void initForReactivation(com.ibm.IManagedServer.IDataObject do)
    throws com.ibm.IManagedServer.IReactivationFailed
{
    theDO = PolicyDOHelper.narrow(do);
}
```

uninitForPassivation() method

No resources need to be released for this method:

```

public void uninitForPassivation()
    throws com.ibm.IManagedServer.IPassivationFailed
{
}

```

syncFromDataObject() method

The business object uses the caching pattern so this method and the companion `syncToDataObject()` method must be implemented. This is where the `initializeState()` method introduced earlier is reused:

```

public void syncFromDataObject()
    throws com.ibm.IManagedServer.ISynchronizationFailed
{
    initializeState();
}

```

syncToDataObject() method

The business object uses the caching pattern so this method and the companion `syncFromDataObject()` method must be implemented. This is where the `initializeState()` method introduced earlier is reused:

```

public void syncToDataObject()
    throws com.ibm.IManagedServer.ISynchronizationFailed
{
    theDO.policyNo(iPolicyNo);
    theDO.amount(iAmount);
    theDO.premium(iPremium);
    theDO.comment(iComment);
    theDO.beneficiaries(beneficiaries);
}

```

Tip: Preserving the robustness of the Component Broker server environment imposes requirements on executing threads that are not met by threads spawned by the new `java.lang.Thread()` method in Java. Therefore, you should avoid creating threads in a server.

Note: The same restrictions and recommendations pertaining to `initForCreation()`, `uninitForDestruction()`, `initForReactivation()`, `uninitForPassivation()`, `syncToDataObject()`, and `syncFromDataObject()` apply to the Java BO versions of these methods. For additional information, see “Implementing `IManagedObject` required methods” on page 118.

Summary of `IManagedObject` methods

All of the methods defined in the `IManagedObject` subclass chosen must be implemented. They have been described in the previous sections. These methods should be called only by the Component Broker server. See “Framework flows” on page 442 for further information.

Implementing the primary key class

A Primary Key class is used by clients of your business object, and also by the Component Broker server infrastructure. Currently the server infrastructure needs its Primary Key to be implemented in C++, while pure Java clients need a Primary Key class implemented in Java.

Java business objects running in the server can use either the pure Java implementation, or can access the C++ class from Java through the interlanguage capabilities of IOM. The second alternative has slightly poorer performance because of the expense of interlanguage calls.

As a result, whenever you implement any Component Broker business object you must supply a C++ Primary Key class. If this class is built with IOM interlanguage bindings, it is sufficient for server Java use. However, the pure Java client requires a second implementation of the Primary Key in pure Java; if you build one of these, you can also use it with Java in the server. In that case, the server environment contains both the C++ implementation, for use by the server infrastructure, and the Java implementation for use by Java business objects.

The Component Broker object builder generates both C++ and Java implementations of Primary Key classes when you use object builder to create the business object.

A C++ Primary Key class example was discussed in “Chapter 5. MOFW - C++ server programming model” on page 103. The equivalent Java class should not be too surprising now that you have seen what the Java business object looked like. The following segment is an example of an IDL for the PolicyKey interface:

```
#include
module XYZCompanyInsuranceKeys
{

    interface PolicyKey : IManagedLocal::IPrimaryKey
    {
        attribute long policyNo;

    };
    #pragma meta PolicyKey localonly
};
```

The following segment illustrates the Java classes implementing the XYZCompany::PolicyKey interface:

```
package XYZCompanyInsuranceKeys;

public class _PolicyKeyImpl // See note 1
    extends com.ibm.IManagedLocal._IPrimaryKeyImpl // See note 2
```

```

        implements PolicyKey                                // See note 3
    {
        int fPolicyNo = 0;                                  // See note 4

        public int policyNo ()
        {
            return fPolicyNo;
        }

        public void policyNo (int policyNo)
        {
            fPolicyNo = policyNo;
        }

        // Methods from Streamable
        public void externalize_to_stream (                  // See note 5
            org.omg.CosStream.StreamIO targetStreamIO)
        {
            {
                targetStreamIO.write_long(fPolicyNo);
            }
        }

        public void internalize_from_stream (
            org.omg.CosStream.StreamIO sourceStreamIO,
            org.omg.CosLifeCycle.FactoryFinder there)
            throws org.omg.CosStream.StreamDataFormatError
        {
            fPolicyNo = sourceStreamIO.read_long();
        }

        public java.lang.String getName()
        {
            return "XYZCompanyInsuranceKeys::PolicyKey";
        }
    } // _PolicyKeyImpl

    package XYZCompanyInsuranceKeys;

    public class PublicKeyHelper                            // See note 6
    {
        public static PublicKey _create()
        {
            return new _PublicKeyImpl();
        }
    }

```

Notes about the example:

1. Java Primary Key classes are intended for use both in the server and in a pure Java client. Because the Java client environment does not support

IOM, this class is built as a pure Java, single-language class. One consequence is that the class name is not required to end with Base like a business object implementation. Another consequence is that references to this object can be used only from Java and cannot be passed successfully as a parameter of an interlanguage call. An attempt to do so generates a run time CORBA::MARSHAL exception.

2. This Java base class provides implementations of several methods. As described in “Chapter 5. MOFW - C++ server programming model” on page 103, the `IKey::isEqualToKey()` and `IKey::isEqualToKeyString()` methods optionally can be overridden, but the supplied implementations should be sufficient for most needs.
3. The `PolicyKey` interface is defined in the file `PolicyKey.java` generated by the IDLC `-suj PolicyKey.idl` command. No other IDLC-generated source files are needed.
4. Following normal Java practice, instance data can be initialized inline, as shown here, or in a separate constructor. In this particular case, the initializer could have been omitted entirely and the Java default zero initialization of instance data accepted.
5. It is essential that the `internalize` and `externalize` methods on the Java and C++ versions of the `Primary Key` class must match, so that instance data can be streamed out of a Java `Primary Key` and then streamed into a C++ one. The Component Broker server infrastructure does this because it operates only with the C++ version.
6. Replace the `Helper` class that the `idlcc` command generated with a simple one like this. Providing a `Helper` class with a static `_create` method allows users of your `Primary Key` class to create instances in a way consistent with other IDL-defined objects.

Implementing the optional copy helper class

You can implement a pure Java copy helper class for use with the `IManagedClient::IHome::createFromCopyString()` method. Just like the pure Java `Primary Key` class, the copy helper can be used in a pure Java client or on the server.

If you do create a Java copy helper, you must also produce a C++ one and install it in the server. The Component Broker infrastructure manipulates copy helpers and requires that they be implemented in C++.

```
module XYZCompanyInsuranceCopy
{
    interface PolicyCopy : IManagedLocal::INonManageable
    {
        attribute long policyNo;
        void changeAmount(in float newAmount)
            raises (InvalidAmount);
    }
}
```

```

        float getAmount();
        attribute string comment;
    };
}

```

The previous example of an IDL appears to be a subset of the Policy business object interface. This is not an accident. The purpose of the copy helper is to hold the data that is used to initialize a new Policy object, so the copy helper needs to have similar attributes. It is also a good idea to duplicate any simple validity checks that the Policy business object's logic would perform. The following Java implementation corresponds to this IDL:

```

package XYZCompanyInsuranceCopy;

public class _PolicyCopyImpl
    extends com.ibm.IManagedLocal._INonManageableImpl
    implements PolicyCopy
{
    public int policyNo()
    {
        return iPolicyNo;
    }

    public void policyNo(int no)
    {
        iPolicyNo = no;
    }

    public void changeAmount(float newAmount)
        throws InvalidAmount
    {
        if(newAmount < 0.0)
            throw new InvalidAmount();
        iAmount = newAmount;
    }

    public float getAmount()
    {
        return iAmount;
    }

    public String comment()
    {
        return iComment;
    }

    public void comment(String comment)
    {
        iComment = comment;
    }

    private int iPolicyNo;
    private float iAmount;
    private string iComment;
}

```

```

public void externalize_to_stream(org.omg.CosStream.StreamIO s)
{
    s.write_long(iPolicyNo);
    s.write_float(iAmount);
    s.write_string(iComment);
}

public void internalize_from_stream(org.omg.CosStream.StreamIO s,
                                     org.omg.CosLifeCycle.FactoryFinder ff)
throws    org.omg.CosLifeCycle.NoFactory,
          org.omg.CosStream.ObjectCreationError,
          org.omg.CosStream.StreamDataFormatError
{
    iPolicyNo = s.read_long();
    iAmount = s.read_float();
    iComment = s.read_string();
}
}

```

The streaming methods do not need to be stream-compatible with those in the Policy business object implementation. In fact, they cannot, because the copy helper does not contain all of the instance data that the business object does. But, as was true for the Primary Key class, it is essential that the streaming methods of the Java implementation be stream-compatible with those of the C++ implementation.

Loading C++ DLLs from Java business objects

The following guidelines were mandatory in previous releases of Component Broker. They are no longer required but programs which use them will continue to work in this release. The guidelines are included here only for illustration purposes and for backwards compatibility.

To access a managed object that may be in an application installed on another host in the network, the Java BO programmer may explicitly load the C++ DLL to access the managed object's C++ components. To do this in object builder, define a private static method on the BO, and call it something like `libLoad`, and have it return a boolean. Then, define a private static attribute, say `ilibLoaded`, and invoke the `libLoad` method as it's initializer.

The content of the `libLoad` method needs to invoke the Java `System.load` function to explicitly specify the path and DLL name to be loaded. Unfortunately, this forces a tight coupling of the Java BO implementation and the directory in which the application is installed. For example, the code generated by object builder will look something like this:

```

static private boolean iLibLoaded = libload();
static private boolean libload()
{
    //Version identifier DCE:12457A40-339B-11d2-B197-0004ACEA9E5A:1
    // Insert Method modifications here
    boolean retval = false;
    try
    {
        System.load("d:\\ntapps\\PolicyApp\\bin\\PolicyC");
        retval = true;
    }
    catch(UnsatisfiedLinkError u)
    {
        System.out.println("ERROR:Caught a load error: " + u);
    }
    catch(Exception e)
    {
        System.out.println(
            "ERROR:Caught an exception loading the C++ DLL: " + e);
    }
    return retval;
    // End Method modifications here
}

```

An alternative solution is to copy the needed DLL into a directory that exists in the path, then the `System.loadLibrary` function can be used, where only the DLL name is specified. This call would look like this:

```
System.loadLibrary("PolicyC");
```

Additional information component creators should know

Use native language support.

Developing the business logic that implements business object interfaces (that is, component interfaces) is done using as much of the Java language as is appropriate for optimally supporting the interfaces. Use native language class libraries to assist when needed. Choices on how to best implement the business logic are yours to make.

Thread usage

Do not create separate threads. The Component Broker server environment in which the components reside is complex. Component Broker must control the threading environment. Creation of additional threads from within business logic could cause unexpected results on the Component Broker server.

Do not use `exit()`

In many cases in error branches of code a user will code an `exit()` to end the program. This technique must not be used in server programs because the server process will be terminated and all application programs will be terminated. A better approach to use in server programs for these conditions is to be returned to the client application.

Where to next?

In order to have a component that can be tested and installed into a Component Broker server, some additional steps must be taken. If you want to provide more advanced features for clients to use and to leverage additional managed object framework features as part of the implementation of your component, then see “Chapter 11. MOFW - Java client programming model - advanced concepts” on page 281. See “Chapter 12. MOFW - Java server programming model - advanced concepts” on page 305 for information on how to add these advanced features to a business object.

Chapter 11. MOFW - Java client programming model - advanced concepts

For further information on Quality of Service Interfaces, see “Expanding the client programming interface” on page 434.

Transactions

Code segments in this section show how a client could use some of the Transaction Services. The example uses Policy objects. The following code segment performs the initialization that is required to use transactions.

```
import org.omg.CosTransactions.*;
org.omg.CosTransactions.Current currentTransaction;
org.omg.CORBA.Object obj;
org.omg.CORBA.ORB orb;
CBSeriesGlobal.Initialize(args, props);
orb = CBSeriesGlobal.ORB();
obj = orb.resolve_initial_references("TransactionCurrent");
currentTransaction = org.omg.CosTransactions.CurrentHelper.
```

You can proceed to find, create, and use a factory as shown in earlier examples. Before attempting to create a Policy managed object, the `begin()` method can be called on the `currentTransaction` pointer to indicate the start of a transaction. When using transactions it is advisable to begin the transaction early in the processing cycle. All method calls that might result in a database access must be always called within the scope of a transaction. For example, methods such as `findByPrimaryKeyString()` and `createFromPrimaryKeyString()` may result in database accesses, it is safer to begin the transaction before making these calls. Beginning transactions early in the processing cycle will hide differences in implementation differences between application adaptors and other specific implementations in later releases.

```
PolicyKey keyVar = PolicyKeyHelper._create();
keyVar.policyNo(55555);
byte[] theKeyString = keyVar._toString();
// Start the transaction now
try {
    currentTransaction.begin();
} catch (java.lang.Exception exception) {
    System.out.println("Error while attempting to begin transaction");
    exception.printStackTrace();
    System.exit(2);
}
/* Assume that the policy home is found */
try {
    obj = policyHome.createFromPrimaryKeyString(theKeyString);
```

```

} catch (com.ibm.IManagedClient.IDuplicateKey exception) {
    // Handle duplicate key exception
} catch (com.ibm.IManagedClient.IInvalidKey exception) {
    // Handle invalid key exception
}
policy = PolicyHelper.narrow(obj);
policy.amount( (float)25000.00);

```

When the managed object has been created, methods can be called on it and its data can be manipulated as desired, all within the scope of the transaction that was started in the previous example. To indicate the termination of the transaction, the commit method is called and the Policy is released.

```

try
{
    currentTransaction.commit(true);
}
catch (java.lang.Exception exception)
{
    exception.printStackTrace();
    System.exit(3);
}

```

The previous segments assume that the data manipulations resulted in desired updates to the data. The commit method completes the current transaction by making these changes to the data permanent. But what if a condition occurred in which the client wanted to end the transaction without any changes?

The following code example shows how the rollback method can be used to terminate a transaction without an update to the data. Two policies are created and have their data set. An if test determines whether the changes are committed or rolled back.

```

org.omg.CORBA.Object obj;
org.omg.CORBA.ORB orb;
com.ibm.IExtendedLifeCycle.FactoryFinder myFinder;
org.omg.CosTransactions.Current currentTransaction=null;
com.ibm.IManagedClient.IHome policyHome=null;
com.ibm.IManagedClient.IManageable mo=null;
byte[] theKeyString;
Policy aPolicy=null;
Policy aPolicy2=null;
CBSeriesGlobal.Initialize(args, props);

try
{
    obj = CBSeriesGlobal.ORB().resolve_initial_references
        ("TransactionCurrent");
    currentTransaction = org.omg.CosTransactions.CurrentHelper.narrow(obj);
    currentTransaction.set_timeout(600);
    obj = CBSeriesGlobal.nameService().resolve_with_string("host
        /resources/factory-finders/host-scope");
    myFinder = com.ibm.IExtendedLifeCycle.FactoryFinderHelper.narrow(obj);

```

```

    obj = myFinder.find_factory_from_string ("Policy.object
        interface/policyB0IMSpecializedFactory.object home");
    policyHome = com.ibm.IManagedClient.IHomeHelper.narrow(obj);
}
catch (java.lang.Exception exception)
{
    exception.printStackTrace();
    System.exit(1);
}
// Start the transaction now
try {
    currentTransaction.begin();
}
catch (java.lang.Exception exception)
{
    exception.printStackTrace();
    System.exit(2);
}
try
{
    PolicyKey theKey = PolicyKeyHelper._create();
    theKey.policyNo(12345);
    theKeyString = theKey._toString();
    mo = policyHome.createFromPrimaryKeyString(theKeyString);
    aPolicy = PolicyHelper.narrow(mo);
    theKey.policyNo(99999);
    theKeyString = theKey._toString();
    mo = policyHome.createFromPrimaryKeyString(theKeyString);
    aPolicy2 = PolicyHelper.narrow(mo);
}
catch (java.lang.Exception exception)
{
    exception.printStackTrace();
    System.exit(3);
}
System.out.println("Enter 'C' for commit, anything else to rollback");
byte[] input = new byte[512];
try
{
    System.in.read(input);
}
catch (java.lang.Exception exception)
{
    exception.printStackTrace();
    System.exit(4);
}
String inputString = new String(input,0);
if(inputString.startsWith("C") || inputString.startsWith("c") )
{
    // End transaction now with updates to data
    try
    {
        currentTransaction.commit(true);
    }
    catch (java.lang.Exception exception)
    {

```

```

        exception.printStackTrace();
        System.exit(5);
    }
}
else
{
    // End transaction now with no changes to data
    try
    {
        currentTransaction.rollback();
    }
    catch (java.lang.Exception exception)
    {
        exception.printStackTrace();
        System.exit(6);
    }
}
}

```

Transactions, exceptions and time-outs

“Transactions” on page 281 provides a basic usage view of the transactional interfaces. This section provides additional guidelines on application structuring in the area of time-outs and exception handling.

Consider the following code example to explain why the code is structured in this way:

```

// get current See note 1
CBSeriesGlobal.Initialize(args, props);
try {
    org.omg.CORBA.Object obj = CBSeriesGlobal.orb().resolve_initial_references
        ("TransactionCurrent");
    currentTransaction = org.omg.CosTransactions.CurrentHelper.narrow(obj);
    currentTransaction.set_timeout(180); // See note 3
    ...
} catch (java.lang.Exception exception) {
    exception.printStackTrace();
    System.exit(1);
}
// you could have a loop starting here...
try {
    try {
        currentTransaction.begin(); // See note 2
    } catch (java.lang.Exception exception) {
        // Begin transaction failed...
    }
    .... do work
    if (businessLogicSaysCommit)
        try {
            currentTransaction.commit(true); // See note 2
        } catch (java.lang.Exception exception) {
            // Commit failed...
        }
    }
else

```

```

    try { // See note 4
        currentTransaction.rollback();
    } catch (java.lang.Exception exception) {
        // Rollback failed ...
    }
} catch (org.omg.CORBA.SystemException exception) { // See note 5

try {
    current.rollback();
}
catch(java.lang.Exception exception) {
    // Rollback failed ...
}
} catch (org.omg.CORBA.UserException exception) { // See note 6
    try {
        current.rollback();
    }
    catch(java.lang.Exception exception) {
        // Rollback failed ...
    }
}
} catch (java.lang.Exception exception) { // See note 7
    try {
        current.rollback();
    }
    catch(java.lang.Exception exception) {
        // Rollback failed ...
    }
}
}

```

What does all of this mean? In the previous example, the numbers in the comments correspond to the following list items:

1. Get the current. This is the same as always. Get it once for performance reasons.
2. Begin and Commit transactions. Bracket units of work in accordance with the requirements of your business logic and domain needs. In other words, if you want to be able to rollback a set of operations, put them in a begin or commit block. The issue is how much time is normally expected between the begin and commit. You must realize that there is some level of resource locking that goes on while there is a transaction in-flight. Shorter transactions will increase throughput, and so on.
3. This is an important line of code. This code signals to the server that this is how long (in seconds) the server should wait before taking matters into its own hands. This means that the server will roll back without any user code having specified rollback. If a time-out is not specified through this method call, a platform-dependent value is used. This value is configurable using the System Management interfaces.
4. Rollback can be issued at any time based on the needs of the business logic. This undoes changes made since the last begin.

5. When a system exception occurs, you should code a `rollback()` in the catch block or some routine that the catch block calls.
6. When a CORBA user exception occurs, you should code a `rollback()` in the catch block or some routine that the catch block calls. This exception block would include exceptions such as `com.ibm.IManagedClient.INoObjectWithKey`.
7. The exceptions that make it to the final catch block, by definition, were thrown from within the client program. Anything that flowed over a wire would have been remapped to a CORBA exception and would have been caught in item number 5. A rollback here will also work properly and release locks.

Using transactions over reference collections

You can encounter a transaction deadlock problem while running more than one instance of a program which in one transaction:

1. Adds a number of records to reference collections (or home collections)
2. Then retrieves the records either directly by creating an iterator on a collection or indirectly through queries on a collection.

This error could be reported either as `CORBA.INTERNAL` or `CORBA.PERSIST_STORE`.

To prevent transaction deadlock problems, split the operation between two transactions. Add records to collections in one transaction and commit the transaction. Then create a second transaction to retrieve the records from the collections.

Session Service



OS/390 Component Broker does NOT support Session Service.

The Session Service defines the notion of a session. It provides a mechanism for grouping a set of operations together as a logical unit. A session is conceptually similar to a transaction as defined and supported by the Transaction Service. Sessions differ from transactions in the following ways:

- A session is defined on an application scope rather than on an individual transaction scope. This provides a mechanism for checkpointing groups of persistent objects for which no externally coordinated transactional update is available.

- Sessions do not define an atomically recoverable commit scope as do transactions; sessions provide some services for checkpointing and clean conclusions of applications but do not provide the same level of application durability as transactions.
- A session can use an application profile. An application profile consists of attributes that define its properties, such as requirements for data concurrency, duration, visibility, update, execution priority, and resource dependencies.

Sessions and transactions can be used together. Transactions can be used within sessions to provide greater levels of durability within applications.

For additional information about the Session Service, see the *WebSphere Application Server Enterprise Edition Component Broker Advanced Programming Guide*.

A simple example

The following example demonstrates how a single-threaded client can begin a session, perform some work, and end the session, checkpointing any non-transactional work that occurs within the session.

```
import org.omg.CORBA.*;
import com.ibm.ISessions.*;
org.omg.CORBA.Object obj;
com.ibm.ISessions.Current sessionCurrent;

// Get the current for this thread of execution
// from the ORB and narrow to the session current.

obj = CBSeriesGlobal.orb().resolve_initial_references("ISessions::Current");
sessionCurrent = ISessions.CurrentHelper.narrow(obj);
```

Set a time limit for all new sessions

The `ISessions.Current` interface has a `setSessionTimeout()` operation that enables the application to set a time limit for all sessions that are subsequently started. The default time limit is zero. That is, sessions can run indefinitely. To set a time limit for all new sessions, invoke the `setSessionTimeout()` operation on the `ISessions.Current` object, passing the new time-out value.

```
// Set the session timeout.

sessionCurrent.setSessionTimeout(100);
```

Begin a session

You begin a session by invoking the `beginSession()` operation on the `ISessions.Current` interface. You should supply a text string representing the name of your application or the application profile under which you want the session

to operate. If the specified profile cannot be found or if you specify an empty string, then a default application profile is used.

The application profile specifies certain expectations about how the session will behave. This information can be used in combination with the capabilities and policies of the running system to produce a set of execution decisions that optimize the total performance and throughput of the system. Once the session has been started you can perform any number of operations on business objects within the session. All operations invoked within the session are performed with the same session context.

```
try {  
  
    // Begin a session context.  
    sessionCurrent.beginSession("LifeInsuranceApplication");  
    // ... do the methods that will be executed under the  
    // session ...  
} catch (com.ibm.ISessions.SESSION_RESET_FORCED exception) {  
    // The session was forced to reset mid-stream, probably the  
    // session timeout tripped, or a session resource  
    // encountered a significant error and had to force the session  
    // reset.  
};
```

End a session

You complete the session by invoking the `endSession()` operation on the `sessionCurrent`. Normally, you specify the *EndModeCheckPoint* end-mode with this operation. This drives all the sessionable resources used within the session to save their state changes persistently, through embedded operations on the underlying data system.

To reset the session, end it without saving any of the changes that occurred since your last checkpoint (or since the beginning of the session if you did not perform any checkpoints). Specify the *EndModeReset* end-mode with the `endSession()` request.

EndModeCheckPoint and *EndModeReset* have no bearing on any transactions issued within the session other than to ensure that the session is terminated before it ends. However, if you encounter severe errors in your processing, you can end the session with *EndModeResetForce*. This forces the session to be reset immediately, including rolling back any outstanding transactions.

```
// End the session context, including checkpointing any activity that  
// occurred during the session.  
  
try  
{
```

```
        sessionCurrent.endSession(  
            com.ibm.ISessions.EndMode.EndModeCheckPoint,true)  
    }  
    // Catch a variety of exceptions.
```

See the *WebSphere Application Server Enterprise Edition Component Broker Advanced Programming Guide* for details.

Other information

Further information on these and other topics regarding the Session Service can be found in the *WebSphere Application Server Enterprise Edition Component Broker Advanced Programming Guide*. For example, the following information is available:

Suspending and resuming sessions

There may be occasions when you want to switch the session under which you are operating. You can do this by suspending the current session and starting a new one. Later, you can resume the original session.

Explicit and implicit propagation of session context

This section describes different techniques of assigning context information for multi-threaded applications.

Checkpoint and reset a session context

Sessions can be checkpointed to save intermediate results to persistent data storage. Sessions can also be reset to revert the state of operations performed within the session.

Registering sessionable resources

An application can introduce resources and explicitly register the resources with a session.

Visibility rules

Sessions can be run concurrently. Visibility rules define how the data interactions between these concurrent sessions are defined.

Queries, iterations and specialized homes

“Chapter 9. MOFW - Java client programming model” on page 231 introduced generic homes as a facility for creating business objects and for finding business objects based on their primary key. This chapter discusses how a client of a business object could use other features that a home could support. If the object provider has built a specialized home, the usage pattern for that home is different than that of a generic home, because a specialized home has

additional methods that can be used by clients. Besides the additional methods available to the client, there is also a change in how this specialized home is found.

Using iterated homes-specific functions

Many times an application needs access to multiple objects. In “Using sets of objects” on page 95, an example of iterating through an arbitrary collection of business objects is presented. This iteration assumes that all the objects are inserted into an `IManagedReferenceCollection`.

If you wanted to examine a group of homogenous objects (for example, every `Claim`) and perform actions on those objects that met certain criteria, the Component Broker programming model extends the concept of iteration to `IHomes`. The following code segment shows the usage of an iterated home.

```
org.omg.CORBA.Object obj;
com.ibm.ExtendedLifeCycle.FactoryFinder myFinder;
com.ibm.IManagedAdvancedClient.IIterableHome policyHome=null;
com.ibm.CollectionsBase.IIterator theIterator=null;

try {
    obj = CBSeriesGlobal.nameService().resolve_with_string
        ("host/resources/factory-finders/host-scope");
    myFinder = com.ibm.ExtendedLifeCycle.FactoryFinderHelper.narrow(obj);
    obj = myFinder.find_factory_from_string(
        "Policy.object interface/policyBOIMSpecializedFactory.object home");
    policyHome =
        com.ibm.IManagedAdvancedClient.IIterableHomeHelper.narrow(obj);
    if ( policyHome == null ) {
        System.out.println("ERROR: Application improperly configured or
            Home for Policies is not iterable.");
        System.exit(1);
    }
} catch (java.lang.Exception exception) {
    exception.printStackTrace();
    System.exit(2);
}
com.ibm.IManagedClient.IManageable aBO;
// Repeat transactional setup.
try {
    currentTransaction.begin();
} catch (java.lang.Exception exception) {
    exception.printStackTrace();
    System.exit(3);
}
try {

    // Create a new iterator on the Claim home
    theIterator = policyHome.createIterator();

    // Loop through the objects in the home
```

```

aBO = theIterator.next();
while ( aBO != null ) {

    // Get the next element
    Policy aPolicy;
    aPolicy = PolicyHelper.narrow(aBO);

    // Do something useful with it.
    System.out.println("Policy no = " + aPolicy.policyNo());
    System.out.println("Policy amount = " + aPolicy.amount());
    aBO = theIterator.next();
}
} catch (java.lang.Exception exception) {
    System.out.println("ERROR: Problem occurred using iterator.");
    try {
        theIterator.remove();
        currentTransaction.rollback();
    } catch (java.lang.Exception anotherException) {
        // Rollback failed ...
    }
}

// After iterating over the entire collection, the iterator is no
// longer needed. Remove it.
try {
    theIterator.remove();
    currentTransaction.commit(true);
} catch (java.lang.Exception exception) {
    exception.printStackTrace();
    System.exit(4);
}

```

In addition to the next() method, ICollectionsBase.IIterator provides the following methods with similar functionality (next() is much more concise): hasMoreElements(), more(), nextElement(), nextOne().

While the previous examples are functionally identical (with identical performance), there is another interface for iterating which is functionally similar, but with different performance characteristics. The nextN() method is similar to the nextOne() method, except that instead of returning only one element, nextN() returns as many elements as you request.

This interface might be useful if, for example, an application wants to display information about a fixed number of business objects at a time (limited by the size of a window on a screen). Here is an example of how to use this interface in the same application as the previous example:

```

org.omg.CORBA.Object obj;
com.ibm.IExtendedLifeCycle.FactoryFinder myFinder;
com.ibm.IManagedAdvancedClient.IIterableHome policyHome=null;
org.omg.CosTransactions.Current currentTransaction=null;
com.ibm.ICollectionsBase.MemberListHolder theBOlist =
    new com.ibm.ICollectionsBase.MemberListHolder();

```

```

com.ibm.ICollectionsBase.IIterator theIterator=null;
com.ibm.IManagedClient.IManageable aB0;

Policy aPolicy;
byte[] input = new byte[512];
boolean continueLoop=true;
try {
    obj = CBSeriesGlobal.orb().resolve_initial_references(
        "TransactionCurrent");
    currentTransaction = org.omg.CosTransactions.CurrentHelper.narrow(obj);
    currentTransaction.set_timeout(600);
    obj = CBSeriesGlobal.nameService().resolve_with_string(
        "host/resources/factory-finders/host-scope");
    myFinder = com.ibm.IExtendedLifeCycle.FactoryFinderHelper.narrow(obj);
    obj = myFinder.find_factory_from_string(
        "Policy.object interface/policyDefaultTransDB2Factory.object home");
    policyHome =
        com.ibm.IManagedAdvancedClient.IIterableHomeHelper.narrow(obj);
    if ( policyHome == null ) {
        System.out.println("ERROR: Application improperly configured or
            Home for Policies is not iterable.");
        System.exit(1);
    }
} catch (java.lang.Exception exception) {
    exception.printStackTrace();
    System.exit(2);
}
try {
    currentTransaction.begin();
} catch (java.lang.Exception exception) {
    exception.printStackTrace();
    System.exit(3);
}
try {

    // Create a new iterator on the Claim home
    theIterator = policyHome.createIterator();

    // Loop through the objects in the home.
    while (continueLoop) {

        // This will return false when the number of items returned is less
        // than the number requested, this will control exiting the while loop
        continueLoop = theIterator.nextN(5, theB0list);

        try {
            for ( int n = 0; n < 5; ++n) {

                // Narrow the next element obtained by nextN()
                aPolicy = PolicyHelper.narrow(theB0list.value[n]);

                // Do something useful with it
                System.out.println("Policy # " + aPolicy.policyNo());
                System.out.println("amount is " + aPolicy.amount());
            }
        }
    }
}

```

```

        // If the number of entries returned is less than the amount
        // requested then will hit this exception while processing
        // the array, this is a good way to stop processing
    } catch (java.lang.ArrayIndexOutOfBoundsException exception) {

        // Do nothing, processed all the entries

    } catch (java.lang.Exception exception) {
        try {
            theIterator.remove();
            currentTransaction.rollback();
        } catch (java.lang.Exception ignoredException) {

            //fall through to error handling

        }
        exception.printStackTrace();
        System.exit(4);
    }

    // Having displayed as many policies as would fit in the 'window',
    // the application now waits for user input before getting the
    // next N policies and displaying them.

    if (continueLoop) {
        System.out.println("press any key (and hit enter) to continue.");
        try {
            System.in.read(input);
        } catch (java.lang.Exception exception) {
            try {
                theIterator.remove();
                currentTransaction.rollback();
            } catch (java.lang.Exception ignoredException) {

                //fall through to error handling

            }
            exception.printStackTrace();
            System.exit(5);
        }
    }
} catch (java.lang.Exception exception) {

    try {
        theIterator.remove();
        currentTransaction.rollback();
    } catch (java.lang.Exception ignoredException) {

        //fall through to error handling

    }
    exception.printStackTrace();
}

```

```

        System.exit(6);
    }

    // After iterating over the entire collection, the iterator is no
    // longer needed. Remove it.

    try {
        theIterator.remove();
        currentTransaction.commit(true);
    } catch (java.lang.Exception exception) {
        exception.printStackTrace();
        System.exit(7);
    }
}

```

The previous examples illustrate several ways to iterate through the objects in a home. To understand iteration better, a few general notes about iterators follow:

- When an iterator is created, it is positioned so as to precede the first element.
- The various types of next methods (`next()`, `nextElement()`, `nextOne()` and `nextN()`) all move the position of the iterator forward, then return the element at the current position.
- There is a `current()` method that returns the element at the current position without moving the position of the iterator forward. Because the initial position of an iterator precedes the first element, it is an error to invoke `current()` prior to invoking one of the next methods.
- There is a `reset()` method which moves the position back to its initial position.
- A home that is iterated is not guaranteed to return its elements in the same order on successive iterations, but it is guaranteed to return each element once only on a given iteration.

Members of Component Broker homes are always accessible through at least one key (the Primary Key), but may or may not be iterable. If the collection is based on or wrappers a relational or object-oriented database, then both keys and iterator can be supported. The same is true for most data back-ends (such as files, IMS DL/1 or CICS File Control). If the collection wrappers a set of applications (such as the Customer Management Application in the “Chapter 2. Personal life insurance application example” on page 31), iteration might not be supported by the application. Therefore, the wrapping collection cannot be iterated. Check with your System Administrator to see if your home supports iteration. Because the `LifeCycle` object service factory finder interface is used to find homes and this interface has no way of telling an iterable home from a generic home, a client must be properly configured if its function depends on access to an iterable home. The client can protect itself against incorrect configuration by verifying that the home returned by the factory finder supports the `IManagedAdvancedClient.IterableHome` interface. A verification example is:


```

policyHome = com.ibm.IManagedAdvancedClient.IIterableHomeHelper.narrow(obj);
if ( policyHome == null )
{
// Insert error message here to be displayed here.
}

```

Using queryable homes-specific functions

Iterating over the business objects in an `IterableHome` is fine for some applications. Specifically, if the goal is to perform some operation on every element in the home, then iteration is a good approach. However, if the reason for iterating is to select a subset of all the business objects in the home (and then work with only that subset), then there is a more efficient way to accomplish the same thing: query.

Here the iteration example is used again; however, instead of iterating against all of the objects in a home, you select only the objects in which you are interested (by evaluating a query), and then you iterate over the subset in which you are interested. Because it is sometimes possible for the query to be performed in the database (without bringing into memory every object in the home), using query offers the possibility of significant performance increases over iterating against every object in the home.

```

org.omg.CORBA.Object obj;
com.ibm.IExtendedLifeCycle.FactoryFinder myFinder;
com.ibm.IManagedAdvancedClient.IQueryableIterableHome policyHome=null;
com.ibm.ICollectionsBase.IIterator theIterator=null;
org.omg.CosTransactions.Current currentTransaction=null;
com.ibm.IManagedClient.IManageable aB0;
Policy aPolicy;

try {
    obj = CBSeriesGlobal.orb().resolve_initial_references(
        "TransactionCurrent");
    currentTransaction = org.omg.CosTransactions.CurrentHelper.narrow(obj);
    currentTransaction.set_timeout(600);
    obj = CBSeriesGlobal.nameService().resolve_with_string(
        ("host/resources/factory-finders/host-scope"));
    myFinder = com.ibm.IExtendedLifeCycle.FactoryFinderHelper.narrow(obj);
    obj = myFinder.find_factory_from_string(
        "Policy.object interface/policyDefaultTransDB2Factory.object home");
    policyHome = com.ibm.IManagedAdvancedClient.
        IQueryableIterableHomeHelper.narrow(obj);

    if ( policyHome == null ) {
        System.err.println("ERROR: Application improperly configured. Home for
            Policies is not queryable.");
        System.exit(1);
    }
} catch (java.lang.Exception exception) {
    exception.printStackTrace();
    System.exit(2);
}

```

```

// Evaluate a query on the Policy home. The results of the
// query are returned in the form of an iterator.
// Set up for transactions

try {
    currentTransaction.begin();
} catch (java.lang.Exception exception) {
    exception.printStackTrace();
    System.exit(3);
}
try {
    theIterator = policyHome.evaluate("policyNo>3;");

    // Loop through the results of the query
    aB0 = theIterator.next();
    while ( aB0 != null ) {

        // Narrow the element obtained by nextOne()
        aPolicy = PolicyHelper.narrow(aB0);

        // Do something useful with it.
        System.out.println("Policy no = " + aPolicy.policyNo());
        System.out.println("Policy amount = " + aPolicy.amount());
        aB0 = theIterator.next();
    }
} catch (java.lang.Exception exception) {

    try {
        theIterator.remove();
        currentTransaction.rollback();
    } catch (java.lang.Exception innerException) {
        innerException.printStackTrace();
        System.exit(4);
    }
}

// After iterating over the entire collection, the iterator
// is no longer needed. Remove it.

try {
    theIterator.remove();
    currentTransaction.commit(true);
} catch (java.lang.Exception exception) {
    exception.printStackTrace();
    System.exit(5);
}

```

The query language used to express the query is SQL with extensions and provisions for being able to query against objects (instead of data only in a database).

Using atomic transactions with query evaluator

The programming model for method duration transactions (called atomic or automatically start a new transaction) occurs when a method is invoked on the object; the method then will be wrapped within a transaction. That is, prior to invoking the method, a transaction will automatically start. When the method returns the transaction will be committed. The programming model has been extended so that creating or finding an object that is in a container (supporting method duration transactions then creation or finding) will also be wrapped in a transaction. The programming model has not been expanded to include the Query Service.

Important: A transaction must have been started prior to using the Query Service. For additional information about the Query Service, see the *WebSphere Application Server Enterprise Edition Component Broker Advanced Programming Guide*.

More on iterators

This information addresses the following:

- How iterators over homes and collections are used
- How query iterators work
- How atomic containers are configured
- Side effect of end transaction

Managed object iterators over home collections, iterators over persistent reference collections, query iterators over persistent managed objects and query data array iterators over persistent objects become invalid at the end of the current transaction in which they were created.

When designing objects that will be configured into an atomic container, the object interface should not have attributes or methods which return to the client any of the above types of iterators.

If the object interface does have such iterators, they can be retrieved by a client only if the client explicitly starts a transaction and retrieves and uses the iterator in the same transaction scope.

Note the following example:

```
interface X
{
    com.ibm.CollectionsBase.IIterator methodX();
}
```

If this interface is implemented by an MO, configured into an atomic container and if my client does not start a transaction, the following statements will result in an exception because the iterator is being used outside the scope of a transaction.

```
com.ibm.ICollectionsBase.IIterator i = X.methodX();
com.ibm.IManagedClient.IManageable mo = X.next();
```

However if the client does the following, it will be valid because the iterator is being retrieved and used in the same transaction scope.

```
currentTransaction.begin();
com.ibm.ICollectionsBase.IIterator i = X.methodX();
com.ibm.IManagedClient.IManageable mo = X.next();
```

Using keyed reference collections

“Using sets of objects” on page 95 describes how to use a Reference Collection (that is, `IManagedCollections.IReferenceCollection`) to maintain references to a set of managed objects. If the elements in a Reference Collection require keyed access (that is, Hashtable-like access), a Keyed Reference Collection can be used instead. Keyed Reference Collections are defined by the `IManagedCollections.IKeyedReferenceCollection` interface.

Keyed and non-keyed Reference Collections have many similarities. Both the `IReferenceCollection` and `IKeyedReferenceCollection` interfaces derive from the `ManagedCollections.ICommonCollection` common base interface that defines a number of methods that apply to both keyed and non-keyed collections. The interfaces include many commonly used methods such as `containsElement()`, `isEmpty()`, `numberOfElements()`, and `createIterator()`. In addition, Keyed Reference Collections are created using the same specialized home (that is, `IManagedCollections::ICollectionHome`) as non-keyed collections. The following code segment illustrates the creation of a Keyed Reference Collection:

```
org.omg.CORBA.Object obj;
com.ibm.IManagedCollections.ICollectionHome kcHome;
com.ibm.IManagedCollections.ICommonCollection cc;
com.ibm.IManagedCollections.IKeyedReferenceCollection kc;
com.ibm.IExtendedLifeCycle.FactoryFinder myFinder;

try {
    obj = CBSeriesGlobal.nameService().resolve_with_string
        ("host/resources/factory-finders/host-scope");
    myFinder = com.ibm.IExtendedLifeCycle.FactoryFinderHelper.narrow(obj);
    obj = myFinder.find_factory_from_string
        ("IManagedCollections::IKeyedReferenceCollection.object
        interface/TransientKeyedReferenceCollectionFactory.object home");

    kcHome = com.ibm.IManagedCollections.ICollectionHomeHelper.narrow(obj);
```

```

    cc = kcHome.createCollection();
    kc = com.ibm.IManagedCollections.IKeyedReferenceCollectionHelper.narrow(cc);
} catch (java.lang.Exception exception) {
    exception.printStackTrace();
}
}

```

This code is identical to that shown in “Using sets of objects” on page 95 to create non-keyed collections except that the `IReferenceCollection` interface name is replaced by `IKeyedReferenceCollection`.

The difference between keyed and non-keyed collections is the way objects are added and accessed. Instead of calling `IReferenceCollection.addElement()`, the `IKeyedReferenceCollection.addElementByString()` method is used to add elements to the collection. The `addElementByString()` method requires two arguments, the object to be added and a stringified key (that is, a `String`) that is used for identifying the object in the collection. The key can be any appropriate subclass of `IManagedLocal.IKey`, either general purpose (for example, `StringKey`), application specific (for example, `PolicyHolderIdentifierKey`), or, for that matter, the primary key (`PolicyHolderPrimaryKey`).

For example, assume a general purpose `IKey` subclass has been defined:

```

interface StringKey : IManagedLocal::IKey
{
    attribute string value;
    #pragma meta StringKey localonly
}

```

Using this key, elements can be added to the Keyed Reference Collection as follows:

```

// Get some policy holders
PolicyHolder policyHolderJohn;
PolicyHolder policyHolderKatherine;

// Create a key object
StringKey theKey = StringKeyHelper._create();
byte[] keyString;

// Add policyHolderJohn to the collection with key "John"
theKey.value("John");
keyString = theKey._toString();
kc.addElementByString(policyHolderJohn, keyString);

// Add policyHolderKatherine to the collection with key "Katherine"
theKey.value("Katherine");
keyString = theKey._toString();
kc.addElementByString(policyHolderKatherine, keyString);

```

An element can be retrieved using the `getElementByString()` method. For example, the following code segment retrieves the policy holder “John”:

```

theKey.value("John");
keyString = theKey._toString();
com.ibm.IManagedClient.IManageable mo = kc.getElementByString(keyString);
PolicyHolder thePolicyHolder = PolicyHolderHelper.narrow(mo);

```

To remove an element from the collection, the `removeElementByString()` method is used:

```

kc.removeElementByString(keyString);

```

Iterating through the elements in a Keyed Reference Collection is done in exactly the same way as for non-keyed collections. In fact, if the collection is referenced using the `IManagedCollection::ICollection` base interface, as shown in the following code segment, the same code can be used to iterate over elements of either keyed or non-keyed collections. Note that if the variable *theCollection* were of type `ICollectionsBase.IMIterable`, the same code segment could be used to iterate over an iterable home.

```

// Get a keyed or non-keyed collection from somewhere
com.ibm.IManagedCollections.ICollection theCollection;
com.ibm.ICollectionsBase.Iterator theIterator =
    theCollection.createIterator();
com.ibm.IManagedClient.IManageable theB0;
PolicyHolder aPolicyHolder;
PolicyBeneficiary aPolicyBeneficiary;

try {

//Loop through the collection. The "next" method advances
//the iterator to the next element (on the first invocation, this
//will advance to the first element) and then return the element
//pointed to by the iterator.

    while ( theIterator.more() ) {
        theB0 = theIterator.next();
        aPolicyHolder = PolicyHolderHelper.narrow(theB0);
        if ( aPolicyHolder != null ) {
            // Send them a bill
        }
        aPolicyBeneficiary = PolicyBeneficiaryHelper.narrow(theB0);

        if ( aPolicyBeneficiary != null ) {
            // Send them a check
        }
    }
} catch (java.lang.Exception exception){

    try {
        theIterator.remove();
    } catch (java.lang.Exception ignoredException){

```

```
        // Fall through to normal error handling
    }
    exception.printStackTrace();
}
```

Each element in the keyed or non-keyed collection is accessed only once and in no defined order. A number of other methods are available on the `IKeyedReferenceCollection` interface:

`containsKeyString()`

Determines if an element with a given key exists.

`getElementKeyString()`

Determines the key of a given element.

`replaceElementWithKeyString()`

Replaces an element (that is associated with a specified key) with another object.

With these and the other Keyed Reference Collection methods, an arbitrary set of keyed references to Component Broker managed objects can be maintained easily.

Conventions and guidelines

This section provides additional information on what is happening on the server when client programs are invoking methods on the business objects. This section does not describe new interfaces but provides additional technical details on what actually occurs. Some of these topics are guidelines or coding techniques that can be useful when interacting with the Component Broker server.

The Component Broker server plays a unique role among application servers: it serves objects. A database server, on the other hand, serves data and a file server serves files. However, the Component Broker server also interacts with database or file servers in cases where these are the chosen options for persistence. In a database server, interaction with the data is direct and the semantics of creating, deleting, and finding are straightforward because you use the methods to the resource manager directly. In the case of an object server, from a programming perspective you encapsulate interaction with a database to get persistence. However, to provide this encapsulation, the decisions about how the database is used are unknown to the clients. Sometimes expectations are not met. The following sections discuss the interaction patterns that the Component Broker has with various resource managers and the coding patterns that you can use to meet various client requirements.

Finding persistent objects

When you want to find an object using the `findByPrimaryKeyString()` method on the `IHome`, Component Broker follows a specific algorithm. This algorithm is:

1. Convert parameter into internal key format.
2. Look in the container cache of the active objects.
3. If not found, go to the database or database cache and look.
4. Return the object reference as soon as it is found.

The programming implications of this are:

- If the object exists, the object reference returned is valid and is ready to use.
- If the object was being cached by the container or the cache manager and had subsequently been deleted by a non-object program, an exception is thrown when the object reference is used.
- The probability of getting an exception when using the returned reference is directly proportional to the amount of deleting that is done by existing non-Component Broker applications that are running concurrently with Component Broker applications.

Creating persistent objects

When you want to create an object using the `createFromPrimaryKeyString()` or other create methods, Component Broker follows a specific algorithm with respect to the creation.

1. Create the managed object and its parts.
2. Put it in the container's cache of active objects.
3. Return the object reference to the client.
4. Wait for transaction commit().
5. When the transaction commits, insert the row in the table.

The implications of this algorithm for you are as follows. If the row already exists in the underlying database, an exception is thrown when the transaction completes.

In many cases the exception is desired and should be planned for accordingly. To alter this behavior for create, you can either:

- Do a `findByPrimaryKeyString()` before issuing the `createFromPrimaryKeyString()`. Remember from the previous discussion that `findByPrimaryKeyString()` looks in the database if necessary to find objects and in this instance would return `notFound`. If there is a slim chance that the object already exists, the performance penalty for the `findByPrimaryKeyString()` might be too high. Select a technique based on the needs of the application.
- Bracket `createFromPrimaryKeyString()` in a transaction by itself. This ensures that the exception for `alreadyExists` is thrown before operations on the newly-created object are started.

These techniques do not reduce the amount of written code but are alternatives for structuring client code.

Chapter 12. MOFW - Java server programming model - advanced concepts

The Component Broker server programming model is based on programming by framework completion. Component Broker introduces its APIs as sets of classes and frameworks. Developers implement their business functions by defining business objects that subclass from Component Broker frameworks and use the frameworks in their implementation.

This chapter provides business object builders with additional options for providing function in business objects. Some of the material in this chapter offers alternative ways of accomplishing some of the tasks described in “Chapter 10. Java server programming model” on page 251.

Some of the topics in this chapter assume the presence of specific application adaptors and other features within Component Broker. In other words, some of the techniques described might increase the cost of porting or targeting business objects and associated applications to other back-end databases and resource managers.

Extending a business object

Chapter 10. Java server programming model presents a basic model for building business objects. Most of the discussion and examples in that chapter center around implementing a new business object interface. Careful examination of the examples shows that there are several meaningful layers of inheritance in the business objects that were presented. Often multiple levels of domain inheritance exist and need to be implemented.

This chapter addresses the addition of a subclass to existing business objects using data object inheritance as the implementation technique. This is the model supported by Component Broker through Object Builder. This is not as simple as adding a single class. This chapter revisits each of the steps used to develop a business object in the context of adding another subclass of domain functionality. These steps are:

1. Developing an interface to the business object.
2. Choosing an inheritance pattern.
3. Implementing the business object methods.
4. Implementing the methods required by the MOFW interfaces.
5. Implementing the key classes.
6. Implementing the copy helper classes.

When you create a child component (that is, a component that inherits behavior or data from another component in your application), the child component objects generally inherit from their equivalent parent objects:

- The child business object file must include the parent business object file.
- The child business object interface must inherit from the parent interface.
- The child key and copy helper can inherit from their equivalents in the parent component, or they can contain selected attributes of the parent interface, without inheriting from the parent key or copy helper.
- The child business object implementation must inherit from the parent implementation.
- The child data object interface must inherit from the parent data object interface.
- The child data object implementation must inherit from the parent data object implementation.
- The child managed object must inherit from the parent managed object.

For data inheritance to work, the type of persistence provided by the parent and child data object implementations must be the same.

Many variations involve extending a business object. The next sections give examples of an extreme case where it inherits as much interface and implementation as possible. There are variations that involve less implementation inheritance that can be extrapolated from the examples given.

Extending business object interfaces

A Java business object can be implemented as a subclass of an existing Java business object, adding extra business methods and state data and overriding or extending the framework method implementations of the base business object class. A Java business object cannot be defined as an extension of a C++ business object.

The first step in building a business object is to construct the interface. In this case, the interface is extended. In the following example, a CarPolicy class extends the Policy.




Figure 52. Inheritance of interface for extended business object

The IDL for CarPolicy should look like the following example.

```

interface CarPolicy : Policy
{
  attribute long year;
  attribute string make;
  attribute string model;
  attribute long serialNumber;
  attribute float collisionDeductible;
  attribute boolean glassCoverage;
  long riskQuotient();
};
  
```

The interface looks like almost any other business object interface. However, because the Policy already inherits from IManagedClient.IManageable, the CarPolicy interface does not need to.

Essential state extensions

Extending the essential state is similar to extending the interface. The data object interface of Policy is extended as shown in the following figure.




Figure 53. Extending the interface for essential state

The IDL for this interface looks like the following example.

```
interface CarPolicyDO : PolicyDO
{
  attribute long year;
  attribute string make;
  attribute string model;
  attribute long serialNumber;
  attribute float collisionDeductible;
  attribute boolean glassCoverage;
  #pragma meta CarPolicyDO localonly ,abstract
};
```

Data object customization and inheritance shows the decisions that must be made when customizing data objects from an implementation perspective.

Choosing an inheritance pattern

The choice of inheritance pattern is based on three concerns:

- Identity: whether parent and child have the same identity (that is, they share the same key)
- Performance trade-offs: whether performance or space efficiency is more important.
- Form of persistence: whether the parent has data to be persisted, and where and how the parent's and child's data is persisted.

If the parent and child have different keys, you should probably use the *Attributes Duplication Pattern*. This means that the child's datastore provides persistence for all of its data, including inherited data. The parent's datastore only provides persistence for instances of the parent, never for instances of the child. If you do not use the overriding persistence pattern, the parent's datastore will have two primary keys: the parent's key for the parent's data, and the child's key for the child's inherited data. It then becomes problematic to determine which data belongs to which object type.

If the parent and child have the same key, you can choose between the *Key Duplication Pattern* and the *Single Datastore Pattern*. The Key Duplication Pattern will generally be more efficient in its use of space (because the persistent objects for each component contain only the data required for that component), and the Single Datastore Pattern will generally provide faster lookup time (because both local and inherited data are mapped to the same persistent object and underlying datastore).

If the parent and child are both persisted in a database, you can compromise between the Key Duplication Pattern and the Single Datastore Pattern, by using the *Single Datastore with Views Pattern*. This pattern uses unique persistent objects for retrieval (the Key Duplication Pattern), and a shared persistent object for all other uses (the shared persistence pattern). This

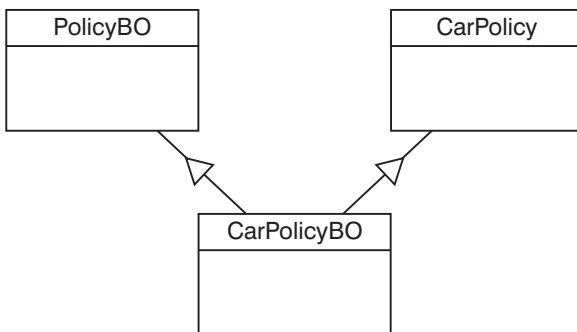
pattern is based on views of the underlying database table, and requires that there be some unique attribute of the child that can be used to select appropriate views of the database.

This example will use the Key Duplication Pattern for illustrations purposes. This is the default pattern supported by Object Builder.

See the “Inheritance” section in the *WebSphere Application Server Enterprise Edition Component Broker Application Development Tools Guide* for further information on the inheritance patterns.

Implement the additional business logic

Next, you must add the implementation of the business logic for the additional methods necessary in the subclass. Introducing another interface for the CarPolicyBO is shown in the following figure and in the following example.



As before, a business object interface is also defined. It is not strictly necessary to explicitly inherit from the `IManagedServer:IWrappable` interface, because `PolicyBO` has already done so:

```
interface CarPolicyBO : CarPolicy, PolicyBO;
```

Use the same pattern (delegating versus caching) in the subclass as was used in the base class. The actual implementation interface has the inheritance of the `PolicyBO` built right into it, as shown in the following example.

In the implementation, bodies for all the new methods must be supplied including set and get methods for the new attributes. Any framework methods whose implementations in the base business object are no longer appropriate must also be replaced:

```
public class _CarPolicyBOBase
    under: _CarPolicyBOBase
    extends _PolicyBOBase
```

```

        implements CarPolicyBO
    {
        public int year()
        public void year(int y)

        public String make()
        public void make(String s)

        public String model()
        public void model(String s)

        public int serialNumber()
        public void serialNumber(int i)

        public float collisionDeductible()
        public void collisionDeductible(float f)

        public boolean glassCoverage()
        public void glassCoverage(boolean b)

        public int riskQuotient()

        private CarPolicyDO iDataObject;
    };

```

This section is focused on business logic. The entire implementation interface is shown in the previous example. How the non-business logic or MOFW framework methods are to be handled is shown in upcoming sections.

The getters and the setters are implemented as they would be in any business object. The only difference in the other methods is that they can choose to access or use state data from the parent class. The following method implementation shows utilization of state data from both the Policy and CarPolicy class.

```

    public long riskQuotient ()
    {
        if ( iGlassCoverage() )
        {
            if ( iYear() > 1960)
                return 1;
            else if ( policyNo() < 1000)
                return 2;
            else
                return 5;
        }
        else if ( amount() > 1000 )
            return 10;
        else
            return 100;
    }

```


In the previous code segment, the accessor or getter methods are used to access the state data that is needed from Policy. This is the most encapsulated way of doing this. However, if the members are declared as protected instead of private in the `_PolicyBOBase` class, then direct access would be possible.

The previous example shows the caching data object case. In the delegating case, the same `CarPolicyDO` would be used to access all of the state data regardless of whether or not it resides in the `CarPolicy` or the `Policy`.

Meet the MOFW IManageable requirements

The following methods are required to be overridden in the simple case:

- `getPrimaryKeyString()` method
- `getHandleString()` method
- `externalize_to_stream()` method
- `internalize_from_stream()` method

Note: `getHandleString()` is not actually required to be overridden in the simple case, unless this object is going to be the target of a one-to-one relationship from another object. The next section discusses considerations for the inheritance case.

getPrimaryKeyString() method

This method implementation in the inheritance case depends on the decision made about the key to be used for the new subclass. That is, what will the key be for the `CarPolicy`? If a new key class is introduced, then this method must be overridden. If the existing key class can be used, that is, use the `PolicyKey` for the `CarPolicy` class, then this method does not need to be overridden and can be inherited directly.

More key classes for help in determining if another key class is needed.

getHandleString() method

A default implementation of this method is provided by the managed object framework. Overriding this method in a subclass of a business object should be done based on the same criteria that are used to determine if an override is needed in the base class.

externalize_to_stream() method

This method needs to be implemented. The general direction is to call the parent class method and add those things which are necessary for the subclass.

```

public void externalize_to_stream(org.omg.CosStream.StreamIO s)
{
    super.externalize_to_stream(s);                // See note
    s.write_long(iDataObject.year());
    s.write_string(iDataObject.make());
    s.write_string(iDataObject.model());
    s.write_long(iDataObject.serialNumber());
    s.write_float(iDataObject.collisionDeductible());
    s.write_boolean(iDataObject.glassCoverage);
}

```

A note about the example

Here is an example where a method implementation that extends rather than replaces the PolicyBO method is supplied. To get this effect, the base class (also known as the superclass) method is invoked using Java's super keyword, which exists for this purpose. In most cases it is appropriate to call the super method first, but sometimes it makes more sense to do it last.

While the previous example is for a caching business object, the pattern for a delegating data object is similar.

internalize_from_stream() method

This method needs to be implemented. The general direction is to call the parent class method and then add those things which are necessary for the subclass.

```

public void internalize_from_stream(org.omg.CosStream.StreamIO s,
                                   org.omg.CosLifeCycle.FactoryFinder ff)
    throws org.omg.CosLifeCycle.NoFactory,
           org.omg.CosStream.ObjectCreationError,
           org.omg.CosStream.StreamDataFormatError
{
    super.internalize_from_stream(s, ff);
    iDataObject.year(s.read_long());
    iDataObject.make(s.read_string());
    iDataObject.model(s.read_string());
    iDataObject.serialNumber(s.read_long());
    iDataObject.collisionDeductible(s.read_float());
    iDataObject.glassCoverage(s.read_boolean());
}

```

While the previous example is for a caching business object, the pattern for a delegating data object is similar.

MOFW requirements - IManagedServer

The following methods must be overridden in the simple case:

- "initForCreation() method" on page 313
- "uninitForDestruction() method" on page 313

- “initForReactivation() method”
- “uninitForPassivation() method” on page 314
- “syncFromDataObject() method” on page 314
- “syncToDataObject() method” on page 314

The next section discusses considerations for the inheritance case.

initForCreation() method

Code this method following the same guidelines that were specified for building business objects independent of this inheritance case. If the subclass introduces additional state data, as the CarPolicy example does, then the data object must be set into a data member of the object. If the subclass does not introduce additional state data, it does not need to save a pointer to the data object that is passed as a parameter. In the delegating case, the subclass might still want to hold a pointer to the data object rather than going through parent class get and set methods. However, regardless of whether the subclass introduces additional state or not, the parent class’ initForCreation() method must be called at the beginning of the subclass’ initForCreation() method.

```
public void initForCreation(com.ibm.IManagedServer.IDataObject inputDO)
    throws com.ibm.IManagedServer.ICreationFailed
{
    super.initForCreation(inputDO);
    IDataObject = CarPolicyDOHelper.narrow(inputDO);
    initializeState();
}
```

uninitForDestruction() method

Implement this method following the guidelines described previously. It is also a good practice to call the parent class’ uninitForDestruction() at the beginning of the subclass’ uninitForDestruction() method if needed.

initForReactivation() method

Code this method following the same guidelines that were specified for building business objects independent of this inheritance case. If the subclass introduces additional state data, as the CarPolicy example does, then the data object must be set into a data member of the object.

If the subclass does not introduce additional state data, it does not need to save a pointer to the data object which is passed as a parameter. In the delegating case, the subclass might still want to hold a pointer to the data object rather than going through parent class get and set methods. However, regardless of whether the subclass introduces additional state or not, the

parent class' `initForReactivation()` method must be called at the beginning of the subclass' `initForReactivation()` method.

```
public void initForReactivation(com.ibm.IManagedServer.IDataObject inputDO)
    throws com.ibm.IManagedServer.IReactivationFailed
{
    super.initForReactivation(inputDO)
    iDataObject = CarPolicyDOHelper.narrow(inputDO);
}
```

uninitForPassivation() method

This method should be implemented following the normal guidelines described earlier. It is also good practice to call the parent class' `uninitForPassivation()` at the beginning of the subclass' `uninitForPassivation()` method if needed.

syncFromDataObject() method

The pattern for implementing this method is similar to that used for the `internalize_from_stream()` method. The parent method must be called first, followed by the code necessary to prime the subclassing business object cache with values from the data object.

This method is also a good place for any initialization logic that is dependent on the presence of an active and usable data object. The act of loading up the business object cache ensures that the state data of the object is ready to be used.

syncToDataObject() method

The pattern for implementing this method is similar to that used for the `externalize_to_stream()` method. The parent method must be called first, followed by the code necessary to push the subclassing business object cache values back into the data object.

It is a good practice to always implement all of these `IManagedServer` methods even if they are not explicitly required based on the particular subclass being introduced. This practice results in consistently generated code which is less prone to error if future changes are made to the subclass that introduces new attributes.

More key classes

The MOFW requires that every managed object have a primary key class associated with it. When extending a business object, there are several possibilities you can use as a primary key class. The business object subclass can:

- Use its base class' primary key class.

- Extend its base class' primary key class.
- Introduce its own primary key class.

The simplest approach is to reuse the existing key. This approach is applicable if the attributes that uniquely identify objects of the subclass are the same ones that identify objects of the base class. This is most likely the case if the subclass introduces no additional state (that is, the subclass only re-implements the base class methods, or introduces new methods). Even if the subclass introduces an additional state beyond that of the base class, this additional state might not contribute anything to object identity.

If the subclass introduces additional state, some of which, combined with the key attributes of the base class, is used to uniquely identify objects of the subclass, then the best approach is to extend the base class' primary key class. When extending the primary key class of a base class, interface and implementation inheritance can be used. Implementation inheritance allows for the reuse of the base class key functionality (specifically, its getters and setters, as well as its streaming code).

Finally, if the attributes that uniquely identify objects of the subclass are neither the same ones as the base class, nor a superset of the base class' key attributes, then the subclass must introduce its own primary key class.

However, if the existence of new state data has altered the way in which the object is uniquely identified, then a new primary key class is necessary.

Note: If the subclass does not use the base class' primary key class, then the subclass' data object needs to be able to handle this extended or new key class.

More copy helper classes

If the subclassing business object introduces additional state data, then a new copy helper class might be useful. The data object needs to be able to handle this new copy helper. The copy helper can inherit interface and implementation from the base class' copy helper, or it can be written from scratch.

Extension summary

In the managed object framework based programming model, there are a number of inheritance activities to follow. Interfaces should be inherited consistently. Implementations should be inherited when the base class' implementation can be reused to some degree. The simplest model is to inherit at all levels of the MOFW architecture and add in additional business logic as necessary. Additional requirements from the MOFW are also added in the new implementation subclass.

Managed object customization and data object customization are also different for business objects that inherit from other business objects. These topics are discussed in “Data object customization and inheritance” on page 389.

Other variations to consider

There are other variations to consider. Some are restrictions and others are tips for leveraging inheritance in the MOFW-based programming environment:

- Do not change data object patterns. It is possible to change data object patterns from caching to delegating at various levels of the data object hierarchy but this adds undue complexity in most cases.
- There might be cases when the subclasser does not know the data object pattern being used by the base class.

Object relationships

Component Broker applications often require persistent relationships between business objects. For instance, in the personal life insurance sample application, a Claim object has a relationship with a Policy object representing the insurance policy against which the claim has been filed. Also, a Policy object has a relationship with Claim objects for pending claims against the insurance policy.

Relationships between objects can be described in many ways. First, there is the cardinality of the relationship. If an object has a relationship to one other object at most, the relationship is considered to be cardinality-1. On the other hand, if an object has a relationship with more than one other object at a time, the relationship is considered to be cardinality-N. In a business object, relationships to other objects are implemented as object references (cardinality-1), or as collections of object references (cardinality-N).

A relationship can also be described as *optional* or *required*. If a relationship is optional, then an object is considered to be in a valid state even when it is not related to (linked to) another object. If the relationship is required, then the object must always be linked to another object. This distinction is often combined with cardinality to form the following combinations:

Class Relationship	Cardinality	Required
An instance of class X is related to 0..1 instances of class Y	-1	No
An instance of class X is related to 1 instance of class Y	-1	Yes
An instance of class X is related to 0..n instances of class Y	-N	No
An instance of class X is related to 1..n instances of class Y	-N	Yes

Finally, a relationship can also be described in terms of ownership. An object which is related to another object might or might not be considered to be the owner of the related object. If the first object is not considered to be the owner of the second object, then the relationship is often referred to as a *uses a* relationship, as in, for example, the first object *uses* the second object. This is sometimes also called an association. On the other hand, if the first object is considered to be the owner of the second object, then the relationship is often referred to as a *has a* relationship, as in, for example, the first object *has* (or contains) the second object. This is sometimes also called an aggregation.

The cardinality-1 or cardinality-N distinction results in much more fundamental differences in the business object code than the optional or required distinction and the uses or has distinction. Cardinality-1 relationships are discussed separately from Cardinality-N relationships.

Cardinality-1 relationships

The following diagram shows an example of a cardinality-1 relationship, a simple link from a Claim object to a Policy object.



Figure 54. Cardinality-1 relationship

The relationship depicted in the previous figure is that of an optional cardinality-1 *uses a* relationship. In a business object interface, a cardinality-1 relationship is declared as a CORBA attribute whose type is a reference to the interface of the target object. For example, an attribute called *thePolicy* on a Claim business object could be used to link to the Policy object as follows:

```
interface Claim : ...
{
    attribute Policy thePolicy;
    ...
};
```

Clients can establish and traverse the link using the attribute set and get methods respectively:

```
Policy aPolicy = ... // Find or create a Policy object
Claim aClaim = ... // Find or create a Claim object

// Set the claim's policy to "aPolicy".
aClaim.thePolicy(aPolicy);

// Get the claim's policy as "somePolicy".
Policy somePolicy = aClaim.thePolicy();
```

Using the Component Broker delegating pattern, the business object implementation of the access methods for thePolicy passes the object pointer to or from the data object:

```
public void thePolicy(Policy policy)
{
    fDataObject.thePolicy(policy);
}

public Policy thePolicy()
{
    return fDataObject.thePolicy();
}
```

Because the relationship is optional it is possible that the data object will return either a valid pointer to a policy object or a null pointer.

To speed up the performance of link access, caching the Policy pointer might be appropriate. The Claim business object implementation class could cache a pointer to the Policy object as shown in the following example:

```
class ClaimBO_Impl ...
{
    public:
        ...
    private:
        ...
        Policy fCachedPolicy;
        ...
}
```

The access methods for thePolicy would now use the cached pointer:

```
public void thePolicy(Policy policy)
{
    fCachedPolicy = policy;
}

public Policy thePolicy()
{
    return fCachedPolicy;
}
```

The general Component Broker business object caching pattern uses the methods syncFromDataObject() and syncToDataObject() respectively to load and flush the cached values. See the description of syncFromDataObject() method and syncToDataObject() method. The implementation of the syncFromDataObject() method calls data object get methods to retrieve the thePolicy pointer as well as all other data contained in the Claim object. The implementation of the syncToDataObject() method calls all the data object set methods. These methods appear as follows:


```

public void syncToDataObject()
{
    ...
    fDataObject.thePolicy( fCachedPolicy);
    ...
}

public void syncFromDataObject()
{
    ...
    fCachedPolicy = fDataObject.thePolicy();
    ...
}

```

Because links can be expensive to compute, and sometimes are not needed, a better pattern for caching object references is to leave them out of the syncTo/FromDataObject methods and instead compute and cache them in the get method on first use. This lazy evaluation approach can be implemented using the following pattern:

```

public void thePolicy(Policy policy)
{
    fCachedPolicy = policy;

    // Synchronize the new Policy in the BO with (to) the DO
    fDataObject.thePolicy(fCachedPolicy);
}

public Policy thePolicy()
{
    // If this is the first access of the Policy, get it from the DO
    if ( fCachedPolicy == null )
        fCachedPolicy = fDataObject.thePolicy();
    return fCachedPolicy;
}

```

Assuming fCachedPolicy is initialized to nil in the constructor, this pattern results in the data object get method being called the first time the Policy is accessed, or not at all if the Policy is set in the same session as the get call. Subsequent accesses return the cached pointer value.

Optional or required cardinality-1 relationships

The previous section discusses how an optional relationship is implemented in a business object. An optional relationship is more flexible than a required relationship, and thus requires less code. However, if the relationship is required, the following diagram represents a required cardinality-1 *uses a* relationship.



Figure 55. Required cardinality-1 "uses a" relationship

If the relationship is truly required, then the link from a Claim to its Policy must exist throughout the life cycle of the Claim. Specifically, a Policy must be linked to a Claim as part of creating a Claim, and the link must be broken when a Claim is removed. The Policy that must be linked to a Claim as part of creation might exist prior to creating the Claim, or it might be created in the process of creating the Claim.

When reviewing a Claim, the link to its Policy must be broken, but the Policy is not necessarily removed. If the claim *uses a* (knows about a) Policy, then it is sufficient to release the Policy when a Claim is removed. However, if the Claim to Policy relationship is one of ownership (that is, a *has a* relationship), then when a Claim is removed, its associated Policy must also be removed.

When a Claim is removed, the link to its Policy must be broken. Changes between an optional and a required relationship depend on one of four different scenarios (discussed over the next several pages). However, there is one change which is common to all four. Under any scenario, if the relationship is required instead of optional, the code for the set method for the *Policy* attribute must make sure that the link is not broken by setting the Policy reference to nil. The following example shows the changes necessary in the case of the delegating pattern. The changes for the caching pattern are analogous.

```

public void thePolicy(Policy policy)
{
    if ( policy != null )
    {
        fDataObject.thePolicy(policy);
    }
}

```

To understand how a link from a Claim to its Policy gets established as part of creating a Claim, it is necessary to recall that there are several ways to create a business object:

- Using a generic home configured for the appropriate type of business object. See *Creating a claim with an existing policy using a generic home* and *Creating a claim with a new policy using a generic home*.
- Using a specialized home developed by the business object provider. See *Creating a claim with an existing policy using a specialized home* and *Creating a claim with a new policy using a specialized home*.

Creating a claim with an existing policy using a generic home: Using a generic home, a business object is created using the `createFromPrimaryKeyString()` method. If an object provider has provided a copy helper class, then a business object can also be created using the `createFromCopyString()` method. However, this has no effect on this discussion because a copy helper class' attributes are defined as being a superset of the primary key class' attributes. The only input to this method is a stringified version of the business object's primary key helper object. This means that the primary key helper class for Claim must contain enough information in it such that the link to its (preexisting) Policy can be established. There are many ways in which this can be done, but the two most straightforward ways are the following:

- The primary key for Claim contains a reference to the Policy object. See "Primary key contains reference to related object".
- The primary key for Claim contains all of the Policy object's primary key attributes. See "Primary key contains key attributes of related object" on page 322.

Primary key contains reference to related object: In this case, the interface of the primary key for Claim would look like the following example:

```
interface ClaimKey : IManagedLocal::IPrimaryKey
{
    attribute long claimNo; // key attribute for Claim
    attribute Policy thePolicy; // associated Policy
}
```

The attribute for the associated Policy object would be implemented as follows:

```
public void thePolicy(Policy thePolicy)
{
    fPolicy = thePolicy;
}

public Policy thePolicy()
{
    return fPolicy;
}
```

For a primary key to flow over the wire from the client to the server, the primary key must be able to externalize and internalize all of its attributes, including those which are necessary for establishing relationships to other objects. In the case of the primary key class for the Claim object, it must externalize and internalize a stringified object reference to the Policy object as follows:

```
public void externalize_to_stream(
    org.omg.CosStream.StreamIO targetStreamIO)
{
    // Insert Method modifications here
}
```

```

        targetStreamIO.write_long(fClaimNo);
        String policyRefString = CBSeriesGlobal.orb().object_to_string(fPolicy);
        targetStreamIO.write_string(policyRefString);
        // End Method modifications here
    }
    public void internalize_from_stream(
        org.omg.CosStream.StreamIO targetStreamIO,
        org.omg.CosLifecycle.FactoryFinder ff)
        throws      org.omg.CosLifecycle.NoFactory,
                   org.omg.CosStream.ObjectCreationError,
                   org.omg.CosStream.StreamDataFormatError
    {
        // Insert Method modifications here
        fClaimNo = sourceStreamIO.read_long();
        String policyRefString = sourceStreamIO.read_string();
        org.omg.CORBA.Object policyRef =
            CBSeriesGlobal.orb().string_to_object(policyRefString);
        fPolicy = PolicyHelper.narrow(policyRef);
        targetStreamIO.write_string(policyRefString);
        // End Method modifications here
    }
}

```

The streaming code in the previous example allows the primary key for Claim to flow from the client application to the Claim home on the server. When the primary key reaches the Claim home on the server, it eventually gets passed to the data object implementation for Claim. The link has been established by invoking the `string_to_object()` method on the ORB while internalizing the Claim primary key from its stream.

Primary key contains key attributes of related object: In the second case, the interface of the primary key for Claim would look like the following example:

```

interface ClaimKey : IManagedLocal::IPrimaryKey
{
    attribute long claimNo; // key attribute for Claim
    attribute long policyNo; // key attribute for Policy
}

```

The attribute for the associated Policy object would be implemented as follows:

```

public void policyNo(long policyNo)
{
    fPolicyNo = policyNo;
}

public long policyNo()
{
    return fPolicyNo;
}

```

Again, in order for a primary key to go from the client to the server, the primary key must be able to externalize and internalize all of its attributes, including those which are necessary for establishing relationships to other objects. In this case, the primary key class for the Claim object must externalize and internalize the key attributes of its related Policy object as follows:

```
public void externalize_to_stream(
    org.omg.CosStream.StreamIO targetStreamIO)
{
    // Insert Method modifications here
    targetStreamIO.write_long(fClaimNo);
    targetStreamIO.write_long(fPolicyNo);
    // End Method modifications here
}

public void internalize_from_stream(
    org.omg.CosStream.StreamIO targetStreamIO,
    org.omg.CosLifeCycle.FactoryFinder ff)
    throws org.omg.CosLifeCycle.NoFactory,
           org.omg.CosStream.ObjectCreationError,
           org.omg.CosStream.StreamDataFormatError
{
    // Insert Method modifications here
    fClaimNo = sourceStreamIO.read_long();
    fPolicyNo = sourceStreamIO.read_long();
    // End Method modifications here
}
```

As with the previous scenario, the streaming code in this example allows the primary key for Claim to flow from the client application to the Claim home on the server. However, that is the end of the similarity. The previous scenario shows that the cardinality-1 relationship is established as part of internalizing the Claim primary key inside the Claim home on the server.

In this scenario, because the Claim primary key class contains the key attributes of the related Policy object, the relationship must be established in the Claim data object implementation. The Claim data object implementation would do so by extracting the Policy key attributes from the Claim primary key, finding a home of Policy objects, and invoking the `findByPrimaryKeyString()` method on the home.

Creating a claim with a new policy using a generic home: In the previous scenario, the application domain specified a constraint that a Claim can only be created for an existing Policy. It is possible that the application domain, while requiring that a relationship between two objects exist, does not require one object to already be created when creating the second. While an insurance company which allows a Policy to be created at the same time as a Claim would probably not stay in business long. That scenario is used here for consistency.

Like the previous scenario, this scenario assumes that the Claim object provider has chosen not to develop a specialized home and instead expects clients of the Claim business object to use a generic home. Because the only input to the `createFromPrimaryKeyString()` method of a generic home is a stringified version of the business object's primary key helper object, this scenario could be implemented in a similar fashion to the previous scenario, where the Claim primary key class contains the key attributes of a Policy. However, the one difference is what happens when the Claim primary key containing information about a Policy reaches the Claim home on the server.

In the previous scenario, where the Claim is being created with an existing Policy, there is no need for the Claim data object implementation to differentiate between when the business object is being created for the first time, and when it is being reactivated after having been previously passivated. In other words, in either case, the Claim data object implementation uses the Policy key attributes from its own key to find the existing Policy object.

In this scenario, however, the Policy is being created during the creation of a Claim. As such, the Claim data object implementation must distinguish between the creation and reactivation of a Claim as follows:

- If a Claim is being created, then a new Policy object must be created.
- If a Claim is being reactivated, then the previously created Policy object must be found.

A data object does not know if the business object is being created or reactivated. However, the business object itself does know if it is being created or reactivated. If the business object is being created the home invokes the `initForCreation()` method on it; if the business object is being reactivated the home invokes the `initForReactivation()` method. In its implementation of the `initForCreation()` method, the Claim business object would do the following:

- Use a factory finder to find a Policy home
- Use the Policy key attributes from the Claim data object to create a new Policy by invoking the `createFromPrimaryKeyString()` method on the Policy home.

In its implementation of the `initForReactivation()` method, the Claim business object would do the following:

- Use a factory finder to find a Policy home.
- Use the Policy key attributes from the Claim data object to find its related Policy by invoking the `findByPrimaryKeyString()` method on the Policy home.

There is an alternate way to implement this scenario. Assuming that the client of a Claim is not required to provide the identity of its Policy, then the

primary key class for Claim does not need to contain any information about the Policy. The primary key class for Claim would look like the following example:

```
interface ClaimKey : IManagedLocal::IPrimaryKey
{
    attribute long claimNo; // key attribute for Claim
}
```

In this case, it is the responsibility of the Claim business object, not the data object, to create the Policy during its own creation. A data object has no way to distinguish between object creation and object reactivation. However, a business object knows it is being created when its home invokes the `initForCreation()` method on it. Thus, the implementation of `initForCreation()` would need to somehow create a Policy object. Because the Claim business object was not provided with any information on how to identify the Policy object, it would probably do one of the following:

- Create a primary key for Policy based on some combination (or function) of its own attributes.
- Create a primary key for Policy based on some random number generator or UUID (Universal Unique Identifier).
- Use a specialized home for Policy (if one was provided).

Because the relationship of Claim to Policy is required, if the Claim business object is unable to create a Policy for some reason, then it should cause its own creation to fail. The following example shows how this would be done:

```
public void initForCreation(com.ibm.IManagedServer.IDataObject inputDO)
    throws com.ibm.IManagedServer.ICreationFailed
{
    // Save the data object for later use
    fDataObject = ClaimDOHelper._narrow(theDO);

    // to create the Claim's Policy somehow
    try
    {
        fCachedPolicy = ...
    }
    catch(java.lang.Exception exception)
    {
        throw new com.ibm.IManagedServer.ICreationFailed();
    }

    // Other initialization...
}
```

Creating a claim with an existing policy using a specialized home: Using a specialized home, it is even easier to establish a link from a Claim to its Policy as part of creating the Claim. With a specialized home, there is no need to add any information about the Policy to the primary key for Claim. Adding information about one object to the primary key of another in the previous

scenarios was done for the sole purpose of establishing the required relationship. The primary key was being used not just to establish unique identity, but also as a vehicle for communicating information about the relationship from the client to the home. This is not what primary key classes were intended for, but was done out of necessity given that a specialized home was not provided.

Unfortunately, doing so introduced the following unwanted side-effect. Using the generic home configured for Claim objects at some later time to find a previously created Claim object required the client to have some information about the Policy too (at least some of the key attributes, if not a reference to the Policy object itself, depending on the scenario). This might not be an acceptable constraint in the application domain. If not, then the solution is to have the Claim object provider provide a specialized home as well.

Using a specialized home, the link between a Claim and its Policy is established in the specialized home itself, as opposed to in the Claim business object or data object. In order to support the required cardinality-1 relationship to a Policy, the specialized home for Claim would introduce new methods for creating a Claim. These new methods would have parameters which would allow the specialized home to establish the link between a Claim and its Policy. The following shows two examples of such create methods:

```
public Claim createClaimWithPolicyRef(long claimNo, Policy policy)
    throws ClaimHome.MissingPolicy
{
    // Before we get too far, let's make sure that we have a valid
    // Policy reference for the required cardinality-1 relationship.

    if ( policy == null )
        throw new ClaimHome.MissingPolicy();
    ClaimKey claimKey = ClaimKeyHelper._create();
    claimKey.claimNo(claimNo);
    byte[] claimKeyString = claimKey._toString();

    // For an inheriting specialized home, pass the claimKeyString
    // to the parent of the specialized home on createFromPrimaryKey-
    // String(). For a delegating specialized home, pass the claim-
    // KeyString to the contained home on createFromPrimaryKeyString().

    Claim newClaim = ...

    // Now that the Claim has been created, establish the link to its
    // Policy using the Policy which was passed in on createClaim().
    newClaim.thePolicy(policy);
    return newClaim;
}
```

Creating a claim with a new policy using a specialized home: This scenario is similar to the previous scenario. However, because in this scenario the

Policy is not created prior to creating a Claim, the `createClaimWithPolicyRef()` method from the previous scenario is not applicable and the implementation of the `createClaimWithPolicyNo()` method is different. Instead of invoking the `findByPrimaryKeyString()` method on the Policy home it must invoke the `createFromPrimaryKeyString()` method as illustrated in the following example:

```
public Claim createClaimWithPolicyNo(long claimNo, long policyNo)
{
    // Assume that in the ClaimHomeB0 initForCreation()
    // the specialized home for Claim objects has used a factory
    // finder to find the home for Policy objects, and saved
    // its reference in 'fPolicyHome'.

    PolicyKey policyKey = PolicyKeyHelper._create();
    policyKey.policyNo(policyNo);
    byte[] policyKeyString = policyKey._toString();

    Policy aPolicy; // Assume fPolicyHome is a specialized home...
    aPolicy = fPolicyHome.createFromPrimaryKeyString(policyKeyString);

    ClaimKey claimKey = ClaimKeyHelper._create();
    claimKey.claimNo(claimNo);
    byte[] claimKeyString = claimKey._toString();

    // For an inheriting specialized home, pass the claimKeyString
    // to the parent of the specialized home on creatFromPrimary-
    // KeyString(). For a delegating specialized home, pass the claim-
    // KeyString to the contained home on createFromPrimaryKeyString().
    Claim newClaim = ...

    // Now that the Claim has been created, establish the link to its
    // Policy using the Policy which was passed in on createClaim().
    newClaim.thePolicy(policy);

    return newClaim;
}
```

"Uses a" and "Has a" cardinality-1 relationships

Now that the differences between optional and required cardinality-1 relationships have been described, especially as they pertain to developing a business object, the *uses a* and *has a* cardinality-1 relationships are described.

Previous sections discuss how *uses a* relationships are represented. The following figure illustrates an optional cardinality-1 *has a* relationship.




Figure 56. Optional cardinality-1 "has a" relationship

The following diagram, on the other hand, illustrates a required cardinality-1 *has a* relationship.




Figure 57. Required cardinality-1 "has a" relationship

A “uses a” relationship means that one object has a reference to another object. In the example, a Claim object has a reference to a Policy object. In fact, there might be more than one Claim object with a reference to the same Policy object and there might also be objects of other types with references to the same Policy object. In any case, any object with a reference to this Policy object is capable of invoking any method on the Policy object’s client interface. This includes the `remove()` method. Of course, things could get chaotic if any object with a “uses a” relationship to the same Policy object were allowed to invoke the `remove()` method on it. To avoid such confusion, Component Broker recommends that an object be removed only by its owner, and that an object be owned only by a single object. All other objects with references to that object should only release the reference.

A “has a” relationship is what is used to represent the concept of ownership. In other words, an object with a “has a” relationship to another object is said to own that object. In the previous two diagrams, a Claim object “has a” reference to a Policy object. Although the relationship would most likely be reversed in the insurance application domain, this scenario continues with the original example instead of introducing two new objects. Thus, a Claim is considered to be the owner of a Policy and as such is responsible for removing the Policy should this become necessary. One case in which it becomes necessary for a Claim to remove the Policy is when the Claim itself is being removed. A business object knows it is being removed when its home invokes the `uninitForDestruction()` method on it. The following example shows how the Claim business object would implement this:

```

public void uninitForDestruction()
{
    // Remove the (owned) Policy object
    fCachedPolicy.remove();
    // Other un-initialization...
}

```

Another case in which it becomes necessary for a Claim to remove the Policy is if the Claim interface has a method that gives its clients the opportunity to request that the Policy be removed. This is similar to the `cancelPolicy()` method in the following example:

```

public void cancelPolicy()
{
    // Remove the (owned) Policy object
    fCachedPolicy.remove();
    // Other clean-up associated with the Policy cancellation
}

```

Of course, the `cancelPolicy()` method makes sense only in an optional *has a* relationship. In any case, if one business object is responsible for the removal of another business object as the result of a *has a* relationship, then the reverse must not also be the case. In other words, in Component Broker, two objects might not have a *has a* relationship with one another. If such a bidirectional relationship is required, then one object must release the other, while the other object removes the first one as described in this section.

Making cardinality-1 relationships persistent

Because object references are really just memory addresses, they cannot be made persistent in that form. Therefore, object references must be converted to other forms that can be made persistent. Because persistence of all attributes, not just object relationships, is implemented in the data object, not the business object, this conversion is described further in “Data object customization for cardinality relations” on page 395.

As previously mentioned, in the Component Broker architecture, a business object’s persistent state is managed by its associated data object. Links (1-1 and 1-N object relationships) are no different. The persistent representation of a link is also typically managed by a data object. However, there might be cases where it makes more sense for the business object to manage the link itself. For instance, if the link can be computed based on some application domain business logic, then it makes sense not to burden the data object with managing the link.

For example, assume that a claim maintains a cardinality-1 link to the insurance policy with which it is associated. Assume as well that the design of the application is such that the primary key of a Claim (the claim number) includes the policy number of its associated Policy, for example, claim

numbers are actually prefixed by their associated policy number. In this situation, implementing the Policy reference involves no additional persistent data. The Policy get method could be implemented directly in the business object as follows:

```
Policy thePolicy()
{
    byte[] pkString;
    // get my primary key

    ClaimPrimarykey myKey = ClaimPrimarykeyHelper._create();
    myKey.fromString(pkString = getPrimarykeyString());

    // extract the policy number from my primary key
    long myPolicyNo = myKey.policyNo();

    // get the policy home somehow (for example, using a factory finder)
    IHome policyHome = ...

    // create and initialize a policy primary key
    PolicyPrimarykey policyKey = PolicyPrimarykeyHelper._create();
    policyKey.policyNo(myPolicyNo);

    // get the policy from the policy home
    com.ibm.IManagedClient.IManageable mo;
    return PolicyHelper.narrow(mo = policyHome.findByPrimarykey(policyKey));
}
```

As shown, this algorithm does not involve an explicit representation of the converted pointer. It is extracted from the key and is based on the application-level knowledge that claim numbers are prefixed by their associated policy number. In this situation the persistent representation of the link is actually maintained as part of another one of the object's attributes. Also, when a link can be computed based on some business logic, the object reference is often considered to be a read-only attribute of the business object.

While it is possible for a business object to manage object references itself, more commonly the business object passes the pointer to the data object and it performs the required conversion. With this approach, every data object that is used with a particular business object is free to use whatever conversion algorithm (and persistent representation) it chooses. This approach allows the business object to remain de-coupled from the persistent storage mechanism and thus possibly be reused in different scenarios with a wide range of back-end data stores.

Cardinality-N relationships

The previous section shows how to link an insurance Policy to a single Claim object. If, instead, the Policy object needs to reference multiple Claims, a cardinality-N relationship is required.




Figure 58. Cardinality-N relationship

There are two approaches for providing a type-safe implementation for a cardinality-N relationship in Component Broker. Both approaches rely on a persistent collection object to manage the references, although one hides the collection in the implementation, while the other exposes it to clients as a first class object.

With the first approach, where the collection is hidden, the interface for adding and removing elements is provided on the Policy object itself. With this approach, the Policy business object interface would provide a set of methods for establishing, deleting, and accessing the links to Claim objects. The interface could be defined as follows:

```

interface Policy : ...
{
    void addClaim(in Claim claim);
    void removeClaim(in Claim claim);
    com.ibm.ICollectionsBase.IIterator listClaims();
    ...
};

```

Clients would then use this interface to add, access, and remove claims, as follows:

```

Policy aPolicy = ...
Claim aClaim1 = ...
Claim aClaim2 = ...

// Add a couple of claims to the policy's set of claim

aPolicy.addClaim(aClaim1);
aPolicy.addClaim(aClaim2);

// Iterate through the policy's set of claims
// Get an iterator for the set of claims
com.ibm.ICollectionsBase.IIterator iter = aPolicy.listClaims();
com.ibm.IManagedClient.IManageable element;

while (element = iter.next())
{
    // iterate through the set of claims
    Claim aClaim = ClaimHelper.narrow(element);
    // do something with (for example, process) "aClaim"
    ...
}

```

```

}

// Remove a claim from the policy's set of claims
aPolicy.removeClaim(aClaim1);

```

As shown, although the underlying implementation of the relationship methods might use a collection object to manage the references, the client program is completely shielded from this fact. The Policy object encapsulates the collection behind the type-safe relationship interface methods, `addClaim()`, `removeClaim()`, and `listClaims()`.

The second approach for implementing the cardinality-N relationship exposes the collection in the client programming model. Instead of referencing multiple Claim objects directly, the Policy object references a single collection object using a cardinality-1 link. The collection, in turn, references the multiple claims.




Figure 59. Cardinality-1 link to a collection object that references claims

This picture actually applies to the previous approach as well, although there the reference collection is hidden in the implementation. One difference is that when the collection is hidden, the Policy interface is the only client interface for adding elements to the collection and therefore it provides static type checking for the elements in the collection. If, on the other hand, the collection is visible to clients, the collection itself must prevent clients from adding objects other than Claims to the relationship.

When using the explicit collection approach for implementing a cardinality-N relationship, the Policy business object interface includes an attribute whose type is a reference to the collection.

```

interface Policy : ...
{
    readonly attribute com.ibm.IManagedCollection.IReferenceCollection claims;
    ...
};

```

The relationship of the collection object to the Policy object can be considered to be a standard cardinality-1 link as described in Cardinality-1 relationships, with one exception, the attribute is read-only. The *readonly* attribute allows clients to add and remove elements in the collection but prevents them from replacing the collection itself. In some situations this latter operation might be

warranted, in which case, the *readonly* tag would be removed. In general, however, you should prevent clients from modifying the reference to the collection.

A client program uses the *claims* attribute access method to access the collection of claims. Methods on the collection can then be called to manipulate and access the actual Claim references.

```
Policy aPolicy = ...
Claim aClaim1 = ...
Claim aClaim2 = ...

// Get the set of claims;
com.ibm.IManagedCollections.IReferenceCollection claims =
    aPolicy.claims();

// Add a couple of claims to the policy's set of claims
claims.addElement(aClaim1);
claims.addElement(aClaim2);

// Iterate through the policy's set of claims

com.ibm.ICollectionsBase.IIterator iter = claims.createIterator();
// get a claims iterator
com.ibm.IManagedClient.IManageable element;

while (element = iter.next())
{
    // iterate through the set of claims
    Claim aClaim = ClaimHelper.narrow(element);
    // do something with (for example, process) "aClaim"
    ...
}

// Remove a claim ("aClaim1") from the policy's set of claims
claims.removeElement(aClaim1);
}
```

Implementing the relationship interface

Depending on the interface requirements and possibly the back-end datastore being used, many different implementations of the Policy/Claim relationship are possible. Some common approaches include implementing the relationship using:

- A simple Non-keyed reference collection. See “Implementing a relationship with a simple reference collection” on page 334.
- A Keyed reference collection. See “Implementing a relationship with a keyed reference collection” on page 335.
- A Home or Collection and some additional information that is used to identify a subset of the entries in the collection. Some examples include a non-unique secondary key supported by the Home or Collection and a

query evaluation string if the Home or Collection is queryable.
See “Subsetting a home or collection” on page 336.

The following sections describe each of these implementation approaches.

Implementing a relationship with a simple reference collection: The relationship collection is exposed in the data object interface as a readonly attribute of type IReferenceCollection:

```
interface PolicyDO : ...
{
    readonly attribute com.ibm.IManagedCollection.
        IReferenceCollection claims;
    ...
};
```

The relationship interface methods, addClaim(), removeClaim(), and listClaims(), would be implemented as shown in the following example:

```
public void addClaim(Claim claim)
{
    com.ibm.IManagedCollections.IReferenceCollection claims = claims();
    claims.addElement(claim);
}

public void removeClaim(Claim claim)
{
    com.ibm.IManagedCollections.IReferenceCollection claims = claims();
    claims.removeElement(claim);
}

public com.ibm.ICollectionsBase.IIterator listClaims()
{
    com.ibm.IManagedCollections.IReferenceCollection claims = claims();
    return claims.createIterator();
}
```

In each of the methods, the method claims is used to return a pointer to the reference collection containing the claims. Each method then delegates to its corresponding method on the reference collection.

The persistent object reference for the reference collection itself can be implemented using any of the design patterns described in “Cardinality-1 relationships” on page 317. Using the lazy evaluation caching pattern in the business object, the claims method would be implemented as follows:

```
public com.ibm.IManagedCollections.IReferenceCollection claims()
{
    // first time?
    if (fCachedClaims == null)
        fCachedClaims = fDataObject.claims();
    return fCachedClaim;
}
```


Before the data object claims method can return a collection, a reference collection must actually be created. This is typically done in the data object claims method the first time it is called. As previously discussed, the reference collection used to implement a relationship may be required to guarantee that the type of elements added to it are of a specific type (for example, Claims). A type-specific reference collection can be created by calling the `createCollectionFor` method on the `IManagedCollections.ICollectionHome` interface:

```
// get the collection home (for example, using a factory finder)
com.ibm.IManagedCollections.ICollectionHome cHome = ...

// create a reference collection that will only hold claims
com.ibm.IManagedCollections.IReferenceCollection rc =
    cHome.createCollec
```

Alternatively, the simpler `createCollection` method can be used when no restriction on the type of collection element is required:

```
rc = cHome.createCollection();
```

Implementing a relationship with a keyed reference collection: A Keyed reference collection can be used to implement relationships where the individual links are accessed using some kind of identifier. For example, if the claims associated with a given insurance policy are identified by, for example, a claim ID, the relationship interface might appear in the Policy business object as follows:

```
interface Policy : ...
{
    void addClaim(in Claim claim, in long id);
    void removeClaim(in long id);
    Claim getClaim(in long id);
    com.ibm.ICollectionsBase.IIterator listClaims();
    ...
};
```

This relationship is most easily implemented using a keyed collection. For example, the data object interface might now include an attribute of type `IKeyedReferenceCollection`:

```
interface PolicyDO : ...
{
    readonly attribute com.ibm.IManagedCollections.
        IKeyedReferenceCollection claims;
    ...
};
```

The `addClaim` method could be implemented as follows:

```
public void addClaim(Claim claim, long id)
{
    NumberKey key = NumberKeyHelper._create();
    key.setValue(id);
```

```

com.ibm.IManagedCollections.IKeyedReferenceCollection claims =
    claims();
byte[] keyString = key._toString();
claims.addElementByString(claim, keyString);
}

```

In this example, a key helper class (see “Chapter 9. MOFW - Java client programming model” on page 231), `NumberKey`, is used to add the element to the collection. The class `NumberKey` wrappers the ID in a key class that is capable of being stringified. When the key is created and initialized, the add operation is delegated to the reference collection.

The `removeClaim` and `getClaim` methods would be implemented similarly. The remaining method, `listClaims`, as well as the `claims` method that is used to access the collection object, would be implemented as shown in the non-keyed reference collection example.

Subsetting a home or collection: This approach uses a home augmented with information that identifies a subset of the objects in the home. This approach is particularly applicable in bottom-up development scenarios where related data already exists in legacy applications.

For example, an RDB-based insurance application might contain two related tables, one for policies and one for claims. The claims registered against a particular policy are identified by a `policy#` column in the claim table. This column, a foreign key in the claim table, is the primary key for the policy table.

POLICY TABLE			CLAIM TABLE		
Policy#	Owner	...	Claim#	Policy#	...
12	"Joann"		87	100	
34	"John"		65	34	
56	"Katherine"		43	101	
78	"Katherine"		21	34	

Figure 60. RDB-based insurance application tables

In Component Broker object space, these tables represent a cardinality-N relationship between a policy object and its associated claims. For example, policy number 34 would have links to claim numbers 21 and 65.

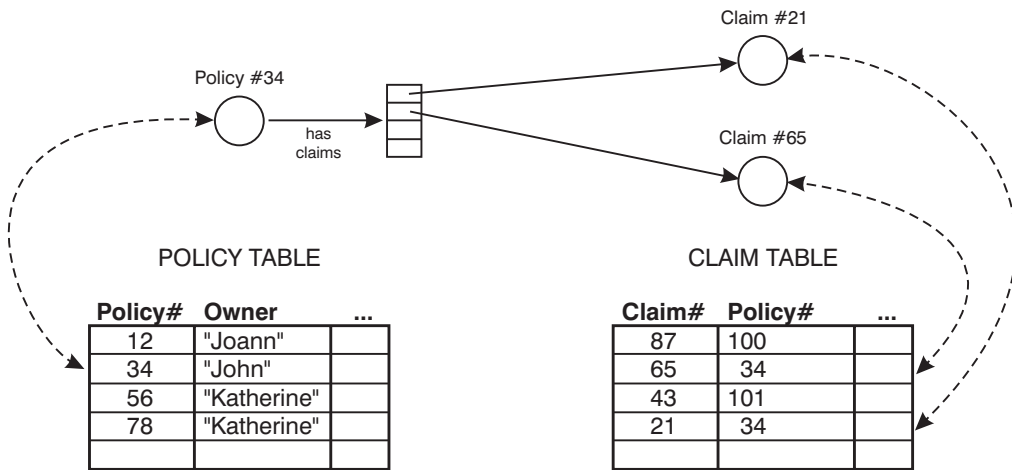


Figure 61. Links between policy and Claims

Normally these links are stored as converted pointers maintained by a persistent reference collection of some type. In this case, the relationship information stored in the claim table itself eliminates the need for a persistent reference collection. The set of claim links for a particular policy can be derived using a query on the claim table. In Component Broker object space, this can be implemented by evaluating an OOSQL query (for example, `policyNo == 34`) on a queryable claim home.

```

com.ibm.IManagedAdvancedClient.IQueryableIterableHome claimHome = ...
com.ibm.ICollectionsBase.IIterator iter = claimHome.evaluate(
    "policyNo=34");

// iterate over the claims
...

```

In this example, the home object would be retrieved using the Naming Service. The home name could be hard coded in the application, available in an environment variable, or stored persistently in the policy table.

The particular pattern used for finding the home might or might not be dictated by legacy requirements.

Because the relationship is derived from other data, adding and removing claims often involve side effects. For example, adding a claim to the reference collection would require a change in state to the claim object (for example, `policy#` field must be updated). If the claim is already in a relationship with another policy, it would be implicitly removed from it as a result of the add operation here. To avoid this kind of change in semantics of the `addClaim` method, the relationship interface could be changed to one that is more semantically consistent with the underlying implementation. For example:

```

interface Policy : ...
{
    Claim createClaim( in long id);
    void deleteClaim(in long id);
    Claim getClaim(in long id);
    com.ibm.ICollectionsBase.IIterator listClaims();
    ...
};

```

By replacing the addClaim and removeClaim operations with createClaim and deleteClaim methods, the relationship semantics map well to the home-based implementation.

Creating specialized homes

Component Broker provides a default IManagedClient.IHome implementation. There might be specializations of this interface specific to an underlying application adaptor. These include IManagedAdvancedClient.IQueryableIterableHome and IManagedAdvancedClient.IIterableHome. This section explains how to extend a home with domain-specific methods for create and find.

There are a number of cases where the usage of IManagedClient.createFromPrimaryKeyString, IManagedClient.createFromCopyString and IManagedClient.findByPrimaryKeyString might not present the optimal interface for clients wishing to interact with the home to create and find business objects.

Java Specialized Homes construction differs from the procedure detailed in “Chapter 7. MOFW - C++ server programming model – advanced concepts” on page 157. The construction also differs from the way conventional Java business objects are built. These differences exist because the base Home implementation provided by Component Broker is written in C++ and there is no corresponding Java implementation for a Java derived class to extend. Since IOM does not support implementation inheritance between languages, the approach used earlier to create a Java CarPolicy cannot be used here.

Instead, for this special case, we simulate a restricted form of inheritance where the Java derived class is not permitted to override arbitrary base class methods, but can only add new methods. If you try to violate the override restriction by coding an override of a method in IManagedClient.IHome into the Java derived class, you will not get any compile time indication that anything is wrong but the resulting class will not behave as you expect: Whenever that method is called on the resulting Home only the C++ method will be called and the Java override will be ignored.

This section describes the process of extending the home and follows the pattern set out in “Extending a business object” on page 305.

- Extending the interface to IHome
- Implement the extended IHome interface
- Overriding specific methods on specialized homes
- Summary of home extension

Extending the interface to IHome

For the example, PolicyHome is an interface that specializes the IHome interface with some specific create and find methods. The goal is to provide methods specific to the insurance policy abstraction. These methods are shown in the IDL in the following example:

```
interface PolicyHome : IManagedClient::IHome
{
    Policy create(in float premium, in float amount)
        raises ( IManagedClient::IInvalidKey,
                IManagedClient::IDuplicateKey,
                IManagedClient::IInvalidCopy );
        // create a policy passing in the attribute values

    Policy defaultCreate()
        raises ( IManagedClient::IInvalidKey,
                IManagedClient::IDuplicateKey );
    Policy createWithNumber(in long policyNo)
        raises ( IManagedClient::IInvalidKey,
                IManagedClient::IDuplicateKey );
    Policy findByPolicyNumber(in long policyNo)
        raises ( IManagedClient::IInvalidKey,
                IManagedClient::INoObjectWKey ); };
```

Details

IManagedClient.IHome is being extended because it is the interface that is supported by all of the application adaptors provided by the server.

Exceptions have also been put on the methods that are introduced. For completeness, creation methods should raise the IManagedClient.IInvalidKey and IManagedClient.IDuplicateKey exceptions. Find related methods should raise the IManagedClient.IInvalidKey and IManagedClient.INoObjectWKey exceptions. Additional exceptions can also be introduced and raised as appropriate.

An alternative exception strategy is to have the specialized home implementations actually handle some of the exceptions. For example, catching the IManagedClient.IDuplicateKey exception on methods where the key is not passed in could be appropriate.

Alternatives to IManagedClient.IHome

Other alternatives for extending homes are IManagedAdvancedClient.IIterableHome and IManagedAdvancedClient.IQueryableIterableHome. This is true only if the additional interface supported by these extended homes is to be a proper superset of that which is available for the particular application adaptor configuration. Not all homes support these IManagedAdvancedClient interfaces.

Implement the extended IHome interface

The IHome itself is a managed object and the development of an extended home should be done like creating any subclass of a managed object. This is described in “Extending a business object” on page 305. This section highlights things that are specific to the IHome interface and extension.

Implementation interface

Now create the implementation of this client interface as though extending a business object implementation called IManagedAdvancedServer.ISpecializedHome:

```
interface PolicyHomeBO : PolicyHome,
                        IManagedAdvancedServer::ISpecializedHome,
                        IManagedServer::IWrappable
{
};
```

The implementation

The implementation of the new create methods involves careful usage of keys, copies, and the basic interface supported by home. All of the create methods end up using createFromCopyString or createFromPrimaryKeyString, passing a key or copy that has been loaded with the proper information to all proper creation of the object. The findByNumber() method follows a similar pattern.

```
public class _PolicyHomeBOBase
    extends com.ibm.IManagedAdvancedServer._IHomeBase
    implements PolicyHomeBO
{
    ...

    private int counter = System.currentTimeMillis(); // See note 1
    private synchronized int getUnique() // See note 2
    {
        return counter++;
    }
}
```

```

        private com.ibm.IManagedAdvancedServer.
            ISpecializedHomeDataObject iDataObject;
        private String userData;
    };

```

Notes on the example:

1. A simple way to get an initial value for a unique ID generator. If something like this is not sufficient and you find you need a more robust mechanism that involves persistent storage, you should implement a normal Java business object to manage that data. This second object can then be located through the Naming Service using a fixed name known to this Home.
2. This method is marked synchronized to ensure that multiple requests running on different threads in the server do not get the same ID value.

It was previously stated that there is no Java implementation to extend, but that appears to be happening in this example. Actually, this Java base class contains only dummy versions of the `IManagedAdvancedServer.ISpecializedHome` methods, which need to be present in order to successfully compile the derived class. Once the Home is installed in the server, the dummy methods are never called because the C++ implementation is used instead.

create() method:

```

public Policy create(float premium, float amount)
    throws com.ibm.IManagedClient.IInvalidKey,
           com.ibm.IManagedClient.IDuplicateKey,
           com.ibm.IManagedClient.IInvalidCopy
{
    PolicyCopy theCopy = PolicyCopyHelper._create();
    theCopy.policyNo(getUnique());
    theCopy.premium(premium);
    theCopy.amount(amount);
    return PolicyHelper.narrow(createFromCopyString(theCopy._toString()));
}

```

This pattern of creating a Copy Helper object, calling `createFromCopyHelper()`, and narrowing the result, is just what a client would write to implement this functionality if `Policy` did not have a specialized Home. This is typically what the extension methods do. That is, they encapsulate client code sequences so the client does not need to contain them anymore.

defaultCreate() method:

```

public Policy defaultCreate()
    throws com.ibm.IManagedClient.IInvalidKey,
           com.ibm.IManagedClient.IDuplicateKey
{

```

```

        PolicyKey pkey = PolicyKeyHelper._create();
        pkey.policyNo(getUnique());
        return PolicyHelper.narrow(createFromPrimaryKeyString(pkey._toString()));
    }

```

createWithNumber() method:

```

public Policy createWithNumber(int policyNo)
    throws com.ibm.IManagedClient.IInvalidKey,
           com.ibm.IManagedClient.IDuplicateKey
{
    PolicyKey pkey = PolicyKeyHelper._create();
    pkey.policyNo(policyNo);
    return PolicyHelper.narrow(createFromPrimaryKeyString(pkey._toString()));
}

```

Note that this method does not perform exception checking. If something other than `IManagedClient:IDuplicate` should be thrown when the number provided as input is already in use, then a try/catch block and associated logic would be required.

findByNumber() method:

```

public Policy findByPolicyNumber(int policyNo)
    throws com.ibm.IManagedClient.IInvalidKey,
           com.ibm.IManagedClient.INoObjectWKey
{
    PolicyKey pkey = PolicyKeyHelper._create();
    pkey.policyNo(policyNo);
    return PolicyHelper.narrow(findByPrimaryKeyString(pkey._toString()));
}

```

Meet MOFW IManageable requirements

Because this extension to the home introduces no new additional methods and no new key for the home itself, `getPrimaryKeyString`, `getHandleString`, `externalize_to_stream`, and `internalize_from_stream` do not need to be implemented.

MOFW Requirements - IManagedObject interfaces

This section includes the following:

initForCreation() method: The `initForCreation` method is only required to call the parent, passing the `dataObject` that comes in as a parameter. Currently, homes are not actually created using this method. Homes are configured onto systems and brought into existence as part of server initialization. This code is not executed; the example is included for completeness.

```

public void initForCreation( com.ibm.IManagedServer.IDataObject theDO)
    throws com.ibm.IManagedServer.ICreationFailed // See note 1
{

```



```

        super.initForCreation(theDO); // See note 2
        iDataObject =
            IManagedAdvancedServer.ISpecializedHomeDataObjectHelper.narrow(theDO);
        userData = iDataObject->getConfigInfo();
    }

```

Notes on example:

1. The following framework methods are given special treatment so that both the methods in the C++ Home base class (IManagedAdvancedServer.IHome_Impl) and the Java methods will be executed:

- initForCreation()
- uninitForDestruction()
- initForReactivation()
- uninitForPassivation()
- syncToDataObject()
- syncFromDataObject()

This special treatment applies only to these methods. The general rule remains: For methods introduced in the specialized Home interface PolicyHome, only the Java versions will be called; for inherited methods, only the C++ versions will be called.

After the Java Home extension is configured into the server environment, calls are routed to the C++ code for methods in IManagedAdvancedServer.IHome and to the Java side for the extension methods introduced in PolicyHome, and to both implementations for those six special case framework methods. This occurs whether the call originates in C++ or Java, and whether the client is another business object or an application on a remote system.

2. You may be confused by this super call because it calls up to the Java base class whose methods, as was previously explained, are dummy methods. However, these calls are here for future compatibility.

initForReactivation: Set the data object using the same pattern as initForCreation. Any other specialized home specific code that needs to execute when a home is activated goes in this method implementation.

```

public void initForReactivation(com.ibm.IManagedServer.IDataObject inputDO)
    throws com.ibm.IManagedServer.IReactivationFailed
{
    super.initForReactivation(inputDO)
    iDataObject = CarPolicyDOHelper.narrow(inputDO);
}

```

uninitForDestruction: All that is necessary is to call the parent.

```

public void uninitForDestruction()
    throws com.ibm.IManagedServer.IDestructionFailed
{
    super.uninitForDestruction();
}

```

uninitForPassivation: Call the parent following the same pattern as uninitForDestruction.

```

public void uninitForPassivation()
    throws com.ibm.IManagedServer.IPassivationFailed
{
    super.uninitForPassivation();
}

```

syncFromDataObject() method: Call the parent following the same pattern as uninitForDestruction.

```

public void syncFromDataObject()
    throws com.ibm.IManagedServer.ISynchronizationFailed
{
    super.syncFromDataObject();
}

```

syncToDataObject() method: Call the parent following the same pattern as uninitForDestruction.

```

public void syncToDataObject()
    throws com.ibm.IManagedServer.ISynchronizationFailed
{
    super.syncToDataObject();
}

```

Keys

The same class used for the MOFW home works here. Homes are found generally using factory finders and are not created. Keys are not used in the programming model. Keys are used in the internal server run time.

Copy helper

The same class used for the MOFW home works here. Homes are not created, and therefore no copy helper is needed or usable based on the life cycle of homes.

Leveraging server-provided essential state extensions

This facility, described in is also available in Java. It allows a single string of read-only data to be associated with each Home; the value is supplied under the name of `userData` in the Object Builder dialog and as an attribute of the same name in the generated DDL file that defines a Home.

This data is accessed by calling the `ISpecializedHomeDataObject.getConfigInfo()` method on the Home's Data Object. Because this is a fairly expensive operation, you should call this function only once and keep a copy of the result in your Home, as demonstrated in the previous example.

Overriding specific methods on specialized homes

In addition to supplying specific methods that allow clients to create and find objects without using keys and copy helper, specialized homes also provide the mechanism for overriding other interfaces supported by the `IManagedClient.IHome` interface. The following specific methods that can be overridden are described below:

- `IManagedClient.IHome.createFromPrimaryKeyString`
- `IManagedClient.IHome.createFromCopyString`
- `IManagedClient.IHome.findByPrimaryKeyString`

The `IManagedClient.IHome.createFromPrimaryKeyString` and `createFromCopyString()` can be overridden to prevent creation of objects of a particular type. There are cases when using implementation inheritance and polymorphism where homes are configured onto a system even when the class is abstract. This is usually to facilitate polymorphic `findByPrimaryKeyString` operations. While `findByPrimaryKeyString()` is desired as a polymorphic operation, the create capabilities may in fact be prohibited. Other reasons include special checking that needs to be done before the actual create is issued. While this can be done with specialized create methods, using the base programming model methods of `createFromPrimaryKeyString()` and `createFromCopyString()` may be desirable in some situations.

Overriding `IManagedClient.findByPrimaryKeyString` is done to facilitate special find logic. This is most common in cases where polymorphic relationships exist between business objects. For example, if an abstract class of "a" has subclasses of "b" and "c", it would be reasonable to try to find an "a" with a given key. While there are no "a" instance because it is abstract, it is reasonable to have an overridden `findByPrimaryKeyString` that would look at all concrete subclasses of "a" to properly execute the find operation. By overriding `findByPrimayKeyString()` instead of using a specialized find method, data objects that deal with polymorphic relationships can work unchanged. They depend on `findByPrimaryKeyString()` as part of the attribute getter implementation. Another reason to override `findByPrimaryKeyString` is to prevent standard access to objects by throwing an exception in this method. While this will work, and is allowed, a side effect is that the objects supported by this home cannot be resolved in object relationships using the standard Home/Key pattern. This is because the `findByPrimaryKeyString()` method will be called during the resolution of these objects. Since this method will now throw an exception the object reference will fail. Using an alternative object

relationship pattern, such as SOR, will solve some of this problem, however ForeignKey relationships would still fail.

The contracts of these methods must be maintained even when they are overridden. The exceptions thrown must still be honored by the specialized home. For example, `IManagedClient.IDuplicateKey` must still be thrown by `create` and `IManagedClient.INoObjectWKey` should be returned. These are integral to much of the mainline programming practice in Component Broker client programs. This contract must be maintained even in the specialized home overridden versions of these methods. New exceptions cannot be introduced on these methods as there is not overriding or overloading of interface specification allowed in IDL.

Thread safety

Special care needs to be taken when writing specialized homes that access and modify static data. This is because homes are more likely to be accessed from many clients concurrently. Because of this they are required to be thread-safe in order to provide correct access and control of static data.

Summary of home extension

Extending a home is much like extending any business object but simpler. Refer to *Create the managed object class and implementation* for more details.

Creating UUID specialized homes

There are cases where a developer may want to use a transient object but does not have a unique value to use for a primary key. Component Broker provides a mechanism called UUID (Universally Unique Identifier) that can be used for this purpose. UUID support is most useful when there are short-lived components that do not need to be found after they are created. The developer can use the provided mechanisms to create a specialized home that can create these components. The Component Broker framework provides a mechanism to generate a UUID for the component. This value must be generated on the server in order to guarantee uniqueness of the value on all Component Broker platforms.

This section describes the process of creating this particular type of home and follows the pattern set out in “Creating specialized homes” on page 338.

Extending the interface

For the example, `AgentUUIDHome` is an interface that specializes the `IHome` interface with some specific `create` methods that generate UUID objects. The

goal, as with any specialized home, is to provide methods specific to the application. These methods are shown in the IDL in the following simple example:

```
interface AgentUUIDHome : IManagedClient::IHome
{
AgentUUID createAgentWithKey(in float commPercent,
                             in float pendingPaycheck,
                             in string<100> agentName )
    raises    (IManagedClient::IInvalidKey,IManagedClient::IDuplicateKey);

AgentUUID createAgentWithCopy(in float commPercent,
                              in float pendingPaycheck,
                              in string<100> agentName )
    raises    (IManagedClient::IInvalidKey,IManagedClient::IDuplicateKey);
};
```

As with any specialized home, you must first decide which IHome interface should be specialized. This seems to have a simple solution: Inherit from the IHome in the managed object framework. Note that there are other home interfaces that support iteration and query. This functionality is not typically required or useful for most UUID implementations.

Details

The details behind this specialized home are the same as Details except for the discussion on queryable and iterable homes. As described above, it is possible to use these classes for as the parent class for UUID specialized homes but it is not typical. The AgentUUIDHomeBO.idl file is shown in the following example:

```
interface AgentUUIDHomeBO : AgentUUIDHome,
                           IManagedAdvancedServer::ISpecializedHome,
                           IManagedServer::IWrappable
{
};
```

The implementation

As with any specialized home, the implementation of the new create methods involves careful usage of keys, copies, and the basic interface supported by home. All of the create methods end up using createFromCopyString() or createFromPrimaryKeyString(), passing a key or copy that has been loaded with the proper information to all proper creation of the object. The unique aspect of creating a UUID specialized home is that the keys and copy helpers will subclass special UUID classes provided by the Component Broker framework that provide special UUID support.

```
public Agent createAgentWithKey(float commPercent,
                                float pendingPaycheck, java.lang.String agentName)
{
byte[] theKeyString;
```

```

com.ibm.IManagedClient.IManageable mo;
Agent newAgent;

// create a UUID key

com.ibm.IManagedAdvancedServer.IUUIDPrimaryKey newKey =
    com.ibm.IManagedAdvancedServer.IUUIDPrimaryKeyHelper_create();
newKey.generateUuid();
theKeyString = newKey._toString();

try {
    mo = createFromPrimaryKeyString(theKeyString);
} catch (java.lang.Exception exception) {
    throw exception;
}
newAgent = AgentHelper.narrow(mo);

if (newAgent == null) {
    throw new com.ibm.IManagedClient.IInvalidKey();
} else {
    newAgent.commPercent(commPercent);
    newAgent.pendingPaycheck(pendingPaycheck);
    newAgent.agentName(agentName);
}
return newAgent;
}
public Agent createAgentWithCopy(float commPercent,
    float pendingPaycheck, java.lang.String agentName)
{
    byte[] theCopyString;
    com.ibm.IManagedClient.IManageable mo;
    Agent newAgent;

    // create a UUID key
    AgentCopy newCopy = AgentCopyHelper._create();
    newCopy.commPercent(commPercent);
    newCopy.pendingPaycheck(pendingPaycheck);
    newCopy.agentName(agentName);
    theCopyString = newCopy->_toString();

    try {
        mo = createFromCopyString(theCopyString);
    } catch (java.lang.Exception exception) {
        throw exception;
    }
    newAgent = AgentHelper.narrow(mo);

    if (newAgent == null) {
        throw new com.ibm.IManagedClient.IInvalidKey();
    }
    return newAgent;
}

```

Meet MOFW IManageable requirements

Because this extension to the home introduces no new additional methods and no new key for the home itself, `getPrimaryKeyString()`, `getHandleString()`, `externalize_to_stream()`, and `internalize_from_stream()` do not need to be implemented.

MOFW Requirements - IManagedObject interfaces

These requirements are the same as any specialized home, see this section “Creating specialized homes” on page 338.

Keys

Keys are not used to create homes, however a special primary key class is used in the home to create instances of components that this home creates. This class is `IManagedAdvancedServer.IUUUIDPrimaryKey`.

Copy helper

Copy helpers are not used to create homes, however a special copy helper class is subclassed by the component that this home creates. This class is `IManagedAdvancedServer::IUUUIDCopyHelperBase`.

Leveraging server-provided essential state extensions

These requirements are the same as any specialized home, see this section “Creating specialized homes” on page 338.

Overriding the standard create interface

These requirements are the same as any specialized home, see this section “Creating specialized homes” on page 338.

Chapter 13. Assembling and installing components on Component Broker/Workstation



The following chapter is platform-dependent and does NOT apply to OS/390 Component Broker. See *OS/390 Component Broker Programming: Assembling Applications* for further information.

The previous chapters in this book have provided direction and guidance on building business objects and client applications. Examples of various techniques used for structuring and implementing business objects have been given. However, many of the details about how these business objects and client applications are transformed from simple artifacts into running applications have been omitted. This chapter addresses some of these topics.

Much of the material in this chapter is for reference only. Component Broker tools generate most of the code and configuration information which is described. The details are here for those seeking a more complete understanding of the Component Broker adaptors and run-time environment. Details about the managed objects which accompany the business logic are provided. Detailed and general flows of common scenarios give some perspective of server and business object interactions.

Container configuration alternatives make up one of the sections in this chapter. This section outlines common container configurations and provides important background for administrators that are looking for a more complete understanding of the options that are available.

Finally, some Component Broker workstation specific client programming interfaces techniques are discussed. These “quality of service” interfaces are specific to the adaptors supported by Component Broker on workstation platforms.

Local only object implementation details

The local only development process includes the following topics:

- “C++ local only” on page 352
- “Java local only” on page 353

C++ local only

This development process applies to all keys and copy helpers. The basic process is as follows:

1. Create the IDL file inheriting from the base class specified by the programming model. This might be `IManagedLocal::IPrimaryKey` (for a primary key) or `IManagedLocal::INonManageable` (for a copy helper) or some other base class.
2. Introduce the appropriate `#pragma` to tell the tools and the world that the bindings that are to be generated do not need to contain things that assist CORBA objects in becoming remote. These are objects that are addressed from within the process. The `#pragma` to be used is: `#pragma meta <interface name> localonly`. The following example code shows how to do this for a `PolicyPrimaryKey` interface:

```
#ifndef PolicyPrimaryKey_idl
#define PolicyPrimaryKey_idl
#include <ILocalOnly.idl>
interface PolicyPrimaryKey : IManagedLocal::IPrimaryKey {
#pragma meta PolicyPrimaryKey localonly
    /* This makes it generate local only bindings */

    /* put your attributes methods here... */
    attribute long policyNo;
};
```

Note: There is an abstract keyword supported by the emitters as well. When combined with `local only`, this keyword suppresses the generation of the implementation stub. The abstract keyword should not be used for key and helper interfaces.

3. Run the emitter against the local only IDL file.

```
idlc -s"uc;hh;ih;ic"
```

This generates the following files:

- `PolicyPrimaryKey.hh` - usage bindings containing `PolicyPrimaryKey` and `PolicyPrimaryKey_Skeleton` class. The `PolicyPrimaryKey` class introduces the `PolicyPrimaryKey::_create()` interface that is implemented later.
- `PolicyPrimaryKey_C.cpp` - implementation of client side usage bindings
- `PolicyPrimaryKey.ih` - implementation class interface
- `PolicyPrimaryKey_I.cpp` - implementation class stubs. Included in this file is the method `classname::_create()`; which should be implemented to new up and return an instance of the proper `_Impl` class (for example, `PolicyPrimaryKey_Impl`).

Table 7 on page 353 contains a summary of the artifacts involved in this activity.

Table 7. Artifacts of the local only development process

Artifact Name (across) ArtifactType (down)	Non-Remoteable Abstraction
IDL Put these under configuration management control.	Abstraction.idl (Use Local Only and Abstract #pragmas to indicate if this is a regular local only or an abstract local only).
.hh (language usage binding) Never modify these – have the makefile generate them.	Abstraction.hh (contains the Abstraction_Skeleton class and the Abstraction C++ class). This also contains the definition of _create() if the local only #pragma is used.
_C.cpp (client side binding) Never modify these – have the makefile generate them.	Abstraction_C.cpp (The #pragma ensures that only the necessary stuff is in here, no remotable and dispatcher-related pieces).
.ih (implementation interface – emit once or enter manually).	Abstraction.ih (defines Abstraction_Impl that inherits only from Abstraction_Skeleton).
_I.cpp (implementation code – emit once or enter manually).	Abstraction_I.cpp (implementation of real logic – the code). When the local only #pragma is used, this file contains an implementation of the Abstraction::_create() method.
_C.obj (from _C.cpp – need rule in makefile).	Abstraction_C.o (binding used by clients).
_I.o (place on servers) (from _I.cpp – need rule in makefile).	Abstraction_I.o
Abstraction.LIB (group one more abstraction into a LIB).	Abstraction_I.o and Abstraction_C.o can get packaged together into one LIB that goes on both the client (C++) and the server. See “Packaging for client and server (VA C++)” on page 406 for more information on how to assemble all of the pieces needed for a managed object.

4. Create a Makefile - Use Table 7 to determine proper packaging strategy.
5. Write the Code - Use the section in this book that maps to the kind of local only object being constructed. The implementation of the code is in the _I.cpp file with the definition of any private variables or methods being done in the .ih file.
6. Run make from the environment with the correct paths set
7. Test and debug.
8. Install.

Java local only

Most of the process is the same for Java. The emitter is IDL-to-Java. The results are .class files. Final packaging might involve .zip or .jar files that go onto the Web server for downloading.

Managed object implementation details

This section describes how the managed object class is constructed. To complete the rest of the “Not ready to use at this point” things from Table 4 on page 129, a new class is introduced. This is called the managed object class. It adds management capabilities to the component. It is called PolicyMO in the example.

However, it does much more than filling out the rest of the things labeled “Not ready to use at this point” in Table 4 on page 129. There are additional methods that are implemented in PolicyMO. These deal specifically with the application adaptor into which a component is installed. In some cases, behavior is added or the business logic is surrounded. In others, there are additional methods that assist in implementing the quality of service provided by the application adaptor.

The example managed object is for use in BOIM-based application adaptors.

From an interface perspective, the managed object mixes together the component’s business object interface with some additional things from the application adaptor in which the component resides.

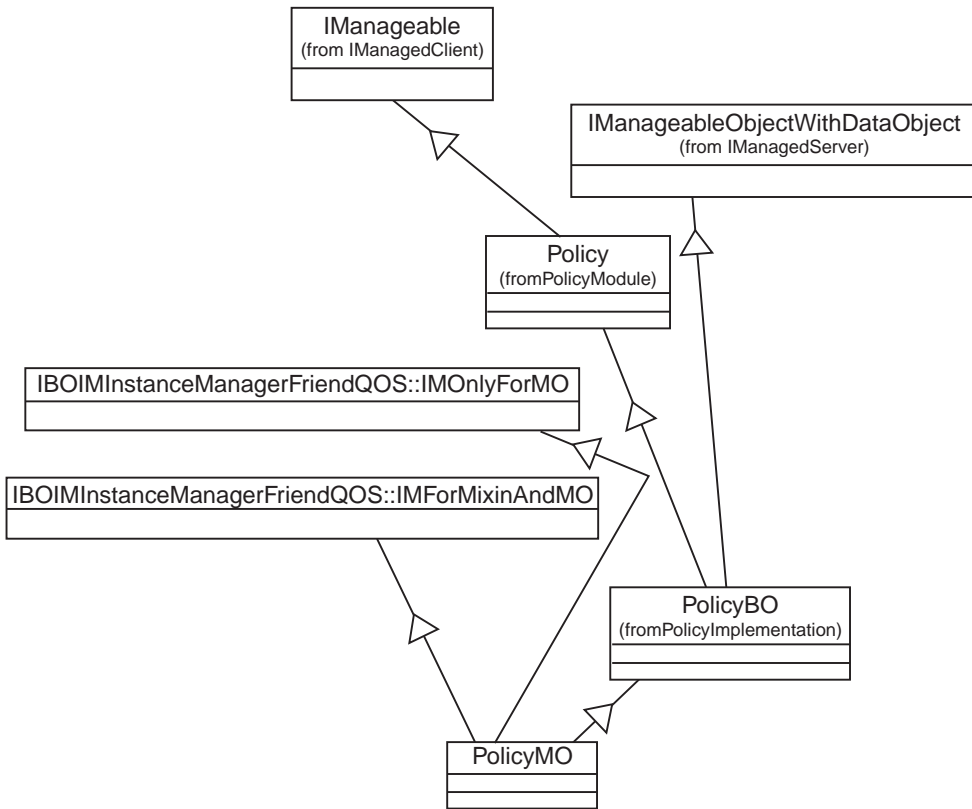


Figure 62. IDL hierarchy view of a managed object

The new abstraction being introduced is the PolicyMO. The IDL for the Policy managed object is as follows:

```

#ifndef _PolicyTransMO_idl
#define _PolicyTransMO_idl
#include <PolicyBO.idl>
#include <IBOIMInstanceManagerFriendQOS.idl>
#include <IBOIMServerFriendQOS.idl>
interface PolicyMO : PolicyBO,
    IBOIMInstanceManagerFriendQOS::IMOnlyForMO,
    IBOIMInstanceManagerFriendQOS::IMForMixinAndMO,
    IBOIMServiceFriendQOS::IMDynamicDispatching
{
};
#endif

```

This is a pretty simple interface that mixes together the IDL for the business object and some additional interfaces that need to be implemented. These additional interfaces are represented by IBOIMInstanceManagerFriendQOS::IMOnlyForMO,

IBOIMInstanceManagerFriendQOS::IMForMixinAndMO and IBOIMServiceFriendQOS::IMDynamicDispatching. They need to be implemented by the managed object. The implementation of these methods in the managed object is most often just a delegation to another object called the mixin.

Implementing the new PolicyMO interface requires inheritance of the PolicyBO_Impl and implementing the other interfaces of the PolicyMO generated by the emitters. The implementation interface is as follows:

```

#ifdef _PolicyMO_ih_included
#define _PolicyMO_ih_included
#ifdef SOMCBNOLOCALINCLUDES
#include <PolicyBO.ih>
#include <PolicyTransMO.hh>
#include <IBOIMInstanceManagerFriendQOS_IMixinImpl.ih>
#else
#include "PolicyBO.ih"
#include "PolicyTransMO.hh"
#include "IBOIMInstanceManagerFriendQOS_IMixinImpl.ih"
#endif

class PolicyTransMO_Impl : public virtual ::PolicyTransMO::Skeleton,
                          public virtual PolicyBO_Impl

{
public:
    PolicyMO_Impl();
    virtual ::CORBA::Float amount ();
    virtual ::CORBA::Void amount (::CORBA::Float amount);
    virtual ::CORBA::Long policyNo ();
    virtual ::CORBA::Float premium ();
    virtual ::CORBA::Void premium (::CORBA::Float premium);
    virtual ::CORBA::Void addBeneficiary ();
    virtual ::CORBA::Void delBeneficiary ();

#ifdef CBS_TRACE_DEBUG
    virtual ::CORBA::Void initForCreation
        (::IManagedServer::IDataObject_ptr theDO);
    virtual ::CORBA::Void initForDestruction();
    virtual ::CORBA::Void initForReactivation
        (::IManagedServer::IDataObject_ptr theDO);
    virtual ::CORBA::Void initForPassivation();
    virtual ::CORBA::Void externalize_to_stream
        (::CosStream::StreamIO_ptr targetStreamIO);
    virtual ::CORBA::Void internalize_from_stream
        (::CosStream::StreamIO_ptr sourceStreamIO,
         ::CosLifecycle::FactoryFinder_ptr there);
#endif
    virtual ::ByteString* getPrimaryKeyString();

    //Methods from Framework

    virtual ::IManagedClient::IHome_ptr getHome();

```

```

virtual ::CosLifeCycle::Key* external_form_id();
virtual ::CosObjectIdentity::ObjectIdentifier constant_random_id();
virtual ::CORBA::Boolean is_identical(
    ::CosObjectIdentity::IdentifiableObject_ptr other_object);
virtual ::CosLifeCycle::LifeCycleObject_ptr copy(
    ::CosLifeCycle::FactoryFinder_ptr there,
    const ::CosLifeCycle::Criteria & the_criteria);
virtual ::CORBA::Void move(::CosLifeCycle::FactoryFinder_ptr there,
    const ::CosLifeCycle::Criteria & the_criteria);
virtual ::CORBA::Void remove();
virtual ::IExtendedObjectIdentity::AbsoluteIdentity*
    get_absolute_identity();

//Application Adaptor Methods

virtual ::CORBA::Void checkpointToDatastore();
virtual ::CORBA::Void refreshFromDatastore();
virtual ::CORBA::Void externalizeKey
    (::IIMFLocalToServer::IKeyStream_ptr keyStream,
    ::IIMF::IContainer_ptr container);
virtual ::IIMF::IContainer_ptr getContainer();
virtual ::CORBA::Void before_completion();
virtual ::CORBA::Void after_completion(
    ::CosTransactions::Status status);

//Special Methods

virtual ::CORBA::Void setMixin(
    ::IIMFLocalToServer::IMixinForDelegatingMO_ptr mixinPtr);
virtual ::IBOIM::InstanceManagerFriendQOS::IMixin_ptr getMixin();
virtual ::CORBA::Void _incrcf();
virtual ::CORBA::Void _decref();
virtual ::CORBA::Void callMethodByName(const char * method_name,
    ::IBOIM ServiceFriendQOS::ArgList & method_arguments,
    ::CORBA::Any*& method_return_value);
virtual CORBA::Object_ptr _localReference();
virtual void _localReference(CORBA::Object_ptr objectPointer);

protected:
    IBOIMInstanceManagerFriendQOS_IMixinImpl mixinPointer;

private:
    PolicyBO_var _localProxy;
};
#endif /* _PolicyTransMO_ih_included */

```

The following are a few kinds of methods in this interface.

- Methods defined in the business object interface (for example, policy). These are the domain specific methods and the get and set methods associated with public attributes.
- Methods defined by the OMG Object Services that are supported by the application adaptor. This also includes extensions to the OMG Object Services interfaces that have been made.

- Methods defined by the BOIM application adaptor interfaces. These are augmentations made to the interface so that application adaptors can manage the components.
- Special Methods. These are necessary to make the infrastructure work correctly. This set of methods is actually implemented directly in the managed object.

Each of these kinds of methods are described in sections, with an example of the implementation pattern that is used.

Business object methods in the managed object implementation

All business domain specific methods which are introduced need to have additional implementation in the managed object. This implementation makes a call to the mixin object before and after calling the real business logic. An example follows:

```

::CORBA::Float PolicyTransMO_Impl::amount()
{
    ::CORBA::Float retval;
    IBOIMExtLocalToServer_IMixinPointerImpl mixinPointer(getMixin());
    CALL_MIXIN_BEFORE2(mixinPointer,"getamount");

    #ifdef CBS_TRACE_DEBUG
        void * trc_handle = BOSS_TRACE_SERVER_START_2(this, "getamount");
    #endif

    //call the real business logic
    retval = PolicyBO_Impl::amount();

    #ifdef CBS_TRACE_DEBUG
        BOSS_TRACE_SERVER_STOP(trc_handle,"getamount");
    #endif

    CALL_MIXIN_AFTER2(mixinPointer, "getamount");
    return retval;
}

```

There are some #ifdef statements in here for enabling remote debugging. These are only included in the compiled code when there is a desire for a version of the managed object that is enabled for debug mode.

The method on the mixinPointer to constMethod (“getamount”) is a special method call. This method will be generated on all attribute getters and on all methods that are marked as “constant method” using the objectBuilder business object interface wizard. This method is processed by the mixin. If every method called in a unit of work is a constMethod, then the datastore will not be unnecessarily updated. This applies only to components that cache data in the business object.

This basic pattern for the implementation of the xxxxxMO_Impl methods applies to all business logic methods.

OMG services methods in the managed object implementation

There is a set of methods defined by the OMG Object Services, for which the mixin object has an implementation. The pattern of implementation for these methods follows:

```
::CORBA::Boolean PolicyTransMO_Impl::is_identical(  
    ::CosObjectIdentity::IdentifiableObject_ptr other_object)  
{  
    IBOIMExtLocalToServer_IMixinPointerImpl mixinPointer(getMixin());  
    return mixinPtr->is_identical(other_object);  
}
```

The mixin object reference holder (returned by `mixin()`) is used to construct a mixin pointer object. The mixin pointer object adjusts a counter in the mixin reference to prevent it from being removed while it is in use. The mixin pointer “->” operator selects the mixin or alternate mixin if the mixin no longer exists.

Application adaptor methods in the managed object implementation

Another set of methods is introduced by the application adaptor framework. The implementation pattern for these is the same as the pattern used for the Object Services. An example implementation follows:

```
void PolicyTransMO_Impl::checkpointToDatastore()  
{  
    IBOIMExtLocalToServer_IMixinPointerImpl mixinPointer(getMixin());  
    mixinPointer->checkpointToDatastore();  
}
```

Special methods in the managed object implementation

The `setMixin()`, `_incref()`, and `_decref()` methods in the managed object have special forms. The `setMixin` method is implemented as follows:

```
void PolicyMO_Impl::setMixin(  
    IIMFLocalToServer::IMixinForDelegatingMO_ptr theMixin)  
{  
    mixin_.setMixin(theMixin);  
}
```

The `setMixin` method stores the pointer to the mixin in the mixin holder (`mixin_`). The holder is used when delegating to the mixin. If this pointer is not set correctly, attempting to use the managed object would likely raise a `CORBA::INV_OBJREF` exception.

The `_incref()` method is implemented as follows:

```

void PolicyMO_Impl::_ineref()
{
    IBOIMExtLocalToServer_IMixinPointerImpl mixinPointer(mixin());
    mixinPointer._ineref();
}

```

The `_decref()` method is similar. The mixin pointer class delegates the `_ineref()` and `_decref()` methods to the mixin, if it exists, or to the ORB otherwise. Thus, the `_ineref()` and `_decref()` methods work correctly even before the `setMixin()` method is called and after the `remove()` method is called.

Handling component augmentation of OMG services methods

In some cases, you may need to augment the behavior of some OMG Services methods in the component. This requires you to manually change the specific managed object method to invoke both the business object and the mixin. The following is an example of the augmented `remove()` method:

```

::CORBA::Void PolicyMO_Impl::remove()
{
    IBOIMExtLocalToServer_IMixinPointerImpl mixinPointer(getMixin());
    mixinPointer->remove();
}

```

The above example is simplified to remove any exception handling that should be done.

Managed objects and specialized homes

In most cases, subclassing is done from a base class that is also created by the object provider. Some cases, however, such as the one described in this section, require extending objects that are part of the Component Broker server.

The interface for the home, used at the managed object customization phase of developing an extended home, is named `IManagedAdvancedServer::ISpecializedHome`. This is located in the `IManagedAdvancedServer.idl` file.

The interface name that is inherited for implementation is `IManagedAdvancedServer_ISpecializedHome_Impl` and is located in the `IManagedAdvancedServer_ISpecializedHome_Impl.ih` files.

Managed objects for specialized homes are very similar to other managed objects except for the interfaces that they inherit. The following IDL is for the specialized home for policy.

```

#ifdef _PolicyHomeMO_idl
#define _PolicyHomeMO_idl

```

```

#include <PolicyHomeB0.idl>
#include <IManagedAdvancedServer.idl>
interface PolicyHomeM0 : PolicyHomeB0,
                        IManagedAdvancedServer::ISpecializedHome
{
};
#endif

```

The implementation is very similar. The standard framework methods that have been previously described for managed objects are required along with the implementation of any business methods defined for that home. In this case:

```

virtual ::Policy_ptr createWithNumber(::CORBA::Long policyNo);
virtual ::Policy_ptr findByPolicyNumber(::CORBA::Long policyNo);

```

Data object implementation details

The PolicyDO which is described in “Chapter 2. Personal life insurance application example” on page 31 is just an interface to the essential state of the component. PolicyDO is not ready to run yet. In order to get that ready to run, PolicyDO must inherit from one of the subclasses of the abstract IDataObject class. While tools generate the implementations of the data objects, the sections that follow describe what is going on under the covers.

The IManagedServer::IDataObject interface and its subclasses provided augmentation of the basic interface to the data which was described by the object provider. Additional methods are available on the data objects that the application adaptors need to use. These interfaces are oriented around what is necessary to interact with the underlying data store. This is encapsulated from the object provider. As much as possible, a key goal of Component Broker is to encapsulate the differences in the underlying data store.

Customization, therefore, is not necessarily the job of the object provider. It is more probably the job of an application or component customizer. The customizer has some familiarity with the domain of the components, but is an expert on the particular environment and resource managers into which the applications and components are to be installed.

Data object customization is a key part of turning a component into a compatible managed object. Depending on the server environment into which the component is installed, different options for the data objects present themselves. This section goes through each of the options available for Component Broker. These are presented starting at the simplest and working towards the more complex data object customizations.

For Component Broker, there are a number of choices for data object customization. These choices are based on the support available in the server. Refer to Component Broker for Windows NT and AIX Online Documentation and *WebSphere Application Server Enterprise Edition Component Broker Application Development Tools Guide* for further information on data object customization.

The choices include:

Persistent Data Object - Static SQL (production use)

This is the choice to make if the component is being customized for installation into an application adaptor that is configured to use static SQL backed by either DB2 on NT or DB2 on MVS. This choice provides data objects with SQL included in their implementations. Optionally, if the component is developed to be queryable (that is, returned as part of the result set of a query), then the data object must be developed to support *query pushdown*. Query pushdown refers to queries on objects in a home that can be done in the database rather than object space. The primary keys in this case are a subclass of the `IManagedClient::IPrimaryKey` base class. For customization activity, see “BOIM data object customization – static SQL” on page 364.

Persistent Data Object -- Cache Service (production use)

An alternative to static SQL that uses the Component Broker Cache Service for improved concurrency, optimistic caching. See the discussion about Cache Service in the *WebSphere Application Server Enterprise Edition Component Broker Advanced Programming Guide* for more details. This implementation replaces the SQL statements in the dataobject with calls to the Cache Service. The business object developed to work with this type of dataobject is expected to be queueable and to use the delegating pattern. For more information, see “BOIM data object customization – Cache Service” on page 371.

Transient Data Object -UUID Key (production use)

This is the choice to make if the component is going to be targeted for an application adaptor that supports transient managed objects. Primary keys in this case are a subclass of the `IManagedAdvancedServer::IUUIDPrimaryKey` base class. This type of data object customization is described in “Transient data object customization – UUID key (production use)” on page 379.

Transient Data Object - (production use)

This choice is similar to the previous choice, except that any domain specific key that ensures uniqueness can be used. This type of data object customization is described in “Transient data object – any key (production use)” on page 381.

Table 8 on page 363 summarizes these data object customization choices and the interfaces they use as part of the customization.

Table 8. Interfaces needed by persistent data objects

Interface Required	Data Object Customization			Notes
	Static SQL data object	Static SQL with query push down	Caching Cache Service data object with query push down	
IBOIMExtLocalToServer:: IQueryable Data Object	No	Yes	Yes	This interface adds the <code>internalizeData()</code> method that must be implemented.
IBOIMExtLocalToServer:: IHomeAwareDataObject	No	Yes	Yes	This interface adds the <code>setHome()</code> method whose implementation is provided by the Component Broker implementation of the <code>IRDBIMExtLocalToServer::ICachingServiceDataObject</code> interface. For the caching DAO data object with query pushdown, this interface is inherited by <code>IRDBIMExtLocalToServer::ICachingServiceDataObject</code> and so does not have to be explicitly inherited.
IRDBIMExtLocalToServer:: IDataObject	Yes	Yes	No	This interface adds the <code>setConnection()</code> method that should be implemented in the persistent object. This persistent object is a companion object to the SQL data object that allows the database specific code, in this case SQL, to be isolated from non-specific data object code.
IRDBIMExtLocalToServer:: ICachingServiceDataObject	No	No	Yes	This interface adds the <code>getMOInterfaceName()</code> , <code>getMapExpression()</code> , and <code>setHome()</code> methods. These method implementations are provided by the interface implementation.

BOIM data object customization – static SQL

If the data object customization is going to target static SQL data objects that run in a DB2 application adaptor, then you should read this section. In this section, the PolicyDO interface is used as the basis for describing the customization necessary.

SQL data object interfaces

The IDL-based picture for SQL-backed BOIM data objects follows:

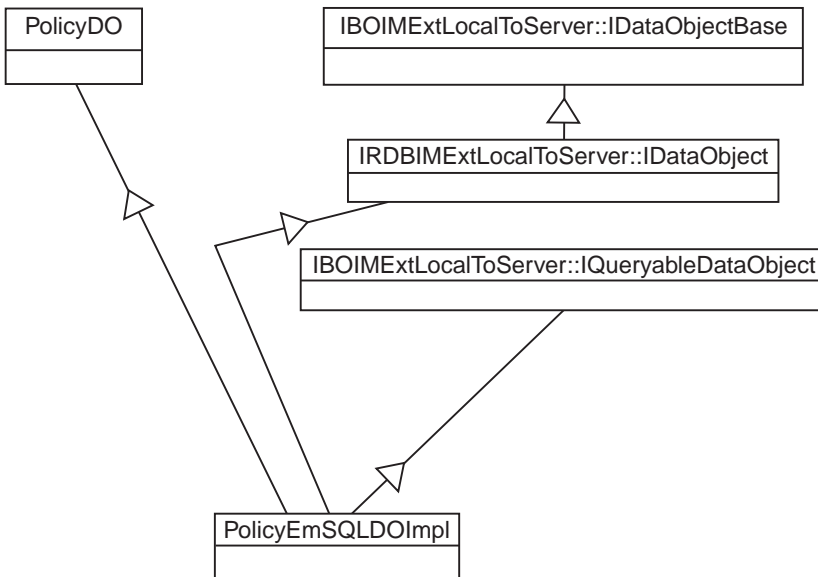


Figure 63. SQL BOIM data object IDL interface view

The example shows the case where the component is also configured to be queryable. (that is, IBOIMExtLocalToServer::IQueryableDataObject should be inherited if the data object is to be used with queryable homes otherwise it should not be inherited.) Here is the IDL for this:

```
#include <PolicyDO.idl>
#include <IRDBIMExtLocalToServer.idl>
#include <IBOIMExtLocalToServer.idl>

interface PolicyEmSQLDOImpl : PolicyDO,
    IRDBIMExtLocalToServer::IDataObject,
    IBOIMExtLocalToServer::IQueryableDataObject
{
    #pragma meta PolicyEmSQLDO localOnly
};
```

In fact there is more implementation inheritance as well. The interfaces added to the basic PolicyDO include those required for all two-level store application adaptors (BOIM) and those specific to the relational database adaptor (RDBIM) This is shown in the implementation view following:

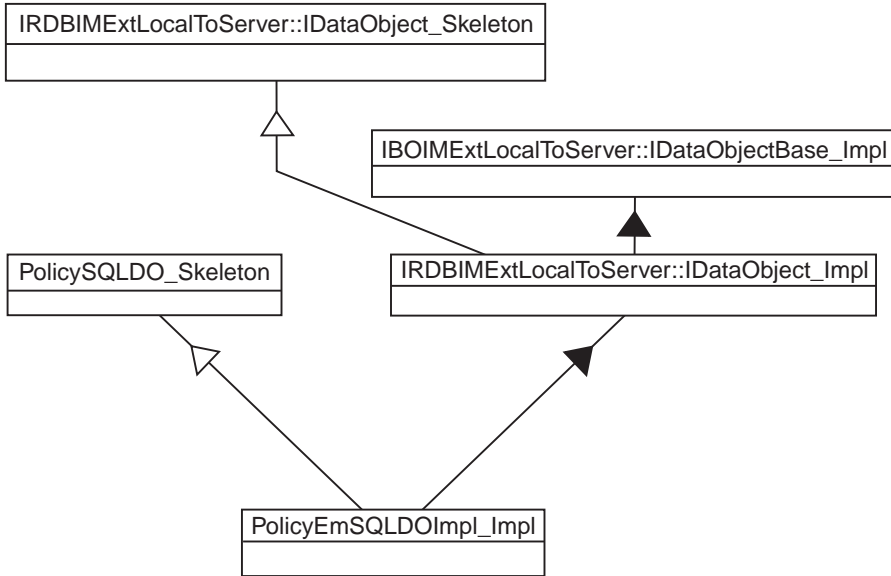


Figure 64. SQL BOIM data object implementation interface inheritance

The IRDBIMExtLocalToServer::IDataObject_Impl in combination with IBOIMExtLocalToServer::IDataObjectBase_Impl provides an implementation for some of the methods that are described in the IBOIMExtLocalToServer::IDataObject and IRDBIMExtLocalToServer::IDataObject interfaces. The following methods have implementations which are inherited:

- string connection()
- void connection(string value)
- void markDirty()
- void clearDirty()
- boolean isDirty()
- void updateToDataStore()
- void insertToDataStore()
- void retrieveFromDataStore()
- void deleteFromDataStore()

These methods should not be overridden.

Given this base, the PolicyEmSQLDO_Impl implementation interface should look like this:

```

#include <PolicyEmSQLPO.hpp>

#ifdefSOMCBNOLOCALINCLUDES
#include <IRDBIMExtLocalToServer.ih>
#include <PolicyEmSQLDOImpl.hh>
#else
#include "IRDBIMExtLocalToServer.ih"
#include "PolicyEmSQLDOImpl.hh"
#endif

class PolicyEmSQLDOImpl_Impl : public virtual ::PolicyEmSQLDOImpl_Skeleton,
public virtual IRDBIMExtLocalToServer_IDataObject_Impl
{
    public:

    //default constructor doimpl
    PolicyEmSQLDOImpl_Impl();

    ::CORBA::Float amount();
    ::CORBA::Void amount( ::CORBA::Float amount);
    ::CORBA::Long policyNo();
    ::CORBA::Void policyNo( ::CORBA::Long policyNo);
    ::CORBA::Float premium();
    ::CORBA::Void premium( ::CORBA::Float premium);

    virtual ::CORBA::Void insert();
    virtual ::CORBA::Void update();
    virtual ::CORBA::Void retrieve();
    virtual ::CORBA::Void del();
    virtual ::CORBA::Void setConnection(const char* dbName);
    virtual ::CORBA::Void internalizeFromPrimaryKey(
        ::IManagedLocal::IPrimaryKey_ptr inKey);
    virtual ::CORBA::Void internalizeFromCopyHelper(
        ::IManagedLocal::INonManageable_ptr inCopy);
    virtual ::CORBA::Void internalizeKeyAttributes(
        ::IIMFLocalToServer::IKeyComponent_ptr keyComp);
    virtual ::CORBA::Void externalizeKeyAttributes(
        ::IIMFLocalToServer::IKeyComponent_ptr & keyComp);
    virtual ::CORBA::Void internalizeData(
        const ::IBOIMLocalToServerMetadata::dataSequence & dataSeq);
    virtual ::CORBA::Boolean verifyKey();

    protected:
    private:
    ::PolicyEmSQLPO iPolicyEmSQLPO;
    ::CORBA::Boolean iKeyValueSet;
};

```

Static SQL data object implementation

The following methods need to be implemented in order for the static SQL data object to work properly in the server.

Framework required method – internalizeFromPrimaryKey: The managed object framework requires the data object to implement the `IManagedServer::IDataObject::internalizeFromPrimaryKey()` method. The *in* parameter is an `IManagedLocal::IPrimaryKey`. This key is how the component is uniquely identified. In the case of the example, the *in* parameter is narrowed to a `PolicyKey` and the policy number is extracted and stored as part of the data object. This method is important because it establishes the linkage between the key and the data object that supports the component being constructed. The `tKeyValueSet` flag should be set to true when valid key values are successfully retrieved from the key.

```

::CORBA::Void PolicyEmSQLD0Impl_Impl::internalizeFromPrimaryKey
    (::IManagedLocal::IPrimaryKey_ptr inKey)
{
    // Insert Method modifications here
    PolicyKey_var iPolicyKey = PolicyKey::_narrow(inKey);
    long iPolicyNoTemp = iPolicyKey->policyNo();
    iKeyValueSet = 1;
    PolicyEmSQLPOKey iPolicyEmSQLPOKey;
    iPolicyEmSQLPOKey.policyNo(iPolicyNoTemp);
    iPolicyEmSQLPO.internalizeFromPrimaryKey(iPolicyEmSQLPOKey);

    // End Method modifications here
}

```

This method is called during construction of objects and during reactivation of objects.

Framework required method – internalizeFromCopyHelper: The managed object framework requires the data object to implement the `IManagedServer::IDataObject::internalizeFromCopyHelper()` method. The *in* parameter is an `IManagedLocal::INonManageable`. This copy contains state data, including the key that shows how the component is uniquely identified. In the case of the example, the *in* parameter is narrowed to a `PolicyCopy` and the data is extracted and stored as part of the data object. In addition, the `tKeyValueSet` flag should be set to true when valid key values are successfully retrieved from the key.

```

::CORBA::Void PolicyEmSQLD0Impl_Impl::internalizeFromCopyHelper(
    ::IManagedLocal::INonManageable_ptr inCopy)
{
    PolicyCopy_var pc = PolicyCopy::_narrow(inCopy);
    tPolicyNo = pc->policyNo();
    // Stores the value of the "policyNo" attribute
    tAmount = pc->amount();
    tPremium = pc -> premium();
    iKeyValueSet = 1;
}

```

This method is called during construction of objects.

Framework required code – create() function: This method is called by the unit test environment or application adaptor when an empty uninitialized data object is required. This happens during component creation and reactivation.

The implementation must have an external entry point so that when the application adaptor loads the DLLs required for a particular component (which happens dynamically), it can access the create() function. The code for the example looks like: See the description in “Framework required code – create() function”.

```
extern "C"
{
    SOMDLLEXPORT IBOIMExtLocalToServer::IDataObject_ptr
    PolicyEmSQLDOImpl_create()
    {
        return new PolicyEmSQLDOImpl_Impl();
    }
}
```

Methods to support attributes – getters: All the attributes need to have a getter method on them. A getter method for one of the attributes of the PolicyDO follows:

```
::CORBA::Float PolicyEmSQLDOImpl_Impl::amount()
{
    // Insert Method modifications here
    ::CORBA::Float iAmountTemp;
    iAmountTemp = iPolicyEmSQLPO.amount();
    return iAmountTemp;

    // End Method modifications here
}
```

Methods to support attributes – setters: All of the attributes need to have a set method on them. A setter method for one of the attributes of the PolicyDO follows:

```
::CORBA::Void PolicyEmSQLDOImpl_Impl::amount(::CORBA::Float amount)
{
    // Insert Method modifications here
    ::CORBA::Float iAmountTemp = amount;
    iPolicyEmSQLPO.amount(iAmountTemp);
    markDirty();

    // End Method modifications here
}
```

One point of interest is the markDirty() method that is run. This indicates to the application adaptor that it is necessary to update the underlying data store at prescribed times.

Additional methods – default constructor: The transient data object implementation has values for each of the attributes that the component is interested in. If the component's business object is caching, the values from the data object are loaded into the business object when the business object methods first require access to state data. In the delegating case, values in the data object are accessed directly, as needed. Either way, when the managed object is originally created, these default constructor values can be used. This ensures that uninitialized data does not get used during the period of time where the key or copy helper is known to the data object, but the data object has not yet been created into or refreshed from the data store. Therefore the values in the data object should be initialized properly in a default constructor. It is required that a default constructor be developed.

```
PolicyEmSQLD0Impl_Impl::PolicyEmSQLD0Impl_Impl ()
: iKeyValueSet(0)
{
}
```

Required method – verifyKey: This method returns a boolean to indicate whether or not the key values in the data object have been acquired and are valid. This method is used by the application adaptor to validate the key before an object reference for the object is built.

```
::CORBA::Boolean PolicyEmSQLD0Impl_Impl::verifyKey ()
{
    //Insert Method modifications here
    return iKeyValueSet;
    //End Method Modifications here
}
```

Required method – externalizeKeyAttributes: This method is used by the application adaptor as an object reference for the component being created. The code in this method writes out the key values into the outKey that is provided on the parameter list.

```
::CORBA::Void PolicyEmSQLD0Impl_Impl::externalizeKeyAttributes
 (::IIMFLocalToServer::IKeyComponent_ptr& keyComp)
{
    // Insert Method modifications here
    long iPolicyNoTemp;
    PolicyEmSQLPOKey iPolicyEmSQLPOKey;
    iPolicyEmSQLPO.externalizeKeyAttributes(iPolicyEmSQLPOKey);
    iPolicyNoTemp = iPolicyEmSQLPOKey.policyNo();
    keyComp->write_long(iPolicyNoTemp);
    // End Method modifications here
}
```

Required method – internalizeKeyAttributes: This method is used by the application adaptor for reactivating objects. This method should pull information out of the inKey provided and set the flag indicating that the key is valid to true. Additional logic may be placed in this method to further verify that the key value is good.

```

::CORBA::Void PolicyEmSQLD0Impl_Impl::internalizeKeyAttributes
    (::IIMFLocalToServer::IKeyComponent_ptr keyComp)
{
    // Insert Method modifications here
    long iPolicyNoTemp = keyComp->read_long();
    iKeyValueSet = 1;
    PolicyEmSQLPOKey iPolicyEmSQLPOKey;
    iPolicyEmSQLPOKey.policyNo(iPolicyNoTemp);
    iPolicyEmSQLPO.internalizeKeyAttributes(iPolicyEmSQLPOKey);
    // End Method modifications here
}

```

Required method – update: This is the code that puts data back into the database. This is .sql code that needs to run through the SQL preprocessor before it is compiled.

```

::CORBA::Void PolicyEmSQLD0Impl_Impl::update ()
{
    // Insert Method modifications here
    iPolicyEmSQLPO.update();
    // End Method modifications here
}

```

Required method – insert: The insert is called when a create is done. The code looks like this:

```

::CORBA::Void PolicyEmSQLD0Impl_Impl::insert ()
{
    // Insert Method modifications here
    iPolicyEmSQLPO.insert();
    // End Method modifications here
}

```

Required method – retrieve: The retrieve gets the data from the database and looks like this:

```

::CORBA::Void PolicyEmSQLD0Impl_Impl::retrieve ()
{
    // Insert Method modifications here
    iPolicyEmSQLPO.retrieve();
    // End Method modifications here
}

```

Required method – del:

```

::CORBA::Void PolicyEmSQLD0Impl_Impl::del ()
{
    // Insert Method modifications here
    iPolicyEmSQLPO.del();
    // End Method modifications here
}

```

Required method – setConnection: This method defines the database to which the SQL in the insert(), del(), retrieve(), and update() methods is directed.

```

::CORBA::Void PolicyEmSQLDOImpl_Impl::setConnection(
    const char* dataBaseName)
{
    // Insert Method modifications here
    iPolicyEmSQLPO.setConnection(dataBaseName);
    // End Method modifications here
}

```

Required method – internalizeData: This method enables query to work on the data object.

```

::CORBA::Void PolicyEmSQLDOImpl_Impl::internalizeData(
    const::IBOIMLocalToServerMetadata::dataSequence & dataSeq)
{
    // Insert Method modifications here
    PolicyEmSQLPOCopy iPolicyEmSQLPOCopy;

    long PolicyEmSQLPO_policyNoTemp;
    if (!((((dataSeq[0])))) >= PolicyEmSQLPO_policyNoTemp))
        PolicyEmSQLPO_policyNoTemp = NULL;
    iPolicyEmSQLPOCopy.policyNo(PolicyEmSQLPO_policyNoTemp);

    double PolicyEmSQLPO_amountTemp;
    if (!((((dataSeq[1])))) >= PolicyEmSQLPO_amountTemp))
        PolicyEmSQLPO_amountTemp = NULL;
    iPolicyEmSQLPOCopy.amount(PolicyEmSQLPO_amountTemp);

    double PolicyEmSQLPO_premiumTemp;
    if (!((((dataSeq[2])))) >= PolicyEmSQLPO_premiumTemp))
        PolicyEmSQLPO_premiumTemp = NULL;
    iPolicyEmSQLPOCopy.premium(PolicyEmSQLPO_premiumTemp);

    iPolicyEmSQLPO.internalizeData(iPolicyEmSQLPOCopy);
    iKeyValueSet = 1;

    // End Method modifications here
}

```

Additional considerations

The static SQL data object could contain the SQL code within the update(), insert(), retrieve(), and delete() methods. It is more convenient to put this code into a C++ helper object that is referred to as the persistent object. The persistent object has many of the same methods as the data object and the data object just deals with the data in the data object and lets the persistent object take care of using SQL.

BOIM data object customization – Cache Service

If the data object customization is going to target Cache Service data objects that run in a DB2 application adaptor, then you should read this section. In this section, the PolicyDO interface is used as the basis to describe the necessary customization.

To create a data object implementation that uses the Cache Service, you must use Object Builder to create the data object. When defining the data object implementation and persistent object, be sure to check the **Use Cache Service** check box.

Managed objects that use the Cache Service must also be configured to use containers with the **Use Cache Service** option.

Cache Service data objects interfaces

The IDL-based picture for Cache Service data objects follows:

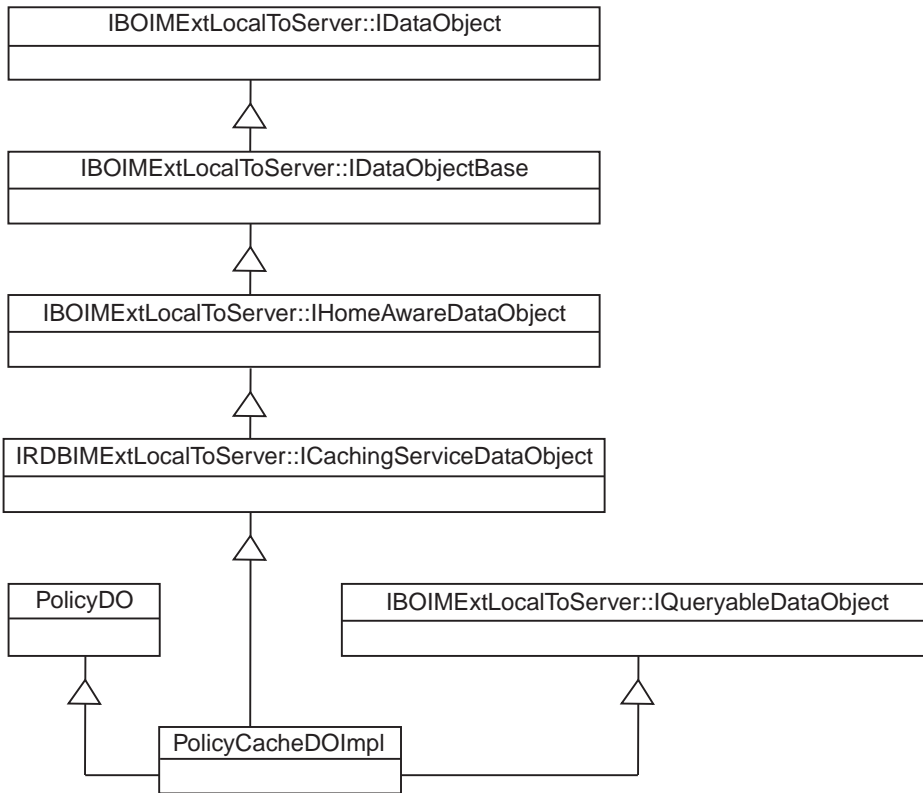


Figure 65. Cache Service data object IDL interface view

The example shows the case where the component is also configured to be queryable. (That is, `IBOIMExtLocalToServer::IQueryableDataObject` should be inherited because the data object is to be used with queryable homes.) The IDL is:

```

#include <PolicyD0.idl>
#include <IRDBIMExtLocalToServer.idl>
#include <IBOIMExtLocalToServer.idl>
interface PolicyCacheDOImpl : PolicyD0,
    IRDBIMExtLocalToServer::ICachingServiceDataObject,
    IBOIMExtLocalToServer::IQueryableDataObject
{
    #pragma meta PolicyCacheDOImpl localonly
};

```

From this, it can be seen that at the IDL level, the difference between the embedded SQL case and the cached DAO case is whether to inherit `IRDBIMExtLocalToServer::IDataObject` or `IRDBIMExtLocalToServer::ICachingServiceDataObject`.

These differences are also apparent from looking at the implementation interfaces and how they interact to provide the basis for the `PolicyCacheDOImpl` class which is used to implement the `PolicyCacheDOImpl` interface shown in the previous figure. The implementation interface inheritance picture follows.

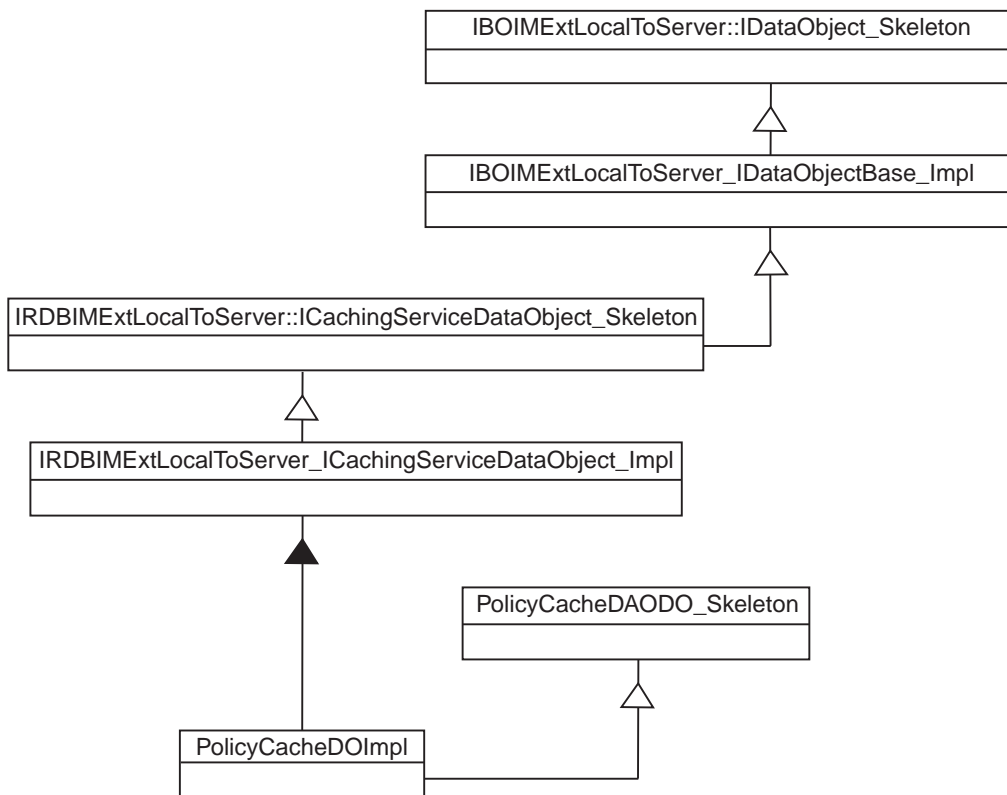


Figure 66. BOIM DO implementation interface inheritance

The `IBOIMExtLocalToServer_IDataObjectBase_Impl` provides an implementation for some of the methods that are described in the `IBOIMExtLocalToServer::IDataObject` interface which is represented in C++ in the previous figure by the `IBOIMExtLocalToServer::IDataObject_Skeleton`. The following methods have implementations which are inherited:

- `string connection()`
- `void connection(string value)`
- `void markDirty()`
- `void clearDirty()`
- `boolean isDirty()`
- `void updateToDataStore()`
- `void insertToDataStore()`
- `void retrieveFromDataStore()`
- `void deleteFromDataStore()`

These methods should not be overridden.

Given this base, the PolicyCacheDOImpl_Impl implementation interface should look like this:

```
#include "PolicyCachePO.hpp"

#ifdefSOMCBNOLOCALINCLUDES
#include <IRDBIMExtLocalToServer.ih>
#include <PolicyCacheDOImpl.hh>
#else
#include "IRDBIMExtLocalToServer.ih"
#include "PolicyCacheDOImpl.hh"
#endif

class PolicyCacheDOImpl_Impl :
public virtual ::PolicyCacheDOImpl_Skeleton,
public virtual IRDBIMExtLocalToServer_ICachingServiceDataObject_Impl
{
public:

//default constructor doimpl
PolicyCacheDOImpl_Impl();

::CORBA::Float amount();
::CORBA::Void amount( ::CORBA::Float amount);
::CORBA::Long policyNo();
::CORBA::Void policyNo( ::CORBA::Long policyNo);
::CORBA::Float premium();
::CORBA::Void premium( ::CORBA::Float premium);

virtual ::CORBA::Void insert();
virtual ::CORBA::Void update();
virtual ::CORBA::Void retrieve();
virtual ::CORBA::Void del();
virtual ::CORBA::Void internalizeFromPrimaryKey(
::IManagedLocal::IPrimaryKey_ptr inKey);
virtual ::CORBA::Void internalizeFromCopyHelper(
::IManagedLocal::INonManageable_ptr inCopy);
virtual ::CORBA::Void internalizeKeyAttributes(
::IMFLocalToServer::IKeyComponent_ptr keyComp);
virtual ::CORBA::Void externalizeKeyAttributes(
::IMFLocalToServer::IKeyComponent_ptr & keyComp);
virtual ::CORBA::Void internalizeData(
const ::IBOIMLocalToServerMetadata::dataSequence & dataSeq);
virtual ::CORBA::Boolean verifyKey();

protected:
private:
::PolicyCachePO iPolicyCachePO;
::CORBA::Boolean iKeyValueSet;
};
```

Most of the methods are the same as the static SQL implementation. Only the methods that differ are described here.

Required method - default constructor:: This method uses a different initialization technique by using the persistent object capabilities of attribute initialization.

```
PolicyTransDOImpl_Impl::PolicyTransDOImpl_Impl()
: iKeyValueSet(0)
{
    policyNo(0);
    amount(0);
    premium(0);
}
```

Required method - internalizeFromPrimaryKey:: This method must handle the exceptions throw by the persistent object and map them to exceptions that are accepted by the framework. Since this method is called for both the create and find scenarios, additional logic is required to either initialize the persistent object or retrieve the already existing values. The isNew() method is used to make this determination.

```
::CORBA::Void PolicyCacheDOImplCache_Impl::internalizeFromPrimaryKey(
::IManagedLocal::IPrimaryKey_ptr inKey)
{
    // Insert Method modifications here
    PolicyKey_var iPolicyKey = PolicyKey::_narrow(inKey);
    long iPolicyNoTemp = iPolicyKey->policyNo();
    iKeyValueSet = 1;
    PolicyCacheDOImplCachePOKey iPolicyCacheDOImplCachePOKey;
    iPolicyCacheDOImplCachePOKey.policyNo(iPolicyNoTemp);
    try
    {
        iPolicyCacheDOImplCachePO.internalizeFromPrimaryKey(
        iPolicyCacheDOImplCachePOKey);
        if (isNew())
        {
            iPolicyCacheDOImplCachePO.create();
            amount(0);
            premium(0);
        }
        else
        {
            iPolicyCacheDOImplCachePO.retrieve();
        }
    }
    catch (IBOIMException::IDataKeyAlreadyExists &dkae)
    {
        throw IManagedClient::IDuplicateKey();
    }
    catch (IBOIMException::IDataKeyNotFound &dknf)
    {
        throw IManagedClient::INoObjectWKey();
    }
    catch (IBOIMException::IDataObjectFailed &dof)
    {
        throw CORBA::PERSIST_STORE(0, CORBA::COMPLETED_NO);
    }
}
```

```

    catch (...)
    {
        throw; // home will handle
    }
    // End Method modifications here
}

```

Required method - `internalizeFromCopyHelper`:: This method must handle the exceptions throw by the persistent object and map them to exceptions that are accepted by the framework.

```

::CORBA::Void PolicyCacheDOImplCache_Impl::internalizeFromCopyHelper(
::IManagedLocal::INonManageable_ptr inCopy)
{ // Insert Method modifications here
    PolicyCopy_var iPolicyCopy = PolicyCopy::_narrow(inCopy);
    long iPolicyNoTemp = iPolicyCopy->policyNo();
    float iAmountTemp = iPolicyCopy->amount();
    float iPremiumTemp = iPolicyCopy->premium();
    iKeyValueSet = 1;
    PolicyCacheDOImplCachePOCopy iPolicyCacheDOImplCachePOCopy;
    iPolicyCacheDOImplCachePOCopy.policyNo(iPolicyNoTemp);
    iPolicyCacheDOImplCachePOCopy.amount(iAmountTemp);
    iPolicyCacheDOImplCachePOCopy.premium(iPremiumTemp);
    try
    {
        iPolicyCacheDOImplCachePO.internalizeFromCopyHelper(
        iPolicyCacheDOImplCachePOCopy);
    }
    catch (IBOIMException::IDataKeyAlreadyExists &dkae)
    {
        throw IManagedClient::IDuplicateKey();
    }
    catch (IBOIMException::IDataKeyNotFound &dknf)
    {
        throw IManagedClient::INoObjectWKey();
    }
    catch (IBOIMException::IDataObjectFailed &dof)
    {
        throw CORBA::PERSIST_STORE(0, CORBA::COMPLETED_NO);
    }
    catch (...)
    {
        throw; // home will handle
    }
    // End Method modifications here
}

```

Required method - `getters`:: Many of the methods are required to handle exceptions that may be thrown by the persistent object and mapping them to exceptions that are supported by the calling framework. For getters the exception that is used as the CORBA standard exception `PERSIST_STORE`.

```

::CORBA::Long PolicyCacheDOImplCache_Impl::amount()
{
    // Insert Method modifications here
}

```

```

        ::CORBA::Float iAmountTemp;
    try
    {
        iAmountTemp = iPolicyCachePO.id();
    }
    catch (CORBA::UserException &cue)
    {
        throw CORBA::PERSIST_STORE(0, CORBA::COMPLETED_NO);
    }
    return iAmountTemp;
    // End Method modifications here
}

```

Required method - setters: Many of the methods are required to handle exceptions that may be thrown by the persistent object and mapping them to exceptions that are supported by the calling framework. For getters the exception that is used as the CORBA standard exception PERSIST_STORE.

```

::CORBA::Void PolicyCacheDOImplCache_Impl::amount( ::CORBA::Float amount)
{
    // Insert Method modifications here
    ::CORBA::Float iAmountTemp = amount;
    try
    {
        iPolicyCache.amount(iAmountTemp);
    }
    catch (CORBA::UserException &cue)
    {
        throw CORBA::PERSIST_STORE(0, CORBA::COMPLETED_NO);
    }
    markDirty();
    // End Method modifications here
}

```

Required method - internalizeKeyAttributes: This method must also handle the exceptions throw by the persistent object and map them to exceptions that are accepted by the framework.

```

::CORBA::Void PolicyCacheDOImplCache_Impl::internalizeKeyAttributes(
::IIMFLocalToServer::IKeyComponent_ptr keyComp)
{
    // Insert Method modifications here
    iKeyValueSet = 1;
    try
    {
        iPolicyCacheDOImplCachePO.internalizeKeyAttributes(keyComp);
    }
    catch (IBOIMException::IDataKeyAlreadyExists &dkae)
    {
        throw IManagedClient::IDuplicateKey();
    }
    catch (IBOIMException::IDataKeyNotFound &dknf)
    {
        throw IManagedClient::INoObjectWKey();
    }
}

```

```

        catch (IBOIMException::IDataObjectFailed &dof)
        {
            throw CORBA::PERSIST_STORE(0, CORBA::COMPLETED_NO);
        }
        catch (...)
        {
            throw; // home will handle
        }
        // End Method modifications here
    }
}

```

Required method - externalizeKeyAttributes:: This method must also handle the exceptions throw by the persistent object and map them to exceptions that are accepted by the framework.

```

::CORBA::Void PolicyCacheDOImplCache_Impl::externalizeKeyAttributes(
::IIMFLocalToServer::IKeyComponent_ptr & keyComp)
{
    // Insert Method modifications here
    try
    {
        iPolicyNoDOImplCachePO.externalizeKeyAttributes(keyComp);
    }
    catch (CORBA::UserException &cue)
    {
        throw IBOIMException::IExternalizeKeyAttributesFailed();
    }
    // End Method modifications here
}

```

Required method - internalizeData:: This method must also handle the exceptions throw by the persistent object and map them to exceptions that are accepted by the framework but is also much simpler due to the support provided by the persistent object.

```

::CORBA::Void PolicyCacheDOImplCache_Impl::internalizeData(
const ::IBOIMLocalToServerMetadata::dataSequence & dataSeq)
{
    // Insert Method modifications here
    iPolicyCacheDOImplCachePO.internalizeData(dataSeq[0]);
    iKeyValueSet = 1;
    // End Method modifications here
}

```

Transient data object customization – UUID key (production use)

This section includes the following topics:

- “Interfaces” on page 380
- “Implementation” on page 380
- “Additional considerations” on page 381

Interfaces

The interface view follows. The `IManagedAdvancedServer::IUUIDDataObject` is the interface that should be supported for data object implementations using UUID keys for transient objects.

graphic to be completed

The implementation interface follows.

graphic to be completed

Implementation

This section includes the following topics:

- “Framework required method – `internalizeFromPrimaryKey`”
- “Framework required method – `internalizeFromCopyHelper`”
- “Framework required code – `create()` function”
- “Methods to support attributes – getters” on page 381
- “Methods to support attributes – setters” on page 381
- “Additional methods – default constructor” on page 381
- “Required method – `verifyKey`” on page 381
- “Required method – `externalizeKeyAttributes`” on page 381
- “Required method – `internalizeKeyAttributes`” on page 381
- “Required method – `update`” on page 381
- “Required method – `insert`” on page 381
- “Required method – `retrieve`” on page 381
- “Required method – `del`” on page 381

Framework required method – `internalizeFromPrimaryKey`: This method should not be implemented. It is implemented by the framework. It works with a `IManagedAdvancedServer::IUUIDPrimaryKey`.

Framework required method – `internalizeFromCopyHelper`: This method must be implemented. It should copy its attributes and call its parent’s `internalizeFromCopyHelper` method which copy the UUID primary key value.

Framework required code – `create()` function: See the description in “Framework required code – `create()` function” on page 368. It is the same for UUID-based data objects.

Methods to support attributes – getters: See the description in “Methods to support attributes – getters” on page 368.

Methods to support attributes – setters: See the description in “Methods to support attributes – setters” on page 368.

Additional methods – default constructor: See the description in “Required method - default constructor:” on page 376.

Required method – verifyKey: This method should not be implemented. The framework takes care of it.

Required method – externalizeKeyAttributes: This method should not be implemented. The framework takes care of it.

Required method – internalizeKeyAttributes: This method should not be implemented. The framework takes care of it.

Required method – update: This method should not be implemented. The framework takes care of it.

Required method – insert: This method should not be implemented. The framework takes care of it.

Required method – retrieve: This method should not be implemented. The framework takes care of it.

Required method – del: This method should not be implemented. The framework takes care of it.

Additional considerations

Using the UUIDO and UUIDKey limits the Programming Model for the components. FindByPrimaryKeyString is not meaningful because the UUIDKey does not contain business logic attributes. UUID support is most useful when there are short-lived components that do not need to be found after they are created.

Transient data object – any key (production use)

If the data object customization is targeting transient data objects that run in a BOIM application adaptor and leverage business object information for making up the key, then you should read this section. In this section, the PolicyDO interface is used as the basis to describe the necessary customization.

BOIM data object interfaces

The IDL-based picture for BOIM data objects follows:

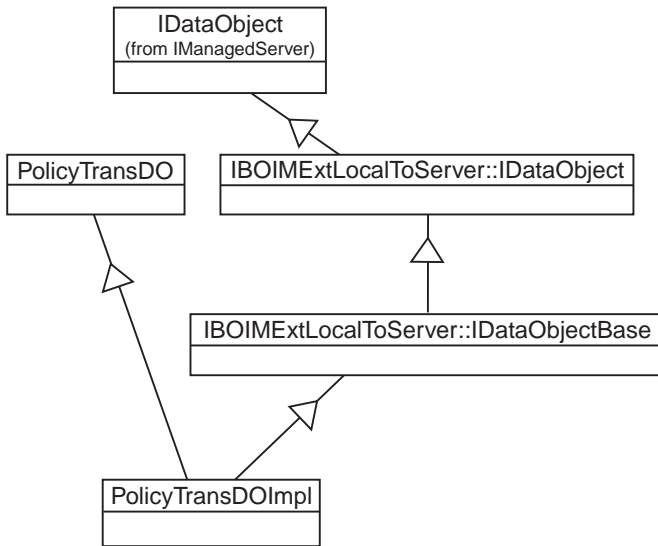


Figure 67. BOIM data object IDL interface view

The example does not show objects capable of being included in queries. The PolicyDO_Impl class implements the PolicyDO interface shown in the previous figure. The implementation interface inheritance picture follows.

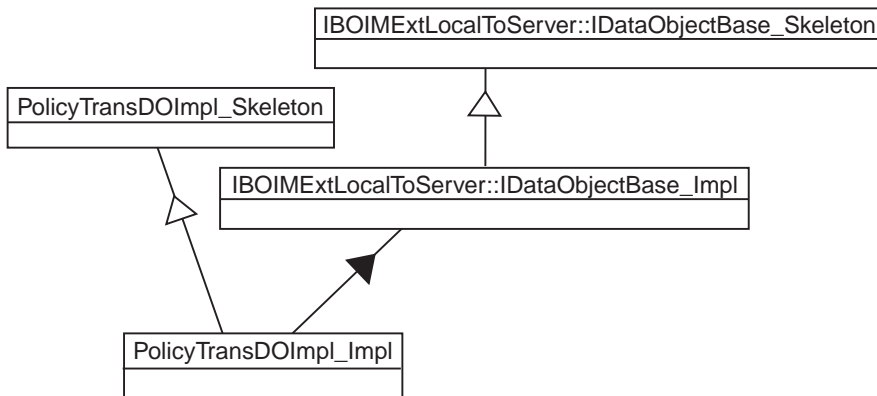


Figure 68. BOIM data object implementation interface inheritance

The `IBOIMExtLocalToServer::IDataObjectBase_Impl` provides an implementation for some of the methods that are described in the `IBOIMExtLocalToServer::IDataObject` interface. The following methods have implementations which are inherited:

- `string connection()`
- `void connection(string value)`
- `void markDirty()`
- `void clearDirty()`
- `boolean isDirty()`
- `void updateToDataStore()`
- `void insertToDataStore()`
- `void retrieveFromDataStore()`
- `void deleteFromDataStore()`
- `initDO()`

These methods should not be overridden.

Given this base, the `PolicyTransDO_Impl` implementation interface should look like this:

```
#ifdef SOMCBNOLocalINCLUDES
#include <IBOIMExtLocalToServer>
#include <PolicyTransDOImpl.hh>
#else
#include "IBOIMExtLocalToServer.ih"
#include "PolicyTransDOImpl.hh"
#endif

class PolicyTransDOImpl_Impl :
public virtual ::PolicyTransDOImpl_Skeleton,
public virtual IBOIMExtLocalToServer_IDataObjectBase_Impl
{
public:

//default constructor doimpl
PolicyTransDOImpl_Impl();

::CORBA::Float amount();
::CORBA::Void amount( ::CORBA::Float amount);
::CORBA::Long policyNo();
::CORBA::Void policyNo( ::CORBA::Long policyNo);
::CORBA::Float premium();
::CORBA::Void premium( ::CORBA::Float premium);

virtual ::CORBA::Void insert();
virtual ::CORBA::Void update();
virtual ::CORBA::Void retrieve();
virtual ::CORBA::Void del();
virtual ::CORBA::Void internalizeFromPrimaryKey(
::IManagedLocal::IPrimaryKey_ptr inKey);
virtual ::CORBA::Void internalizeFromCopyHelper(
::IManagedLocal::INonManageable_ptr inCopy);
```

```

virtual ::CORBA::Void internalizeKeyAttributes(
    ::IIMFLocalToServer::IKeyComponent_ptr keyComp);
virtual ::CORBA::Void externalizeKeyAttributes(
    ::IIMFLocalToServer::IKeyComponent_ptr & keyComp);
virtual ::CORBA::Boolean verifyKey();

protected:
private:
    ::CORBA::Float iAmount;
    ::CORBA::Long iPolicyNo;
    ::CORBA::Float iPremium;
    ::CORBA::Boolean iKeyValueSet;
};

```

BOIM data object implementation

The following methods need to be implemented in order for the BOIM data object to work properly on the server.

Framework required method – internalizeFromPrimaryKey: See the description in “Framework required method – internalizeFromPrimaryKey” on page 367. It is the same for BOIM-based data objects.

In addition, the `keyValueSet` flag should be set to true when valid key values are successfully retrieved from the key.

Framework required method – internalizeFromCopyHelper: See the description in “Framework required method – internalizeFromCopyHelper” on page 367. It is the same for BOIM-based data objects.

In addition, the `iKeyValueSet` flag should be set to true when valid key values are successfully retrieved from the key.

Framework required code – create() Function: See the description in “Framework required code – create() function” on page 368. It is the same for BOIM-based data objects.

Methods to support attributes – getters: See the description in “Methods to support attributes – getters” on page 368. It is the same for BOIM-based data objects.

Methods to support attributes – setters: All of the attributes need to have a set method on them. A setter method for one of the attributes of the PolicyDO follows.

```

::CORBA::Void PolicyTransDOImpl_Impl::amount( ::CORBA::Float amount)
{
    // Insert Method modifications here
}

```

```

        iAmount = amount;
        markDirty();
        // End Method modifications here
    }

```

The difference from the transient case is where the `markDirty()` method is run. This indicates to the application adaptor that it is necessary to update the underlying datastore at prescribed times.

Additional methods – default constructor: The default constructor needs to initialize the attributes and synchronize the dataobject with the persistent object. This is best done by calling the setters for these attributes. However, a side effect of calling the setters is that the 'dirty' bit is set. Leaving this bit set would cause inappropriate updates to the datastore when the business object is passivated. To solve this problem the `clearDirty()` method is called at the end of the method. Also, the `iKeyValueSet` attribute is set to a value that indicates the primary key attributes have not yet been set.

```

PolicyEmSQLDOImpl_Impl::PolicyEmSQLDOImpl_Impl ()
:iKeyValueSet(0)
{
    policyNo(0);
    amount(0);
    premium(0);
    clearDirty();
}

```

Required method – verifyKey: This method returns a boolean to indicate whether or not the key values in the data object have been acquired and are valid. This method is used by the application adaptor to validate the key before it builds reference data for the object.

```

::CORBA::Boolean PolicyTransDOImpl_Impl::verifyKey()
{
    // Insert Method modifications here
    return iKeyValueSet;
    // End Method modifications here
}

```

Required method – externalizeKeyAttributes: This method is used by the application adaptor as it creates reference data for the object reference that is being created. The code in this method writes out the key values into the outKey that is provided on the parameter list.

```

::CORBA::Void
PolicyTransDOImpl_Impl::externalizeKeyAttributes(
    ::IIMFLocalToServer::IKeyComponent_ptr & keyComp)
{
    // Insert Method modifications here
    keyComp->write_long(iPolicyNo);
    // End Method modifications here
}

```

Required method – internalizeKeyAttributes: This method is used by the application adaptor for reactivating objects. This method should pull information out of the `inKey` provided and set the flag indicating that the key is valid to true. Additional logic may be placed in this method to further verify that the key value is good.

```
::CORBA::Void
PolicyTransDOImpl_Impl::internalizeKeyAttributes(
    ::IIMFLocalToServer::IKeyComponent_ptr keyComp)
{
    // Insert Method modifications here
    iPolicyNo = keyComp->read_long();
    iKeyValueSet = 1;
    // End Method modifications here
}
```

Required methods – del, insert, retrieve, and update: For true transient objects, the `del()`, `insert()`, `retrieve()`, and `update()` methods are empty. For situations where the transient object adaptor is being used to manage the object references, but not the persistent state data, these methods contain code that removes, creates, retrieves, or updates a new entry in the persistent store.

Summary of DataObject customization

Table 9 on page 387 enumerates the general strategy for each data object method or data object method type as it applies to each of the kinds of data object customization that are possible using Component Broker.

Table 9. Summary of data object customization methods.

Method	Persistent data object static SQL	Persistent Cache Service data object Cache	Transient (production) - UUID	Transient - Any Key (production)
IManagedServer::IDataObject::internalizeFromPrimary	Required, uses local cache	Required, talks to Cache Service data object	Implemented by framework	Required, uses local cache
IManagedServer::IDataObject::internalizeFromCopyHelper	Required, uses local cache	Required, talks to Cache Service data object	Required	Required, uses local cache
_create	Required	Required	Required	Required
Attribute getters	Required, uses local cache	Required, talks to Cache Service data object	Required	Required, uses local cache
Attribute setters	Required, uses markDirty(), uses local cache	Required, talks to Cache Service data object	Required, uses markDirty(), uses local cache	Required, uses markDirty(), uses local cache
Default constructor	Required	Required	Required	Required
IBOIMExtLocalToServer::IDataObject::connection()	Implement by framework	Implemented by framework	Implemented by framework	Implemented by framework
IBOIMExtLocalToServer::IDataObject::connection(string)	Implemented by framework	Implemented by framework	Implemented by framework	Implemented by framework
IBOIMExtLocalToServer::IDataObject::markDirty()	Implemented by framework	Implemented by framework	Implemented by framework	Implemented by framework
IBOIMExtLocalToServer::IDataObject::isDirty()	Implemented by framework	Implemented by framework	Implemented by framework	Implemented by framework
IBOIMExtLocalToServer::IDataObject::updateToDataStore()	Implemented by framework	Implemented by framework	Implemented by framework	Implemented by framework
IBOIMExtLocalToServer::IDataObject::insertToDataStore	Implemented by framework	Implemented by framework	Implemented by framework	Implemented by framework
IBOIMExtLocalToServer::IDataObject::retrieveFromDataStore	Implemented by framework	Implemented by framework	Implemented by framework	Implemented by framework

Table 9. Summary of data object customization methods. (continued)

Method	Persistent data object static SQL	Persistent Cache Service data object Cache	Transient (production) - UUID	Transient - Any Key (production)
IBOIMExtLocalToServer::IDataObject::deleteFromDataStore	Implemented by framework	Implemented by framework	Implemented by framework	Implemented by framework
IBOIMExtLocalToServer::IDataObject::verifyKey()	Required	Required	Implemented by framework	Required
IBOIMExtLocalToServer::IDataObject::externalizeKeyAttributes	Required	Required, talks to Cache Service data object	Implemented by framework	Required
IBOIMExtLocalToServer::IDataObject::internalizeKeyAttributes	Required	Required, talks to Cache Service data object	Implemented by framework	Required
IBOIMExtLocalToServer::IDataObject::update()	Required, has SQL	Required, null implementation Cache Service data object handles	Implemented by framework	Required
IBOIMExtLocalToServer::IDataObject::insert()	Required, has SQL	Required, null implementation Cache Service data object handles	Implemented by framework	Required
IBOIMExtLocalToServer::IDataObject::retrieve()	Required, has SQL	Required, null implementation	Implemented by framework	Required
IBOIMExtLocalToServer::IDataObject::del()	Required, has SQL	Required, does markDelete on Cache Service data object	Implemented by framework	Required
IRDBIMExtLocalToServer::setConnection	Required	n/a	n/a	n/a
IBOIMExtLocalToServer::IQueryableDataObject::internalizeData()	Optional	Optional	Optional	Optional
IBOIMExtLocalToServer::IDataObject::initDO()	Optional	Optional	Optional	Optional

Data object data management patterns

In the managed object framework, decisions are made about how to deal with the data object. This is done when choosing between `IManagedServer::IManagedObjectWithCachedDataObject` and `IManagedServer::IManagedObjectWithDataObject`. Programming style influences the choice. Performance and other design considerations also come into play.

Data objects, like business objects, also do a level of “caching”. This is not a programming style issue. It is related to the component’s underlying store, or lack thereof, for which the data object is the abstraction. Transient objects definitely need to have some cache in the data object as that is the only way that they can work correctly. Data objects that are abstractions over persistent storage, however, have choices to make about the caching pattern that is chosen.

Data object customization and inheritance

It is necessary to inherit the parent data object interface. It is not necessary and may not even be desirable to inherit implementation of the base class data object. However, if you are inheriting a data object implementation, then parent methods need to be called for `internalizeFromPrimaryKey`, `internalizeFromCopyHelper`, `verifyKey`, and `externalizeKeyAttributes`. Call the parent method first, before executing the subclass function. Other methods may also need to call parent methods. This varies based upon the kind of data object being used.

The following example assumes that the `PolicyEmSQLDOImpl` interface is inherited. (See “Transient data object – any key (production use)” on page 381 for details on the BOIM data object type of data object customization.)

CarPolicy BOIM data object interfaces

The IDL-based picture for the CarPolicy BOIM data object follows:




Figure 69. CarPolicy BOIM data object IDL interface view

The PolicyEmSQLDOImpl_Impl class implements the policy interface as shown in Figure 69 on page 389. The implementation interface inheritance is shown in Figure 70.

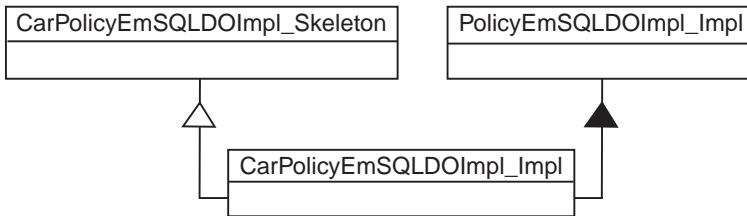


Figure 70. BOIM DO implementation interface inheritance

Given this base, the CarPolicyEmSQLDOImpl_Impl implementation interface should look like this:

```

#include "CarPolicyEmSQLPO.hpp"

#ifdef SOMCBNOLOCALINCLUDES
#include <PolicyEmSQLDOImpl.ih>
#include <CPEmSQLDOImpl.hh>
#else
#include "PolicyEmSQLDOImpl.ih"
#include "CPEmSQLDOImpl.hh"
#endif

class CarPolicyEmSQLDOImpl_Impl : public virtual::CarPolicyEmSQLDOImpl_Skeleton,public virtual PolicyEmSQLDOImpl_Impl
{
public:
    //default constructor doimpl
    CarPolicyEmSQLDOImpl_Impl();

    ::CORBA::Long year();
    ::CORBA::Void year( ::CORBA::Long year);
    char* make();
    ::CORBA::Void make(const char* make);
    char* model();
    ::CORBA::Void model(const char* model);
    ::CORBA::Long serialNumber();
    ::CORBA::Void serialNumber( ::CORBA::Long serialNumber);
    ::CORBA::Float collisionDeductible();
    ::CORBA::Void collisionDeductible( ::CORBA::Float collisionDeductible);
    ::CORBA::Boolean glassCoverage();
    ::CORBA::Void glassCoverage( ::CORBA::Boolean glassCoverage);
    ::CORBA::Long policyNo();
    ::CORBA::Void policyNo( ::CORBA::Long policyNo);

    virtual ::CORBA::Void insert();
    virtual ::CORBA::Void update();
    virtual ::CORBA::Void retrieve();
    virtual ::CORBA::Void del();
  
```



```

virtual ::CORBA::Void setConnection(const char* dataBaseName);
virtual ::CORBA::Void internalizeFromPrimaryKey(
    ::IManagedLocal::IPrimaryKey_ptr inKey);
virtual ::CORBA::Void internalizeFromCopyHelper(
    ::IManagedLocal::INonManageable_ptr inCopy);
virtual ::CORBA::Void internalizeKeyAttributes(
    ::IMFLocalToServer::IKeyComponent_ptr keyComp);
virtual ::CORBA::Void externalizeKeyAttributes(
    ::IMFLocalToServer::IKeyComponent_ptr & keyComp);
virtual ::CORBA::Void internalizeData(
    const ::IBOIMLocalToServerMetadata::dataSequence & dataSeq);
virtual ::CORBA::Boolean verifyKey();

protected:
    ::CarPolicyEmSQLPO iCarPolicyEmSQLPO;

private:
    ::CORBA::Boolean iKeyValueSet;
};

```

CarPolicy BOIM data object implementation

The following methods need to be implemented for the data object to work properly on the server.

Framework required method – internalizeFromPrimaryKey: The *policyNo* attribute needs to be internalized. The *keyValueSet* flag should be set to true when valid key values are successfully retrieved for both attributes.

```

::CORBA::Void
CarPolicyEmSQLDOImpl_Impl::internalizeFromPrimaryKey(
    ::IManagedLocal::IPrimaryKey_ptr inKey)
{
    // Insert Method modifications here
    PolicyKey_var iPolicyKey = PolicyKey::_narrow(inKey);
    long iPolicyNoTemp = iPolicyKey->policyNo();
    iKeyValueSet = 1;
    PolicyEmSQLPOKey iPolicyEmSQLPOKey;
    CarPolicyEmSQLPOKey iCarPolicyEmSQLPOKey;
    iPolicyEmSQLPOKey.policyNo(iPolicyNoTemp);
    iCarPolicyEmSQLPOKey.policyNo(iPolicyNoTemp);

    iPolicyEmSQLPO.internalizeFromPrimaryKey(iPolicyEmSQLPOKey);
    iCarPolicyEmSQLPO.internalizeFromPrimaryKey(iCarPolicyEmSQLPOKey);

    // End Method modifications here
}

```

Framework required method – internalizeFromCopyHelper: The framework required method *internalizeFromCopyHelper()* should first be called using *PolicyEmSQLDOImpl_Impl*. Then copy the attributes from the copy helper

object to those stored in the CarPolicyEmSQLDOImpl_Impl. In addition, the keyValueSet flag should be set to true when valid key values are successfully retrieved from the key.

```

::CORBA::Void
CarPolicyEmSQLDOImpl_Impl::internalizeFromCopyHelper(
    ::IManagedLocal::INonManageable_ptr inCopy)
{
    // Insert Method modifications here

    CarPolicyCopy_var iCarPolicyCopy = CarPolicyCopy::_narrow(inCopy);
    long iYearTemp = iCarPolicyCopy->year();
    ::CORBA::String_var iMakeTemp = iCarPolicyCopy->make();
    ::CORBA::String_var iModelTemp = iCarPolicyCopy->model();
    long iSerialNumberTemp = iCarPolicyCopy->serialNumber();
    float iCollisionDeductibleTemp = iCarPolicyCopy->collisionDeductible();
    ::CORBA::Boolean iGlassCoverageTemp = iCarPolicyCopy->glassCoverage();
    float iAmountTemp = iCarPolicyCopy->amount();
    long iPolicyNoTemp = iCarPolicyCopy->policyNo();
    float iPremiumTemp = iCarPolicyCopy->premium();
    iKeyValueSet = 1;
    PolicyEmSQLPOCopy iPolicyEmSQLPOCopy;
    CarPolicyEmSQLPOCopy iCarPolicyEmSQLPOCopy;
    iPolicyEmSQLPOCopy.amount(iAmountTemp);
    iPolicyEmSQLPOCopy.policyNo(iPolicyNoTemp);
    iCarPolicyEmSQLPOCopy.policyNo(iPolicyNoTemp);
    iPolicyEmSQLPOCopy.premium(iPremiumTemp);
    iCarPolicyEmSQLPOCopy.year(iYearTemp);
    {
        DB2VarCharMap maketemp;
        CARPOLICYEMSQLPO_MAKEDB2VARCHAR maketemp2;
        maketemp.maxLen = 2000;
        maketemp.dataP = maketemp2.data;
        if (iMakeTemp)
            DB2MappingHelper::stringToVarChar(iMakeTemp, maketemp);
        else
            maketemp.length = 0;
        maketemp2.length = maketemp.length;
        iCarPolicyEmSQLPOCopy.make(maketemp2);
    }
    {
        DB2VarCharMap modeltemp;
        CARPOLICYEMSQLPO_MODELDB2VARCHAR modeltemp2;
        modeltemp.maxLen = 2000;
        modeltemp.dataP = modeltemp2.data;
        if (iModelTemp)
            DB2MappingHelper::stringToVarChar(iModelTemp, modeltemp);
        else
            modeltemp.length = 0;
        modeltemp2.length = modeltemp.length;
        iCarPolicyEmSQLPOCopy.model(modeltemp2);
    }
    iCarPolicyEmSQLPOCopy.serialNumber(iSerialNumberTemp);
    iCarPolicyEmSQLPOCopy.collisionDeductible(iCollisionDeductibleTemp);
    iCarPolicyEmSQLPOCopy.glassCoverage(iGlassCoverageTemp);
}

```

```

        iPolicyEmSQLPO.internalizeFromCopyHelper(iPolicyEmSQLPOCopy);
        iCarPolicyEmSQLPO.internalizeFromCopyHelper(iCarPolicyEmSQLPOCopy);

        // End Method modifications here
    }

```

Framework required code – create() function: See the description in “Framework required code – create() function” on page 368. It is the same for this data object.

Methods to support attributes – getters: All of the attributes need to have a getter method for them. Methods that support getting attributes work in the usual way. For example, the getter for make would be:

```

char* CarPolicyEmSQLD0Impl_Impl::make()
{
    // Insert Method modifications here
    ::CORBA::String_var iMakeTemp;
    {
        DB2VarCharMap maketempR;
        maketempR.maxLen = 2000;
        maketempR.length = iCarPolicyEmSQLPO.make().length;
        maketempR.dataP = (char *)iCarPolicyEmSQLPO.make().data;
        iMakeTemp = new char[2001];
        DB2MappingHelper::varCharToString(maketempR, iMakeTemp);
    }
    return CORBA::string_dup(iMakeTemp);
    // End Method modifications here
}

```

Methods to support attributes – setters: All of the attributes need to have a setter method for them. Methods that support setting attributes work in the usual way. For example, the setter for make would be:

```

::CORBA::Void CarPolicy_Impl::make(const char* make)
{
    fMake = make;
    markDirty();
}

```

Additional methods – default constructor: The implementation for this method is similar to other data objects.

```

CarPolicyEmSQLD0Impl_Impl::CarPolicyEmSQLD0Impl_Impl()
:iKeyValueSet(0)
{
    year(0);
    make(CORBA::string_dup(""));
    model(CORBA::string_dup(""));
    serialNumber(0);
    collisionDeductible(0);
    policyNo(0);
    clearDirty();
}

```

Required method – externalizeKeyAttributes: The implementation for this method is similar to other data objects, except that it first calls `externalizeKeyAttributes` on the parent data object.

```

::CORBA::Void
CarPolicyEmSQLDQImpl_Impl::externalizeKeyAttributes(
    ::IIMFLocalToServer::IKeyComponent_ptr & keyComp)
{
    // Insert Method modifications here
    long iPolicyNoTemp;
    PolicyEmSQLPOKey iPolicyEmSQLPOKey;
    CarPolicyEmSQLPOKey iCarPolicyEmSQLPOKey;
    iPolicyEmSQLPO.externalizeKeyAttributes(iPolicyEmSQLPOKey);
    iCarPolicyEmSQLPO.externalizeKeyAttributes(iCarPolicyEmSQLPOKey);
    iPolicyNoTemp = iPolicyEmSQLPOKey.policyNo();
    iPolicyNoTemp = iCarPolicyEmSQLPOKey.policyNo();
    keyComp->write_long(iPolicyNoTemp);
    // End Method modifications here
}

```

Required method – internalizeKeyAttributes: The implementation for this method is similar to other data objects, except that it first calls `internalizeKeyAttributes` on the parent data object. In addition, all keys need to be verified that the values are good.

```

::CORBA::Void
CarPolicyEmSQLDQImpl_Impl::internalizeKeyAttributes(
    ::IIMFLocalToServer::IKeyComponent_ptr keyComp)
{
    // Insert Method modifications here
    long iPolicyNoTemp = keyComp->read_long();
    iKeyValueSet = 1;
    PolicyEmSQLPOKey iPolicyEmSQLPOKey;
    CarPolicyEmSQLPOKey iCarPolicyEmSQLPOKey;
    iPolicyEmSQLPOKey.policyNo(iPolicyNoTemp);
    iCarPolicyEmSQLPOKey.policyNo(iPolicyNoTemp);
    iPolicyEmSQLPO.internalizeKeyAttributes(iPolicyEmSQLPOKey);
    iCarPolicyEmSQLPO.internalizeKeyAttributes(iCarPolicyEmSQLPOKey);
    // End Method modifications here
}

```

Required method —del: The `del()` method is called when the object is removed. Both of the persistent objects need to be called in this method.

```

::CORBA::Void CarPolicyEmSQLDQImpl_Impl::del()
{
    // Insert Method modifications here
    iPolicyEmSQLPO.del();
    iCarPolicyEmSQLPO.del();
    // End Method modifications here
}

```

Required method —insert: The `insert()` method is called when a create is done. Both of the persistent objects need to be called in this method.

```

::CORBA::Void CarPolicyEmSQLD0Impl_Impl::insert()
{
    // Insert Method modifications here
    iPolicyEmSQLPO.insert();
    iCarPolicyEmSQLPO.insert();
    // End Method modifications here
}

```

Required method —retrieve: The `retrieve()` method gets data from the database. Both of the persistent objects need to be called in this method.

```

::CORBA::Void CarPolicyEmSQLD0Impl_Impl::retrieve()
{
    // Insert Method modifications here
    iPolicyEmSQLPO.retrieve();
    iCarPolicyEmSQLPO.retrieve();
    // End Method modifications here
}

```

Required method —update: The `update()` method is called when data is put into the database. Both of the persistent objects need to be called in this method.

```

::CORBA::Void CarPolicyEmSQLD0Impl_Impl::update()
{
    // Insert Method modifications here
    iPolicyEmSQLPO.update();
    iCarPolicyEmSQLPO.update();
    // End Method modifications here
}

```

Required method —setConnection: The `setConnection()` method is called when initializing connections. Both of the persistent objects need to be called in this method.

```

::CORBA::Void CarPolicyEmSQLD0Impl_Impl::setConnection(
    const char* dataBaseName)
{
    // Insert Method modifications here
    iPolicyEmSQLPO.setConnection(dataBaseName);
    iCarPolicyEmSQLPO.setConnection(dataBaseName);
    // End Method modifications here
}

```

Data object customization for cardinality relations

The data object implementation for a component that contains a reference to another component requires getters and setters that are more complex than those that implement attributes with simple mappings to backend resource managers such as SQL tables. This section describes the reference mapping patterns, what the purpose of each pattern is, and how best to apply the patterns to specific relationship situations.

Top-down versus bottom-up relations

The component's business object interface defines references, but does not know how the relationship is implemented. This is a key part of the encapsulation provided by the managed object framework. In this section, implementation strategies and how they appear on the data object implementation are introduced.

There are two basic strategies that are applied to implementing object references that appear in the business object:

Top-down

This approach allows alteration or definition of a database schema that underlies a component that has references to other components.

Bottom-up

This approach implies preservation of an existing schema.

While there are times when the object resolution approach characterized as bottom-up can be used for new top-down applications, the inverse is not true. The two strategies apply equally well to cardinality-1 and cardinality-n types of relationships.

Top-down customizations

The top-down approach is characterized by the presence of a string in the database table of the containing object. This string contains information which the data object can use to produce the object reference required by the business object interface. Data object getter methods take on the general form of:

```
Claim_ptr PolicyDO_Impl::currentClaim()
{
    // retrieve the stored value for persistent store
    ...
    // convert the retrieved value into a pointer and return it
    ...
}
```

The Data object setter methods take on the general form of:

```
::CORBA::Void PolicyDO_Impl::currentClaim(Claim_ptr claim)
{
    // map the claim pointer into a storable form
    ...
    // invoke and/or notify the application adaptor regarding
    // the change to persistent data
    ...
}
```

The implementation of these methods is variable in the following dimensions:

- Mapping or conversion design pattern dimension
- Persistent storage dimension

CORBA provides conversion interfaces that assist in the mapping. These interfaces are `string_to_object` and `object_to_string`. The string form of an object reference (often called a stringified object reference, or SOR), must stand alone and may be quite large, because it must contain enough information to locate and materialize a remote object. It is, however, a simple way to get the string representation of an object that can then be stored persistently and later used to bring an object reference back to life.

Conceptually, storing a stringified reference sounds reasonable. However, more efficient mechanisms with dimensions that possess different levels of robustness are also possible. These patterns take advantage of abstractions introduced by the managed object framework. The patterns are generally preferable to SORs for the following reasons:

- They require less storage to implement.
- They are based on the applications object model, instead of the objects physical location in the network, making them more robust and allowing them to be maintained more easily than SORs. For example, moving a home from one container to another would not break the home/key reference (this pattern is described later), but would break an outstanding SOR.

There are many different combinations of CORBA and Component Broker abstractions that could be used to map object references. The following useful patterns have been identified by the programming model:

Stringified Object Reference

Stores the stringified object reference as a variable length string. It is the simplest form in structure and the largest in size. Multiple references to the same object or other objects in the same application adaptor environment store duplicate prefix data.

Object Name

Stores a Naming Service name of an object as a variable length string. This string is much shorter than the SOR. Only objects that are named in the Naming Service can be referenced using this pattern.

Home Name and Key

Used for objects that are not named in the Naming Service. Instead of the object name, it stores the name of the object home and its primary key within the home. The stored representation for this pattern is a pair of variable length strings: the home name and the stringified primary key.

Queryable Collection Name and Query String

Used for objects that are contained in a queryable home or named collection. It stores a pair of variable length strings: the collection name and a query string that uniquely returns the object.

This is not an exhaustive list. Many other patterns and variants of these patterns are possible. The intent here is to describe some of the more generally applicable patterns.

Bottom-up customizations

In the bottom-up case, in addition to knowing that the business object interface has a getter and a setter, there is also a known value in an existing database table or other resource to which the data object is mapped. The value in the existing resource manager is not a mapped object reference as described in “Top-down customizations” on page 396. It is most probably a foreign-key, a value that can be used, along with other Component Broker system function, to render the object reference which is being requested in the upper level (business object) interface. The methods take on the same form as in the top-down case excepting this additional restriction of the value used to do the mapping.

Conceptually, there are a number of patterns that can be used to map the foreign-key into the object reference and back again:

FindByPrimaryKeyString

Using this pattern, a value (or values) from the underlying resource manager is used to formulate a key. The home for the kind of object being found is determined. The key is used to do a `findByPrimaryKeyString` on the object. The resulting reference is returned back to the business object. On the setter side, the object reference primary key is extracted and stored for later use when the containing object is again activated and the referred to object is requested. A basic structure is shown in the following segments:

```
Claim_ptr PolicyDO_Impl::currentClaim()
{
    // retrieve the stored foreign key from the PolicySQL table
    // create a claim key
    ClaimKey_var theKey = Claim::_create();
    // set foreign key from table into primary key for claim
    theKey->claimNo(n);
    // find the object through its home
    IManagedClient::IManageable_var aMgbl;
    ByteString_var thekeyString = theKey->toString();
    aMgbl = claimHome->findByPrimaryKeyString(thekeyString);
    return Claim::narrow(aMgbl);
    ...
}
```



```

::CORBA::Void PolicyDO_Impl::currentClaim(Claim_ptr claim)
{
    // map the claim pointer into a storable form
    // n is part of a struct used to talk to the database
    n=claim->claimNo;
    // invoke and/or notify the application adaptor regarding
    // the change to persistent data
    markDirty();
    ...
}

```

The above segments have exception handling and NULL checks removed to allow a more simplified view. They are also incomplete in how they deal with the conditionality aspect of the relationship. If, for example, there is always total ownership of the referenced object, then the setter has to consider removing the referenced object.

The other issue is that of finding the home that is used in the getter method. This is discussed in later sections.

Query This pattern leverages the Query Service. In this case, the foreign-key is used as the basis to formulate a query into the table that contains the data for the referred to objects. The result set can then be used to return values to the business object interface. This pattern works for both 1-to-1 and 1-to-n relationships with various exception handling required.

The following segment shows conceptually what happens in the 1-to-1 case:

```

Policy_ptr ClaimDO_Impl::policy()
{
    PolicyKey_var iPolicyKey;
    iPolicyKey=PolicyKey::_create();
    // set the foreign key into a policy key
    iPolicyKey->policyNo(iClaimPO.policyNo());
    ::ByteString_var iPolicyKeyString = iPolicyKey->toString();

    // find the home to use
    char buf[128];
    sprintf(buf, "/host/resources/servers/%s/collections/%s"
            , serverName(), "Insurance::Policy");
    CORBA::Object_var obj = nameService()->resolve_with_string(buf);
    iPolicyHome=IManagedClient::IHome::_narrow(obj);

    // find the object using the key
    temp = iPolicyHome->findByPrimaryKeyString(*iPolicyKeyString);
    return Policy::_narrow(temp);
}

::CORBA::Void ClaimDO_Impl::policy(Policy_ptr policy)
{
    // get the primary key string from the in object

```

```

::ByteString_var iPolicyKeyString =
    policy->getPrimaryKeyString();

// make a key
PolicyKey_var iPolicyKey = PolicyKey::_create();

// set string into the key
iPolicyKey->fromString(*iPolicyKeyString);

// get number from key and tell PO about it
iPolicyPolicyNoFK = iPolicyKey->policyNo();
iClaimPO.policyNo(iPolicyPolicyNK);

```

In the example above, the getter method creates the reference to return by issuing `findByPrimaryKeyString` with the key that is created from the foreign key value. On the setter, the foreign key determined by examining the key that comes with a reference is extracted and pushed back down to the database. The value for the `Home Insurance::Policy` represents the value that is used by client programs for the factory finder and is part of the management information that is found in the Home Image under the name of managed object interface.

The general patterns that can be used for top-down and bottom-up relationships are described in previous sections. The following sections detail what is done for each specific combination.

1-1 relationships

For the top-down case, using the handle-pattern that is supported by Object Builder is recommended. The handle concept of Component Broker encapsulates the actual pattern that is used to store the object reference. From this perspective, each 1-to-1 relationship in the top-down case stores a handle. Using the handle pattern requires that some decisions be made about the nature of the handle.

When a component that will be referred to by other components is constructed, the handle-pattern that it supports must be decided. If no choice is made, then the handle implementation that encapsulates a stringified object reference is the default. This means that it is essential that the stringified object reference pattern be selected on Object Builder whenever constructing relationships to this component from other components. If, when constructing the referenced component, the `managedObjectName` or `Home Name and Key` handles are selected, then the corresponding mapping pattern must be used when making relationships to these objects.

The home name and key pattern should be used when appropriate. This allows creation of database tables that hold shorter handles than is possible

with the stringified reference version of handles. For example, the home name and key can generally be stored in 200 bytes (plus or minus 20 bytes) but the SOR version of handles can require more than 1000 bytes.

An optimization for this pattern that can be used when the object being linked is always in a specific known home, is to store only the key string in the database and to hard code the home name (for example, with an installation-time settable environment variable) in the access methods. This reduces the storage overhead per link tremendously but can be used only to link to objects in the same home. Attempts to set the link to point to an object of the same type in another home cannot be statically checked and would therefore need to fail at run time.

Alternatively, a factory finder can be used to implement this pattern. This variant of the pattern also eliminates the home name from the stored representation by using a factory finder instead of the Naming Service to retrieve the home object in the get method. This approach can be used only if the factory finder can be guaranteed to always return the same home object (for example, when there is known to be one and only one home for a certain managed object type).

The following figure summarizes what happens at the object (top of figure) and database level (bottom of figure) to make this relationship happen.

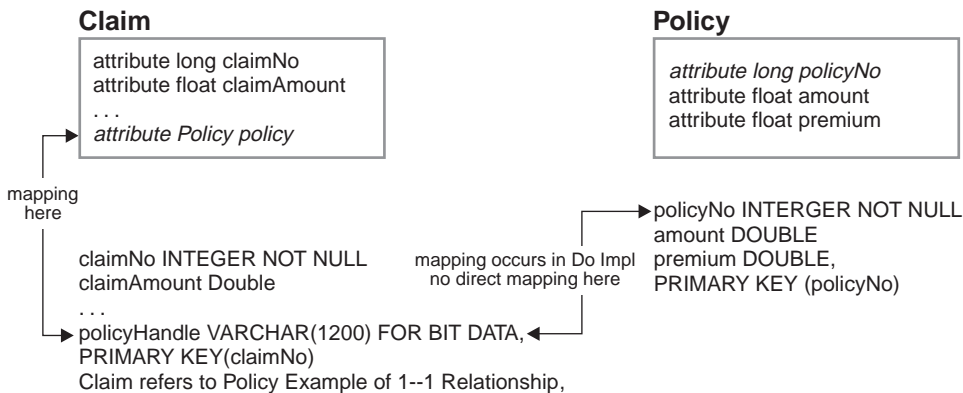


Figure 71. Claim refers to Policy example of 1-1 relationship (top down)

For the bottom-up case, Object Builder supports the foreign key pattern suggested and described previously. This is the recommended pattern.

The following figure summarizes what happens at the object (top of figure) and database level (bottom of figure) to make this relationship happen.

Claim

```
attribute long claimNo
attribute float claimAmount
.....
attribute Policy policy
```

Policy

```
attribute long policyNo
attribute float amount
attribute float premium
```

```
claimNo INTEGER NOT NULL
claimAmount DOUBLE
```

```
.....
policyNo INTEGER,
PRIMARY KEY (claimNo)
```

*foreign key
mapping here*

```
policyNo INTEGER NOT NULL
amount DOUBLE
premium DOUBLE,
PRIMARY KEY (policyNo)
```

Figure 72. Claim refers to Policy example of 1-1 relationship (bottom up)

1-n relationships

For the top-down case, two patterns are supported by Object Builder for 1-n (one to-many) relationships. Conceptually they involve either keeping a reference collection of the object relationships or resolving the object relationships using the Query Service.

The reference collection method is useful if there is no query predicate from which a results set can be calculated. This option creates a reference collection of object relationships. A handle to this reference collection is stored in the component that contains the reference to the objects. This reference collection is then accessed whenever the members of the object relationship are used.

The query solution for 1-n (described following) is also useful in some top-down scenarios.

For the bottom-up case, the solution for 1-n relationships is to use the Query Service. Objects are retrieved using a code segment similar to the following:

```
::IManagedCollections::IIterator_ptr PolicyDOImple::myClaims()
{
    char buf[1024];
    sprintf(buf2, "/host/resources/servers/%s/query-evaluators/default",
            serverName());
    ::CORBA::Object_var obj = nameService()->resolve_with_string(buf2);
    IExtendedQuery::QueryEvaluator_var qe =
        IExtendedQuery::QueryEvaluator::_narrow(obj);
    ICollectionsBase::IIterator_var qIter;
    IExtendedQuery::MemberList_var members;
    char buf[1-24];
```

```

sprintf(bufm,"select a from %s a where
    a.\policy\..\policyNo=d%";,"Insurance::Policy",policyNo());
qe->evaluate_to_iterator(buf,NULL,NULL,NULL,-,members,qIter);
return ICollectionsBase::IIterator::_duplicate(qIter);
}

```

Object Builder currently supports a pattern similar to this. The difference is that the code to do the query is actually located in the business object implementation. This is a tactical statement. Either way, the general concept of the query code being used to resolve the relationship is the same.

This option allows query to be used to determine membership. This is required in cases where there are other legacy programs that could affect membership in any given relationship. This also has the advantage of using less storage in the component that has a relationship to many other objects of a different type.

Summarizing relationships implementations

The patterns described above are summarized in the following table.

	Top-Down	Bottom-Up
Cardinality-1	Store the reference <ul style="list-style-type: none"> • Uses handles support to store a stringified version of the object reference 	Foreign key <ul style="list-style-type: none"> • Uses findByPrimaryKeyString to locate referenced object
Cardinality-n	Reference collection <ul style="list-style-type: none"> • Uses handles support to store a reference to a collection that contains the "object relationships" 	Foreign key <ul style="list-style-type: none"> • Uses query service to return the "object relationships" • Requires a 1-to1 reference to implement the other way

Additional customizations

Additional data object customizations are possible. If the customization options that are optimal for a given relationship are not supported directly, it is possible to alter the mapping patterns that are used in such a way that Object Builder round-tripping is still preserved. This is done through the use of a mapping helper.

Mapping helpers: A mapping helper is a class that contains mapping methods. Mapping methods provide the conversion between the attribute types of the two objects. You can either use the mapping helpers provided by Object Builder, or you can define your own. Object Builder provides the default mapping helper (the class and its methods) in the following cases:

- When a stringified object reference (SOR) of the data object is mapped to a persistent object attribute of type VARCHAR.
- When a data object attribute of type string is mapped to a persistent object attribute of type VARCHAR. A data object attribute of type string is normally mapped to a persistent object attribute of C++ string type, for example, a string of length 20 is mapped to char[21].

Object Builder does not provide the default mapping between complex data types (any, wchar, and wstring and types defined as constructs, that include typedefs, structures, and unions) and DB2 database types. You must provide your own helper class for these mappings.

You cannot use a mapping helper to map many data object attributes to either one or many persistent object attributes; however you can use one to map many data object attributes to one persistent object attribute. Refer to *WebSphere Application Server Enterprise Edition Component Broker for Windows NT and AIX Online Documentation* and *WebSphere Application Server Enterprise Edition Component Broker Application Development Tools Guide* for further information on mapping helpers.

Examples of mapping helper usage: Object Builder integrates the usage of Mapping Helpers into the data object implementation files as part of the customization of the application. Following is an example of a VARCHAR-to-string conversion:

```

::CORBA::String_var iASStringTemp = aString;
{
    DB2VarCharMap aStringtemp;
    MAPPINGHELPERPO_ASTRINGDB2VARCHAR aStringtemp2;
    aStringtemp.maxLen = 10;
    aStringtemp.dataP = aStringtemp2.data;
    if (iASStringTemp)
        DB2MappingHelper::stringToVarChar(iASStringTemp, aStringtemp);
    else
        aStringtemp.length = 0;
    aStringtemp2.length = aStringtemp.length;
    iMappingHelperPO.aString(aStringtemp2);
}

```

Following is a similar case to stringToByteString() mapping helper. The stringToWstring() needs to allocate enough size for *wszData* before the conversion. Therefore you should add:

```
wszData = ::CORBA::wstring_alloc(iRc);
```

in your mapping helper as follows:

```

// Count the number of wide characters to expect in the
// result wide string.
iRc = mbstowcs(NULL, szData, 0);

```

```

if (iRc < 0)           // This behavior, when the conversion is not
wscpy(wszData, L ""); // possible, is user dependent.
else
{
    cerr << "mbstowcs(szData)" << szData << endl;
    wszData = ::CORBA::wstring_alloc(iRc);
    iRc2=mbstowcs(wszData, szData, iRc);
    cerr << "mbstowcs(wszData)" << wszData << endl;
}

```

Example scenario - createFromPrimaryKeyString

This section describes a `createFromPrimaryKeyString` scenario, giving special attention to the framework calls that the BOIM application adaptor would make to various pieces of the component implementation. This example should help tie together many of the component implementation details shown in the previous sections of this chapter. The scenario assumes that the component was constructed from the `IManagedObjectWithCachedDataobject` base class (that is, it is caching data in the BO). Some of the application adaptor implementation details have been purposely left out, and the exact flow of these calls may change with releases. This example is only intended to show how a typical application adaptor might interact with various pieces of the component implementation.

1. The home receives a `createFromPrimaryKeyString` method request from a client.
2. The home creates a copy of the Primary Key for the component.
3. The home calls `fromString()` method on the Primary Key, passing in the key bytestring received from the client. This causes `internalize_from_stream()` method to run on the Primary Key so that it can load its key data from the client input.
4. The home calls `toString()` method on the Primary Key. This causes `externalize_to_stream()` method to run on the Primary Key so that it can write its key data back into a bytestring. These steps of reading the key from a bytestring into the Primary Key object, and then writing it back to a bytestring assure that the key data has gone through any code page conversion required.
5. The home creates the data object for the component.
6. The home calls `isNew()` method on the data object to signal that a new component is being created (that is, as opposed to reactivating an existing component).
7. The home calls `internalizeFromPrimaryKey()` method on the data object, passing in the converted key string from step 4.
8. The home calls `verifyKey()` method on the data object to make sure that all key data has been initialized.

9. The home creates the managed object for the component.
10. The home creates a KeyComponent for the component.
11. The home calls externalizeKeyAttributes() method on the data object, passing in the KeyComponent. The data object fills the component's key data into the KeyComponent.
12. The home creates the mixin object for the component.
13. The home calls setMixin() method on the managed object, passing in the mixin.
14. The home calls initForCreation() method on the managed object, passing in the data object. This allows the managed object to perform any necessary initialization tasks.
15. The home creates a second KeyComponent for the component.
16. The home calls buildKeyComponent() method on the mixin, passing in the second KeyComponent.
17. The mixin calls externalizeKeyAttributes() method on the data object so that the data object may fill the component's key data into the KeyComponent.
18. The home calls beginMOAssemblyActivation() method on the mixin.
19. The mixin calls syncToDataObject() method on the managed object so that the component can move any cached data in the business object to the data object.
20. The mixin calls insert() method on the data object so that the component has an opportunity to write its state data to a backing store.
21. The home registers the component with the appropriate container.
22. The home returns a proxy for the component to the client.

Packaging the application

This section describes how to pull together all of the pieces and package them for installation on the server or various Component Broker clients.

This section recommends a packaging strategy for managed objects. Alternate packaging possibilities exist. The Policy example is used here as the example.

Packaging for client and server (VA C++)

This section describes packaging recommendations for components.

Library packaging

The Component Broker server uses libraries created with the C++ compiler that is specific to either Component Broker for Windows NT, AIX, and Solaris. On Windows NT, libraries implies dynamic link library. On the AIX and Sun

platforms, library implies a shared library. Libraries that represent component implementations are dynamically loaded by the server.

The general strategy for component libraries is to have two libraries. One of the libraries contains items needed on clients (that is, clients running VisualAge C++) and on the server, while a separate library created for the component implementation on the server. The following example packages only one component into the library, although multiple components can be put into each library. For simplicity, the example uses only Policy. There is a library for both client and server named `polycy` and a server-side library named `polcys`.

The following table shows the `.o` files that must be placed into one of the two LIBs that are created.

Note: In this book, “.o” means that it is an object file targeted for a LIB/DLL. An “.obj” is something that is targeted for an .exe. For example, a `client.cpp` might generate a `client.obj` that would go into a `client.exe` while all of the components associated with a managed object are `.o`’s targeted for LIBs/DLLs.

Table 10. Object files targeted for a LIB or DLL

Object File	LIB/DLL	Notes
PolicyKey_C	polycy	Bindings for the key class.
PolicyKey_I	polycy	Implementation for the key class.
PolicyCopy_C	polycy	Bindings for the copy helper.
PolicyCopy_I	polycy	Implementation for the copy helper.
Policy_S	polycy	Bindings to be used by clients for the business object. Note: This is an <code>_S</code> and not an <code>_C</code> so that the same LIB/DLL can work on both the client and the server. The <code>_S</code> has all the function of the <code>_C</code> , plus the ability to dispatch on the server. The overhead of the <code>_S</code> versus the <code>_C</code> is minimal.
PolicyBO_S	polcys	Bindings for business object implementation.
PolicyBO_I	polcys	Implementation of business logic.
PolicyDO_C	polcys	Data object interface bindings.
PolicyDOBOIM_C	polcys	Data object implementation bindings.
PolicyDOBOIM_I	polcys	Data object implementation code.
PolicyMO_S	polcys	Managed object bindings.
PolicyMO_I	polcys	Managed object implementation.

Client programs should be able to link the policy.lib into their programs (.exe or .dll) and successfully communicate with components on the server. Components that wish to directly use references to policies should link to policy.lib.

See systems management and application installation information about distributing and installing these DLL files on clients and on servers.

Create functions for dynamic DLL loading

The server dynamically loads the DLLs that contain the component function. Each component is outfitted with a set of external functions that allow the server to create the various pieces of a component. The following table summarizes the requirement for these create functions.

Table 11. Function summary

Class	Function Name	Implementation	Return Type
DataObject_Impl (for example, Policy_Impl)	Policy_create	return new Policy_Impl();	IBOIMManagedObjectCustomization::IDataObject_ptr
ManagedObject_Impl (for example, PolicyMO_Impl)	PolicyMO_create	return new PolicyMO_Impl();	IManagedServer::ImangedObject_ptr
Key_Impl (for example, PolicyKey_Impl)	PolicyKey_create	return PolicyKey::_create();	IManagedLocal::IPrimaryKey_ptr
CopyHelper_Impl (for example, PolicyCopy_Impl)	PolicyCopy_create	return PolicyCopy::_create	IManagedLocal::INonManageable_ptr

The general form of these functions is:

```
extern "C"
{
    _declspec(dllexport) return type functionname()
    {
        return implementation;
    }
}
```

Notice that the two shaded boxes for the local only objects are using the same _create function that is used by clients when they are making one of these objects. The others are made only by the server and thus do not need to have similar _create methods. Do not get the _create methods that some local only objects have confused with the create functions that are used by the server.

Exposing interfaces to clients

Given the .LIB files shown previously, the next step is to list and determine which interfaces are needed by clients of the policy class. Pure clients need access to the following interfaces:

PolicyKey.hh

To create and use the key class.

PolicyCopy.hh

To create and use the copy helper class.

Policy.hh

To interact with Policy components that are on the server.

Exposing interfaces to component builders

If an object provider wishes to have a reference to a component from another component, then the interfaces described previously are sufficient. If however, the component provider wants to extend or override the component and create a new component as described in “Extending a business object” on page 157, then some additional interfaces are necessary. These are listed next:

- All of the IDL files associated with the Policy would be needed to make a subclass of policy.
- By implication of the previous item, all of the .hh files associated with Policy would be needed to make the bindings for the new class.
- Furthermore, all of the .ih files associated with Policy would be needed if the implementation of the new business object was to make full use of the implementation of Policy.

Packaging the DLL for the ActiveX Visual C++ Client

Packaging for a pure Visual C++ client could be as simple as taking the policyc DLL and rebuilding it for the Microsoft Visual C++ client, and providing the same interfaces to clients. For a minimum footprint client, the Policy_S could be replaced with Policy_C as the Microsoft version of policyc is not used on the server.

To get from a pure Microsoft Visual C++ client to a full ActiveX/COM enabled client, COM wrappers must be created. Run the idl2com tool on those interfaces that are to be exposed to the ActiveX/COM programmer. The idl2com tool produces a series of files that include all of the source code necessary for the COM wrapper for a particular IDL file, as well as a makefile for building a DLL which contains the COM wrapper.

By default, the makefile compiles all COM wrapper source, and all _C files needed by the COM wrapper. Because it is possible that any method inherited

by Policy from other classes could be invoked by a COM wrapper user, all `_C` files for any inherited classes are also compiled into the ActiveX/COM wrapper DLL. This allows the COM wrapper to expose all methods and attributes defined by Policy, as well as those which are inherited from other classes. Once this DLL is built, it must be registered in the Windows system registry, so that the DLL that implements the COM wrapper can be found. Running `regsvr32 dllname.dll` against the produced DLL, when it has been placed into its final destination directory, accomplishes this task.

A set of basic LIBs must be linked to in order to get the client DLLs or EXEs made properly. This list includes:

- `somororm.lib` - ORB
- `sompmcim.lib` - Programming Model
- `somosa1m.lib` - Object Services
- `somax00m.lib` (the ActiveX client run time)

Packaging the Java client code

Packaging for the Java Client involves taking the IDL for Policy, PolicyKey and PolicyCopy and running them through the IDL-to-Java emitter. This generates the necessary bindings. These bindings can then be packaged into a `.zip` file or exist as `.class` files in the CLASSPATH of the clients.

Assembling and installing Java components

The next step is to generate and build the extra objects necessary to adapt a component to a server environment and to a particular persistent store, then to package and install all the necessary pieces into a server.

In Component Broker, most of the extra code that is required is C++ code and the Component Broker tools generate almost all of it. As a result, the discussion earlier in this chapter applies with few or no changes when the component's business object is implemented in Java. In addition, there are some extra steps only in the Java business object case.

This section details the differences in the procedures and code described elsewhere in this chapter when they are applied to a Java business object. As was the case previously, it is not normally necessary for a component developer to understand the tool-generated code. It is described in case unusual circumstances require debugging.

Create the C++ client and server bindings

Previously Java has been used exclusively, but in order to adapt a Java business object to the Component Broker server environment it is necessary to create some C++ code. If you are using Object Builder to create the component, it generates a make utility file that does all this for you.

The first group of required C++ files are the bindings generated from the IDL that has already been produced. These are created using the `idlC` command with the `-shh;uc;sc` option. In the case of the Policy example this command would be used on all of `Policy.idl`, `PolicyKey.idl`, `PolicyCopy.idl`, `PolicyBO.idl`, and `PolicyDO.idl`. It would also be repeated for the IDL created later in this chapter for the `PolicyMO` class and the specialized data object.

Create the managed object class and implementation

This Object Builder-generated class differs only slightly from that described earlier in this chapter. Its IDL is identical, and there are only two differences in the implementation:

- Inheritance
- Construction and destruction

Inheritance: The `PolicyMO_Impl` class constructed for a C++ business object inherits from the `PolicyBO_Impl` C++ class, but when the business object is implemented in Java there is no `PolicyBO_Impl`. Instead, the inheritance is from a generated IOM C++ proxy class that delegates method calls to the Java business object:

```
class PolicyMO_Impl : public virtual ::PolicyMO_Skeleton,
                    public virtual Object_SOMProxyRemotable,
                    public virtual PolicyBO_SOMProxy
```

The `Object_SOMProxyRemotable` base class is added to resolve some ambiguous overrides of methods introduced in `CORBA::Object`, and does not contribute any new function.

Inside the business object methods of the `PolicyMO_Impl`, this same renaming happens for each call to the parent BO business logic:

```
::CORBA::Float PolicyMO_Impl::amount()
{
    ::CORBA::Float retval;
    IBOIMExtLocalToServer_IMixinPointerImpl mixinPointer(mixin());
    CALL_MIXIN_BEFORE(mixinPointer);

    #ifdef CBS_TRACE_DEBUG
        void *trc_handle = BOSS_TRACE_SERVER_START_2(this, "name");
    #endif

    retval = PolicyBO_SOMProxy::amount();

    #ifdef CBS_TRACE_DEBUG
        BOSS_TRACE_SERVER_STOP(trc_handle, "name");
    #endif

    CALL_MIXIN_AFTER(mixinPointer);
    return retval;
}
```

Construction and destruction: The PolicyMO_Impl for a Java business object has a more complex constructor and destructor than in the C++ case. The constructor calls through the IOM interlanguage run time in order to create the Java business object and to connect it to the C++ managed object:

```
static SOMRef* _PolicyMO_Impl_helper()
{
    SOMException e = { 0,0 };
    SOMRef* sr;
    SOMClassRef* scr = SOM_FindClass("PolicyBO", &e);
    if(scr == NULL || (sr=scr->NewObject(&e)) == NULL)
    {
        throw CORBA::NO_IMPLEMENT(0, CORBA::COMPLETED_NO);
    }
    return sr;
}

PolicyMO_Impl::PolicyMO_Impl()
: Object_SOMProxy(_PolicyMO_Impl_helper()) , mixin_(this)
// See note 1
{
    m_target_type_id = (char*)Policy_RID;
}

PolicyMO_Impl::~PolicyMO_Impl()
// See note 2
{
    if(m_somref2 != m_somref && m_somref != NULL)
    {
        SOMRef* tmp = m_somref;
        m_somref = m_somref2;
        setWrapper(this);
        tmp->Destroy();
    }
}
```

Notes about the example:

1. A helper routine creates the corresponding Java business object and returns an interlanguage handle that is used to tie together the C++ managed object and Java business object.
2. The objects are eventually untied in the managed object destructor, allowing them to be separately destroyed.

Data object customization

Data objects created for use with a Java business object are identical to those created for a C++ business object, and the same variety of customization choices is available. For more information, see “Data object implementation details” on page 361.

Marking data object interfaces: Interfaces of customized data objects, such as the PolicyEmSQLDO interface illustrated earlier in this chapter are marked

with the abstract meta information pragma in their IDL. This was done to suppress the declaration of a static `_create()` method, which is not needed.

In Component Broker, this generates incorrect bindings, and such a data object is not usable with a Java business object. Interfaces that do have corresponding implementations must not be marked abstract. Instead, if the function of the static `_create()` method is not needed, you can just define a `_create` that returns `NULL`.

Generating server-specific Java classes

Two Java classes are used on the server to wrap the Java business object and adapt it's component to run in the server environment. Both classes are generated from the business object IDL by running the `idlcc` command with the `-sbj` option:

```
idlcc -sbj PolicyBO.idl
```

The command generates the `_PolicyBOWrapper` and `_PolicyBOImpl` classes. `_PolicyBOWrapper` extends `_PolicyBOBase` and is used by IOM to dispatch calls from C++ to the business logic. `_PolicyBOImpl` extends `_PolicyBOWrapper` and delegates calls from Java code over to the C++ managed object class, which is then able to add manageability and quality of service factors when the component is called from Java, just as when it is called from C++.

Generating other C++ classes

If you have not already done so, a C++ primary key class must now be generated and compiled. The server infrastructure requires this, even if the component's business object is implemented in Java. "Chapter 5. MOFW - C++ server programming model" on page 103 describes how to do this.

If the component makes use of a copy helper class, you also need to generate and compile a C++ version of the class.

In both cases, it is essential that the `internalize_from_stream()` and `externalize_to_stream()` methods of the C++ and Java versions be consistent with one another. If the Java version reads and writes the policyNo as a `CORBA::Long`, followed by the amount as a `CORBA::float`, the C++ version must do the same.

Debugging Java code running on the server

As initially delivered, Component Broker does not contain a Java debugger capable of debugging code running in a server. There are two options available for server debugging.

The `System.out.println()` standard Java method can be used to produce debugging output in the server console log. A `_PolicyBOBase` class can be quickly modified to add or delete these statements, recompiled, and replaced, and the server restarted to use it, without regenerating or replacing any of the other Java classes or C++ DLLs.

A second option is to use the `jdb` debugger supplied as part of the JavaSoft Java Development Kit. This requires that you first use the Component Broker System Manager User Interface to enable the debug option in the server. The procedure to do this follows:

1. In System Manager, from the **View** menu, select **User Level > Super User**. This causes the System Manager window to display more detailed information.
2. Open the Host Images folder and expand the host image that corresponds to the name of the server.
3. Expand Server Images and find the server image where the application resides. Right-click this image.
4. From the pop-up menu, select **Edit**.
5. Select the **Main** notebook tab.
6. Change the *debug enabled* attribute to yes.
7. Select the **ORB** tab.
8. Change the **request timeout value** to 0.
9. Select the **Log Controls** tab.
10. Change the **Console Disposition** to "Console".
11. Click **OK** to exit the server image notebook.

You are now ready to start the Component Broker daemon, name server, and application server. To do so, follow these steps:

1. From the Host Images folder, find the host image that corresponds to the name of the server computer.
2. Right-click this host image.
3. From the pop-up menu, select **Activate**. This starts the communication daemon and the name server.
4. Monitor the Action Console window for completion status. When activation is complete, right-click to select the application server.
5. From the pop-up menu, select **Run Immediate**. Monitor the Action Console window for completion status.

Note: If the server fails to start, it is possible that the `CLASSPATH` environment variable may not be configured to include the path to your JDK installation's `classes.zip` file. To correct the `CLASSPATH` and restart the server, perform the following steps:

1. Logon to the user ID specified for Systems Management.
2. Change the CLASSPATH environment variable to include the path to the classes.zip file either in the **System Variable** or in the **User Variable** section of the environment setting (**System Variable** is recommended to ensure that other users logging on also pick up the CLASSPATH environment variable changes).
3. Stop the CBCConnector service.
4. Restart the CBCConnector service.
5. Reactivate the applications using the System Manager User Interface.

You can now run a client program that uses the Java-based components. The Java environment initializes as soon as the first Java class is used in the server, and at that time a single line of output appears in the server console window:

```
Agent password=password
```

Where *password* is a random sequence of letters and digits. You can now start the jdb debugger:

```
jdb -host hostname -password password
```

Where *hostname* is the Internet-style name of the server host (for example, server1.ibm.com) and *password* is the agent password. The server continues to run throughout this procedure, so you should set breakpoints in the Java business object or to explicitly halt the server using the jdb command set. See the JDK documentation on the jdb debugger for more information.

The managed object for a Java specialized home

The managed object class for a Java specialized home has the responsibility for routing method calls to the correct implementation. Calls to methods from the IManagedAdvancedServer::ISpecializedHome interface are passed to a C++ implementation, while calls to extension methods are forwarded to Java. To accomplish this, the managed object class, which is implemented in C++, inherits from both the C++ base implementation and from several IOM proxy classes:

```
class PolicyHomeMO_Impl : public virtual ::PolicyHomeMO_Skeleton,
    public virtual Object_SOMProxyRemotable,
    public virtual PolicyHomeBO_SOMProxy,
    public virtual IManagedAdvancedServer_ISpecializedHome_Impl,
    public virtual IManagedServer::IWrappable_SOMProxy
```

The second last of these is the managed object class for the base C++ home implementation, and provides implementations of the IManagedClient::IHome interface as well as all the usual framework methods. The two surrounding it, PolicyHomeBO_SOMProxy and IManagedServer::IWrappable_SOMProxy,

forward their respective methods over to the Java side. Because many of the framework methods are implemented in more than one of these classes, the managed object class `PolicyHomeMO_Impl` has to override all methods and explicitly call up to the correct base class or classes.

As a result, `PolicyHomeMO_Impl` has a total of 50 methods. Only a representative sample is shown here. The constructor and destructor, and all of the framework methods that delegate to the mixin object are no different than for any Java business object, so they are not discussed.

Some of the more interesting managed object methods correspond to methods implemented in the C++ home class

`IManagedAdvancedServer_ISpecializedHome_Impl`. These are delegated to that base class:

```
IManagedClient::IManageable_ptr PolicyHomeMO_Impl::findByPrimaryString(
    const ::ByteString & key)
{
    IManagedClient::IManageable_var temp;
    IBOIMExtToLocalServer_IMixinPointerImpl mixinPointer(mixin());

    CALL_MIXIN_BEFORE2(mixinPointer, "findByPrimaryString" );

    temp= IManagedAdvancedServer_ISpecializedHome_Impl::findByPrimaryString(
        key);

    CALL_MIXIN_AFTER(mixinPointer);

    return IManagedClient::IManageable::_duplicate(temp);
}
```

This managed object method only calls up to the C++ base class, and ignores any alternate definition of the function that may have been provided in Java. Other managed object methods, though, correspond to methods introduced in the specialized `PolicyHome` interface, and they are delegated through the IOM proxy class over to the Java class `_PolicyHomeBOBase`:

```
::Policy_ptr PolicyHomeMO_Impl::default_create()
{
    ::Policy_var retval;
    IBOIMExtLocalToServer_IMixinPointerImpl mixinPointer(mixin());
    CALL_MIXIN_BEFORE2(mixinPointer, "findSomething");

    #ifdef CBS_TRACE_DEBUG
        void *trc_handle = BOSS_TRACE_SERVER_START_2(this, "default_create",
            "Java");
    #endif

    // call the real business logic
    retval = PolicyHomeBO_SOMProxy::default_create();

    #ifdef CBS_TRACE_DEBUG
```

```

        BOSS_TRACE_SERVER_STOP(trc_handle, "default_create");
    #endif

    CALL_MIXIN_AFTER(mixinPointer);

}

```

Finally, the following framework methods are special cased so that both the C++ and Java implementations are called:

- `initForCreation()`
- `uninitForDestruction()`
- `initForReactivation()`
- `uninitForPassivation()`
- `syncToDataObject()`
- `syncFromDataObject()`

This ensures that both the C++ and Java portions of the composite implementation will be able to correctly initialize and shut down, and allows them both access to the home's data object:

```

::CORBA::Void PolicyHomeMO_Impl::initForReactivation(
    ::IManagedServer::IDataObject_ptr theDO)
{
    #ifdef CBS_TRACE_DEBUG
        BOSS_TRACE_CREATE(this, "PolicyHome");
    #endif

    IManagedAdvanceServer_ISpecializedHome_Impl::initForReactivation(theDO);

    PolicyHomeBO_SOMPProxy::initForReactivation(theDO);

}

```

Enabling additional clients

Component Broker provides special support for the clients listed in the previous sections. For additional clients such as Orbix, the appropriate set of Component Broker IDL files would have to be run through the emitter of the client ORB. Then the IDL files for the business object (policy) would need to be run through the emitter of the client ORB.

Note that clients not provided by IBM do not have facilities for handling `INonManageable` Objects. This means that alternative means of client access are necessary. Using stringified object references to components, creating specialized homes or wrapping homes appear as the most preferable options for providing access to other clients.

Configuring the CBCConnector run time

This section provides information about the configuration options offered by the CBCConnector run time. The application development process produces artifacts that provide initial settings for most of these parameters, but many developers may wish to learn more. This information is also useful for administrators who want to do application performance tuning. The following topics are covered:

- Introduction to container configuration parameters
- Typical settings for container configuration parameters
- Summary of supported container configurations
- Configuring application adaptors - RDB
- Configuring homes

Advanced information about tuning the CBCConnector run time to suit the specific requirements of your application can be found in “CBCConnector Tuning” topic of the “Performance Monitoring and Tuning” chapter in the *WebSphere Application Server Enterprise Edition Component Broker Planning, Performance and Installation Guide*.

Introduction to container configuration parameters

This section provides an overview of the container configuration parameters. Instructions are provided for setting each parameter using Object Builder wizards. The active settings for these parameters may be viewed from the System Management Interface. Note that any change made to a container parameter by editing its image using the System Management Interface is not guaranteed to be permanent. A subsequent model activation may overwrite the change. Permanent changes must be made using Object Builder. The application artifacts need to be regenerated and reinstalled.

Description

The description parameter gives you an ability to provide meaningful information about each of the containers used by your applications. This information is helpful for problem determination. It is included in several activity log messages.

- In Object Builder this is configured on the Name page of the Container wizard by filling in text for the **Description** field.
- From the System Management GUI this parameter may be viewed by displaying the *description* attribute of the container.

Memory management policies

Memory management policies allow you to control when an application adaptor brings your components into and out of the server's memory. This is referred to as reactivation and passivation, respectively. These policies affect both the amount of memory your application server will require and the performance of your applications.

Passivation policy: Each container is configured to apply one of the following four passivation schemes to the components it manages.

1. Passivate after checkpoint.
 - This option is only applicable to components which use no services. Selecting this option will cause these components to be taken out of memory whenever the client application explicitly checkpoints the component. Explicit checkpoint of a component involves narrowing to the `IBOIMManagedObjectFriendQOS::IMMMixin` interface so that the `checkpointToDatastore()` method may be invoked.
 - In Object Builder this is configured on the Service page of the Container wizard by selecting **Use No Services** and selecting the **Passivate a component after checkpoint** check box.
 - From the System Management GUI this parameter may be viewed by examining the *memory management policy* attribute of the container.
2. Passivate at end of transaction.
 - This option is applicable to components which use the RDB Transaction Service, the PAA transaction service, or the MQAA Transaction Service. Selecting this option will cause components, which are not actively in use by any other transaction, to be taken out of memory when the current transaction commits or rolls back. This is accomplished in the client code by narrowing to the `TransactionCurrent` interface so that the `commit()` method or `rollback()` method may be invoked.
 - In Object Builder this is configured on the Service page of the Container wizard by selecting either the **Use RDB Transaction Service**, **Use PAA Transaction Service**, or **MQAA Transaction Service** radio button and selecting the **Passivate a component at end of a transaction** check box.
 - From the System Management GUI this parameter may be viewed by examining the *memory management policy* attribute of the container.
3. Passivate at end of session
 - This option is only applicable to components which use the PAA Session Service. Selecting this option will cause these components to be taken out of memory at the end of each session. This is accomplished in the client code by narrowing to the `ISessions::Current` interface so that the `endSession()` method may be invoked.

- In Object Builder this is configured on the Service page of the Container wizard by selecting **Use PAA Session Service** and selecting the **Passivate a component at end of a session** check box.
 - From the System Management GUI this parameter may be viewed by examining the *memory management policy* attribute of the container.
4. Never passivate.
- This option is applicable to all components. Selecting this option will cause the components to always stay in memory. This option must never be used for non-shared components. The application adaptors supporting non-shared components make new copies of these components on a per transaction or per session basis. The application server's memory will be quickly exhausted if these copies are not allowed to passivate. Please read "CBCConnector Tuning" topic in the "Performance Monitoring and Tuning" chapter of the *WebSphere Application Server Enterprise Edition Component Broker Planning, Performance and Installation Guide* for more information.
 - In Object Builder this is configured on the Service page of the Container wizard by clearing the **Passivate a component** check box.
 - From the System Management GUI this parameter may be viewed by examining *memory management policy* attribute of the container.

Number of managed objects in container: Each container is configured with an estimate of the number of components it will hold. This parameter represents the number of active component instances the container will have in memory at a given point in time. The value of this parameter is used by the container to improve the efficiency of its hashing mechanism.

- In Object Builder this is configured on the Name page of the Container wizard by setting a value in the **Number of Components** box.
- From the System Management GUI this parameter may be viewed by examining the *cache Size* attribute of the container.

Data caching policies

Data caching policies configure the container so that the application adaptor understands the data synchronization requirements of your components. These policies tell the container about decisions made when your components were developed. These decisions affect both the amount of memory your application server will require and the performance of your applications.

Cache data in business object: Selecting this option notifies the application adaptor that a copy of your component's essential state data is kept in the business object.

In Object Builder this is configured on the Data Access Patterns page of the Container wizard by selecting **Caching** in the Business Object area.

From the System Management GUI this parameter may be viewed by examining the *data cached in business object* attribute of the container.

Cache data in data object: Selecting this option notifies the application adaptor that a copy of your component's essential state data is kept in the data object in addition to the copy which may already kept in an associated persistent object.

In Object Builder this is configured on the Data Access Patterns page of the Container wizard by selecting **Local copy** in the Data Object area.

From the System Management GUI this parameter may be viewed by examining the *data cached in data object* attribute of the container.

Use the Cache Service: Selecting this option causes a copy of the backing store data for your components to be managed by the Component Broker Cache Service. This service improves performance for many application scenarios by avoiding unnecessary database accesses.

In Object Builder this is configured on the Data Access Patterns page of the Container wizard by selecting **Cache Service** in the Data Object area.

From the System Management GUI this parameter may be viewed by examining the *use cache service* attribute of the container.

Transaction policy

The transaction policy configures the container so it understands the transactional semantics required by your components. This policy only applies to containers which have been configured to use either the RDB Transaction Service, the PAA Transaction Service, or the MQAA Transaction Service.

Behavior for methods called outside a transaction: Each container is configured to apply one of the following three semantics whenever an attempt is made to run a method on your component and there is no active transaction (that is, neither the client nor the server portions of the application have explicitly started a transaction).

1. Throw exception

- This option causes the application adaptor to return an exception to the caller without running the requested method on your component.
- In Object Builder this is configured on the Service Details page of the Container wizard by selecting **Throw an exception and abandon the call** in the Behavior for Methods Called Outside a Transaction area.
- From the System Management GUI this parameter may be viewed by examining the *default transaction policy* attribute of the container.

2. Begin a new transaction
 - This option causes the application adaptor to begin a new transaction prior to running the requested method on your component, and commit or roll back the transaction after the method has run.
 - In Object Builder this is configured on the Service Details page of the Container wizard by selecting the **Start a new transaction and complete the call** radio button in the Behavior for Methods Called Outside a Transaction area.
 - From the System Management GUI this parameter may be viewed by examining the *default transaction policy* attribute of the container.
3. Ignore condition
 - This option is currently not supported by Component Broker. If this option is selected, the application adaptor will behave as if Throw exception was configured.
 - In Object Builder this is configured on the Service Details page of the Container wizard by selecting **Ignore the condition and complete the call** in the Behavior for Methods Called Outside a Transaction area. The model checker function of Object Builder will flag this as an inappropriate configuration.
 - From the System Management GUI this parameter may be viewed by examining the *default transaction policy* attribute of the container.

Session policies

The session policies configure the container so it understands the sessional semantics required by your components. These policies only apply to containers which have been configured to use the PAA Session Service.

Behavior for methods called outside a session: Each container is configured to apply one of the following two semantics whenever an attempt is made to run a method on your component and there is no active session (that is, neither the client nor the server portions of the application have explicitly started a session).

1. Throw exception
 - This option causes the application adaptor to return an exception to the caller without running the requested method on your component.
 - In Object Builder this is configured on the Service Details page of the Container wizard by selecting **Throw an exception and abandon the call** in the Behavior for Methods Called Outside a Session area.
 - From the System Management GUI this parameter may be viewed by examining the *session policy* attribute of the container.
2. Ignore condition

- This option is currently not supported by Component Broker. If this option is selected, the application adaptor will behave as if Throw exception was configured.
- This option is not available from Object Builder.
- From the System Management GUI this is parameter may be viewed by examining the *session policy* attribute of the container.

Session priority: This option affects when the application adaptor will make certain framework calls to sessionable components. This parameter cannot be set from Object Builder. This parameter can be viewed from the System Management GUI by examining the *session priority* attribute of the container. The value of this parameter should not be changed from its default setting.

Security policy

The security authorization parameter gives you the ability to specify whether security checks should be made on the components in the container before methods are implemented. This parameter applies to all component types.

- This parameter is not available in Object Builder. The default value of “yes” is always applied.
- From the System Management GUI this parameter may be viewed by examining the *Enable fine authorization level* attribute of the container. This parameter has no affect on the application adaptor unless the server attribute called *authorization level* is set to *fine*.

Serialization policy

The component serialization parameter gives you the ability to specify if access to the methods of your component should be serialized by the CBCConnector run time to provide thread-safety. This parameter only applies to non-transactional, transient components. This function is provided to satisfy a portion of the entity bean component contract required to support enterprise bean components.

- This parameter is not available in Object Builder. The default value of “no” is always applied.
- From the System Management GUI this parameter can be viewed by examining the *Enable managed object serialization* attribute of the container.

Termination policy

The termination policy allows you to configure the container so that it will automatically checkpoint the essential state data for your components when the application server process terminates. This policy only applies to containers which have been configured to use no services.

- In Object Builder this is configured on the Service page of the Container wizard by selecting **Use No Services** and selecting the **Passivate a component after checkpoint** check box.
- From the System Management GUI this parameter may be viewed by examining the *termination policy* attribute of the container.

Component type

The component type parameter tells the container whether your components have persistent or transient state data.

- In Object Builder the setting for this parameter is determined by the combination of the type of service selected on the Service page of the Container wizard, and the data object type selected for your component. For example, certain components, such as homes are always persistent. Transactional components may be transient or persistent, depending on which type of data object implementation is chosen.
- From the System Management GUI this parameter may be viewed by examining the *hold persistent objects* attribute of the container.

Object reference type

The object reference type parameter tells the container whether to generate persistent or transient references for your components. All components with persistent state data need to have persistent references so the application adaptors can reactivate and passivate the components, as needed. Components with transient state data may select either transient or persistent references. Components with transient references should be created, used, and then destroyed. They cannot be reactivated so passivation is equivalent to destroying them.

- In Object Builder this is configured on the Service page of the Container wizard by selecting or clearing the **Enable persistent references** check box.
- From the System Management GUI this parameter may be viewed by examining the *persistent references* attribute of the container.

State

The state parameter describes status about your container. The value for this parameter is filled in by the CBConnector run time. It can be displayed from the System Management GUI, but cannot be changed. This parameter is not available in Object Builder.

Workload management

In Object Builder there is a Container wizard page which asks the developer to make decisions about whether the container will be workload managed.

This input does not directly configure any parameters on the container. For more information about configuring workload management see the “Workload Management” chapter in the *WebSphere Application Server Enterprise Edition Component Broker Advanced Programming Guide*.

Typical settings for container configuration parameters

This section describes typical container configurations that support the eight most commonly used component implementations provided by Component Broker. These eight encompass the bulk of the components used by most applications. The information is organized by data object implementation type. As described in “Introduction to container configuration parameters” on page 418, each of these parameters may be set from Object Builder wizards and viewed from the System Management interface. This section is written using the System Management interface terms for these parameters.

The following are supported for all component implementations:

- Both business object data caching patterns (Delegating and Caching)
- Both data object caching patterns (Delegating and Local Copy)
- Both security policies (Enable/Disable Fine Authorization Level Checking)

With Component Broker Workstation all component implementations are supported on each of the deployment platforms, except for Transactional Components for MQ systems, which may only be deployed on Windows NT.

Transient components

The following discusses the transient components typically used by applications.

1. Shared, Non-Transactional Components with Transient data

- These components do not have persistent state data or require transactional coordination. They are either stateless, or have shared state data which may be simultaneously accessed through multiple clients without requiring serialization. These components are created, used, and then removed by the client application. They always stay in memory.
- Containers for these components are typically configured as follows.

Table 12. Shared, transient, non-transactional container configuration

Memory management policy	never passivate
Use Caching Service	no
Default transaction policy	no transaction
Session policy	no session
Managed object serialization	no
Termination policy	no check point

Table 12. Shared, transient, non-transactional container configuration (continued)

Hold persistent objects	no
Persistent references	no

2. Shared, Transactional Components with Transient data

- These components do not have persistent state data. They have shared state data which is protected against simultaneously updates through multiple transactions. Serialization is accomplished by the BOIM application adaptor obtaining a concurrency lock. This lock works on a transactional basis. It is locked when the component is first used in a transaction, and unlocked when the transaction commits or rolls back. Other transactions are forced to wait for their turn to access the component. The transaction policy is set so that the BOIM application adaptor will automatically start and commit a transaction for each method if the client application fails to start a transaction. These components are created, used, and then removed by the client application. They always stay in memory.
- Containers for these components are typically configured as follows.

Table 13. Shared, transient, transactional container configuration

Memory management policy	never passivate
Use Caching Service	no
Default transaction policy	begin transaction
Session policy	no session
Managed object serialization	no
Termination policy	no check point
Hold persistent objects	no
Persistent references	no

Basic components

The following discusses the basic components typically used by applications.

1. Shared, Transactional Components with Embedded SQL access to DB/2 data

- These components have persistent state data in DB/2 databases which is accessed by embedded SQL. Their shared state data is protected against simultaneously updates through multiple transactions. Serialization is accomplished by the DB/2 application adaptor obtaining a concurrency lock. This lock works on a transactional basis. It is locked when the component is first used in a transaction, and unlocked when the transaction commits or rolls back. Other transactions are forced to wait for their turn to access the component. The transaction policy is set so that the DB/2 application adaptor will return an exception if the client

application fails to start a transaction. These components are passivated at the end of each transaction, unless another transaction is waiting to use them.

- Containers for these components are typically configured as follows.

Table 14. Shared, DB/2, transactional, ESQL container configuration

Memory management policy	passivate at end of transaction
Use Caching Service	no
Default transaction policy	throw exception
Session policy	no session
Managed object serialization	no
Termination policy	no check point
Hold persistent objects	yes
Persistent references	yes

2. Transactional Components with Cache Service access to DB/2 data

- These components have persistent state data in DB/2 databases which is accessed by the Cache Service. Each transaction is given its own copy of the component by the DB/2 application adaptor. Any state data update conflicts that occur are resolved by the Cache Service. When conflicts occur, the Cache Service picks a winning transaction and forces the losing transaction to rollback. Multiple transactions can have read access to these components without causing conflicts. The transaction policy is set so that the DB/2 application adaptor will return an exception if the client application fails to start a transaction. These component copies must be passivated at the end of each transaction because they are not shared by multiple transactions.
- Containers for these components are typically configured as follows.

Table 15. DB/2, transactional, Cache Service container configuration

Memory management policy	passivate at end of transaction
Use Caching Service	yes
Default transaction policy	throw exception
Session policy	no session
Managed object serialization	no
Termination policy	no check point
Hold persistent objects	yes
Persistent references	yes

3. Transactional Components with Cache Service access to Oracle data

- These components have persistent state data in Oracle databases which is accessed by the Cache Service. Each transaction is given its own copy

of the component by the Oracle application adaptor. Any state data update conflicts that occur are resolved by the Cache Service. When conflicts occur, the Cache Service picks a winning transaction and forces the losing transaction to rollback. Multiple transactions can have read access to these components without causing conflicts. The transaction policy is set so that the Oracle application adaptor will return an exception if the client application fails to start a transaction. The component copies must be passivated at the end of each transaction because they are not shared by multiple transactions.

- Containers for these components are typically configured as follows.

Table 16. Oracle, transactional, Cache Service container configuration

Memory management policy	passivate at end of transaction
Use Caching Service	yes
Default transaction policy	throw exception
Session policy	no session
Managed object serialization	no
Termination policy	no check point
Hold persistent objects	yes
Persistent references	yes

4. Transactional Components to access procedural systems

- These components have persistent state data which is stored in procedural back end systems like CICS or IMS. Each transaction is given its own copy of the component by the CICS and IMS application adaptor. Any state data update conflicts that occur are resolved by the EAB Cache. When conflicts occur, the Cache Service picks a winning transaction and forces the losing transaction to rollback. Multiple transactions can have read access to these components without causing conflicts. The transaction policy is set so that the CICS and IMS application adaptor will return an exception if the client application fails to start a transaction. The component copies must be passivated at the end of each transaction because they are not shared by multiple transactions.
- Containers for these components are typically configured as follows.

Table 17. CICS and IMS, transactional container configuration

Memory management policy	passivate at end of transaction
Use Caching Service	no
Default transaction policy	throw exception
Session policy	no session
Managed object serialization	no

Table 17. CICS and IMS, transactional container configuration (continued)

Termination policy	no check point
Hold persistent objects	yes
Persistent references	yes

5. Sessional Components to access procedural systems

- These components have persistent state data which is stored in procedural back end systems like CICS or IMS. Each session is given its own copy of the component by the CICS and IMS application adaptor. There is no cross-session protection of the component data. The component's state data is shared. The session policy is set so the CICS and IMS application adaptor will return an exception if the client application fails to start a session. The component copies must be passivated at the end of each session because they are not shared by multiple sessions.
- Containers for these components are typically configured as follows.

Table 18. CICS and IMS, sessional container configuration

Memory management policy	passivate after end of session
Use Caching Service	no
Default transaction policy	no transaction
Session policy	throw exception
Managed object serialization	no
Termination policy	no check point
Hold persistent objects	yes
Persistent references	yes

6. Shared, Transactional Components for MQ systems

- These components have persistent state data which is stored on an MQSeries queue. Unlike other basic components, MQ Transactional components are deleted (and their state data destroyed) when they are retrieved from the persistent data store (that is, the queue). Their shared state data is protected against simultaneously updates through multiple transactions. Serialization is accomplished by the MQ application adaptor obtaining a concurrency lock. This lock works on a transactional basis. It is locked when the component is first used in a transaction, and unlocked when the transaction commits or rolls back. Other transactions are forced to wait for their turn to access the component. The transaction policy is set so that the application adaptor will return an exception if the client application fails to start a transaction. These components are passivated at the end of each transaction.

- Containers for these components are typically configured as follows.

Table 19. Shared, MQAA, transactional container configuration

Memory management policy	passivate at end of transaction
Use Caching Service	no
Default transaction policy	throw exception
Session policy	no session
Managed object serialization	no
Termination policy	no check point
Hold persistent objects	yes
Persistent references	yes

Summary of supported container configurations

The table below provides a summary of all the container configurations that are currently supported by the CBConnector run time. The “typical” container configurations explained in Typical settings for container configuration parameters are in **bold** print in Table 20 on page 431. Although the CBConnector run time is very flexible and supports many other “non-typical” container configurations, care must be taken when using these options. Please refer to the *WebSphere Application Server Enterprise Edition Component Broker Planning, Performance and Installation Guide* for more information on “non-typical” container configurations.

Table 20. Summary of supported container configurations

	Memory management policy	Use Caching Service	Default transaction policy	Session policy	Managed object serialization	Termination policy	Hold persistent objects	Persistent references
Shared, Transient, Non-transactional	never passivate	no	no transaction	no session	no	no check point	no	no
	never passivate	no	no transaction	no session	yes	no check point	no	no
	never passivate	no	no transaction	no session	yes or no	no check point	no	yes
	never passivate	no	no transaction	no session	yes or no	check point	no	yes or no
	passivate after check point	no	no transaction	no session	yes or no	check point	no	yes
Shared, Transient, Transactional	never passivate	no	begin transaction	no session	no	no check point	no	no
	never passivate	no	begin transaction	no session	no	no check point	no	yes
	never passivate	no	throw exception	no session	no	no check point	no	yes or no
	passivate at end of transaction	no	begin transaction	no session	no	no check point	no	yes or no
	passivate at end of transaction	no	throw exception	no session	no	no check point	no	yes or no
Transient, Transactional	passivate at end of transaction	yes	begin transaction	no session	no	no check point	no	yes
	passivate at end of transaction	yes	throw exception	no session	no	no check point	no	yes

Table 20. Summary of supported container configurations (continued)

	Memory management policy	Use Caching Service	Default transaction policy	Session policy	Managed object serialization	Termination policy	Hold persistent objects	Persistent references
Shared, DB/2, Transactional, ESQL	passivate at end of transaction	no	throw exception	no session	no	no check point	yes	yes
	never passivate	no	begin transaction	no session	no	no check point	yes	yes
	passivate at end of transaction	no	throw exception	no session	no	no check point	yes	yes
	never passivate	no	begin transaction	no session	no	no check point	yes	yes
DB/2, Transactional, Caching Service	passivate at end of transaction	yes	begin transaction	no session	no	no check point	yes	yes
	passivate at end of transaction	yes	throw exception	no session	no	no check point	yes	yes
Oracle, Transactional, Caching Service	passivate at end of transaction	yes	begin transaction	no session	no	no check point	yes	yes
	passivate at end of transaction	yes	throw exception	no session	no	no check point	yes	yes
CICS and IMS, Transient, Transactional	passivate at end of transaction	no	begin transaction	no session	no	no check point	no	yes
	passivate at end of transaction	no	throw exception	no session	no	no check point	no	yes
CICS and IMS, Transactional	passivate at end of transaction	no	begin transaction	no session	no	no check point	yes	yes
	passivate at end of transaction	no	throw exception	no session	no	no check point	yes	yes
CICS and IMS, Transient, Sessional	passivate after end of session	no	no transaction	throw exception	no	no check point	no	yes

Table 20. Summary of supported container configurations (continued)

	Memory management policy	Use Caching Service	Default transaction policy	Session policy	Managed object serialization	Termination policy	Hold persistent objects	Persistent references
CICS and IMS, Sessional	passivate after end of session	no	no transaction	throw exception	no	no check point	yes	yes
Shared, MQAA, Transactional	passivate at end of transaction	no	throw exception	no session	no	no check point	yes	yes
	never passivate	no	throw exception	no session	no	no check point	yes	yes

Note: Selections in **bold** print indicate the “typical” container configurations.

Configuring homes

Homes that can be queried and iterated can be configured for components that are using embedded SQL data objects or using data objects that use the Cache Service. Components that use any other data object cannot be configured into homes that can be queried or iterated, or into specialized homes that can be queried or iterated.

The four default homes that are provided by Component Broker are as follows:

BOIMHomeOfNotRegHomes

This home is used for system objects and should not be used for components. Objects that are stored in this home are not registered in the name space and cannot be found using factory finding techniques.

BOIMHomeOfRegHomes

This home is for components that are not workload managed and do not have queryable or iterable interfaces. This home should be used for standard components if they are not configured for WLM and are not configured for Query.

BOIMHomeOfRegQIHomes

This home is for components that are not workload managed and are configured for Query. This home should be used for those components that match this criteria.

BOIMHomeOfRegWLMHomes

This home is for components that are workload managed. Components that have been configured for WLM support must be stored in these homes.

Object Builder allows configuration of components into the appropriate home based on their configuration criteria. Only those homes that match the configuration of the components are allowed as selections on the appropriate configuration windows.

Expanding the client programming interface

In “Chapter 4. MOFW C++ client programming model” on page 85 there was one interface described that could be used to access the business logic functions of a managed object. That is the easiest way to get started and should be used when possible. Object providers, when they are building components, need to subclass from the appropriate interfaces in managed object framework and follow a set of rules for building components. The client interface that has been discussed so far is the one that is the first subclass of IManageable or one of its Component Broker Frameworks provided dependents. The figure below shows this interface.

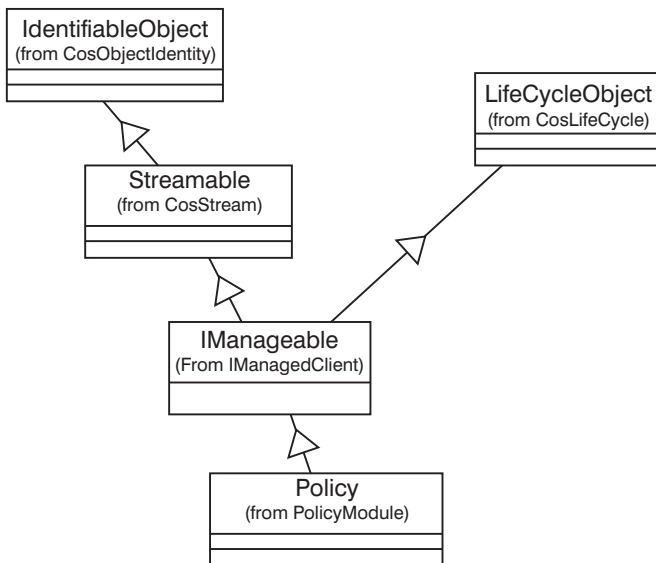


Figure 73. Primary client programming interface

This figure shows the interface of a managed object. There is much to this interface that is described in other chapters. The important thing to recognize

here is that the Policy interface introduces the business logic methods and provides access to other methods that can be useful.

Application adapter quality of service interfaces

Each component that is constructed by the object provider is installed and run in a particular application adaptor or container. These are Component Broker server concepts that deal with how resource managers are used beneath the managed object framework. The managed object framework provides a consistent programming model for clients and for object providers, encapsulating the details of the underlying resource managers when possible and practical. However, there are cases where accessing the additional capabilities afforded by a particular implementation of the Component Broker server application adaptor may be desirable. To accommodate this, a different client programming interface is available to you. The intent is to call this the client *Quality Of Service* (QOS) interface. Some clients are considered to be "friends" and might have access to additional methods. Friends are probably other components playing the role of clients.

The following interfaces fall into this category:

- IBOIMManagedObjectQOS::IMMMixin
- IBOIMManagedObjectFriendQOS::IMMMixin

These interfaces collect a number of interfaces which are supported by the BOIM container. IBOIMManagedObjectQOS::IMMMixin brings in CosTransactions::TransactionalObject; IBOIMManagedObjectFriendQOS::IMMMixin brings in the checkpointToDataStore() and refreshFromDataStore() methods.

Normally, a client program deals with a client interface by specifying it when the object is found in the IHome, retrieving from some other collection, creating, or reviving from a stringified object reference.

```
Policy_var myPolicy;
ByteString_var theKeyString;
IManagedClient::IManageable_var mptr;
IManagedClient::IHome_var myPolicyHome;

// Get the primary key (string) of the object in "theKeyString"
// Get the home configured for Policy objects in "myPolicyHome"

mptr = myPolicyHome->findByPrimaryKeyString(theKeyString);
myPolicy = Policy::_narrow(mptr);

// invoke business logic or IManageable methods on "myPolicy"
myPolicy->amount(25000.00);
```

Narrow the QOS interface to access specific methods introduced by that interface:

```

IBOIMManagedObjectFriendQOS::IMMixin_ptr myPolicyAsMOFQOS;
ByteString_var theKeyString;
IManagedClient::IManageable _var mptr;
IManagedClient::IHome_var myPolicyHome;

// Get the primary key (string) of the object in "theKeyString"
// Get the home configured for Policy objects in "myPolicyHome"
mptr = myPolicyHome->findByPrimaryKeyString(theKeyString);
myPolicyAsMOFQOS =
    IBOIMManagedObjectFriendQOS::IMMixin::_narrow(mptr)

// now use a method from this quality of service interface
myPolicyAsMOFWQOS->checkpointToDatastore();

```

The `checkpointToDatastore()` method is used as an example of a method that is available on the QOS interface. Exactly which methods are available on this interface depends on the particular Component Broker installation and the container in which the managed object resides.

If you already have an object reference to a Policy and want to get at the QOS interface, use the following code segment to narrow it down.

```

myPolicy->....; // run some business logic
myPolicy->....; // run some more business logic

// Now ManagedObjectFriendQOS capability is needed.
myPolicyAsMOFQOS =
    IBOIMManagedObjectFriendQOS::IMMixin::_narrow(mypolicy);
// now use a method from this quality of service interface
myPolicyAsMOFWQOS->checkpointToDatastore();

```

While the previous code-segments are understandable, they leave you saddled with the responsibility of maintaining two references to a single managed object. The first reference, `myPolicy`, has the business logic and `IManageable` interfaces, while the second, `myPolicyAsMOFQOS`, has the special QOS interface that relates to the quality of service available on this particular type of managed object.

An interface that combines the QOS quality of service interface with the business logic interface can be introduced. This makes it possible for clients to code to one interface for the duration of an application. A code segment of the IDL follows:

```

#include <IBOIMManagedObjectFriendsQOS.idl>
#include <Policy.idl>

interface PolicyMOFQOS: Policy, IBOIMManagedObjectFriendQOS::IMMixin
{
}

```

If the `PolicyMOFQOS` interface exists and is usable, then the previous code segments can be simplified as follows:

```

PolicyMOFQOS_ptr myPolicy;
ByteString *theKeyString;
IManagedClient::IManageable_ptr mptr;
IManagedClient::IHome_var myPolicyHome;

// Get the primary key (string) of the object in "theKeyString"
// Get the home configured for Policy objects in "myPolicyHome"

mptr = myPolicyHome->findByPrimaryKeyString(*theKeyString);
myPolicy = PolicyMOFQOS::_narrow(mptr);

// invoke business logic or IManageable methods on "myPolicy"
myPolicy->amount(25000.00);

// invoke any ManagedObjectFriendQOS methods that you want to ...
myPolicy->checkpointToDataStore();

```

You should make a decision for your applications and remain consistent. If you expect to use QOS interfaces rarely, then choose the option implied by the first set of code segments. This has the advantage of letting the application code be as independent as possible from the underlying implementation of the components. Only surfacing the QOS type when it is needed means that changes to the methods supported on this interface have minimal impacts on the application code. This interface should not often change. One reason to change is the movement of a particular component from one container to another. If application code is dependent on a specific container's quality of service interface, it must change. The design challenge is to minimize and isolate these usages to be able to deal with change easily when it occurs.

The advantage of using the combined interface (PolicyMOFQOS in the example) exclusively is one of housekeeping and simplicity. When declared, it provides complete access to all methods associated with the object. Narrowing or casting from one to the other and keeping track of two sets of references to the same object is not necessary.

Using QOS interfaces for non-transactional support

The server supports the concept of transactions in a number of ways. Containers are configured so that all objects within them have the same transactional semantics. Besides the variations that leverage the Transaction Service, one option exists in which there is no transactional capabilities.

In cases where objects live in a container that does not use transactions, use the following general programming model:

1. Find or create objects.
2. Make changes to those objects.
3. Use QOS interface to `checkpointToDataStore`.
4. If you want a refresh, then do `refreshFromDataStore`.

Behind the scene - an overview of application adaptors

An application adaptor provides a place for the managed objects of components, much like a database system provides a home for data or records. An application adaptor is similar to an object-oriented database; it is responsible for providing systems capabilities (for example, identity, caching, and persistence) for its managed objects. By providing such capabilities, an application adaptor provides a certain “quality of service” to its managed objects. Different application adaptors may differ in the trade-off between cost and their quality of service.

An application adaptor provides the following benefits:

- It provides a higher level of abstraction than the Object Services interfaces for the systems capabilities that it provides. This allows the flexibility that is very important for providing efficient support for objects.
- It enables Object Services to be packaged in a more integrated and efficient manner than is possible with Object Services that are provided separately.
- It can delegate many of its Object Services to existing resource managers, such as database systems, that typically provide systems capabilities in a very integrated way.
- It provides a boundary for administrative functions.

A specific type of application adaptor can only handle one type of legacy store. This means that a DB2 application adaptor provides support for managed objects that are persistent in DB2 tables. However, a different type of application adaptor would be required for managed objects that are stored through CICS.

Behind the scene - the BOIM application adaptor

Business object application adaptor (BOIM) is the term used to identify IBM's first application adaptor. BOIM fulfils the following needs:

- As an application adaptor it forms a framework for applications
- It is an implementation of the application adaptor framework

The BOIM application adaptor is targeted at a two-level store implementation model for servers and a single-level store model for clients.

An object is fully functional only when in memory; the object returns to a dormant state when removed from memory. The BOIM application adaptor provides memory management schemes for bringing objects into memory and removing them from memory. Such schemes mean that clients do not need to perform special memory management tasks. For example, the BOIM retrieves objects from the database or disk transparently to the clients.

The environment for managed objects supported by the BOIM application adaptor is geared toward having each data object provide its own access to the underlying backend permanent data store.

Adapting applications

The simplest applications to convert to the environment targeted by the BOIM application adaptor are those that have the following characteristics:

- The application is not widely distributed. That is, it is built with a traditional client-server model.
- The application has basic backend access requirements, for example, accessing files or having all the activity for an entire transaction within one database and server process.

Managed object assembly

The managed object assembly pattern supported by the BOIM application adaptor has a delegated mixin, a delegated data object, and a managed object that extends the business object implementation.

The BOIM application adaptor provides interfaces (methods) to the managed object without requiring the business object to provide implementation for those methods.

Managed object

Adapting a component to the BOIM application adaptor requires development of a managed object and a function that instantiates a managed object. The basic design of the managed object is to give it the appearance of being the component. Developing the managed object requires support for the following types of interfaces:

- An interface used by the BOIM application adaptor to communicate with the managed object instead of the component
- The business interfaces implemented in the business object
- The interface for which the BOIM application adaptor provides the implementation
- The interface for which no implementation is provided by the BOIM application adaptor but for which the environment requires a default implementation for which no implementation is provided by the BOIM application adaptor

Business object

The BOIM application adaptor supports business objects that support either of the two interfaces `IManagedServer::IManagedObjectWithDataObject` or `IManagedServer::IManagedObjectWithCachedDataObject` from the Component

Broker programming model. The interface supported by the business object is detected by the BOIM application adaptor and no special configuration is required.

Data object

The BOIM application adaptor provides a framework for data object development by providing a base class with implementation. All data objects are expected to extend this implementation and provide the implementation of the methods that it has defined.

The life cycle of components

Because the BOIM application adaptor is based on a two-level store, there is more to the life cycle of a component than creation and deletion. There are the following additional states for a component:

A passivated state

There is no representation of the component in memory, however the state data for the component is saved.

An activated state

There is no representation of the component in memory.

Qualities of service

The BOIM provides the following qualities of service:

Sharing Data with the Database

The BOIM application adaptor provides control information that helps data objects manage their synchronization with the database.

Objects are synchronized with the underlying data store on the following occasions:

- When the object is created, its corresponding data needs to be inserted into the underlying data store.
- When the object is removed, its corresponding data needs to be deleted from the underlying data store.
- When the data in the object becomes stale, the data needs to be retrieved from the underlying data store.
- When the underlying data store becomes stale, the underlying data store needs to be updated.

Simplifying Interaction with the Data Store

The BOIM application adaptor simplifies interaction with the underlying data store in the following ways:

- A base class for data objects that provides a check to ensure that the key is valid before invoking any of the database access commands

- A performance improvement by updating the database only if data has been changed

Memory Management

The BOIM application adaptor determines when objects are activated in memory and removed from memory.

Tolerance of Programming Errors

The BOIM application adaptor tolerates programming errors within applications.

Locking

One very useful quality of service provided by the BOIM application adaptor is an ability to have multiple copies of a single object. This allows multiple clients access the same object simultaneously by giving each client its own copy of the object. However, the fact that there are multiple copies is hidden from the client and each client thinks it has the only copy of the object. Another useful quality of service provided by the application adaptors is obtaining the locks in the database. All of this sounds great because locks are obtained and components can be used simultaneously and yet do not need to be thread safe, so what could be better? All of this does come with a price. There is a possibility of creating applications that result in dead locks as they are accessed by multiple clients.

The following coding practices will help prevent these dead lock situations from occurring.

Identify and label those methods that are read only. Especially when data is cached in the data object. When a method is not identified as read only, then a write lock is obtained as the transaction is committed. This can lead to multiple clients all trying to upgrade locks from read locks to write locks yet not being able to because of the others holding the read locks.

Delegate all data management to the data object (do not cache data in the business object). When the management of the data is delegated to the data object then it can detect when values are changed and only upgrade the lock when necessary.

Ensure that there is recovery code in the client to cover the case where the transaction will not commit. When a dead lock occurs in the database, or a time out is detected in the Component Broker run time, then all but one of the transactions involved in the deadlock should be rolled back. Leaving one to succeed and the others to fail. Those that failed should try again.

Use the optimistic mode of the cache. This will detect clashes immediately rather than wait for a timer to expire to determine that a deadlock condition exists.

Framework flows

The following samples show the order in which the framework invokes the `IManagedObject` methods for a given scenario.

Create – Business Object has Cached Data Object

1. Framework calls `internalizeFromPrimaryKey` or `internalizeFromCopyHelper` on data object.
2. Framework calls `initForCreation` on business object passing data object.
3. Framework calls `syncToDataObject` on business object.
4. Framework tells “DataObject” to insert.
5. Object is created and can now be used by client.

Create – Business Object has Data Object (Delegating)

1. Framework calls `internalizeFromPrimaryKey` or `internalizeFromCopyHelper` on data object.
2. Framework calls `initForCreation` on business object passing data object.
3. Framework tells “DataObject” to insert.
4. Object is created and can now be used by client.

Reactivation – Business Object has Cached Data Object (first touch)

1. Framework calls `internalizeFromPrimaryKey` or `internalizeFromCopyHelper` on data object.
2. Framework calls `initForReactivation` on business object passing data object.
3. Object reference returned to client.
4. Client invokes a method.
5. Framework tells “DataObject” to retrieve.
6. Framework calls `syncFromDataObject` on business object.

Reactivation – Find Scenario with Cached Data Object

1. Framework calls `internalizeFromPrimaryKey` or `internalizeFromCopyHelper` on data object.
2. Framework calls `initForReactivation` on business object passing data object.
3. Framework tells “DataObject” to retrieve.
4. Framework calls `syncFromDataObject` on business object.
5. Object reference return to client.
6. Client invokes a method.

Remove Object

1. Framework calls `uninitForDestruction` on business object.
2. Framework tells “DataObject” to `deleteFromDataStore`.
3. Client can no longer use the object reference.

Passivation – Business Object with Cached Data Object

1. Framework calls `syncToDataStore` on business object.

2. Framework tells “DataObject” to update.
3. Framework calls `uninitForPassivation`.

Passivation – Business Object with Data Object

1. Framework tells “DataObject” to update.
2. Framework calls `uninitForPassivation`.

Appendix A. Artifacts produced in building objects

The following table summarizes the artifacts of the development process for object components (interface, business, and managed objects).

Table 21. Artifacts of the development process for object components

Artifact Type	Artifact Name		
	Abstraction Interface	AbstractionBusiness Object	AbstractionManaged Object
IDL	Abstraction.idl	AbstractionBO.idl	AbstractionMO.idl (Object Builder builds these automatically for BOIM managed objects that customers write.)
.hh (language usage binding) Never modify these – have the makefile generate them.	Abstraction.hh (contains the Abstraction_Skeleton class and the Abstraction C++ class)	AbstractionBO.hh (contains the AbstractionBO_Skeleton class and the Abstractionbusiness object C++ class)	AbstractionMO.hh (contains the AbstractionMO_Skeleton class and the AbstractionManaged object C++ class)
_C.cpp (client side binding) Never modify these – have the makefile generate them.	Abstraction_C.cpp	AbstractionBO_C.cpp	AbstractionMO_C.cpp
_S.cpp (server side bindings) These #include the _C.cpp. Never modify these – have the makefile generate them.	Abstraction_S.cpp	Abstraction_S.cpp	Abstraction_S.cpp
.ih (implementation interface - emit once or enter manually)	Not needed	AbstractionBO.ih (defines AbstractionBO_Impl that inherits only from AbstractionBO_Skeleton) (and inherits from IManageable_Impl)	AbstractionMO.ih (defines AbstractionMO_Impl that inherits from AbstractionBO_Impl and AbstractionMO_Skeleton) (Object Builder builds these automatically for BOIM managed objects that customers write.)

Table 21. Artifacts of the development process for object components (continued)

Artifact Type	Artifact Name		
	Abstraction Interface	AbstractionBusiness Object	AbstractionManaged Object
_I.cpp (implementation code - emit once or enter manually)	Not needed	AbstractionBO_I.cpp (implementation of real logic – the code)	AbstractionMO_I.cpp (delegator, traffic cop) (Object Builder builds these automatically for BOIM managed objects that customers write.)
_C.obj (from _C.cpp – need rule in makefile)	Abstraction_C.o (primary binding used by remote clients. See note.)	Not needed	Not needed
_S.obj (placed on servers; from _S.cpp – need rule in makefile)	Abstraction_S.o	AbstractionBO_S.o	AbstractionMO_S.o
_I.obj (placed on servers; from _I.cpp – need rule in makefile)	Not possible	AbstractionBO_I.o	AbstractionMO_I.o
CLIENT.LIB (group one or more abstractions into an LIB)	contributes its _C.o only (see note)	no contribution	no contribution
ASERVER.LIB (group one or more abstractions into an LIB)	contributes _S.o	contributes _S.o and _I.o	contributes _S.o and _I.o
<p>Note: If the same DLL is to be used on both client and server (which makes sense in the VisualAge for C++ case), then using the Abstraction_S.o for the client DLL allows this same DLL to be used on both client and server.</p>			

Appendix B. Interface Definition Language

The interface to a class of objects contains the information that a caller must know to use an object, specifically, the names of its attributes and the signatures of its operations. The interface is described in a formal language independent of the programming language used to implement the object's operations. The formal language used to define object interfaces is the Interface Definition Language (IDL), standardized by CORBA.

The implementation of a class of objects (that is, the procedures that implement operations and the variables used to store an object's state) is written in the implementor's preferred programming language (for example, C++ or Java).

A completely implemented class definition consists of the following parts:

- An IDL specification of the interface to instances of the class: the interface definition file (or IDL file).
- Method procedures written in the implementor's language of choice: the implementation file(s).

The IDL compiler takes as input an object interface definition file (the IDL file) and produces binding files that make it convenient to implement and use objects that support the defined interface within a particular programming language.

Note: Component Broker is based on CORBA Version 2.0. All IDL used with Component Broker must be CORBA 2.0-compliant without IDL extensions.

IDL name scoping

The IDL file forms a naming scope (or scope). Modules, interface statements, structures, unions, operations, and exceptions form nested scopes. An identifier can only be defined once in a particular scope. Identifiers can be redefined in nested scopes.

Names can be used in an unqualified form within a scope, and the name will be resolved by successively searching the enclosing scopes. Once an unqualified name is defined in an enclosing scope, that name cannot be redefined.

Fully qualified names are of the form:

scope-name::identifier

For example, operation name `meth` defined within interface `Test` of module `M1` would have the fully qualified name:

`M1::Test::meth`

A qualified name is resolved by first resolving the *scope-name* to a particular scope, *S*, and then locating the definition of *identifier* within that scope. Enclosing scopes of *S* are not searched.

Qualified names can also take the form:

::identifier

These names are resolved by locating the definition of *identifier* within the outermost name scope.

Every name defined in an IDL specification is given a global name, constructed as follows:

- Before the IDL Compiler scans the IDL file, the name of the current root and the name of the current scope are empty. As each module is encountered, the string `::` and the module name are appended to the name of the current root. At the end of the module, they are removed.
- As each interface, struct, union, or exception definition is encountered, the string `::` and the associated name are appended to the name of the current scope. At the end of the definition, they are removed. While parameters of an operation declaration are processed, a new unnamed scope is entered so that parameter names can duplicate other identifiers.
- The global name of an IDL definition is then the concatenation of the current root, the current scope, a `::`, and the local name for the definition.

The names of types, constants, and exceptions defined by base interfaces are accessible in a derived interface. References to these names must be unambiguous. Ambiguities can be resolved by using a scoped name (prefacing the name with the name of the interface that defines it, and the characters `::`, as in *base-interface::identifier*). Scope names can also be used to refer to a constant, type, or an exception name defined by a base interface but redefined by a derived interface.

Type and constant declarations

IDL specifications may include type declarations and constant declarations as in C and C++, with the restrictions and extensions described below. IDL supports the following declarations.

Integral types

IDL supports only the integral types short, long, unsigned short, and unsigned long, which represent the following value ranges:

- short $-2^{15} .. (2^{15})-1$
- long $-2^{31} .. (2^{31})-1$
- unsigned short $0 .. (2^{16})-1$
- unsigned long $0 .. (2^{32})-1$

Floating point types

IDL supports the float and double floating-point types. The float type represents the IEEE single-precision floating-point numbers; double represents the IEEE double-precision floating-point numbers.

Since returning floats and doubles by value may not be compatible across Microsoft Windows[®] compilers, client programs should return floats and doubles by reference.

Character type

IDL supports a char type, which represents an 8-bit quantity. The ISO Latin-1 (8859.1) character set defines the meaning and representation of graphic characters. The meaning and representation of null and formatting characters is the numerical value of the character as defined in the ASCII (ISO 646) standard. Unlike C/C++, type char cannot be qualified as signed or unsigned. (The octet type, below, can be used in place of unsigned char.)

Boolean type

IDL supports a boolean type for data items that can take only the values zero (FALSE) and one (TRUE).

Octet type

IDL supports an octet type, an 8-bit quantity guaranteed not to undergo conversion when transmitted between a client and server process. The octet type can be used in place of the unsigned char type.

Any type

IDL supports an any type, which permits the specification of values of any IDL type. Conceptually, an any consists of a value and a TypeCode that represents the type of the value. The TypeCode class provides functions for obtaining information about an IDL type.

Constructed types (struct, union, enum)

In addition to the above basic types, IDL also supports three constructed types: struct, union, and enum. The structure and enumeration types are specified in IDL just as they are in C and C++, with the following restrictions:

- Unlike C/C++, recursive type specifications are allowed only through the use of the sequence template type (see below).
- Unlike C/C++, structures, discriminated unions, and enumerations in IDL must be tagged. For example, `struct { int a; ... }` is an inappropriate type specification (because the tag is missing). The tag introduces a new type name.
- In IDL, constructed type definitions need not be part of a typedef statement; furthermore, if they are part of a typedef statement, the tag of the struct must differ from the type name being defined by the typedef. For example, the following are valid IDL struct and enum definitions:

```
struct myStruct {
    long x;
    double y;
};
/* defines type name myStruct */
enum colors { red, white, blue };
/* defines type name colors */
```

The following IDL definitions are **not** valid:

```
typedef struct myStruct {
    /* NOT VALID */
    long x;
    /* Tag myStruct is the same */
    double y;
    /* as the type name below; */
} myStruct;
/* myStruct has been redefined */
typedef enum colors { red, white, blue } colors;
/* NOT VALID */
```

Union type

IDL also supports a union type, which is a cross between the C union and switch statements. The syntax of a union type declaration is as follows:

```
union identifier switch (switch-type) { case+ }
```

The *identifier* following the union keyword defines a new legal type. (Union types may also be named using a typedef declaration.)

The *switch-type* specifies an integral, character, boolean, or enumeration type, or the name of a previously defined integral, boolean, character or enumeration type.

Each *case* of the union is specified with the following syntax:

```
case-label+ type-spec declarator;
```

Where

- Each *caselabel* has one of the following forms:

case *const-expr*:

default: The *const-expr* is a constant expression that must match or be automatically castable to the *switch-type*. A default case can appear no more than once.

- *type-spec* is any valid type specification.
- *declarator* is an identifier or an array declarator (such as, `foo[3][5]`).

Template types (sequences and strings)

IDL defines two template types not found in C and C++: sequences and strings. A sequence is a one-dimensional array with two characteristics: an optional maximum size (specified at compile time) and a length (determined at run time). Sequences permit passing unbounded arrays between objects. Sequences are specified as follows:

- `sequence < simple-type [, positive-integer-const] >`
where *simple-type* specifies any valid IDL type, and the optional *positive-integer-const* is a constant expression that specifies the maximum size of the sequence (as a positive integer).

A string is similar to a sequence of type `char`. It can contain all possible 8-bit quantities except NULL. Strings are specified as follows:

- `string [< positive-integer-const >]`
where the optional *positive-integer-const* is a constant expression that specifies the maximum size of the string (as a positive integer, which does not include the extra byte to hold a NULL as required in C/C++).
- CORBA does not prescribe specific rules on how to process blanks contained within strings. Thus, Component Broker needs help in determining whether keys “ABC” and “ABC” (Insert three trailing blanks after the second “ABC”.) refer to the same or different managed objects. To aid with this decision, Component Broker offers multiple semantic choices for processing string attributes when they are defined on a business object. When multiple string attributes exist within a single business object, mixing and matching of the various semantics will be allowed. The three semantic choices are:
 - CORBA - this is consistent with the support provided through release 1.3 of Component Broker. No rules are defined for processing blanks within these strings.
 - Trailing blanks stripped - object builder will generate code that removes trailing blanks.

- Pad with blanks to fixed length - object builder will generate code that adds trailing blanks to some bounded string length.

Customers are discouraged from using the CORBA semantics for strings to be used as key attributes. Customers must follow similar semantic rules as enforced by our generated code for any string extensions they implement.

Arrays

Multidimensional, fixed-size arrays can be declared in IDL as follows:

```
identifier { [ positive-integer-const ] }+
```

where the *positive-integer-const* is a constant expression that specifies the array size, in each dimension, as a positive integer. The array size is fixed at compile time.

Object types

The name of the interface to a class of objects can be used as a type name. For example, if an IDL specification includes an interface declaration (described below) for a class (of objects) C1, then C1 can be used as a type name within that IDL specification. When used as a type, an interface name indicates a reference to an object that supports that interface. An interface name can be used as the type of an operation argument, as an operation return type, or as the type of a member of a constructed type (a struct, union, or enum). In all cases, the use of an interface name indicates a reference to (as opposed to an instance of) an object that supports that interface.

Constants

Constants are declared in IDL just as in C++, except that the type of the constant must be a valid IDL type. IDL Constant declarations take the following form:

```
const <const-type> identifier=<constant-expression>;
```

The *const-type* must be a valid IDL integer, char, boolean, floating point, string, or user-defined type name. The *identifier* is the name of the constant being defined. The *constant-expression* is a constant expression as in C/C++, and can include the usual C/C++, unary and binary operators (|, ^, &, >>, <<, +, -, *, /, %, ~), parentheses for controlling operator precedence, literal values (integer, string, character, and floating point) as in C/C++, and the boolean literal values TRUE and FALSE.

Interface declarations

The IDL specification for a class of objects must contain a declaration of the interface these objects will support. When objects are implemented using classes, the interface name is used as a class name as well. In addition to the interface name and its base interface names, an interface indicates new methods (operations), and any constants, type definitions, and exception structures that the interface exports. An interface declaration has the following syntax:

```
interface interface-name [: base-interface1, base-interface2, ...]
{
    constant declarations (optional)
    type declarations (optional)
    exception declarations (optional)
    attribute declarations (optional)
    operation declarations (optional)
};
```

The base-interface names specify the interfaces from which *interface-name* is derived. Parent-interface names are required only for the immediate base interface(s). Each base interface must have its own IDL specification (which must be `#included` in the IDL file). A base interface cannot be named more than once in the interface statement header.

In general, an interface header must precede any subsequent references to . For a discussion of multiple interface statements, see [Multiple IDL Interfaces and Modules](#).

The topics listed below describe the various declarations/statements that can be specified within the body of an interface declaration. The order in which these declarations are specified is usually optional, and declarations of different kinds can be intermixed. Although all of the declarations/statements are listed above as optional, in some cases using one of them may mandate another. For example, if an operation raises an exception, the exception structure must be defined beforehand. In general, types, constants, and exceptions, as well as interface declarations, must be defined before they are referenced, as in C/C++.

- “Attribute declarations” on page 457
- “Comments” on page 460
- “Type and constant declarations” on page 448
- “Operation declarations” on page 454

Constant, type, and exception declarations within an interface

The form of a constant, type, or exception declaration within the body of an interface declaration is the same as described in “Interface declarations” on page 453. Constants and types defined within an interface are transferred by the IDL compiler to the binding files it generates for that interface.

Types, constants, and exceptions defined in a base interface are accessible to the derived interface. References to them, however, must be unambiguous. Ambiguities can be resolved by using a fully scoped name, (prefacing a name with the name of the interface that defines it) separated by the characters “::” as illustrated below:

```
MyBaseInterface::myType
```

A leading “::” can be used to fully-qualify a reference starting from the outermost name scope.

The derived interface can redefine any of the type, constant, and exception names that were inherited. The derived interface cannot, however, redefine attributes or operations. To refer to a constant, type, or exception “name” defined by a base interface and redefined by “interface-name,” use the “parent-name::name” syntax.

Operation declarations

Operation declarations define the interface of each operation introduced by the interface. (An IDL operation is typically implemented by a method in the implementation programming language. Hence, the terms operation and method are often used interchangeably.) An operation declaration is similar to a C++ virtual function definition:

```
[ oneway ] type-spec identifier ( parameter-list ) [ raises-expr ] [ context-expr ] ;
```

where

identifier is the name of the operation.

type-spec is any valid IDL type, except a sequence, or the keyword void, indicating that the operation returns no value. (Although the return type cannot be a sequence, it can be a user-defined type that is a sequence.)

Unlike C and C++ procedures, operations that do not return a result must specify void as their return type.

The remaining syntax of an operation declaration is elaborated in the following subtopics.

"oneway" keyword

The optional oneway keyword specifies that when a caller invokes the operation, no reply is expected or received. The invocation semantics of a oneway operation are best-effort, which does not guarantee delivery of the call. Best-effort implies that the operation will be invoked at most once. A oneway operation must not have any output parameters and must have a return type of void. A oneway operation also must not include a raises expression (see below).

If the oneway keyword is not specified, then the operation has *at-most-once* invocation semantics if an exception is raised, and it has *exactly-once* semantics if the operation succeeds. This means that an operation that raises an exception has been implemented zero or one times, and an operation that succeeds has been implemented exactly once.

Parameter list

The parameter-list contains zero or more parameter declarations for the operation, delimited by commas. (The target object for the operation is not explicitly specified as an operation parameter in IDL.) If there are no explicit parameters, the syntax "(" must be used, rather than "(void)". A parameter declaration has the following syntax:

```
{ in | out | inout } type-spec declarator
```

where *type-spec* is any valid IDL type (except a sequence), and *declarator* is an identifier or an array declarator. Although the type of a parameter cannot be a sequence, it can be a user-defined type that is a sequence.

The required in|out|inout directional attribute indicates whether the parameter is to be passed from caller to callee (in), from callee to caller (out), or in both directions (inout). The following are examples of valid operation declarations:

```
short meth1(in char c, out float f);
oneway void meth2(in char c);
float meth3();
```

An operation's implementation should not modify an in parameter. If a change must be made by the implementation, the implementation should copy the parameter and only modify the copy.

If an operation raises an exception, the values of the return result and the values of the out and inout parameters (if any) are undefined.

"raises" expression

The optional raises expression in an IDL operation declaration indicates which exceptions the operation may raise. A raises expression is specified as follows:

```
raises ( identifier1, identifier2, ... )
```

where each *identifier* is the name of a previously defined exception. In addition to the exceptions listed in the raises expression, an operation may also signal any of the standard exceptions. Standard exceptions, however, should not appear in a raises expression. If no raises expression is given, then an operation can raise only the standard exceptions. "Exception declarations" on page 457 contains further information on defining exceptions and the list of standard exceptions.

"context" expression

The optional context expression (context-expr) in an operation declaration indicates which elements of the caller's context the operation's implementation may consult. A context expression is specified as follows:

```
context ( identifier1, identifier2, ... )
```

where each *identifier* is a string literal made up of alphanumeric characters, periods, underscores and asterisks. The first character must be alphabetic, and an asterisk can only appear as the last character, where it serves as a wildcard matching any characters. If convenient, identifiers may consist of period-separated valid identifier names, but that form is optional.

The Context is a special object that is specified by the CORBA standard. It contains a property list: a set of property-name/string-value pairs that the caller can use to store information about its environment that operations may find useful. It is used in much the same way as environment variables. It is passed as an additional parameter to operations that are defined as context-sensitive in IDL.

The context expression of an operation declaration in IDL specifies which property names the operation uses. If these properties are present in the Context object supplied by the caller, they will be passed to the object implementation, which can access them through the interface of the Context object.

The argument that is passed to the operation having a context expression is a Context object, not the names of the properties. The caller must create a Context object and use the interface of the Context object to set the context properties. The caller then passes the Context object in the operation

invocation. The CORBA standard allows properties in addition to those in the context expression to be passed in the Context object.

Attribute declarations

Declaring an attribute as part of an interface is equivalent to declaring one or two accessor operations: one to retrieve the value of the attribute (a get or read operation) and (unless the attribute specifies readonly) one to set the value of the attribute (a set or write operation).

Attributes are declared as follows:

```
[ readonly ] attribute type-spec declarators;
```

where:

type-spec specifies any valid IDL type (except a sequence).

declarators is a list of identifiers, delimited by commas. An array declarator cannot be used directly when declaring an attribute, but the type of an attribute can be a user-defined type that is an array. Although the type of an attribute cannot be a sequence, it can be a user-defined type that is a sequence. The optional `readonly` keyword specifies that the value of the attribute can be accessed but not modified. (In other words, a `readonly` attribute has no set operation.) Below are examples of attribute declarations, which are specified within the body of an interface statement:

```
interface Goodbye: Hello
{
    void sayBye();
    attribute short xpos;
    attribute char c1, c2;
    readonly attribute float xyz;
};
```

Attributes are inherited from base interfaces. An inherited attribute name cannot be redefined to be a different type.

Exception declarations

IDL specifications can include exception declarations, which define data structures to be returned when an exception occurs during the execution of an operation. A name is associated with each type of exception. Optionally, a struct-like data structure for holding error information can also be associated with an exception. Exceptions are declared as follows:

```
exception identifier
{
    member*
};
```

The *identifier* is the name of the exception, and each *member* has the following form:

```
type-spec declarators ;
```

The *type-spec* is a valid IDL type specification and *declarators* is a list of identifiers or array declarators, delimited by commas. The members of an exception structure should contain information to help the caller understand the nature of the error. The exception declaration can be treated like a struct definition: whatever you can access in an IDL struct, you can access in an IDL exception. Unlike a struct, an exception can be empty, meaning the exception is just identified by its name.

If an exception is returned as the outcome of an operation, the exception *identifier* indicates which exception occurred. The values of the members of the exception provide additional information specific to the exception. “Operation declarations” on page 454 describes how to indicate that a particular operation may raise a particular exception.

The following is an example showing the declaration of a BAD_FLAG exception:

```
exception BAD_FLAG
{
    long ErrCode; char Reason[80];
};
```

In addition to user-defined exceptions, there are several predefined exceptions for system run-time errors. The standard exceptions as prescribed by CORBA are subclasses of CORBA::SystemException. These exceptions correspond to standard run-time errors that may occur during the execution of any operation (regardless of the list of exceptions listed in the operation’s IDL specification).

Each of the standard exceptions has the same structure: an error code (to designate the subcategory of the exception) and a completion status code. For example, the NO_MEMORY standard exception has the following definition:

```
enum completion_status
{
    COMPLETED_YES, COMPLETED_NO, COMPLETED_MAYBE
};
exception NO_MEMORY
{
    unsigned long minor;
    completion_status completed;
};
```

The "completion_status" value indicates whether the operation was never initiated (COMPLETED_NO), if the operation completed its execution prior to the exception (COMPLETED_YES), or if the operation's completion status is indeterminate (COMPLETED_MAYBE).

IDL syntax

This section describes the syntax of the Interface Definition Language (IDL), as specified by the CORBA standard. This section describes the syntax and semantics of IDL using the following conventions:

- bold** Indicates Literals (such as keywords).
- italics* Indicate user-supplied elements.
- { } Groups related items together as a single item.
- [] Encloses an optional item.
- * Indicates zero or more repetitions of the preceding item.
- + Indicates one or more repetitions of the preceding item.
- | Separates alternatives.
- _ Within a set of alternatives, an underscore indicates the default, if defined.

IDL is a formal language used to describe object interfaces. An IDL definition specifies, for a class of objects, what methods (operations) are available, their return types, and their parameter types. For this reason, we often speak of an IDL specification for a class (as opposed to simply an object interface).

IDL generally follows the same lexical rules as C and C++. Exceptions to C++ lexical rules include:

- IDL uses the ISO Latin-1 (8859.1) character set.
- White space is ignored except as token delimiters.
- C and C++ comment styles are supported.
- IDL supports standard C/C++ preprocessing, including macro substitution, conditional compilation, and source file inclusion.
- Identifiers (user-defined names for operations, attributes, instance variables, and so on) are composed of alphanumeric and underscore characters (with the first character alphabetic) and can be of arbitrary length, up to an operating-system limit of about 250 characters.
- Identifiers must be spelled consistently with respect to case throughout a specification.
- Identifiers that differ only in case yield a compilation error.

- Within a particular name scope, there is a single name space for all identifiers, regardless of their type. For example, using the same identifier for a constant and an interface name within the same name scope yields a compilation error.
- Integer, floating point, character, and string literals are defined as in C and C++.

The terms listed in Table 22 are reserved keywords and may not be used otherwise. Keywords must be spelled using upper- and lower-case characters exactly as shown in the table. For example, “void” is correct, but “Void” yields a compilation error.

Table 22. Reserved keywords for IDL

any	default	FALSE	oneway	read-only
attribute	double	float	out	sequence
boolean	enum	in	raises	short
case	exception	inout	unsigned	string
char		interface	union	struct
const		long	void	switch
context		module		TRUE
		Object		
		octet		typedef

A typical IDL specification for a single interface, residing in a single IDL file, has a form which includes the following specifications listed.

- “Interface declarations” on page 453 (optional)
- “Exception declarations” on page 457 (optional)
- “Include directives” on page 461 (optional)
- “Type and constant declarations” on page 448 (optional)

The order is unimportant, except that interface names must be declared (or forward referenced) before they are referenced.

For more information on the CORBA standard for IDL, see *The Common Object Request Broker: Architecture and Specification* Object Management Group, Inc. (OMG) and x/Open .

Comments

IDL supports both C and C++ comment styles. The characters “//” start a line comment, which finishes at the end of the current line. The characters “/*”

start a block comment that finishes with `*/`. Block comments do not nest. The two comment styles can be used interchangeably.

Because comments appearing in an IDL specification may be transferred to the files that the IDL Compiler generates, and because these files are often used as input to a programming language compiler, avoid using characters that are not generally allowed in comments of most programming languages. For example, the C language does not allow `*/` to occur within a comment, so its use is to be avoided, even when using C++ style comments in the IDL file.

IDL also supports throw-away comments. They may appear anywhere in an IDL specification. Throw-away comments start with the string `//#` and end at the end of the line. Use throw-away comments to comment out portions of an IDL specification.

Include directives

The IDL specification for an interface normally contains `#include` statements that tell the IDL compiler where to find the interface definitions (the IDL files) for each of the interface's parent (direct base) interfaces and for other referenced types.

The file `orb.idl` can be included to access IDL types defined by the CORBA specification that are not IDL keywords.

As in C and C++, if a filename is enclosed in angle brackets (`[]`), the search for the file begins in system-specific locations. If the filename is in double quotation marks (`""`), the search for the file begins in the current working directory, before searching the system-specific locations.

In addition to the `#include` directive, other preprocessor directives can be used in IDL.

Pragma directives

Component Broker supports the following pragmas.

- `localonly`
- `localonly abstract`
- `cpponly`
- `init`
- `ID`
- `Prefix`
- `version`

localonly pragma

This Component Broker unique pragma supports the generation of bindings for objects that are known to be local (not distributed). This pragma may occur at any point in the IDL file following the definition or forward declaration of the designated interface.

The syntax is:

```
#pragma meta interface-name localonly
```

The IDL interface identified by *interface-name* is treated by generated bindings as strictly local to the caller's process. No calls to the CORBA ORB occur when invoking the operations defined in this interface. *interface-name* may be a simple name of an interface in the current scope or a fully- or partially-qualified interface name. The interface must be previously defined or forward declared when the pragma statement is encountered.

localonly abstract pragma

This Component Broker unique pragma is like the localonly pragma, but it signifies an abstract function that cannot be instantiated. These types of interfaces are used to just define interfaces.

The syntax is:

```
#pragma meta interface-name localonly abstract
```

cpponly pragma

This Component Broker unique pragma suppresses the generation of IOM interlanguage bindings.

The syntax is:

```
#pragma meta interface_name cpponly
```

In the default case, without this pragma, two sets of bindings are produced:

- The standard CORBA C++ bindings suitable for use with the ORB component.
- IOM bindings suitable for interlanguage interaction.

Without this pragma, only the standard CORBA C++ bindings are produced.

init pragma

This Component Broker unique pragma specifies a function to use to initialize newly created objects.

The syntax is:

```
#pragma meta method-name init
```

This pragma allows the IDL to specify the name of a function to be used to initialize the newly created method. When this pragma is not used, the emitters produce a `_create()` function that takes no parameters and does no initialization after the new object is created.

For example, if the IDL contains:

```
interface A
{
    void initFunction(int);
};
#pragma meta A::initFunction init
```

the C++ class `A` that implements interface `A` will have a `_create()` function that takes an `int` parameter (because `initFunction` takes an `int`). Also, the code inside `_create(int)` creates a new instance of class `A` and then call `initFunction(int)` on the newly created object, passing along its `int` parameter.

ID pragma

This CORBA-defined pragma overrides the default `RepositoryID` for an IDL entity.

The syntax is:

```
#pragma ID scoped-name literal-string
```

which sets the `RepositoryID` of *scoped-name* to *literal-string* instead of the default `RepositoryID`.

Prefix pragma

This CORBA-defined pragma sets the `RepositoryID` prefix

The syntax is:

```
#pragma prefix string
```

which sets the current prefix used in generating OMG IDL format `RepositoryIDs`. The specified prefix applies to `RepositoryIDs` generated after the pragma until the end of the current scope is reached or another prefix pragma is encountered.

version pragma

This CORBA-defined pragma sets the RepositoryID version number.

The syntax is:

```
#pragma version scoped-name major.minor
```

which uses the *major.minor* as the version number for RepositoryID of the *scoped-name*.

Multiple IDL interfaces and modules

A single IDL file can define multiple interfaces. When a file defines two or more interfaces that reference one another, forward declarations can be used to declare the name of an interface before it is defined. This is done as follows:

```
interface interfaceName ;
```

The actual definition of the interface for *interfaceName* must appear later in the same IDL file.

If multiple interfaces are defined in the same IDL file, they can be grouped into modules, by using the following syntax:

```
module moduleName { definition+ };
```

where each definition is a type declaration, constant declaration, exception declaration, interface statement or nested module statement. Modules are used to scope identifiers.

Alternatively, multiple interfaces can be defined in a single IDL file without using a module to group the interfaces. Whether a module is used for grouping multiple interfaces or not, the languages bindings produced from the IDL file will include support for all of the defined interfaces.

The idlc command

Creates usage and implementation bindings for interfaces described in IDL files.

The Interface Definition Language Compiler (`idlc`) command compiles one or more files containing CORBA 2.0-compliant IDL statements, and optionally produces generated language bindings appropriate to each named input file.

The syntax for the `idlc` command is:

```
idlc [options] <filename>...
```

Where:

[options]

See Options for the idlc command for a complete list of the options, their usage, and their restrictions.

<filename>

<filename> may be specified without a file name extension; if no file name extension is supplied, it is assumed to be ".idl". The wildcard character "*" is permitted to appear once in the non-path portion of the file name. For example, the following are acceptable ways to refer to the file "xyz.idl" in directory "E:\idl\src":

```
E:\idl\src\xyz.idl
E:\idl\src\xyz
E:\idl\src\*.idl (this may refer to additional files as well)
E:\idl\src\x*.idl (all files starting with x)
xyz.idl (if E:\idl\src is the current directory)
xyz
x*
```

When all specified input files are compiled, the idlc command returns a value of zero if no errors were detected; otherwise, a non-zero value is returned.

If no emitters are specified for the idlc command, then only the syntax of the named files is checked, and any errors reported. (See the discussion of the -e option (IDL Command options) for how to specify emitters.) When a compilation error (but not a warning) is detected for a particular input file, the emit phase for that file is skipped.

For additional information on emitted files, see Emitted file names.

Options for the idlc command

Options for the idlc command are preceded with a dash (-) character and may be specified individually or run together. For example, -p -v -V or -pvV is acceptable.

Some options accept an argument. These options must either be specified individually or as the last option in a run-together grouping (for example, -p -m tie or -pm tie). The space between this type of option and its argument is optional. For example, either -mtie or -m tie is equally acceptable.

All options are case-sensitive, even on platforms where file names are not case-sensitive.

Table 23 on page 466 describes each available option:

Table 23. idlc command options

Option	Description
-d <directory-name>	Specifies the directory in which to place emitted output files and directories. If none is specified, the default is the current directory.
-V	Shows the version number of the idlc command.
-v	Specifies verbose mode. This shows all internal commands (and their arguments) issued by the idlc command.
-?	Writes a brief description of the idlc command syntax to standard output.
-h	Synonymous with -? .
-D <define-expression>	Predefines a preprocessor variable for the IDL compiler.
-I <include-directory>	Adds a directory to the list of directories used by the IDL compiler to find #include files. In addition to the -I option, the IDLC_INCLUDE environment variable can be used to specify a list, with <include-directory> names separated by the PATH separator character.
-i <file-name>	Specifies the name of a file to be compiled that does not have the .idl extension. The <file-name> should not have an implicit .idl suffix added to its name.
-p	Used as a shorthand for -D__PRIVATE__ .

Table 23. *idlc* command options (continued)

Option	Description
-e <emit-list>	<p>Specifies a list of emitters to run. Emitters are specified with a short 2- or 3-character designator. Each emitter in the list should be separated from the others with a colon (:) or semicolon (;) character. Valid emitter names are:</p> <p>hh Produces C++ usage bindings. If no modifiers are present, bindings with support for remotable cross-language operation are produced. The <i>cpponly</i>, <i>localonly</i>, and <i>somthis</i> modifiers cause specialized bindings to be produced (see <code>-m<name[=value]></code>).</p> <p>sc Produces a C++ skeleton for the Basic Object Adapter of the ORB. If no modifiers are present, bindings with support for remotable cross-language operation are produced. The <i>cpponly</i>, <i>localonly</i>, and <i>somthis</i> modifiers cause specialized bindings to be produced (see <code>-m<name[=value]></code>).</p> <p>uc Produces local implementations needed by the C++ usage bindings. If no modifiers are present, bindings with support for remotable cross-language operation are produced. The <i>cpponly</i>, <i>localonly</i>, and <i>somthis</i> modifiers cause specialized bindings to be produced (see <code>-m<name[=value]></code>).</p> <p>ih Produces a C++ implementation header. If the <i>mo</i> modifier (see <code>-m<name[=value]></code>) is specified, bindings that support Component Broker managed objects are produced; otherwise, pure CORBA C++ bindings, suitable for use with a standalone ORB, are produced.</p> <p>ic Produces a Component Broker C++ managed object implementation template. The <i>mo</i> modifier has the same effect as the <i>ih</i> emitter.</p> <p>uj Produces the cross-language Java usage bindings.</p> <p>sj Produces a Java implementation skeleton.</p>

Table 23. idlc command options (continued)

Option	Description
-e <emit-list> (continued)	<p>bj Creates files needed to support business objects written in Java. The files are</p> <ul style="list-style-type: none"> • <code>_<i><interface_name></i>Wrapper.java</code> that replaces <code>_<i><interface_name></i>ImplBase.java</code> • <code>_<i><interface_name></i>Impl.java</code> that is the implementation-side proxy for the C++ managed object associated with the Java business object. <p>ir Updates the CORBA Interface Repository with the interfaces in this compilation unit.</p> <p>The idlc command looks for emitters specified by the -e or -s flags and looks for an <emit-list> in an environment variable named IDLC_EMIT. If no <emit-list> can be found from any source, no emitters are run, but the IDL compiler is invoked to check for syntax errors in the input files.</p>
-s <emit-list>	Synonymous with -e.
-m <name[=value]>	<p>Specifies an output modifier. A modifier may be given as a name or a name=value expression. The emitters are sensitive to the following modifiers:</p> <p>LINKAGE=<value> Used to insert customized C++ linkage modifiers into the generated bindings.</p> <p>notconsts Eliminates the generation of C++ TypeCode constants and overloaded any operators.</p> <p>tie Generates “tie-style” bindings that assume delegation rather than inheritance.</p> <p>cpponly Suppresses the production of cross-language bindings and produces standard CORBA C++ bindings suitable for use with a standalone ORB. cpponly affects the bindings produced by the hh, sc, and uc emitters.</p> <p>localonly Generates bindings that can only be used to access a local object for all of the most-derived interfaces in the IDL file. For an alternative mechanism that also offers finer control over the affected interfaces, see “localonly” Pragmas.</p> <p>orbadapter Generates C++ bindings that allow the C++ ORB to dispatch Java implementations.</p>

Table 23. *idlc* command options (continued)

Option	Description
-m <name[=value]> (continued)	<p>IRforce Forces the IR emitter to destroy objects already present in the IR with the same name as in the IDL being produced.</p> <p>dllname=<value> Puts NT import/export specifications into classes contained in the DLL named by <value>.</p> <p>preInclude=<file-name> Adds the line: #include <file-name></p> <p>to the .hh file, just before the line that includes corba.h.</p> <p>postInclude=<file-name> Adds the line: #include <file-name></p> <p>just before the end of the .hh file.</p>
-J	<p>Passes options through to the Java interpreter used internally. For example:</p> <pre data-bbox="467 852 583 874">-J" -mx32m"</pre> <p>sets the heap size for the interpreter to 32M.</p>

Emitted file names

The *idlc* command process IDL files and produces output files that contain language-specific usage and implementation bindings for the IDL interface. Each emitter (see *-emit-list* in Options for the *idlc* command) produces one or more output files. The rules used to generate the names of these output files are described in the following sections.

C++ emitters

The names of the generated output files are derived from the file name of the corresponding IDL file. For a file named *filestem.idl*, the following list of output files may be emitted when the *idlc* command is run. The list contains the emitter and its corresponding output file name.

```

hh    filestem.hh
sc    filestem_S.cpp
uc    filestem_C.cpp
ih    filestem.ih
    
```

ic `filestem_I.cpp`

Java emitters

The Java emitters produce multiple output files as defined by the CORBA Java bindings. These file names are unrelated to the name of the corresponding IDL file, and depend instead on the name of the IDL elements being mapped. In keeping with the required consistency between Java package names and the directory structure where Java source files reside, subdirectories of the current directory (or the directory specified by the `-d` command line option) are created dynamically to hold the generated .java files.

IDLC_OPTIONS environment variables

Any idlc command line option can be specified in the environment by adding the option to the string named IDLC_OPTIONS environment variable. Options specified in the IDLC_OPTIONS variable are treated as if they were keyed on the command line before any of the actual command line options. For example, if:

```
IDLC_OPTIONS="-m cpponly -mdllname=mydll"
```

and the command line is:

```
idlc -ehh idlfile
```

the result is the same as if the IDLC_OPTIONS variable was not set and the command line was:

```
idlc -m cpponly -mdllname=mydll -ehh idlfile
```

The IDL-to-Java compiler

The IDL-to-Java compiler generates Java bindings for a given IDL file.

Quick reference

The command to invoke the IDL-to-Java code compiler has the general form:

```
java com.ibm.idl.toJava.Compile [options] source_IDL
```

where *source_IDL* is the name of a file that contains IDL definitions, and [options] is any combination of the options listed in "Compilation options". Options may appear in any order, but must precede the IDL file specification.

Compilation options

Invoke the compiler without any arguments to view the following options:

-d This is equivalent to the following line in an IDL file: `#define symbol`.

-emitAll

Emit all types, including those found in #include files. By default, only those types found in *idl file* are emitted.

-fside Defines what bindings to emit. *side* is one of `client`, `server`, `all`, `serverTie`, and `allTie`. Assumes `-fclient` if the flag is not specified.

-i *include_path*

By default, the current directory is scanned for included files. This option adds another directory.

-keep If a file to be generated already exists, do not overwrite it. By default, it is overwritten.

-m Generate information to be included in a make description file; output goes to `.u` file. By default, this information is not generated.

-pkgPrefix *type package*

Wherever *type* is encountered, ensure it resides within *package* in all generated files. *type* is a fully qualified, Java-style name.

-sep *string*

Only valid with `-m`. Replace the file separator character with *string* in the file names listed in the `.u` file.

-td *target_directory*

Emit bindings to *target_directory* rather than to the current directory.

-v, -verbose

Verbose mode. By default, no messages are output unless there are errors.

-version

Display version information.

The sections that follow provide complete instructions for using each option along with tips about when to use them.

Emitting client and server bindings

To generate the Java bindings for an IDL file named `My.idl`, set the current working directory to that containing `My.idl` and issue the following command:

```
java com.ibm.idl.toJava.Compile My.idl
```

This command generates client-side bindings only and is equivalent to:

```
java com.ibm.idl.toJava.Compile -fclient My.idl
```

Client-side bindings include all generated files except the Skeleton. If you wish to generate server-side bindings for `My.idl`, issue the command:

```
java com.ibm.idl.toJava.Compile -fserver My.idl
```

This command generates all client-side bindings plus an inheritance-model Skeleton (ImplBase). Currently, server-side bindings include all generated files, even the Stub. Thus, the command above is currently equivalent to each shown below:

```
java com.ibm.idl.toJava.Compile -fclient -fserver My.idl
java com.ibm.idl.toJava.Compile -fall My.idl
```

The compiler generates inheritance-model Skeletons by default. Given an interface `My` defined in `My.idl`, the compiler generates Skeleton `_MyImplBase.java`. You provide the implementation for `My`, which must extend `_MyImplBase`.

Specifying an alternate location for emitted files

By default, the compiler outputs bindings to the directory from which it was invoked (the current directory). To direct the output to another directory, specify the target directory immediately following the `-td` flag. The target directory may be absolute or relative. For example, to direct the output to directory `/my/bindings` while compiling `My.idl`, you would invoke the compiler with the following command:

```
java com.ibm.idl.toJava.Compile -td /my/bindings My.idl
```

Similarly, if `/my` is the current directory, you could direct the output to `/my/bindings` by issuing the command:

```
java com.ibm.idl.toJava.Compile -td ./bindings My.idl
```

Specifying alternate locations for include files

If `My.idl` included another idl file, `MyOther.idl`, the compiler assumes that `MyOther.idl` resides in the local directory. If it resides in directory `/includes`, for example, you would invoke the compiler with the following command:

```
java com.ibm.idl.toJava.Compile -i /includes My.idl
```

If `My.idl` also included `Another.idl` that resided in `/moreIncludes`, then you would invoke the compiler as:

```
java com.ibm.idl.toJava.Compile -i /includes -i /moreIncludes My.idl
```

You can begin to see that if you have a number of places where included files may come from, the command will become long and unmanageable. So there is another means of indicating to the compiler where to search for included files. This technique is very similar to the idea of an environment variable. You must create a file called `idl.config` in a directory that is listed in your `CLASSPATH`. Inside of `idl.config` you must provide a line of the following form:

```
includes=/includes;/moreIncludes
```

The compiler take the first version of the file it locates and read in its includes list. Note that in this example, the separator character between the two directories is a semicolon (;). This separator character is platform dependent: On NT it is a semicolon, on AIX it is a colon, and so on.

Note: Some platforms will fail when issuing a long command line. If the command line to invoke the compiler becomes too long, use the `idl.config` file.

Emitting bindings for include files

By default, only those interfaces, structs, and so on, that are defined in the `idl` file on the command line have the Java bindings generated for them. The types defined in included files are not generated. For example, assume the following two `idl` files:

```
My.idl

#include MyOther.idl
interface My
{
};

MyOther.idl

interface MyOther
{
};
```

The following command will only generate bindings for types within `My`:

```
java com.ibm.idl.toJava.Compile My.idl
```

To generate bindings for all of the types in `My.idl` and all of the types in files that `My.idl` includes (in this example, `MyOther.idl`), use the following command:

```
java com.ibm.idl.toJava.Compile -emitAll My.idl
```

There is a caveat to the default rule. `#include` statements which appear at the global scope are treated as described. These `#include` statements can be thought of as `import` statements. `#include` statements which appear within some enclosing scope are treated as true `#include` statements, meaning that the code within the included file is treated as if it appeared in the original file and, therefore, Java bindings are emitted for it. Here is an example:

```
My.idl

#include MyOther.idl
interface My
{
  #include Embedded.idl
};
```

```

MyOther.idl

interface MyOther
{
};

Embedded.idl

enum E {one, two, three};

```

Running the following command:

```
java com.ibm.idl.toJava.Compile My.idl
```

will generate the following list of Java files:

```

./MyHolder.java
./MyHelper.java
./_MyStub.java
./MyPackage
./MyPackage/EHolder.java
./MyPackage/EHelper.java
./MyPackage/E.java
./My.java

```

Notice that `MyOther.java` was not generated because it is defined in an import-like `#include`. But `E.java` was generated because it was defined in a true `#include`. Notice also that since `Embedded.idl` was included within the scope of the interface `My` it appears within the scope of `My` (that is, in `MyPackage`).

If the `-emitAll` flag were used in the previous example, all types in all included files would be emitted.

Inserting package prefixes

Say you work for a company called ABC and that company has constructed the following IDL file:

```

Widgets.idl

module Widgets
{
  interface W1 {...};
  interface W2 {...};
};

```

Running this file through the IDL-to-Java compiler will place the Java bindings for `W1` and `W2` within the package `Widgets`. But there is an industry convention that states that a company's packages should reside within a package named `com.company name`. The `Widgets` package is not good enough.

To follow the convention, it should be `com.abc.Widgets`. To place this package prefix onto the `Widgets` module, implement the following:

```
java com.ibm.idl.toJava.Compile -pkgPrefix Widgets com.abc Widgets.idl
```

You should be aware that, if you have an IDL file which includes `Widgets.idl`, the `-pkgPrefix` flag must appear on that command as well. If it does not, then your IDL file will be looking for a `Widgets` package rather than a `com.abc.Widgets` package.

If you have a number of these packages that require prefixes, it might be easier to place them into the `idl.config` file described above. Each package prefix line should be of the form:

```
PkgPrefix.type=prefix
```

So the line for the above example would be:

```
PkgPrefix.Widgets=com.abc
```

Emitting makefiles and specifying the path separator character

When the Java bindings will be compiled using a makefile, it can become tedious to build the makefile by hand. There are two arguments to the IDL-to-Java compiler which help in building the makefile.

```
java com.ibm.idl.toJava.Compile -m My.idl
```

This will generate, besides the usual bindings, file `My.u` which will contain the following lines:

```
MyHelper.java: My.idl
My.java: My.idl
MyHolder.java: My.idl
MyPackage/E.java: Embedded.idl
MyPackage/EHelper.java: Embedded.idl
MyPackage/EHolder.java: Embedded.idl
_MyStub.java: My.idl
```

```
MyHelper.java \
My.java \
MyHolder.java \
MyPackage/E.java \
MyPackage/EHelper.java \
MyPackage/EHolder.java \
_MyStub.java
```

If you are building a makefile that will run on multiple platforms, the slash `'/'` character will not necessarily be the file separator character. Perhaps the build environment has a special variable for the file separator character. If this variable were `$(Sep)`, then the compiler can place this in place of the slash in `My.u` with the following command:

```
java com.ibm.idl.toJava.Compile -m -sep \$(Sep) My.idl
```

So that My.u now contains:

```
MyHelper.java: My.idl
My.java: My.idl
MyHolder.java: My.idl
MyPackage$(Sep)E.java: Embedded.idl
MyPackage$(Sep)EHelper.java: Embedded.idl
MyPackage$(Sep)EHolder.java: Embedded.idl
_MyStub.java: My.idl
```

```
MyHelper.java \
My.java \
MyHolder.java \
MyPackage$(Sep)E.java \
MyPackage$(Sep)EHelper.java \
MyPackage$(Sep)EHolder.java \
_MyStub.java
```

Defining symbols before compilation

You may wish to define a symbol for compilation that is not defined within the idl file, perhaps to include debugging code in the bindings. The command

```
java com.ibm.idl.toJava.Compile -d MYDEF My.idl
```

is the equivalent to putting the line '#define MYDEF' inside My.idl itself.

Preserving pre-existing bindings

If the Java binding files already exist, this argument will keep the compiler from overwriting them. The default is to generate all files without considering if they already exist. If you've customized those files (which you should not do unless you are very comfortable with their contents), then the `-keep` option is very useful. The command

```
java com.ibm.idl.toJava.Compile -keep My.idl
```

emits all client-side bindings that do not already exist.

Viewing progress of compilation

The IDL-to-Java compiler will generate status messages as it progresses through its phases of execution. Use the `-v` (or `-verbose`) option to activate this "verbose" mode:

```
java com.ibm.idl.toJava.Compile -v My.idl
```

By default the compiler does not operate in verbose mode.

Displaying version information

To display the build version of the IDL-to-Java compiler, specify the `-version` option on the command-line:

```
java com.ibm.idl.toJava.Compile -version
```

Version information also appears within the bindings generated by the compiler. Any additional options appearing on the command-line are ignored.

The idl2com Command

The idl2com Command Creates usage and implementation bindings for interfaces described in IDL files.

The Interface Definition Language Compiler-to-COM (idl2com) command compiles one file containing CORBA 2.0-compliant IDL statements, and produces generated language bindings appropriate to the named input file.

Quick reference

The syntax of the idl2com command is:

```
idl2com [options] <-g GUID_VALUE> <filename>
```

Where:

[options]

See Options for the idl2com command for a complete list of the options, their usage, and their restrictions.

<-g GUID_VALUE>

-g is a mandatory parameter. See Options for the idl2com command for a complete description of this parameter.

<filename>

<filename> is the name of a file containing IDL definitions. It is a mandatory parameter and must appear last. It must be specified with a file name extension. For example, the following are acceptable ways to refer to the file "Policy.idl" in directory "E:\idl\src":

```
E:\idl\src\Policy.idl  
Policy.idl (if E:\idl\src is the current directory)
```

If no parameters are specified, idl2com will write its syntax and options to standard output.

When the specified IDL file is compiled, the idl2com command returns a Java exception if an error has occurred. Warnings may also be given. For instance, if the IDL contains a constant of type float or double, idl2com issues a warning statement to standard output and continues processing. The result is that the definition for the constant in the .bas file will be lacking, otherwise the bindings will be complete.

For additional information on emitted files, see *Emitted file names*.

Options for the `idl2com` command

Options to the `idl2com` command may be specified in any combination. The `-g` option is mandatory, all other options are optional. The options may appear in any order. The IDL filename is required and must appear last. A maximum of 9 parameters can be specified, including the IDL filename. If more than 9 parameters need to be specified, the user can resort to invoking Java directly. Examine the contents of `idl2com.bat` to see how to do this. If no parameters are specified, `idl2com` will write its syntax and options to standard output.

With the exception of the IDL filename, options to the `idl2com` command are preceded with a dash (-) character and must be specified individually. For example, `-emitAll -keep -v`. Some options accept an argument. The space between this type of option and its argument is mandatory. For example, `-d DEBUG`. All options are case-sensitive, even on platforms where file names are not case-sensitive.

Table 24 describes each available option:

Table 24. idl2com command options

Option	Description
<code>-d <symbol></code>	This is equivalent to the following line in an IDL file: <code>#define <symbol></code>
<code>-emitAll</code>	Emit all types, including those found in included files. The default is to emit only the types that are part of the IDL being processed.
<code>-g <GUID></code>	The GUID seed to be used for GUID generation in the registry file format of: <code>xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx</code> . This value is generated by using the <code>guidgen.exe</code> utility included with Microsoft Visual C++. See <i>Generating interfaces using the idl2com command</i> for more information.
<code>-i <include path></code>	By default, the current directory is scanned for included files. This option adds another directory. Multiple <code>-i <include path></code> options may be specified.
<code>-keep</code>	If a file to be generated already exists, do not overwrite it. By default it is overwritten.
<code>-v</code>	Verbose mode. Default is non-verbose mode.

Generating interfaces using the `idl2com` command

When an IDL file is processed by the `idl2com` command, a unique GUID seed value (`-g` parameter) is required. This GUID is used to register within the Windows system registry the various interfaces, etc produced by `idl2com`. Many interfaces may be contained within the IDL, and `idl2com` uses the provided GUID parameter as a starting point for the to-be registered

interfaces. If multiple items produced from the IDL file need to have GUIDs, `idl2com` increments using the first 8 digits (AE3E2131 in the example below) as needed. `idl2com` will use all eight of those digits when incrementing and will fail if it hits "FFFFFFFF". It will not roll over to 00000000 since this is the start of reserved GUID ranges held by Microsoft. For example, if the IDL file results in three interfaces being registered, the following GUIDs would be used:

- AE3E2131-C6DE-11d0-92AF-08005ACE818D
- AE3E2132-C6DE-11d0-92AF-08005ACE818D
- AE3E2133-C6DE-11d0-92AF-08005ACE818D

This is important to know because all registered items need to be unique within the registry. Use the `guidgen.exe` program included with Microsoft Visual C++ to provide the `-g` parameter, but be careful not to conflict with the internally-generated values that `idl2com` will use based on the `-g` parameter. To avoid this potential conflict, you need to exit and restart `guidgen.exe` prior to generating a new GUID for another run of `idl2com`. Be sure to include a space between the `-g` and the GUID value on the `idl2com` command line.

Emitted file names

The `idl2com` command processes an IDL file and produces output files that contain language-specific usage and implementation bindings for the IDL interface. The list of generated files is dependent on the contents of the IDL. However, the following is a general list of what can be expected given an IDL file called `Policy`:

Table 25. `idl2com` emitted files

File Name	Content
<code>Policy.odl</code>	the Object Description file for the IDL
<code>Policy.def</code>	definition file
<code>Policy.bas</code>	Visual Basic file containing any constants that were defined in the IDL
<code>Policy.mak</code>	makefile
<code>Policy.rc</code>	resource definition
<code>*.cpp</code>	a set of implementation files which will vary based on the contents of the IDL
<code>*.h</code>	a set of header files which will vary based on the contents of the IDL

The names of the generated `.cpp` and `.h` files vary depending on the contents of the IDL.

Data type restrictions

The `idl2com` command has the following limitations:

- IDL constants of data type `float` or `double` are not supported.

- IDL types long long, unsigned long long, long double and fixed are not supported.

idl2com generated makefile

As part of the idl2com processing, a makefile is automatically generated. The makefile is named <idlfilename>.mak. In most cases it will be complete. If however, the developer extends the code in a way that is not reflected within the IDL, the makefile may require some additional hand coding. Also, if the IDL makes use of other IDL whose bindings are linked into a different .dll, the developer will have to modify the generated .mak file to add the other .LIB to the list of .LIBs.

Appendix C. CORBA programming

This appendix includes the following C++ topics:

- “C++ bindings”
- “Name scoping and modules in the C++ bindings” on page 508
- “C++ bindings for interfaces” on page 508
- “Storage management and `_var` types” on page 512
- “C++ client bindings” on page 520
- “C++ server bindings” on page 521
- “C++ binding restrictions” on page 523

C++ bindings

CORBA 2.0 specifies standard forms by which client C++ code can manipulate data whose types are described using IDL. C++ bindings that support these forms are termed *compliant* and client code that uses (only) these forms is termed *conformant*. The bindings in Component Broker are compliant.

C++ bindings for constants

Constants can be defined within the IDL either of the following ways:

- Within the module or interface
- Globally

Within the module

An IDL constant declaration contained within a module or an interface is mapped to a static constant data member of the C++ class to which the module or interface is mapped. For example, consider the following IDL:

```
module M
{
    const string name = "testing";
};
```

After compiling the client bindings, a C++ programmer could use the following expression to denote the previous name constant:

```
M::name
```

Globally

Globally-defined constants are mapped to static data local to a single compilation unit. For example, if the following IDL appeared globally, un-nested within a module or interface:

```
const string name = "testing";
```

The code that includes the corresponding .hh file refers to the name constant using the expression name.

CORBA types and business objects

Most of the CORBA types are straightforward and can be easily used in business objects. Other CORBA types are more difficult to use but are still useful if care is taken.

Basic types

The basic C++ types are mapped directly into CORBA types. These include:

- Boolean
- Char
- Double
- Enum (enumerations)
- Float
- Long
- LongLong (long long)
- Octet (hexadecimal)
- Short
- Struct
- UShort (unsigned short)
- ULong (unsigned long)
- ULongLong (unsigned long long)
- WChar (wide character)

All types are scoped to the class CORBA and must be declared accordingly. They are used transparently to C++ and are straightforward. For example:

```
CORBA::Short aShortvariable;  
...  
aShortVariable = 12;  
...
```

Types and object references

Other CORBA types are more complex to use because they return object references to the caller. It is the responsibility of the caller to manage these object references and their associated memory. There are two facilities provided by CORBA to do this:

A_var This is most frequently used by client code because it is a *smart* pointer and automatically releases its object reference when it is deallocated or when assigned a new object reference. This is the most straightforward and safest approach to managing these types.

Note: You should avoid declaring C++ **Static** variables as `_var`. The `_var` holds a reference to an object. During process termination, this object could reference another object that was removed before termination processing is completed for this Static type. As a result, the `_var` could reference an inappropriate address or null pointer and thereby cause bad termination.

A_ptr This is a pointer type and provides the most basic object reference, which has similar semantics to a standard C++ pointer.

The CORBA types that return object references include:

- Any
- Array
- Sequence
- String
- Union
- WString (wide string)

C++ bindings for data types

The following are C++ bindings for data types:

- “Any type”
- “Array types” on page 489
- “Atomic data types” on page 492
- “Enums” on page 493
- “Sequence types” on page 493
- “Strings” on page 498
- “Struct types” on page 499
- “Union types” on page 501
- “Using WStrings” on page 503

Any type

The purpose of the IDL “any” type is to encapsulate data of some arbitrary IDL type. The C++ bindings provide a C++ class named `CORBA::Any` that provides this functionality. A `CORBA::Any` object encapsulates a `void*` pointer and a `CORBA::TypeCode` object that describes the thing pointed to by the `void*`.

The Any type can be used with many of the CORBA types and is useful when different types can be used that are unknown to the receiver of the data or as a common storage mechanism for passing a variety of types. It is used easily with many of the CORBA types but has a unique method of redirection operators for setting and retrieving data.

The following types are handled in this manner:

- Double
- Enumerations
- Float
- Long
- Short
- ULong
- UShort
- Unbounded Strings
- Object References

```
::CORBA::Any anything;  
anything <<= (::CORBA::Long) 123456;  
::CORBA::Long anythingStart = 123456;  
::CORBA::Long anythingLongResult = 0;  
policyVar->anything(anything);  
  
::CORBA::Any_var anythingResult_var(policyVar->anything());  
::CORBA::Any anythingResult(anythingResult_var);  
anythingResult >>= anythingLongResult;  
if ( anythingStart != anythingLongResult)  
{  
    cout << "Anything not set" << endl;  
    return 1;  
}  
else  
{  
    cout << "Anything set correctly..." << endl;  
}
```

There are also specialized structures provided for the following types for conversion with Any:

- Boolean
- Char
- Octet
- String

The data in an Any object is initialized and accessed using insertion (<<=) and extraction (>>=) operators defined by the C++ bindings. These operators are provided (using overloading) by CORBA::Any for each primitive data type, and are provided by the generated C++ bindings for each user-defined IDL type. As a result, there is usually no need to indicate a typecode when

inserting or extracting data from a CORBA::Any (although the CORBA::Any class does provide methods for manipulate the data using an explicit TypeCode).

Types that cannot be distinguished by C++ overloading are inserted into and extracted from Any's using special *wrapper* classes. These wrapper classes are not transparent to the application; the application must explicit create and use them when inserting or extracting ambiguous types into or from Any's. For primitive IDL types that do not map to distinct C++ types (boolean, octet, and char), the wrapper classes are defined within the CORBA::Any scope; they are called from_boolean, to_boolean, from_octet, to_octet, from_char, and to_char. For information on the scope see "IDL name scoping" on page 447. Because bounded strings cannot be distinguished in C++ from unbounded strings, CORBA::Any provides the from_string and to_string wrapper classes, for inserting/extracting bounded strings. For extracting object references from Any's as the base CORBA::Object type, CORBA::Any provides a to_object wrapper class.

For application-specific arrays, the bindings provide a special *forany* class, for inserting/extracting the array into/from an Any. For example, given the following IDL array definition:

```
typedef long LongArray[4][5];
```

the emitted bindings define the following:

```
typedef CORBA::Long LongArray[4][5];
typedef CORBA::Long LongArray_slice[5];
typedef LongArray_slice* LongArray_slice_vPtr;
typedef const LongArray_slice* LongArray_slice_cvPtr;
class LongArray_forany
{
public:
    LongArray_forany();
    LongArray_forany(LongArray_slice*, CORBA::Boolean nocopy=0);
    LongArray_forany(const LongArray_forany&);
    LongArray_forany & operator= (LongArray_slice*);
    LongArray_forany & operator= (const LongArray_forany &);
    ~LongArray_forany();
    const LongArray_slice& operator[] (int) const;
    const LongArray_slice& operator[] (CORBA::ULong) const;
    LongArray_slice & operator[] (int);
    LongArray_slice & operator[] (CORBA::ULong);
    operator LongArray_slice_cvPtr () const;
    operator LongArray_slice_vPtr& ();
};
void operator<=<(Any&, const LongArray_forany &);
CORBA::Boolean operator>=>(Any&, LongArray_forany &);
```

Note that the nocopy optional parameter of the *_forany*'s second constructor indicates whether the *_forany* makes a copy of the input array or assumes

ownership of it. (The default is for the `_forany` to assume ownership of the input array; that ownership will then be transferred to the `Any` when the `_forany` is inserted into the `Any`.)

To determine what kind of data is in `Any`, invoke the `type` method on a `CORBA::Any` to access a `TypeCode` that describes the data it holds. Alternatively, you can simply try to extract data of a particular type from the `Any`; the extraction operator returns a boolean to indicate success. If the extraction operation fails, the `Any` doesn't hold data of the type you tried to extract.

A `CORBA::Any` object always owns the data that its `void*` points to, and deletes (or releases) it when the `Any` is given a new value or deleted. The only question is whether this data is a copy of the data that was inserted into the `Any`. When primitives (including strings and enums) are inserted, a copy is made, and a copy is returned when the data is extracted.

For non-primitive (constructed) data, extraction from an `Any` always updates a pointer (owned by the caller) so that it points to the data owned by the `Any`. The caller should not, therefore, free this data or reference it after the `Any` has been given a new value or deleted. For constructed IDL type `T`, the emitted bindings define the following extraction operator:

```
CORBA::Boolean operator>>=(Any&, T*&);
```

When a reference to constructed data is inserted into an `Any` (when the C++ syntax looks as if you are inserting a value instead of a pointer) a copy is made. In this case, the caller retains ownership of the original data. For example, for constructed type `T` and interface `I`, the emitted bindings define the following insertion operators, which copy (or, in the case of object references, `_duplicate`) the inserted value):

```
void operator<<=(Any&, const T&);  
void operator<<=(Any&, T_ptr);
```

When a pointer to constructed data is inserted into an `Any`, as when using the following insertion operators emitted for type `T` and interface `I`:

```
void operator<<=(Any&, T*);  
void operator<<=(Any&, T_ptr*);
```

The `Any` takes ownership of the constructed type's top-level storage only; however, the `Any` makes no copy of the top-level storage or any embedded storage. All further use of the pointer that was inserted is forbidden; the `Any` now owns it and is free to delete it at any time. The next time data is inserted into the `Any`, or when the `Any` is destroyed, the `Any` deletes the previously-inserted pointer. However, if the constructed type consists of multiple dynamically-allocated regions of memory, only the top-level storage is deleted. (The `Any` deletes arrays using a single array delete; other

constructed types are deleted using a single, normal delete.) Further, the top-level storage is deleted as a void*, rather than its true type, which means that the constructed type's destructor will not be run. Due to these restrictions, insertion by pointer of constructed types into an Any should be used with caution.

In summary, when extracting data from an Any, the caller does own the data for primitive types, but does not own the data for constructed types. When inserting data into an Any, the caller retains ownership of the data for primitive types, for constructed types inserted by value, and for storage embedded within constructed types inserted by pointer. The caller does not retain ownership of the top-level contiguous storage for a constructed type inserted into an Any by pointer.

The following is an example that illustrates the previously discussed aspects of CORBA::Any usage. The IDL for this example appears immediately below. It defines a struct and an array that will be inserted into an Any.

```
Module M
{
  Struct S
  {
    string str;
    long lng;
  };
  typedef long long1[2][3];
}
```

A C++ program illustrating Any insertion and extraction appears below:

```
#include <stdio.h>
#include <any_C.cpp>
main()
{
  CORBA::Any a; // the Any that we'll be using
  // test a long
  long l = 42;
  a <<= l;
  if (a.type()->equal(CORBA::_tc_long))
  {
    long v;
    a >>= v;
    printf("the any holds a long = %d\n", v);
  }
  else
  printf("failure: long insertion\n");
  // test a string
  char *str = "abc";
  a <<= str;
  if (a.type()->equal(CORBA::_tc_string))
  {
    char *ch;
    a >>= ch;
  }
}
```

```

    printf("the any holds the string = %s\n", ch);
    delete ch;
    a >>= ch;
    printf(" the any still holds the string = %s\n", ch);
    delete ch;
}
else
printf("failure: string insertion\n");
// test a bounded string -- note we don't use a typecode here
char *bstr = "abcd";
char *rstr;
a <<= CORBA::Any::from_string(bstr, 6);
if (a >>= CORBA::Any::to_string(rstr,6))
printf("the any holds a bounded string<6> = %s\n", rstr);
else
printf("failure: bounded string insertion\n");

// test a user-defined struct
M::S *s1 = new M::S;
char *saveforlater = CORBA::string_dup("abc");
s1->str = saveforlater;
s1->lng = 42;
a <<= s1; // insertion by pointer
if (a.type()->equal(_tc_M_S))
{
    M::S *sp;
    a >>= sp;
    printf("the any holds an M::S = {%s, %d}\n", sp->str, sp->lng);
}
else
printf("failure: struct insertion by pointer\n");
M::S s2;
s2.str = CORBA::string_dup("def");
s2.lng = 23;
a <<= s2; // note: this deletes *s1, but not saveforlater
printf("saveforlater still = %s\n", saveforlater);
CORBA::string_free(saveforlater);
if (a.type()->equal(_tc_M_S))
{
    M::S *sp;
    a >>= sp;
    printf("the any holds an M::S = {%s, %d}\n", sp->str, sp->lng);
}
else
printf("failure: struct insertion by value\n");
M::S_var s3 = new M::S;
s3->str = CORBA::string_dup("ghi");
s3->lng = 96;
a <<= *s3;
if (a.type()->equal(_tc_M_S))
{
    M::S *sp;
    a >>= sp;
    printf("the any holds an M::S = {%s, %d}\n", sp->str, sp->lng);
}
}

```

```

else
printf("failure: struct insertion by ref to value\n");
// test an array
M::long1_var llv = M::long1_alloc();
for (i=0;i<2;i++)
for (j=0;j<3;j++)
llv[i][j] = (i+1)*(j+1);
a <<= M::long1_forany(llv);
if (a.type()->equal(_tc_M_long1))
{
M::long1_forany ll1s;
a >>= ll1s;
printf("the any holds the array: ");
for (i=0;i<2;i++)
for (j=0;j<3;j++)
printf("%d ",ll1s[i][j]);
printf("\n");
}
else printf("failure: array insertion\n");
}

```

Output from the above program is:

```

the any holds a long = 42
the any holds a string = abc
the any still holds a string = abc
the any holds a bounded string<6> = abcd
the any holds an M::S = {abc, 42}
saveforlater still = abc
the any holds an M::S = {def, 23}
the any holds an M::S = {ghi, 96}
the any holds the array: 1 2 3 2 4 6

```

Array types

An IDL array type is mapped to the corresponding C++ array definition. There is also a corresponding `_var` type. For example, given the following IDL definition:

```
typedef long LongArray [4][5];
```

The C++ bindings provide the following definitions:

```

typedef CORBA::Long LongArray[4][5];
typedef CORBA::Long LongArray_slice[5];
typedef LongArray_slice* LongArray_slice_vPtr;
typedef const LongArray_slice* LongArray_slice_cvPtr;
class LongArray_var
{
public:
LongArray_var();
LongArray_var(LongArray_slice*);
LongArray_var(const LongArray_var&);
LongArray_var & operator= (LongArray_slice*);
LongArray_var & operator= (const LongArray_var &);

```

```

    ~LongArray_var();
    const LongArray_slice& operator[] (int) const;
    const LongArray_slice& operator[] (CORBA::ULong) const;
    LongArray_slice & operator[] (int);
    LongArray_slice & operator[] (CORBA::ULong);
    operator LongArray_slice_cvPtr () const;
    operator LongArray_slice_vPtr& ();
};
LongArray_slice * LongArray_alloc();
void LongArray_free (LongArray_slice*);
LongArray_slice * LongArray_dup (const LongArray_slice*);

```

As shown above, array mappings provide alloc, dup, and free functions (for allocating, copying, and freeing array storage), as static member functions of the class within which the array type name is scoped. The alloc function dynamically allocates an array, which can be later freed using the free function. The dup function dynamically allocates an array and copies the elements of an existing array into it. A NULL pointer can be passed to the free function. None of these functions throws exceptions.

The type of the pointer returned from LongArray_alloc is LongArray_slice*. The C++ bindings define array “slice” types for all arrays declared in IDL. The reason is that using the name LongArray in a program doesn’t denote the array LongArray; rather, it denotes a pointer to the array. For historical reasons (related to the fact that arrays are not an actual data type in C and C++) the type of this pointer has one less array dimension than the array LongArray. Thus, the bindings for LongArray include the following typedef:

```
typedef string LongArray_slice[5];
```

Hence, LongArray_slice* is the correct type for a pointer to an array of IDL type LongArray.

As with structs and sequences, arrays use special auxiliary classes for automatic storage management of String and object reference elements. The auxiliary classes for Strings and object references manage the storage just as the associated _var classes do.

When assigning a value to an array element that is an object reference, the assignment operator will automatically release the previous value (if any). When assigning an object reference pointer to an array element, the array assumes ownership of the pointer (no _duplicate is done), and the application should no longer access the pointer directly. (If this is not the desired behavior, then the caller can explicitly _duplicate the object reference before assigning it.) However, when assigning to an object reference array element from a _var object or from another struct, union, array, or sequence member (rather than from an object reference pointer), a _duplicate is done automatically.

For an array of Strings, when assigning a value to an element or deleting the array, any previously held (non-null) `char*` is automatically freed. As when assigning to `String_vars`, assigning a `char*` to a string element does not make a copy, but assigning a `const char *` or another struct/union/array/sequence String element does make a copy. One should never assign a string literal (for example, "abc") to a String array element without an explicit cast to "const char*". When assigning a `char*` that occupies static storage (rather than one that was dynamically allocated), the caller can use `CORBA::string_dup` to duplicate the string before assigning it.

The following is an example that involves multidimensional arrays, and `array_vars`, from the IDL snippet immediately below:

```
typedef string s2_3[2][3];
typedef string s3_2[3][2];
```

The code that exercises the C++ arrays that correspond to the above IDL is shown below. Notice that in the following example:

- There is no need to explicitly use slice types when working with the `array_var` types, because the bindings declare the pointer held by an `array_var` type using the appropriate slice type.
- At the end, the program explicitly frees the storage pointed to by `s2_3p` (using an array delete operator), but does not do this for `s3_2v`, because its pointer is deleted when the destructor for `s3_2v` is implemented. (This is the purpose of the `_var` types.)

```
#include <arr_C.cpp>
#include <stdio.h>
main()
{
    int i,j;
    char id[40];
    // create arrays
    s2_3_slice* s2_3p = s2_3_alloc();
    s3_2_var s3_2v = s3_2_alloc();
    // load the arrays
    for(i=0; i<2; i++)
    {
        for (j=0; j<3; j++)
        {
            sprintf(id, "s2_3 element [%d][%d]",i,j);
            // Use string_dup when assigning a char*
            // if you don't want the array to own the original:
            s2_3p[i][j] = CORBA::string_dup(id);
        }
    }
    for(i=0; i<3; i++)
    {
        for (j=0; j<2; j++)
        {
            sprintf(id, "s3_2_var element [%d][%d]",i,j);
```

```

        // Use string_dup when assigning a char*
        // if you don't want the array to own the original:
        s3_2v[i][j] = CORBA::string_dup(id);
    }
}
// print the array contents
for(i=0; i<2; i++)
{
    for (j=0; j<3; j++)
    {
        printf("%s\n", s2_3p[i][j]);
    }
}
for(i=0; i<3; i++)
{
    for (j=0; j<2; j++)
    {
        printf("%s\n", s3_2v[i][j]);
    }
}
delete [] s2_3; // needed to prevent a storage leak.
// Nothing is needed for s3_2v, because
// it is a _var type.
}

```

Output from the above program is:

```

s2_3 element [0][0]
s2_3 element [0][1]
s2_3 element [0][2]
s2_3 element [1][0]
s2_3 element [1][1]
s2_3 element [1][2]
s3_2_var element [0][0]
s3_2_var element [0][1]
s3_2_var element [1][0]
s3_2_var element [1][1]
s3_2_var element [2][0]
s3_2_var element [2][1]

```

Atomic data types

The atomic IDL data types (long, short, unsigned long, unsigned short, float, double, char, boolean, and octet) are mapped into types defined in `corba.h`, nested within the CORBA scope. See “IDL name scoping” on page 447 for more information. The first letter of the mapped type is capitalized. For example, to introduce and initialize a local variable named `Myvar` whose type corresponds to the IDL type named `long`, a C++ programmer could employ the following expression:

```
CORBA::Long Myvar = 1;
```

The mapping for the IDL boolean type (CORBA::Boolean) defines only the values 0 and 1. The unsigned long and unsigned short IDL types are mapped to CORBA::ULong and CORBA::UShort, respectively.

Enums

An IDL enum is mapped to a corresponding C++ enum. For example, given the following IDL:

```
module M
{
    enum Color
    {
        red, green, blue
    };
};
```

A C++ programmer could introduce a local variable of the corresponding C++ type and initialize it with the following code:

```
{
    M::Color MYCOLOR = M::red;
}
```

The enumeration constant `red` is not denoted using the expression `M::Color::red`. For this reason, names of enumeration constants must be carefully chosen.

Sequence types

An IDL sequence type is mapped to a C++ class that behaves like an array with a current length (how many elements have been stored) and a maximum length (how much storage is currently allocated). The array indexing operator `[]` is used to read and write sequence elements. (Indexing begins at zero.) It is the programmer's responsibility to check the current sequence length or maximum to prevent accessing the sequence beyond its bounds. The length and maximum of the sequence are not automatically increased to accommodate new elements; the programmer must explicitly increase them.

The maximum length of a bounded sequence is implicit in the sequence class's type and cannot be changed. The initial maximum length of an unbounded sequence is set to zero by the default constructor, or can be initialized by the programmer using a non-default constructor. Setting the maximum of an unbounded sequence using the non-default constructor causes storage to be allocated for the specified number of sequence elements.

Sequence classes provide an overloaded member function `length` that either returns or sets the length of the sequence. Setting the length of an unbounded sequence to a value larger than the current maximum causes the sequence to

allocate new storage of the required size, copy any previous sequence elements to the new storage, free the old storage (if any), and reset the maximum to the new length. Sequence classes also provide `allocbuf` and `freebuf` member functions for explicitly allocating/freeing the sequence's storage buffer. Decreasing a sequence's length does not cause any storage to be deallocated, but any orphaned sequence elements are no longer accessible, even if the sequence length is subsequently increased.

Sequences may or may not manage (own) the storage that contains their elements, and the elements themselves. By default, a sequence manages this storage, but a *release* constructor parameter allows client programmers to request otherwise (when passing in a buffer explicitly allocated using the `allocbuf` function).

The following IDL:

```
typedef sequence s1; // unbounded sequence
```

is mapped to the following C++ sequence class:

```
class s1
{
public:
    s1(); // default constructor
    s1(CORBA::ULong max); // "max" constructor
    s1(CORBA::ULong max, CORBA::ULong length,
        T* data, CORBA::Boolean release=0);
    // "data" constructor
    s1(const s1&); // copy constructor
    s1 &operator= (const s1&); // assignment operator
    ~s1(); // destructor
    CORBA::ULong maximum() const;
    CORBA::ULong length() const;
    void length(CORBA::ULong len);
    T& operator[] (CORBA::ULong index);
    const T& operator[] (CORBA::ULong index) const;
    static T* allocbuf(CORBA::ULong nelems);
    static void freebuf(T* data);
};
```

The default constructor sets the length and maximum to zero. (For a bounded sequence, the default constructor sets the maximum to the sequence bounds and allocates storage for the maximum number of elements, which the sequence owns.)

The "max" constructor sets the initial sequence maximum and allocates a storage buffer for the specified number of sequence elements, which the sequence owns. The length of the sequence is initialized to zero. (This method is not available for bounded sequences.)

The “data” constructor sets the initial length and maximum of the sequence, as well as its initial contents. (For bounded sequences, the maximum cannot be set by the “data” constructor.) The input storage should match the specified sequence maximum. Ownership of the input storage is indicated by the “release” parameter. Passing `release=1` specifies that the storage was allocated using `s1::allocbuf`, that the sequence should delete the storage and the sequence elements when the sequence is deleted or when the storage needs to be reallocated, and that the caller will not directly access the storage after the call (since the sequence is free to delete it at any time). In general, sequences constructed with `release=0` should not be passed as inout parameters, because the callee must assume that the sequence owns the sequence elements.

The copy constructor creates a new sequence with the same maximum and length as the input sequence and copies the sequence elements to storage that the sequence owns. The assignment operator performs a deep copy, releasing the previous sequence elements if necessary. It behaves as if the destructor were run, followed by the copy constructor.

The destructor destroys each of the sequence elements (from zero through `length-1`), if the sequence owns the storage.

The `allocbuf` function allocates enough storage for the specified number of sequence elements; the return value can then be passed to the “data” constructor. Each sequence element is initialized using its default constructor; string elements are initialized to `NULL`; object reference elements are initialized to `nil` object references. `NULL` is returned if storage cannot be allocated for any reason. If ownership of the allocated buffer is not transferred to a sequence using the “data” constructor with `release=1`, the buffer should be subsequently freed using the `freebuf` function. The `freebuf` function insures that each sequence element’s destructor is run (or, for strings, that `CORBA::string_free` is called, or for object references, that `CORBA::release` is called) before the buffer is deleted. The `freebuf` function ignores `NULL` pointers passed to it. Neither `allocbuf` nor `freebuf` throw CORBA exceptions.

As with structs, sequences that manage their elements use special auxiliary classes for automatic storage management of String and object reference sequence elements. These auxiliary classes manage Strings and object references just as the associated `_var` classes do.

For a storage-managing sequence whose elements are object references, when assigning a value to an element, the assignment operator will automatically release the previous value (if any). When assigning an object reference pointer to such a sequence element, the sequence assumes ownership of the pointer (no `_duplicate` is done), and the application should no longer access the pointer directly. (If this is not the desired behavior, then the caller can

explicitly `_duplicate` the object reference before assigning it to the sequence element.) However, when assigning to such an object reference sequence element from a `_var` object or from another struct, union, array, or sequence (rather than from an object reference pointer), a `_duplicate` is done automatically.

For a storage-managing sequence whose elements are Strings, when assigning a value to such an element or deleting the sequence, any previously held (non-null) `char*` is automatically freed. As when assigning to `String_vars`, assigning a `char*` to a string element doesn't make a copy, but assigning a `const char *`, a `String_var`, or another struct/union/array/sequence `String` member does automatically make a copy. Thus, one should never assign a string literal (such as "abc") to such an element without an explicit cast to `const char*`. When assigning a `char*` that occupies static storage (rather than one that was dynamically allocated), the caller can use `CORBA::string_dup` to duplicate the string before assigning it.

There is a corresponding `_var` type defined for every sequence class. The `_var` type for a sequence provides an overloaded operator[] that forwards the operator to the underlying sequence.

Following is an example that illustrates loading and accessing the elements of a sequence. This example illustrates a recursive sequence (whose entries are structs of the same type that contain the sequence). The IDL for the example is shown below:

```
struct S
{
    long sf1;
    sequence sf2;
};
typedef sequence Sseq;
```

The following is an example program that creates and loads a sequence of type `Sseq` and then prints out its contents.

```
#include <seq_C.cpp>
#include <stdio.h>
main()
{
    int i,j;
    Sseq seq;          // create an Sseq
    seq.length(3);    // set length of seq to 3
    for (i=0; i<3; i++) { // index the three S structs in seq
        seq[i].sf1 = i;    // place a number in the i-indexed struct
        seq[i].sf2.length(i+1); // set length of the sequence in
        // the i-indexed struct
        for (j=0; j<i+1; j++) // index the i+1 S structs in the sequence
            // in the i-indexed struct
            seq[i].sf2[j].sf1 = (i+1)*10+j; // place a number in
            // the j-indexed struct
    }
}
```

```

}
// OK. Print out what we have created!
printf("seq = (%d sequence elements)\n", seq.length());
for (i=0; i<3; i++)
{
    printf("    struct[%d] = {\n", i);
    printf("        sf1 = %d\n", seq[i].sf1);
    printf("        sf2 = (%d sequence elements)\n",
        seq[i].sf2[j].length());
    for (j=0; j<i+1; j++)
    {
        printf("            struct[%d] = \n",j);
        printf("                sf1 = %d\n", seq[i].sf2[j].sf1);
        printf("                sf2 = (%d sequence elements)\n",
            seq[i].sf2[j].sf2.length());
        printf("            }\n");
    }
    printf("    }\n");
}
}

```

Note that the above program never explicitly constructs any data of type S, even though the sequences contain structs of this type. The reason is that when a sequence buffer is allocated, default constructors are run for each of the buffer elements. So, when the above program sets the length of a sequence of S structs (either at the top level for the seq variable, or for the sf2 field of an S struct in seq), the resulting buffer is automatically populated with default structs of type S.

The output from the above program is:

```

seq = (3 sequence elements)
  struct[0] = {
    sf1 = 0
    sf2 = (1 sequence elements)
      struct[0] = {
        sf1 = 10
        sf2 = (0 sequence elements)
      }
  }
  struct[1] = {
    sf1 = 1
    sf2 = (2 sequence elements)
      struct[0] = {
        sf1 = 20
        sf2 = (0 sequence elements)
      }
      struct[1] = {
        sf1 = 21
        sf2 = (0 sequence elements)
      }
  }
  struct[2] = {
    sf1 = 2
  }

```

```

sf2 = (3 sequence elements)
  struct[0] = {
    sf1 = 30
    sf2 = (0 sequence elements)
  }
  struct[1] = {
    sf1 = 31
    sf2 = (0 sequence elements)
  }
  struct[2] = {
    sf1 = 32
    sf2 = (0 sequence elements)
  }
}

```

Strings

The mapping for strings is provided by `corba.h`, within the CORBA scope. See “IDL name scoping” on page 447 for more information. The user-visible types are `CORBA::String` and `CORBA::String_var`. `CORBA::String` is a typedef name for `char*`. The `CORBA::String_var` class performs storage management of a dynamically allocated `CORBA::String`. The following functions are for dynamic allocation/deallocation of memory to hold a `String`:

- `CORBA::string_alloc`
- `CORBA::string_free`
- `CORBA::string_dup`

A `String_var` object behaves as a `char*`, except that when it is assigned to, or goes out of scope, the memory it points to is automatically freed by `CORBA::string_free`. When a `String_var` is constructed or assigned from a `char*`, the `String_var` assumes ownership of the string and the caller should no longer access the string directly. (If this is not the desired behavior, as when the `char*` occupies static storage, the caller can use `CORBA::string_dup` to copy the `char*` before assigning it.) When a `String_var` is constructed or assigned from a `const char*`, another `String_var`, or a `String` element of a struct, union, array, or sequence, an automatic copy of the source string is done. The `String_var` class provides subscripting operations to access the characters within the embedded string.

C++ compilers don’t treat a string literal (such as “abc”) as a `const char*` upon assignment; given both a `const` and a non-`const` assignment operator, the compiler will choose the non-`const` operator. As a result, when assigning a string literal to a `String_var`, no copy of the string into dynamically allocated memory is made; the pointer “owned” by the `String_var` will point to memory that cannot be freed. Thus, string literals should not be assigned to a `String_var` without an explicit cast to `const char*`.

Some examples using `String_var` objects are:

```

// first some supporting functions for the examples
char* f1()
{
    return "abc";
}
char* f2()
{
    char* s=CORBA::string_alloc(4);strcpy(s,"abc");return s;
}
// then the examples
void main()
{
    CORBA::String_var s1;
    if (0) s1 = f1();// Wrong!! The pointer can't be freed and
    // no copy is done.
    if (0) s1 = "abc"; // Also wrong, for the same reason.
    const char* const_string = "abcd"; // *const_string can't be changed
    s1 = const_string; // OK. A copy of the string is made because
    // it is const, and the copy can be freed.
    CORBA::String_var s3 = f2();// OK. no copy is made, but f2
    // returns a string that can be freed
    CORBA::String_var s4 = CORBA::string_alloc(10); // also OK. no copy
    s4 = s3; // s4 will use string_free followed by string_dup
    long l4 = strlen(s4); // l4 will receive 3
    long l1 = strlen(s1); // l1 will receive 4
    if (l4 >= l1)
        strcpy(s4,s1); // OK, but only because of the condition.
    // note that s4's buffer only has size=4.
    s4 = const_string; // OK. s4 will use string_free followed by
    // string_dup. The copy is made because String_vars
    // must reference a buffer that can be modified.
}
// The s1, s3 and s4 destructors run successfully, freeing their buffers

```

Struct types

An IDL struct type is mapped to a corresponding C++ struct whose field names correspond to those in the IDL declaration, and whose field types support access and storage of the C++ types corresponding to the IDL struct field types. Dynamically allocated storage used to hold such a C++ struct must be allocated and freed using the C++ new and delete operators.

When a new struct is created, the default constructor for each of its fields implements. Object reference fields are initialized to nil references, and String fields are initialized to NULL. When the struct is deleted (or goes out of scope), the destructor for each of its fields implements. The (default) copy constructor performs a deep copy, including duplicating object references; the (default) assignment operator acts as the destructor followed by the copy constructor.

The actual types of the fields in the C++ struct to which an IDL struct is mapped may be auxiliary classes for the purpose of storage management. In

particular, String and object reference field types are auxiliary classes that manage Strings and object references in the same way that the associated `_var` classes do. Although client code should not depend on the names of these auxiliary classes, the client code does need to know that struct fields containing Strings and object references are managed by these auxiliary classes.

When assigning a value to a struct field that is an object reference, the assignment operator for the struct field will automatically release the previous value (if any). When assigning an object reference pointer to a struct member, the struct member assumes ownership of the pointer (no `_duplicate` is done), and the application should no longer access the pointer directly. (If this is not the desired behavior, then the caller can explicitly `_duplicate` the object reference before assigning it to the struct member.) However, when assigning to an object reference struct member from a `_var` object or from another struct, union, array, or sequence member (rather than from an object reference pointer), a `_duplicate` is done automatically.

When assigning a value to a struct field that is a String, or when the struct is deleted or goes out of scope, any previously held (non-null) String is automatically freed. As when assigning to `String_vars`, assigning a `char*` to a String field does not make a copy, but assigning a `const char*`, a `String_var`, or another struct/union/array/sequence String member does automatically make a copy. One should never assign a string literal (for example, `"abc"`) to a String struct member without an explicit cast to `"const char*"`. When assigning a `char*` that occupies static storage (rather than one that was dynamically allocated), the caller can use `CORBA::string_dup` to duplicate the string before assigning it.

As with all constructed types, a `_var` type is provided for managing an instance of the C++ struct that corresponds to an IDL struct. When assigning one struct's `_var` to another, the receiving `_var` deletes its current pointer (thus running all contained destructors), and creates a new struct to hold the assignment result, which is initialized using copy constructors for each of the contained fields. Thus, for example, if the source struct has an object reference field, the struct `_var` assignment will automatically duplicate this reference.

The IDL that follows is used in the succeeding example, which shows both correct and incorrect ways to create and manipulate the corresponding C++ struct and the corresponding `_var` type :

```
Interface A
{
    struct S
    {
```

```

        string f1;
        A      f2;
    };
};

```

The following code illustrates both correct and incorrect ways to create and manipulate the corresponding C++ struct and the corresponding `_var` type.

```

{
    A::S_var sv1 = new A::S;
    A::S_var sv2 = new A::S;
    // sv1->f1 = "abc"; -- Wrong! f1 can't free this pointer later
    sv1->f1 = CORBA::string_alloc(20);
    A_ptr a1 = // get an A somehow
    A_ptr a2 = // get an A somehow
    sv1->f2 = a1; // a1 still has ref cnt = 1
    sv2->f1 = CORBA::string_alloc(20);
    sv2->f2 = a2; // a2 still has ref cnt = 1
    sv1 = sv2; // This runs copy ctors, and increments a2's ref cnt.
    // Also, a1's ref count is decremented.
    sv1->f1 = sv2->f1;
}

```

Union types

Union fields are not directly accessible to C++ programmers. Instead, the C++ mapping for IDL unions defines a class that provides accessor methods for the union discriminator and the corresponding union fields. The union discriminator accessor is named `_d`. The union field accessors are named using the IDL union field names and are overloaded to allow both reading and writing.

Setting a union's value using a field accessor automatically sets the discriminator, and releases the storage associated with the previous value, if any. It is an error for an application to attempt to access the union's value through an accessor that does not match the current discriminator value. It is also an error for the application to use the discriminator modifier method to implicitly switch between difference union members.

Unions with implicit default members (those that do not have an explicit default case and do not list all possible values of the discriminator as cases) provide a `_default` method, for setting the discriminator to a legal default value. This method causes the union's value to be composed only of the legal default value, since there is no explicit default member in this case.

A `_var` type is defined, for managing a pointer to a union in dynamically allocated memory.

To illustrate the C++ bindings for IDL unions, consider the following IDL:

```

module A
{
    interface X
    {
    };
    union U switch (long)
    {
        case 1: long u1;
        case 2: string u2;
        case 3: X u3;
    };
};

```

The following code illustrates usage of the C++ bindings corresponding to the previous IDL:

```

{
    X_ptr x = // get an X somehow
    A::U_var uv = new A::U;
    uv.u2((const char*) "testing"); // sets the discriminator to 2
    // and copies the string
    if (u.d() == 2) // the condition evaluates to true
    u.u1(23); // frees the string, and sets the discriminator to 1
    if (u.d() == 1) // the condition evaluates to true
    u.u3(x); // duplicates x and sets discriminator to 3
}

```

The default constructor of a union class does not initialize the discriminator or the union members, so the application must initialize the union before accessing it. The copy constructor and assignment operator perform deep copies. The assignment operator and destructor release all storage owned by the union.

With respect to memory management, accessor and modifier methods for union members work similarly to those for struct members. Modifier methods make a deep copy of their input when passed by value (for simple types) or by reference (for constructed types). Accessor methods that return a non-const reference can be used by the application to update a union member's value, but only for struct, union, sequence, and any members.

The modifier method for a string union member makes a copy when given a `const char*` or a `String_var`, but not when given a `char*`. As shown in the example above, a string literal should not be assigned to a union without an explicit `"const char*"` cast. The accessor method for a string union member returns a `const char*`, therefore the string union member cannot be modified. (This is done to prevent the string union member from being assigned to a `String_var`, resulting in memory management errors.)

The modifier method for an object reference union member always duplicates the input object reference and releases the previous object reference value, if

any. The accessor method for an object reference union member does not duplicate the returned object reference, because the union retains ownership of it.

The accessor method for an array union member returns a pointer to the array slice. The application can thus read or write the union-member array elements using subscript operators. If the union member is an anonymous array (one without an explicit type name), the union defines (typedefs) the slice type, by concatenating a leading underscore and appending “_slice” to the union member name.

Using WStrings

The WString type provides support for wide strings. It is fairly comparable to using strings except for type declarations and assignments:

```
#include <wctr.h> // For WChar and WString support
...
const wchar_t* wcomments = L"This policy looks pretty good...";
wchar_t* wcommentsResult=::CORBA::wstring_alloc(wcslen(wcomments));
::CORBA::WString_var wcommentsResult_var(wcommentsResult);
policyVar->wcomments(wcomments);

if (!wcscmp(wcommentsResult_var, wcomments) )
{
    cout << "Wcomments not set" << endl;
    return 1;
}
else
{
    cout << "Wcomments set correctly..." << endl;
}
wcommentsResult = policyVar->wcomments();
```

Exceptions

The preferred coding practice for handling errors in C++ and Java is by using Exceptions. The programming model and CORBA support this coding practice using the standard try and throw logic of exception handling. Handling exceptions is a critical part of the programming model. The exceptions that are thrown must be understood and handled appropriately by application developers.

Note: See a standard C++ book for further information regarding exceptions and their usage.

No matter how much care an object provider takes in implementing a business object, there are times when things go wrong. In these cases, a business object might need to throw an exception to the client to give the client the opportunity to recover from the error.

Which exceptions to use

CORBA exceptions are used to communicate between business objects and client applications. Specific rules must be followed regarding which CORBA Exceptions to use. There are several abstract CORBA exception classes defined:

- CORBA::Exception
- CORBA::UserException
- CORBA::SystemException

CORBA::Exception: This is the abstract class that is the base of all CORBA exceptions. Because this class is abstract, it is never thrown. However, it can be used in catch blocks to process all CORBA exceptions in one block.

CORBA::UserException: This is the abstract class for all CORBA user exceptions and is a subclass of CORBA::Exception. This class should be used as the base class of all user-defined exception classes. The contents of these classes have no special format. Methods that throw these classes must declare their usage in IDL using the raises clause.

CORBA::SystemException: This is the abstract class for all CORBA standard exceptions and is a subclass of CORBA::Exception. These exceptions can be thrown by any method regardless of the interface specification. Standard exceptions cannot be listed in raises expressions, therefore whether or not an interface throws a system exception is unknown. This means you should be prepared to handle standard exceptions on all method calls. Each standard exception includes a minor code to provide more detailed information. This field is used by Component Broker components. The definition of the minor code values is included in the online documentation.

Note: CORBA standard exceptions are a predefined list of exceptions. These can be thrown from any method. CORBA has defined the class that provides this support as CORBA::System Exception. See *The Common Object Request Broker: Architecture and Specification* for further information regarding CORBA exceptions.

Throwing exceptions

A business object might wrapper existing logic which might not be written in C++ or might not use the exception paradigm. These business objects must convert the existing exceptions or error return codes to CORBA exceptions that can be returned to the client program.

Any non-CORBA exception thrown by the business object is automatically mapped to CORBA::UNKNOWN by the framework. This does not provide specific information to the client and severely limits the error recovery

capability of the client program. These C++ exceptions should be mapped to appropriate CORBA exceptions by the business object.

Catching exceptions

It is a requirement to handle exceptions in client programs. Remember that any method might throw a standard exception, therefore an exception can be thrown by these methods at any time – even if there are no exceptions declared in the raises clause of that method. The default behavior for uncaught exceptions is to terminate that process. If this happens, suspect an uncaught exception first. The exact style of how or what exceptions are caught depends on what the client application does for error recovery but there are some good general rules to follow:

- Perform as specific error recovery as makes sense. By proper structuring of catch clauses specific error recovery can be done.
- Check for most specific exceptions first, most general exceptions last.
- Make use of information that is available in the exception. All CORBA exceptions support the `.id()` method that returns the exception identifier. System exceptions also provide `.minor()` and `.completed()` methods which return the minor code and completion status respectively.

A simple client code example:

```
try
{
    // Some real code goes here
    foo.boo();
}
// Catch any specific User exceptions defined for the method in the
// 'raises' clause
catch (IManagedClient::INoObjectWKey &nowk)
{
    // Process the error, more specific recovery could be made here
    // because the specific error is known
}
// Catch and process any other specific User exceptions
...
// Catch any other User exceptions defined for the method in the
// 'raises' clause
catch (CORBA::UserException &ue)
{
    // Process any other User exceptions. Use the .id() method to
    // record or display useful information
    cout << "Caught a User Exception: " << ue.id() << endl;
}
// Catch any System exceptions defined for the method in the
// 'raises' clause
catch (CORBA::SystemException &se)
{
    // Process any System exceptions. Use the .id(), and .minor()
    // methods to record or display useful information
```

```

    cout << "Caught a System Exception: " << ue.id() << ": " << ue.minor() << endl;
}
catch (...)
{
    // Process any other exceptions. This would catch any other C++
    // exceptions and should probably never occur
    cout << "Caught an unknown Exception" << endl;
}

```

Specific standard exceptions cannot be caught individually. If processing individual standard exceptions is required it can be done within the `CORBA::SystemException` catch block and using the `.id()` method.

Commonly used CORBA interfaces

References have already been made to interfaces that are not directly defined by the programming model that is provided by Component Broker, for example `_narrow()`. This is because the programming model is built on top of CORBA and relies on the interfaces that are already defined by CORBA. This section describes the most commonly used CORBA interfaces that Component Broker developers will use.

For further information on these and other operations defined by these interfaces, see the appropriate section in the Object Request Broker section of the *WebSphere Application Server Enterprise Edition Component Broker Programming Reference*.

CORBA class interfaces

The CORBA interface also provides some class operations that are commonly used. These are used like a C++ class reference (e.g. `CORBA::is_nil(somePointer);`)

is_nil This operation returns a boolean that indicates if the input object reference is nil. This is useful for many operations involving object references, including those operations that do not throw exceptions when they fail - for example `CORBA::Object::_narrow()`.

release

This operation releases resources associated with an object or pseudo-object reference. This operation may or may not perform a C++ delete operation. A reference count is used by this operation and `CORBA::Object::_duplicate()`. When the reference count reaches zero then the appropriate delete operations are performed. Care must be taken when using the `release` and `_duplicate` operations to ensure that objects are neither inadvertently deleted or are leaked. Alternatively use the `_var` technique described in the next section.

string_dup

This operation copies a string. The resulting string should be subsequently freed using the CORBA::string_free operation or assign the string to a _var variable which will free the string appropriately. Strings and wide strings, unlike the other basic CORBA types, have associated allocated memory. Care must be taken when using these variables. A common example is when returning a string from an operation.

CORBA::Object interfaces

_duplicate

This operation duplicates an object reference. This is particularly useful when passing references to objects to resolve memory ownership issues. For every _duplicate that is performed on an object an equal number of release() must also be performed for proper memory management. An alternative to the _duplicate() and release() logic is to use _var support as described in the next section of this chapter.

_is_a This operation is used to determine whether an object reference supports a given IDL interface. If the object supports the interface the _narrow operation can be successfully performed.

_is_equivalent

This operation is used to determine whether two object references refer to the same object.

_narrow

This operation is used to narrow a more generic interface to a more specific interface. This operation will return an empty pointer without throwing an exception if the interface cannot be narrowed to the requested type. Care must be taken to check the returned value before using it.

_nil This operation returns a nil CORBA::Object. This object could be used for comparison operations.

_non_existent

This operation determines whether an object reference refers to a valid object. This will result in verification of the object reference only, no other operations are performed on the requested object.

CORBA::ORB interfaces

object_to_string

This operation converts an object reference to an external form that can be stored for later use or exchanged between processes. The string_to_object operation can be used to reconstruct the object reference.

string_to_object

This operation converts an stringified object reference to an reconstructed object reference. The `object_to_string` operation must have been used to create the input stringified data.

Note: Although `object_to_string` is the way to save object references for future usage, the returned data should only be used with `string_to_object` to reconstruct that object reference. Do not use the string for comparing equivalence of object references.

The string is an IOR which can be composed of several pieces of data - some the ORB generates and some contributed by other components. The `object_to_string` operation may return different values at different times because various Object Services may be adding information to this IOR.

Name scoping and modules in the C++ bindings

IDL scoped names are mapped to C++ scopes as follows.

- In the IBM C++ bindings, IDL modules are, by default, mapped to C++ classes of the same name. If the programmer using the bindings `#defines _USE_NAMESPACE` before including the bindings, then the bindings map the IDL module to a C++ namespace of the same name. IDL definitions occurring within a module are mapped to corresponding C++ definitions within the C++ module class or namespace.
- IDL interfaces are mapped to C++ classes. All IDL constructs defined within an interface are mapped to corresponding C++ definitions within the C++ interface class.
- Every use in IDL of a C++ keyword (such as "class") is mapped into the same identifier with a leading underscore.

C++ bindings for interfaces

The CORBA 2.0 C++ client bindings define a variety of C++ types corresponding to a single IDL interface. Specifically, an IDL interface `I` is mapped to C++ types with the following four names: `I`, `I_ptr`, `IRef`, `I_var`. The types named `I` and `I_var` are classes. The types `I_ptr` and `IRef` are unspecified by CORBA, but are required to name the same type; these types are the C++ form for an object reference. (The use of the `IRef` types will be removed by the CORBA 2.0 specification.)

The class `I` is referred to as the interface class corresponding to IDL interface named `I`; the C++ mappings of the typedefs, operations, and constants defined within the IDL interface `I` appear publicly within the C++ interface class `I`. For example, an IDL operation that accepts an in parameter of interface type `I` is

mapped to a C++ virtual member function of the class named I that has a parameter of type I_ptr. Similarly, an IDL operation in I that returns data of type I is mapped to a C++ member function in the class I that returns data of type I_ptr.

As with other user-defined IDL types, the I_var type is used to assist storage management. Specifically, an I_var type holds an I_ptr and can be used as if it were an I_ptr. When a I_var type is assigned a new value or when it goes out of scope, it releases the I_ptr it is holding at that time.

The CORBA 2.0 specification prohibits CORBA-compliant applications from:

- Explicitly creating an instance of an interface class, as in:

```
I my_instance; // NOT ALLOWED!  
I_ptr my_instance = new I; // NOT ALLOWED!
```
- Declaring a pointer (I*) or a reference (I&) to an interface class.

Instead, the I_ptr, IRef, and I_var types must be used to hold object references, and object references can only be created (by client applications) by invoking methods that return object references. The interface class (I) is used by client applications only as a name scope.

IDL operations defined in (or inherited by) interface I are invoked in C++ using the arrow (->) operator on either an I_ptr, IRef, or I_var type.

Nil object references of type I_ptr can be obtained using a static member function of I called _nil(). Operations cannot be invoked on nil object references. The CORBA::is_nil function is the only CORBA-conformant way to determine whether a given object reference is nil. CORBA::release can be invoked on a nil object reference, but is not needed. The _duplicate and _narrow functions defined by the C++ bindings can be given a nil object reference.

In the IBM C++ bindings, the CORBA-prescribed types are implemented as follows:

1. The interface class for I is derived using virtual inheritance from the interface classes for I's IDL parents. When I has no IDL parents, its interface class is derived using virtual inheritance from CORBA::Object. Types, constants, and operations declared within the I interface are mapped to types, constants, and member functions declared within the corresponding interface class.
2. The object reference types I_ptr and IRef are typedef'd to I* (for example, an I_ptr points to an object of type I). However, CORBA 2.0 specifies that treating an I_ptr as a C++ pointer (e.g., using conversion to void*, arithmetic and relational operators, test for equality) is not conformant, although this is not enforced by the bindings.

3. An instance of I addressed is called a proxy, and is created by a proxy factory object of class ProxyFactory. For each interface I, the bindings define an ProxyFactory class, and provide a global instance of this class with the name _ProxyFactory.
4. Nil object references are represented as NULL pointers (but CORBA 2.0 conformant applications should not assume so, and should instead use the _nil() and is_nil() functions to manipulate nil object references).
5. The I_var class introduces an instance variable of type I_ptr. The purpose of an I_var object is to handle release operations on the I_ptr that it holds.
6. An auxiliary class I_SeqElem class is used to return sequence elements, and is similar to the I_var class. It is returned from array access operations on an IDL type sequence. An I_SeqElem is different from an I_var in that it must honor the release setting of the sequence from which it is selected (that is, it only owns the object that its I_ptr references if it was taken from a sequence that owns its buffer storage).

Managing object references

The mapping for interface I defines a static member function named _duplicate that takes as input an object reference of type I_ptr and returns an object reference of type I_ptr (potentially the same reference, when reference counting is employed, as is the case with IBM C++ bindings). The CORBA::release function indicates that the caller will no longer access the object reference, and the resources associated with the object reference can be deallocated. (In the IBM C++ bindings, an object reference is only deleted when its reference count falls to zero, that occurs only if CORBA::release is called for each _duplicate or _narrow performed on the object reference.)

Duplicating an object reference using _duplicate is analogous to copying a string before transferring ownership of it, and releasing an object reference is analogous to deleting a string that is no longer needed. Unlike strings, object references cannot be directly copied or deleted by the client programmer; object references are managed by the ORB and can only be duplicated or released by the application.

Widening object references

If interface A is a (direct or indirect) base of interface B, the following assignments do not require an explicit C++ cast.

- B_ptr to an A_var
- B_ptr to A_ptr
- B_ptr to Object_var
- B_ptr to Object_ptr
- B_var to A_ptr
- B_var to Object_ptr

B_var cannot be assigned to A_var or a compile-time error occurs. To assign B_var to A_var:

- Use B::_duplicate on B_var to create B_ptr.
- Assign B_ptr to A_var.

Narrowing object references

The mapping for an interface I defines a static member function named `_narrow` that takes as input an object reference of any type (for example, an `Object_ptr`) and returns an object reference of type `I_ptr`. If the referenced object (the actual implementation object corresponding to the proxy addressed by the input object reference) does not support the I interface, the result is `NULL`; otherwise, the `I_ptr` addresses an object that also supports the I interface. In the case where the proxy addressed by the input argument does not support interface I and the actual implementation object does, the `I_ptr` returned by `I::_narrow` addresses a different proxy object than the input argument.

The `_narrow` static member function does an implicit `_duplicate` of the input argument. Therefore, the caller is responsible for releasing both the object reference input to `_narrow` and the return result.

Narrowing to a C++ implementation

Given an interface pointer to an object, it is sometimes useful to narrow to the implementation pointer of the object. For example, given interface I, the C++ implementation hierarchy for I might look like:



You might want to convert a pointer to I into a pointer to `I_Impl`.

There is no CORBA-prescribed mechanism for this conversion. Within the confines of the C++ language, dynamic cast can be used. Since Component Broker does not require the C++ compiler to support dynamic cast, a second mechanism is provided, with the virtual `_narrow_impl()` method defined in `CORBA::Object`:

```
class Object
{
    ...
    virtual void *_narrow_impl(const char* impl_name = NULL)
```

```

    {
        return NULL;
    }
};

```

The default implementation of this method in CORBA::Object returns NULL, so that implementations that do not support `_narrow_impl()` need not worry about its existence. Implementations that support `_narrow_impl()` should override it, and provide an alternate implementation. The `impl_name` parameter is optional, and may be used by the implementation to perform sanity checks.

For the example with interface I, I_Impl might look like:

```

class I_Impl: virtual public I_Skeleton ...
{
    ...
    void *_narrow_impl(const char* impl_name)
    {
        if (impl_name != NULL && !strcmp(impl_name, "I_Impl"))
            return this;
        else
            return NULL;
    }
}

```

A user of this service would do the following to issue the narrow:

```

I_ptr i = ...; /* get the IDL ptr somehow */
I_Impl * ii = (I_Impl *) i->_narrow_impl("I_Impl");
if ( ii == NULL )
{
    // problem
}
else
{
    // Use ii
    ...
}

```

Storage management and `_var` types

The C++ bindings try to make the programmer's storage management responsibility as easy as possible. One aspect of this is the "`_var`" types. For each user-defined structured IDL type T (struct, union, sequences, and arrays) and for interfaces, the bindings generate both a class T and a class T`_var`. The classes CORBA::String and CORBA::Any also have corresponding CORBA::String`_var` and CORBA::Any`_var` classes.

The essential purpose of a `_var` object is to hold a pointer to dynamically allocated memory. A `_var` object can be used as if it were a pointer to the IDL type for which it is named; special constructors, assignment operators, and conversion operators make this work in a way that is invisible to programmers. The memory pointed to by a `_var` object is always considered to be owned (managed) by the `_var` object, and when the `_var` object is deleted, goes out of scope, or is assigned a new value, it deletes (or, in the case of an object reference, releases) the managed memory.

A typical `_var` object is declared by a programmer as an automatic (stack) variable within a code block, and is then used to receive an operation result or is passed to an operation as an out parameter. Later, when the code block is exited, the `_var` object destructor runs and its managed memory is deleted (or, for object references, released).

When a pointer (rather than another `_var` object or struct/union/array/sequence element) is assigned to (or used to construct) a `_var` object, this pointer should point to dynamically allocated memory, because the `_var` object does not make a copy; it assumes ownership of the pointer and will later delete it (or, for object references, release it). The single exception is that pointers to const data can be assigned to a `_var` object. When this occurs, the `_var` object dynamically allocates new memory and copies the const data into the new memory. A pointer assigned to a `_var` object must not be “owned” by some other data structure, and the pointer should not be subsequently used by the application except by the `_var` object.

The default constructor for a `_var` type loads the contained pointer with NULL. You must assign a value to a `_var` object created by a default constructor before invoking methods on it, just as you must assign a value to a pointer variable before invoking methods on it.

The copy constructor and `_var` assignment operator of a `_var` type perform a deep copy of the source data. The copy is later deallocated (or released, in the case of object references) when the `_var` is destroyed or when it is assigned a new value.

The following is the typical form for a `T_var` class, emitted for an IDL—constructed data type named `T`:

```
class T_var
{
    public:
        T_var ();
        T_var (T*);
        T_var (const T_var&);
        ~T_var ();
        T_var &operator= (T*);
};
```

```

    T_var &operator= (const T_var &);
    T * operator-> const ();
    ...
};

```

Argument passing considerations for C++ bindings

Rules must be observed when passing parameters to a CORBA object implementation. These rules are dependent on a combination of the IDL type of the argument and the argument mode (in, inout, out or return value), and must be followed to:

- Ensure the required access authority.
- Prevent memory leaks.
- Ensure that the allocation and deallocation of memory is performed consistently.

C++ type mapping for argument passing

The type used to pass the parameters of a method signature is dependent on the IDL type and the directionality of the parameter (in, inout, out, or return value). The rules are dictated by CORBA OMG IDL to C++ mapping. These mapping rules are captured in and enforced by the header files produced when an IDL interface description is compiled. Some rules cannot be enforced by the bindings. For example, parameters that are passed or returned as a pointer type (T*) or reference to pointer(T*&) should not be passed or returned as a null pointer. Memory management responsibilities cannot be enforced by the bindings. Client (caller) and implementation (callee) programmers must understand and implement according to these rules. The argument type mappings are discussed in the following paragraphs and summarized in the table, Table 26 on page 515. Memory management responsibilities are discussed in "Storage management responsibilities for arguments" on page 516.

For primitive types and enumerations, the type mapping is straightforward. For in parameters and return values the type mapping is simply the C++ type representation (abbreviated as "T" in the text that follows) of the IDL specified type. For inout and out parameters the type mapping is a reference to the C++ type representation (abbreviated as "T&" in the text that follows).

For object references, the type mapping uses `_ptr` for in parameters and return values and `_ptr&` for inout and out parameters. That is, for a declared interface A, an object reference parameter is passed as type `A_ptr` or `A_ptr&`. The conversion functions on the `_var` type permit the client (caller) the option of using `_var` type rather than the `_ptr` for object reference parameters. Using the `_var` type may have an advantage in that it relieves the client (caller) of the responsibility of deallocating a returned object reference (out parm or

return value) between successive calls. This is because the assignment operator of a `_ptr` to a `_var` automatically releases the embedded reference.

The type mapping of parameters for aggregate types (also referred to as complex types) are complicated by when and how the parameter memory is allocated and deallocated. Mapping in parameters is straightforward because the parameter storage is caller allocated and read only. For an aggregate IDL type `t` with a C++ type representation of `T` the in parameter mapping is `const T&`. The mapping of out and inout parameters is slightly more complex. To preserve the client capability to stack allocate fixed length types, OMG has defined the mappings for fixed-length and variable-length aggregates differently. The inout and out mapping of an aggregate type represented in C++ as `T` is `T&` for fixed-length aggregates and as `T*&` for variable-length aggregates.

Table 26. Basic argument and result passing

Data	Type In	Inout	Out	Return
short	short	short&	short&	short
long	long	long&	long&	long
unsigned short	ushort	ushort&	ushort&	ushort
unsigned long	ulong	ulong&	ulong&	ulong
float	float	float&	float&	float
double	double	double&	double&	double
boolean	boolean	boolean&	boolean&	boolean
char	char	char&	char&	char
wchar	wchar	wchar&	wchar&	wchar
octet	Octet	Octet&	Octet&	Octet
enum	enum	enum&	enum&	enum
object reference ptr	objref_ptr	objref_ptr&	objref_ptr&	objref_ptr
struct, fixed	const struct&	struct&	struct&	struct
struct, variable	const struct&	struct&	struct*&	struct*
union, fixed	const union&	union&	union&	union
union, variable	const union&	union*&	union*&	union*
string	const char*	char*&	char*&	char*
wstring	const char*	char*&	char*&	char*
sequence	const sequence&	sequence&	sequence*&	sequence*
array, fixed	const array	array	array	array slice*
array, variable	const array	array	array slice*&	array slice*

Table 26. Basic argument and result passing (continued)

Data	Type In	Inout	Out	Return
any	const any&	any&	any*&	any*

For an aggregate type represented by the C++ type T, the T_var type is also defined. The conversion operations on each T_var type allows the client (caller) to use the T_var type directly for any directionality, as opposed to using the required form of the T type (T, T& or T*&). The emitted bindings define the operation signatures in terms of the parameter passing modes shown in Table 27, and the T_var types provide the necessary conversion operators to allow them to be passed directly.

Table 27. T_var argument and result passing

Data Type	In	Inout	Out	Return
object referenc var	const object_var&	objref_var&	objref_var&	objref_var
struct_var	const struct_var&	struct_var&	struct_var&	struct_var
union_var	const union_var&	union_var&	union_var&	union_var
string_var	const string_var&	string_var&	string_var&	string_var
sequence_var	const sequence_var&	sequence_var&	sequence_var&	sequence_var
sequence_var array_var	const array_var&	any_var&	any_var&	any_var

For parameters that are passed or returned as a pointer type (T*) or reference to pointer(T*&) the programmer should not pass or return a null pointer. This cannot be enforced by the bindings.

Storage management responsibilities for arguments

Table 28 on page 517 summarizes the storage access and allocation responsibilities for argument passing. As an overall requirement when allocating and deallocating argument storage, the storage allocation rules for the specific type must be followed. Specifically, for strings, sequences, and arrays or for aggregate types composed of these types, the associated memory allocation and deallocation functions must be used. For string types this means the use of string_alloc(), string_dup(), and string_free(), for sequence types this means the use of allocbuf() and freebuf() and for arrays this means the use of T_alloc(), T_dup() and T_free(). The memory deallocation responsibilities of the client can be minimized by stack allocation and the use of the _var types when that is possible. When an argument is passed or returned as a pointer type, a NULL pointer value should never be passed or returned.

Table 28. Argument storage responsibilities

Data Type	Inout	Out	Return
short	1	1	1
long	1	1	1
unsigned short	1	1	1
unsigned long	1	1	1
float	1	1	1
double	1	1	1
boolean	1	1	1
char	1	1	1
octet	1	1	1
enum	1	1	1
object reference pointer	2	2	2
struct, fixed	1	1	1
struct, variable	1	3	3
union, fixed	1	1	1
union, variable	1	3	3
string	4	3	3
sequence	5	3	3
array, fixed	1	1	6
array, variable	1	6	6
any	5	3	3

in parameters

The caller (client) must allocate the input parameters. The callee (implementation) is restricted to read access. The caller is responsible for the eventual release of the storage. Primitive types and fixed-length aggregate types may either be heap allocated or stack allocated. By their nature, variable-length aggregates cannot be completely stack allocated.

inout parameters

For inout parameters, the caller provides the initial value and the callee may change that value. For primitive types and fixed-length aggregates this is straight forward. The caller provides the storage and the callee overwrites the storage on return. For variable-length aggregates the size of the contained data provided on input may differ than the size of the contained data provided on output. Therefore, the callee is required to deallocate any input contained data that is being replaced on output with callee allocated data. For object references,

the caller provides an initial value: if the callee reassigns the value the callee must first release the original input value. The callee assumes or retains ownership of the returned parameters and must eventually deallocate or release them.

out parameters

For primitive types and fixed-length aggregate types, the caller allocates the storage for the out parameter and the callee sets the value. For variable-length aggregate types, the caller allocates a pointer and passes it by reference and the callee sets the pointer to point to a valid instance of the parameter's type. For object references the caller allocates storage for the `_ptr` and the callee sets the `_ptr` to point to a valid instance.

Because a pointer to an array in C++ must actually be represented as a pointer to the array element type, CORBA defines an `array_slice` type, where a slice is an array with all the dimensions of the original except the first. The output parameter is typed as a reference to an `array_slice` pointer. The caller allocates the storage for the pointer and the callee updates the pointer to point to a valid instance of an `array_slice`.

The caller assumes or retains ownership of the output parameter storage and must eventually deallocate it or, in the case of object references, release it.

return values

For primitive types and fixed-length aggregate types, the caller allocates the storage for the return value and the callee returns a value for the type. For variable-length aggregate types, the caller allocates a pointer and the callee returns a pointer to an instance of the type. For object references the caller allocates storage for the `_ptr` and the callee returns a `_ptr` that points to a valid object instance.

As a pointer to an array in C++ must actually be represented as a pointer to the array element type, the `array_slice` type is used for returning array values. The caller allocates storage for a pointer to the `array_slice` and the callee returns a pointer to a valid instance of an `array_slice`.

The caller assumes or retains ownership of the storage associated with returned values and must eventually deallocate it or, in the case of object references, release it.

For definitions of the numerical values in Table 28 on page 517, refer to Table 29 on page 519.

Table 29. Argument passing cases

Case	Description
1	Caller allocates all necessary storage, except that which may be encapsulated and managed within the parameter itself. For inout parameters, the caller allocates the storage but need not initialize it, and the callee sets the value. Function returns are by value.
2	Caller allocates storage for the object reference. For inout parameters, the caller provides an initial value; if the callee wants to reassign the inout parameter, it will first call CORBA:release on the original input value. To continue to use an object reference passed in as an inout, the caller must first duplicate the reference. The caller is responsible for the release of all out and return object references. Release of all object references embedded in other structures is performed automatically by the structures themselves.
3	The callee sets the pointer to point to a valid instance of the parameter's type. For returns, the callee returns a similar pointer. The callee is not allowed to return a null pointer in either case. In both cases, the caller is responsible for releasing the returned storage. To maintain local/remote transparency, the caller must always release the returned storage, regardless of whether the callee is located in the same address space as the caller or is located in a different address space. Following the completion of a request, the caller is not allowed to modify any values in the returned storage: in order to do so, the caller must first copy the returned instance into a new instance, then modify the new instance.
4	For inout strings, the caller provides storage for both the input string and the char* pointing to it. Since the callee may deallocate the input strings and reassign the char* to point to new storage to hold the output value, the caller should allocate the input string using string_alloc() . The size of the out string is therefore not limited by the size of the in string. The caller is responsible for deleting the storage for the out using string_free() . The callee is not allowed to return a null pointer for an inout, out or return value.
5	For inout sequences and any's , assignment or modification of the sequence or any may cause deallocation of owned storage before any reallocation occurs, depending upon the state of the Boolean release parameter with which the sequence or any was constructed.
6	For out parameters, the caller allocates a pointer to an array slice, which has all the same dimensions of the original array except the first, and passes the pointer by reference to the callee. The callee sets the pointer to point to a valid instance of the array. For returns, the callee returns a similar pointer. The callee is not allowed to return a null pointer in either case. In both cases, the caller must always release the returned storage, regardless of whether the callee is located in the same address space as the caller or is located in a different address space. Following completion of a request, the caller is not allowed to modify any values in the returned storage: in order to do so, the caller must first copy the returned array instance into a new array instance, then modify the new instance.

C++ client bindings

In these bindings, the client usage picture for the IDL types declared in the file `T.idl` appears as follows. Bold lines enclose files that are generated from IDL. Double lines enclose files that would normally be produced by a programmer or development tool.

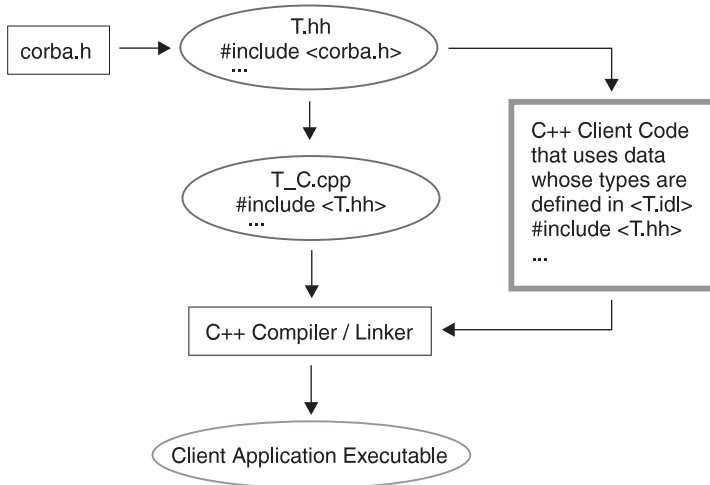


Figure 74. Client usage picture for IDL types

The `corba.h` header file defines the C++ mappings for primitive IDL data types and other types required by the bindings, within a scope called CORBA. For more information about the scope see “IDL name scoping” on page 447. These types are implemented in a shared library that can be linked with a client application. A client application is created by compiling/linking emitted bindings and client code to produce an executable file.

The C++ bindings for the IDL types defined in the file `T.idl` are represented by a set of declarations in the emitted `T.hh` header file. The classes declared in `T.hh` that support client code are implemented by the code emitted into a corresponding `T_C.cpp` implementation file. The pair of files `T.hh` and `T_C.cpp` thus collectively provide the client bindings for `T`. (To minimize the number of generated files, some types used by servers are also declared in `T.hh`).

In general, the C++ bindings map non-primitive IDL types to C++ classes that implement constructors, destructors, assignment operators, and other functionality. Auxiliary classes are also sometimes defined, such as classes to automate storage management for array elements, sequence elements, and structure and union fields. The names of these auxiliary classes are not

specified by CORBA, because specially-designed conversion operators and copy constructors hide their existence from client code. These classes are not of interest to programmers that use the bindings.

C++ server bindings

To allow IDL interfaces to be implemented in C++, server-side bindings are emitted. The resulting classes work with the client bindings. The following figure illustrates the module structure of the server-side bindings, assuming that interface T is declared in the file T.idl.

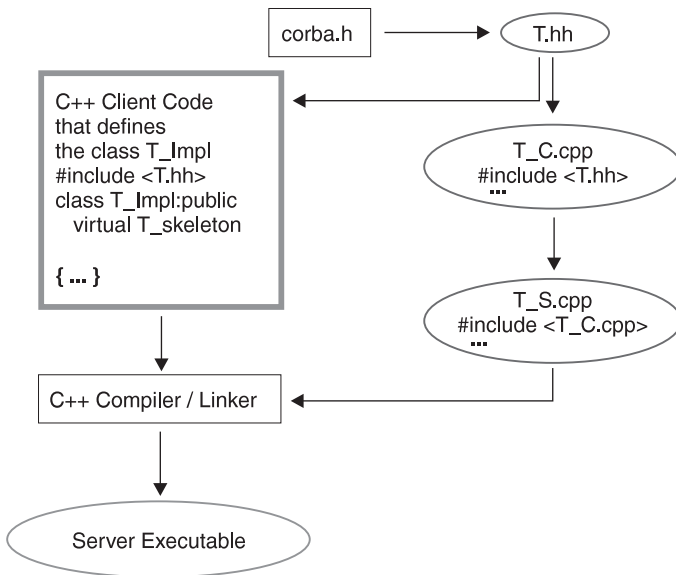


Figure 75. Module structure of server-side bindings

The differences between this figure and the client-side figure presented in Client C++ Bindings are:

- An emitted T_S.cpp file that provides server-side implementation bindings is compiled.
- Server-side C++ code (written by a programmer) defines the implementation for the operations introduced and inherited by the T interface.

The T_S.cpp file provides an implementation for the class T_Dispatcher. This class inherits from the dispatcher classes corresponding to T's parents. An instance of this class contains a T_ptr that addresses the T_Impl object upon which it will dispatch operation invocations. Each target object (for example, each T_Impl instance) exported by a server must have a corresponding

dispatcher object, whose purpose is to receive a CORBA::Request object, determine what method is being invoked, stream the method arguments out into local variables, invoke the method on the target object, then stream the results back into the request so these can be returned to the caller.

The target object for a T_Dispatcher is an instance of the T_Impl class, which subclasses from (at least) the T_Skeleton class defined by the implementation bindings (in the file T.hh). The T_skeleton class inherits from the T interface class and the skeleton classes corresponding to T's parents. As a result, T_skeleton inherits all the methods that T_Impl must support. Furthermore, this is done in a way that forces T_Impl to indeed provide implementations for all of these methods.

Take notice of the fact that the class name T_Impl is entirely arbitrary. The implementation class may have any name. Also note that the implementation class is not nested within any of the C++ classes that might be used to provide nesting scopes corresponding to IDL modules within which the interface T is defined. Thus, naming conflicts are a concern. A simple solution is to use underscores to concatenate module names with the name of the implemented interface. For example, if the interface T is defined within module M, then the implementation class name M_T_Impl can be used.

If the programmer responsible for T_Impl desires, the implementations (and supporting instance data) for any or all of T's parents can be inherited from their implementation classes, using C++ inheritance. Alternatively, T_Impl can provide its own implementation for the operations inherited into T. The image below graphically illustrates these options from an IDL snippet, using a dotted inheritance line to show optional C++ inheritance.

```
interface A
{
    ...
};
interface B : A
{
    ...
};
```

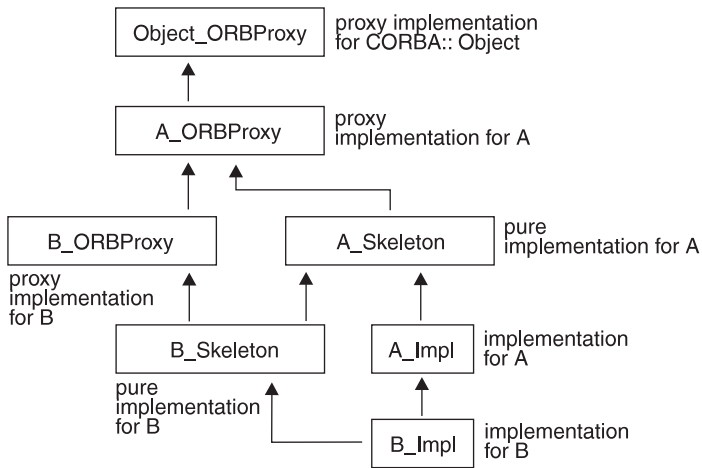


Figure 76. Inherited implementation

Figure 76 focuses on C++ classes. The term *pure* is based on the use of this word in C++ to describe virtual functions that have no implementation (denoted in C++ by assigning a 0 to the name of the virtual function). Classes with pure virtual functions cannot be instantiated. Therefore, the skeleton classes require a subclass to provide complete implementations for all virtual functions in an interface. The dotted line in the previous figure indicates one way that B_Impl can provide implementations for the pure virtual functions inherited from A_Skeleton (using B_Skeleton). One way of viewing the skeleton classes is that they “turn off” proxy behavior and require subclasses to explicitly provide an alternative (non-proxy) implementation.

C++ binding restrictions

When a forward reference to an interface appears within an IDL module, the IDL compiler issues an error message if the referenced interface is not defined within the module. When a similar unresolved forward reference appears at global (file) scope, a warning is issued that indicates the bindings being emitted won’t include a mapping for the undefined interface. For information on the scope see “IDL name scoping” on page 447. The assumption is that the interface will be defined by other bindings than those being currently generated. This approach supports IDL files with mutually-referential interfaces (as long as they appear at global scope). The following example that illustrates how to organize the IDL files for such cases:

```

// file foo.idl
#ifndef foo_idl
#define foo_idl
interface Foo; // declare Foo so bar.idl can refer to it

```

```

#include <bar.idl>
interface Foo
{
    Bar foo1(); // notice the use of Bar
};
#endif // foo_idl
// file bar.idl
#ifndef bar_idl
#define bar_idl
interface Bar; // declare Bar so foo.idl can refer to it
#include <foo.idl>
interface Bar
{
    Foo bar1(); // notice the use of Foo
};
#endif // bar_idl

```

Due to problems inherent to the CORBA 2.0 mapping for C++, there are currently two known limitations with respect to handling legal CORBA 2.0 IDL. The compiler provides informative error messages in these two cases, and indicates that C++ bindings cannot be generated. The cases are:

- The C++ bindings map most IDL data types to C++ classes contained within a nesting scope provided by another C++ class. However, it is not legal to define a nested C++ class (or any other type) that has the same name as a containing C++ class. Thus, for example, the following IDL cannot be mapped to useful C++ bindings:

```

module X
{
    interface X ...;
    // or struct X ... ;
    // or union X ...;
    // or typedef sequence < > X;
    // ...
};

```

- The C++ bindings map attributes into overloaded C++ accessor functions whose name is the attribute name. As a result, for example, the following IDL will not map to useful C++ bindings (because Y's l method interferes with the inherited mapping for X's attribute). If Y's method took any arguments, there would not be a problem, because of C++ overloading. The compiler indicates an error only when C++ overloading won't distinguish inherited accessors from newly introduced methods (or vice versa).

```

interface X
{
    attribute long l;
}
interface Y : X
{
    long l();
};

```

Appendix D. Java ORB properties

The default behavior of the Java ORB can be modified through the use of various properties. These properties may be set in a number of ways, and a specific precedence order applies. There are a set of properties mandated by OMG, as well as a larger set which are IBM specific. Those properties defined by OMG have the prefix "org.omg.CORBA". All of IBM's properties have the prefix "com.ibm.CORBA".

Property setting mechanisms

Properties of the Java ORB can be set using any of the following mechanisms. The precedence order is *1 to n*, meaning that a property set through mechanism 3 will override the same property set by mechanism 2.

1. ClientStyleImageURL - This is actually an indirect setting of the Java ORB's properties through a file containing properties. See com.ibm.CORBA.ClientStyleImageURL in Table 30 on page 526 for details.
2. System properties - This includes all properties set on the command line using the -Dcom.ibm.CORBA.<property> option.
3. Programmatically through the Properties object passed to ORB.init().
4. Command line arguments set using the -ORB<property> format.
5. HTML applet parameters for applets.

Mechanisms 4 and 5 are mutually exclusive. Mechanism 4 can be used in applications. Mechanism 5 can be used in applets.

You may have noticed that the list above defines that the ClientStyleImageURL is the first source for properties, yet its name is a property itself. Therefore, the configuration file cannot be read until all the other properties are read since its value (the URL) may be overridden. This conflict is remedied by parsing and applying each properties mechanism and reading the configuration file at the very end. Any properties set within the configuration file are added to the Java ORB only if they do not already exist, because if they already exist, the values from the configuration file would have been overridden.

Details of Java ORB properties

The following table provides details of Java ORB properties.

Table 30. Java ORB properties

Property	Purpose	Range of Values	Default Value	Required/Optional
org.omg.CORBA.ORBClass	Defines the ORB implementation			Required
org.omg.CORBA.ORBSingletonClass	Defines the ORBSingleton implementation			Required Note: When using CBA Series Global, this property is set internally and cannot be overridden.
com.ibm.CORBA.BootstrapHost	The hostname of the machine on which initial server contact for this client resides.	The value of this property is a string. This string can be a host name or the IP address (ex. 9.5.88.112).	If this property is not set the local host is retrieved by calling one of the following methods: For applications: InetAddress. getLocalHost(). getHostAddress(); For applets: <applet>.getCodeBase(). getHost();	Optional
com.ibm.CORBA.BootstrapPort	The port of the machine on which the initial server contact for this client is listening.	0 to 2147483647 (Java max int)	900	Optional

Table 30. Java ORB properties (continued)

Property	Purpose	Range of Values	Default Value	Required/Optional
com.ibm.CORBA.TunnelAgentURL	The URL of the servlet used to support HTTP tunneling.	The format of the value of this property is a properly formed URL, i.e.: http://w3.mycorp.com:81/servlet/com.ibm.CORBA.services.IIOPTunnelServlet	For applets: http://<applethost:port>/servlet/com.ibm.CORBA.services.IIOPTunnelServlet There is no default for applications.	Required if using HTTP tunneling. See: com.ibm.CORBA.ForceTunnel for more information.
com.ibm.CORBA.ForceTunnel	Controls how the Java client ORB attempts to use tunneling.	“always” : Do HTTP tunneling right away, do not try TCP connections first. “never” : Never try HTTP tunneling, disable the feature. In this case, if TCP connection fails, a CORBA system exception COMM_FAILURE is thrown. “whenrequired” : This is the default value. Try TCP connection first. If that fails, use HTTP tunneling.	“whenrequired”	Optional

Table 30. Java ORB properties (continued)

Property	Purpose	Range of Values	Default Value	Required/ Optional
com.ibm.CORBA.CommTrace	This property is used to turn on wire tracing. Every incoming and outgoing GIOP message will be output to system.err.	"true", "false"	"false"	Optional

Table 30. Java ORB properties (continued)

Property	Purpose	Range of Values	Default Value	Required/Optional
com.ibm.CORBA.ClientStyleImageUrl	This property defines the location of a file containing a set of properties for configuring the Java ORB. This is useful when you have multiple instances of the ORB (perhaps multiple client applications or applets) that you would like to configure identically. For more information see "Administering Clients" in the <i>WebSphere Application Server Enterprise Edition Component Broker for Windows NT and AIX System Administration Guide</i>	The format of the value of this property is a properly formed URL, i.e.: http://w3.mycorp.com/JavaClients.config	None	Optional
com.ibm.CORBA.requestTimeout	Defines the number of seconds to wait before timing out on a Request message	1 to 2147483647 (Java max int) and infinity (0)	180	Optional

Table 30. Java ORB properties (continued)

Property	Purpose	Range of Values	Default Value	Required/Optional
com.ibm.CORBA.locateRequestTimeout	Defines the number of seconds to wait before timing out on a LocateRequest message	1 to 2147483647 (Java max int) and infinity (0)	180	Optional
Note: The URL strings above have been split due to the size of the table. The string continues on the following line with no space.				

This is not an exhaustive list of Java ORB Properties. It is a list of the properties which the Java ORB is aware of at the time of this writing. The Java ORB properties are designed such that any property whose prefix is "com.ibm.CORBA." is considered a valid property and, if seen, will be stored and accessible through ORB.getProperty (String name). Command line property syntax is "-ORB<PropertyName>" This is translated into "com.ibm.CORBA.<PropertyName>".

Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS DOCUMENT "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OR CONDITIONS OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the document. IBM may make

improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

For Component Broker:

IBM Corporation
Department LZKS
11400 Burnet Road
Austin, TX 78758
U.S.A.

For TXSeries:

IBM Corporation
ATTN: Software Licensing
11 Stanwix Street
Pittsburgh, PA 15222
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available

sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples may include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks and service marks

The following terms are trademarks or registered trademarks of the IBM Corporation in the United States, other countries, or both:

AFS	IMS
AIX	MQSeries
AS/400	MVS/ESA
CICS	OS/2
CICS OS/2	OS/390
CICS/400	OS/400
CICS/6000	PowerPC
CICS/ESA	RISC System/6000
CICS/MVS	RS/6000
CICS/VSE	S/390
CICSplex	Transarc
DB2	TXSeries
DCE Encina Lightweight Client	VSE/ESA
DFS	VTAM
Encina	VisualAge
IBM	WebSphere

Microsoft, Windows, Windows NT, and the Windows 95 logo are trademarks of Microsoft Corporation in the United States and/or other countries.

Oracle and Oracle8 are registered trademarks of the Oracle Corporation in the United States and/or other countries.

UNIX is a registered trademark of The Open Group in the United States and/or other countries licensed exclusively through X/Open Company Limited.

OSF and Open Software Foundation are registered trademarks of the Open Software Foundation, Inc.

* HP-UX is a Hewlett-Packard branded product. HP, Hewlett-Packard, and HP-UX are registered trademarks of Hewlett-Packard Company.

Orbix is a registered trademark and OrbixWeb is a trademark of IONA Technologies Ltd.

Sun, SunLink, Solaris, SunOS, Java, all Java-based trademarks and logos, NFS, and Sun Microsystems are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and/or other countries.

Some of this documentation is based on material from Object Management Group bearing the following copyright notices:

Copyright 1995, 1996 AT&T/NCR
Copyright 1995, 1996 BNR Europe Ltd.
Copyright 1991, 1992, 1995, 1996 by Digital Equipment Corporation
Copyright 1996 Gradient Technologies, Inc.
Copyright 1995, 1996 Groupe Bull
Copyright 1995, 1996 Expersoft Corporation
Copyright 1996 FUJITSU LIMITED
Copyright 1996 Genesis Development Corporation
Copyright 1989, 1990, 1991, 1992, 1995, 1996 by Hewlett-Packard Company
Copyright 1991, 1992, 1995, 1996 by HyperDesk Corporation
Copyright 1995, 1996 IBM Corporation
Copyright 1995, 1996 ICL, plc
Copyright 1995, 1996 Ing. C. Olivetti &C.Sp
Copyright 1997 International Computers Limited
Copyright 1995, 1996 IONA Technologies, Ltd.
Copyright 1995, 1996 Itasca Systems, Inc.
Copyright 1991, 1992, 1995, 1996 by NCR Corporation
Copyright 1997 Netscape Communications Corporation
Copyright 1997 Northern Telecom Limited
Copyright 1995, 1996 Novell USG
Copyright 1995, 1996 02 Technolgies
Copyright 1991, 1992, 1995, 1996 by Object Design, Inc.
Copyright 1991, 1992, 1995, 1996 Object Management Group, Inc.
Copyright 1995, 1996 Objectivity, Inc.
Copyright 1995, 1996 Oracle Corporation
Copyright 1995, 1996 Persistence Software

Copyright 1995, 1996 Servio, Corp.
Copyright 1996 Siemens Nixdorf Informationssysteme AG
Copyright 1991, 1992, 1995, 1996 by Sun Microsystems, Inc.
Copyright 1995, 1996 SunSoft, Inc.
Copyright 1996 Sybase, Inc.
Copyright 1996 Taligent, Inc.
Copyright 1995, 1996 Tandem Computers, Inc.
Copyright 1995, 1996 Teknekron Software Systems, Inc.
Copyright 1995, 1996 Tivoli Systems, Inc.
Copyright 1995, 1996 Transarc Corporation
Copyright 1995, 1996 Versant Object Technology Corporation
Copyright 1997 Visigenic Software, Inc.
Copyright 1996 Visual Edge Software, Ltd.

Each of the copyright holders listed above has agreed that no person shall be deemed to have infringed the copyright in the included material of any such copyright holder by reason of having used the specification set forth herein or having conformed any computer software to the specification.

WHILE THE INFORMATION IN THIS PUBLICATION IS BELIEVED TO BE ACCURATE, THE OBJECT MANAGEMENT GROUP, AND THE COMPANIES LISTED ABOVE MAKE NO WARRANTY OF ANY KIND WITH REGARDS TO THIS MATERIAL INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. The Object Management Group and the companies listed above shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.



This software contains RSA encryption code.

Other company, product, and service names may be trademarks or service marks of others.

Index

Special Characters

(IDL), Interface Definition Language

- attribute declarations 457
- emitted file names 469
- exception declarations 457
- IDL syntax 459
- idlc command 464
- IDL_OPTIONS environment variable 470
- interface declarations 453
- multiple IDL interfaces and modules 464
- name scoping 447
- operation declarations 454
- options for the idlc command 465
- reserved keywords for IDL 460
- type and constant declarations 448

(MOFW) client programming model 59, 85

A

ActiveX Client Programming Model 211

ActiveX Client Programming Model: Basic Tasks 216

- concurrency control 228
- creating a managed object 221
 - Creating a New Object - Create from Copy 223
 - Creating a New Object — Create From Key 221
- finding a managed object 217
 - bound in the naming service 218
 - by methods on held objects 221
 - using the PrimaryKey helper class 219
- initializing the CB client environment 216
- releasing and deleting objects 228
- remembering your favorite objects 227
- using a managed object 221
- using sets of objects 225
- When References Explode 228

ActiveX Client View of Component Broker Applications 211

Advanced Concepts, MOFW Client Programming Model

- conventions and guidelines 133
- create_object() method 133
- expanding the client programming interface 133
- queries, iterations and specialized homes 133
- session service 133
- transactions 133
- using keyed reference collections 133

Agent 38

application development, object-oriented

- bottom-up 4, 34
- combined approaches 10
- meet in the middle 8, 33
- top-down 3, 32

application example, personal life insurance

- application model 33
- design model 34
- object model 32

Artifacts Produced in Building Objects 445

Assembling and Installing Business Objects 351

- application adaptor methods 359
- augmentation of OMG services methods 360
- business object methods 358
- Create the Managed Object Class and Implementation 354
- OMG services methods 359
- special methods 359
- specialized homes 360

B

Beneficiary 39

bottom-up application development 4, 34

business logic, life insurance example

- Create Customer 47
- Modify Beneficiary 48
- ModifyPolicy 47

business logic, life insurance example (*continued*)

- Process Claim 48

Business Object, Extending a

- choosing an inheritance pattern 160
- essential state extensions 159
- Extending a Business Object Interface 158
- extension summary 168
- implement the additional business logic 161
- MOFW IManageable requirements 163
- MOFW requirements — IManaged server 165
- more copy helper classes 167
- more key classes 167
- other variations to consider 168

Business Object Attributes 58

Business Object Basics 55, 129

- Artifacts Table for Business Objects 55, 129
- summary 55

Business Object State 57

Business Objects, Assembling and Installing

- application adaptor methods 359
- augmentation of OMG services methods 360
- business object methods 358
- Create the Managed Object Class and Implementation 354
- OMG services methods 359
- special methods 359
- specialized homes 360

C

C++ 13

C++ bindings 481

C++ Bindings for Data Types 483

C++ Bindings for Interfaces 508

C++ CORBA Programming 481

- C++ binding restrictions 523
- C++ bindings for data types 483
- C++ bindings for interfaces 508
- C++ client bindings 520
- C++ server bindings 521
- CC++ bindings 481

- C++ CORBA Programming 481
 - (continued)*
 - CORBA types and business objects 482
 - exceptions 503
 - name scoping and modules in the C++ bindings 508
 - storage management and `_var` types 512
 - storage management responsibilities for arguments 516
- C++ local only development process 352
- Cardinality Relations, Data Object Customization for
 - Bottom-Up Customizations 398
 - Cardinality-1 Relationships 400
 - Cardinality-N Relationships 402
 - enabling additional clients 417
 - mapping helpers 403
 - packaging for client and server (VA C++) 406
 - create functions for dynamic DLL loading 408
 - DLL packaging 406
 - exposing interfaces to business object builders 409
 - exposing interfaces to clients 409
 - packaging the DLL for the ActiveX visual C++ client 409
 - packaging the Java client code 410
 - summarizing relationships implementations 403
 - Top-Down Customizations 396
 - Top-Down Versus Bottom-Up Relations 396
- circular references 65
- Claim 45
- Client Programming Interface 434
 - QOS interfaces for non-transactional support 434, 437
 - Quality of Service Interface 434, 435
- Client Programming Model: Basic Tasks 86
- Client Programming Model — Advanced Concepts, MOFW
 - conventions and guidelines 133
 - `create_object()` method 133
- Client Programming Model — Advanced Concepts, MOFW
 - (continued)*
 - expanding the client programming interface 133
 - queries, iterations and specialized homes 133
 - session service 133
 - transactions 133
 - using keyed reference collections 133
- Client View of Component Broker Applications 60
- Coding Tips for proper CORBA Memory Management 100
- combined development approaches 10
- conventions, programming C++ 13
 - Interface Definition Language 12
 - Java 13
 - naming conventions 14
 - other 13
- Conventions and Guidelines 151
 - creating persistent objects 151, 152
 - finding persistent objects 151, 152
- Copy Helpers Sharing Opportunities 75
- CORBA::Object Interfaces 507
- CORBA class interfaces 506
- CORBA Programming, C++
 - C++ binding restrictions 523
 - C++ bindings for data types 483
 - C++ bindings for interfaces 508
 - C++ client bindings 520
 - C++ server bindings 521
 - CC++ bindings 481
 - CORBA types and business objects 482
 - exceptions 503
 - name scoping and modules in the C++ bindings 508
 - storage management and `_var` types 512
 - storage management responsibilities for arguments 516
- CORBA Types and Business Objects 482
- CosTransactions module 81
- Create Customer 47
- Create From Key, Creating a New Object —
 - Factory 92
 - factory finder 92
- `create_object()` method 153
- creating a managed object 92
- Creating a New Object — Create from Copy 94
- Creating a New Object — Create From Key 92
 - Factory 92
 - factory finder 92
- Creating Specialized Homes
 - extending the interface to IHome 192
 - alternatives to IManagedClient IHome 193
 - details 193
- Implement the Extended IHome Interface 194
 - copy helper 199
 - implementation 195
 - Implementation Interface 194
 - keys 199
 - leveraging server provided essential state extensions 199
 - MOFW IManageable Requirements — implementation 197
 - MOFW Requirements — IManagedObject Interfaces 198
 - summary of home extension 201
- Customer 40
- customization, data object
 - additional considerations 371, 403
 - BOIM data object customization — cache service 371
 - BOIM data object customization — static SQL 364
 - BOIM data object implementation 384
 - BOIM data object interfaces 382
 - data object data management patterns 389
 - interfaces needed by persistent data objects 363
 - SQL data object interfaces 364
 - static SQL data object implementation 366
 - summary of data object customization 386

- customization, data object
(*continued*)
 - transient data object-any key
(production use) 381
 - transient data object
customization – UUID key
(production use) 379
- Customization and Inheritance, Data Object
 - Additional Methods – Default
Constructor 393
 - CarPolicy BOIM Data Object
Implementation 391
 - CarPolicy BOIM Data Object
Interfaces 389
 - Framework Required Code –
create() Function 393
 - Framework Required Method –
internalizeFromCopyHelper 391
 - Framework Required Method –
internalizeFromPrimaryKey 391
 - Methods To Support Attributes –
Getters 393
 - Methods To Support Attributes –
Setters 393
 - required method 394
 - Required Method —del 394
 - Required Method
—insert 394
 - Required Method
—retrieve 395
 - Required Method
—setConnection 395
 - Required Method
—update 395
 - Required Method –
externalizeKeyAttributes 394
 - Required Method –
internalizeKeyAttributes 394
- Customization for Cardinality
Relations, Data Object
 - Bottom-Up Customizations 398
 - Cardinality-1 Relationships 400
 - Cardinality-N Relationships 402
 - enabling additional clients 417
 - mapping helpers 403
 - packaging for client and server
(VA C++) 406
 - create functions for dynamic
DLL loading 408
 - DLL packaging 406
 - exposing interfaces to
business object
builders 409
- Customization for Cardinality
Relations, Data Object (*continued*)
 - packaging for client and server
(VA C++) 406 (*continued*)
 - exposing interfaces to
clients 409
 - packaging the DLL for the
ActiveX visual C++ client 409
 - packaging the Java client
code 410
 - summarizing relationships
implementations 403
 - Top-Down Customizations 396
 - Top-Down Versus Bottom-Up
Relations 396
- D**
 - data object 108
 - data object customization 361
 - additional considerations 371,
403
 - BOIM data object customization –
cache service 371
 - BOIM data object customization –
static SQL 364
 - BOIM data object
implementation 384
 - BOIM data object interfaces 382
 - data object data management
patterns 389
 - interfaces needed by persistent
data objects 363
 - SQL data object interfaces 364
 - static SQL data object
implementation 366
 - summary of data object
customization 386
 - transient data object-any key
(production use) 381
 - transient data object
customization – UUID key
(production use) 379
 - Data Object Customization and
Inheritance 389
 - Additional Methods – Default
Constructor 393
 - CarPolicy BOIM Data Object
Implementation 391
 - CarPolicy BOIM Data Object
Interfaces 389
 - Framework Required Code –
create() Function 393
 - Framework Required Method –
internalizeFromCopyHelper 391
 - Data Object Customization and
Inheritance 389 (*continued*)
 - Framework Required Method –
internalizeFromPrimaryKey 391
 - Methods To Support Attributes –
Getters 393
 - Methods to Support Attributes –
Setters 393
 - required method 394
 - Required Method —del 394
 - Required Method
—insert 394
 - Required Method
—retrieve 395
 - Required Method
—setConnection 395
 - Required Method
—update 395
 - Required Method –
externalizeKeyAttributes 394
 - Required Method –
internalizeKeyAttributes 394
 - Data Object Customization for
Cardinality Relations 395
 - Bottom-Up Customizations 398
 - Cardinality-1 Relationships 400
 - Cardinality-N Relationships 402
 - enabling additional clients 417
 - mapping helpers 403
 - packaging for client and server
(VA C++) 406
 - create functions for dynamic
DLL loading 408
 - DLL packaging 406
 - exposing interfaces to
business object
builders 409
 - exposing interfaces to
clients 409
 - packaging the DLL for the
ActiveX visual C++ client 409
 - packaging the Java client
code 410
 - summarizing relationships
implementations 403
 - Top-Down Customizations 396
 - Top-Down Versus Bottom-Up
Relations 396
 - Design Tips for Business
Objects 106
 - developing a business object 103
 - Developing an Interface to the
Business Object 103

- development, object-oriented application
 - bottom-up 4, 34
 - combined approaches 10
 - meet in the middle 8, 33
 - top-down 3, 32
- E**
- example, personal life insurance application
 - application model 33
 - design model 34
 - object model 32
- Extending a Business Object
 - choosing an inheritance pattern 160
 - essential state extensions 159
 - Extending a Business Object Interface 158
 - extension summary 168
 - implement the additional business logic 161
 - MOFW IManageable requirements 163
 - MOFW requirements — IManaged server 165
 - more copy helper classes 167
 - more key classes 167
 - other variations to consider 168
- F**
- finding a managed object 87
- Finding a Managed Object Bound in the Naming Service 88
- Finding a Managed Object by Methods on Held Objects 91, 506
- Finding a Managed Object Using the PrimaryKey Helper Class 89
- finding persistent objects 152
- G**
- Guidelines, Conventions and creating persistent objects 151, 152
- finding persistent objects 151, 152
- I**
- IManageable Method
 - Implementations Table 117
- IManageable Methods 117
- IManageable Required Methods 113
- IManagedObject Methods 121, 270, 272
- IManagedObject Required Methods 118
- Implementing Business Object Methods and Attributes 109
- Inheritance, Data Object Customization and
 - Additional Methods – Default Constructor 393
 - CarPolicy BOIM Data Object Implementation 391
 - CarPolicy BOIM Data Object Interfaces 389
 - Framework Required Code – create() Function 393
 - Framework Required Method – internalizeFromCopyHelper 391
 - Framework Required Method – internalizeFromPrimaryKey 391
 - Methods To Support Attributes – Getters 393
 - Methods to Support Attributes – Setters 393
 - required method 394
 - Required Method —del 394
 - Required Method —insert 394
 - Required Method —retrieve 395
 - Required Method —setConnection 395
 - Required Method —update 395
 - Required Method – externalizeKeyAttributes 394
 - Required Method – internalizeKeyAttributes 394
- Initialization requirements 237
- Initializing the Client Environment 86
- Installing Business Objects, Assembling and
 - application adaptor methods 359
 - augmentation of OMG services methods 360
 - business object methods 358
 - Create the Managed Object Class and Implementation 354
 - OMG services methods 359
 - special methods 359
 - specialized homes 360
- insurance application example, personal life
 - application model 33
 - design model 34
 - object model 32
- Interface, Client Programming
 - QOS interfaces for non-transactional support 434, 437
 - Quality of Service Interface 434, 435
- interface definition language 12
- Interface Definition Language (IDL) 447
 - attribute declarations 457
 - emitted file names 469
 - exception declarations 457
 - IDL syntax 459
 - idlc command 464
 - IDLC_OPTIONS environment variable 470
 - interface declarations 453
 - multiple IDL interfaces and modules 464
 - name scoping 447
 - operation declarations 454
 - options for the idlc command 465
 - reserved keywords for IDL 460
 - type and constant declarations 448
- Iterations and Specialized Homes, Queries,
 - using iterated homes — specific function 142
 - using queryable homes — specific functions 142, 145
- iterators 147
- J**
- Java 13
- Java Client Programming Model 231
 - Basic Tasks 235
 - Component Broker Applications 231
 - creating a new object 241
 - Finding a Managed Object 237
 - Initializing the Component Broker Client Environment 235
 - Java exception handling 247
 - Preparing Managed Objects for Remote Access 234
 - releasing and deleting objects 246
 - remembering your favorite object 245
 - using a managed object 241
 - using sets of objects 243
 - when references explode 246

- Java local only development process 353
- Java Server Programming Model
 - business object methods 262
 - 'this' references in business objects 266
 - managing memory 265
 - reference scoping 266
- CosStream:: Streamable::
 - externalize_to_stream 269
- CosStream:: Streamable::
 - internalize_from_stream 269
- developing an interface to the business object 253
- IManegeable::
 - getPrimaryKeyString 268
- IManagedClient:: IManegeable::
 - getHandleString 268
- IManagedServer::
 - IManagedObject
 - initForCreation 271
 - uninitForDestruction 271
- IManagedServer::
 - IManagedObjectWithDataObject
 - initForReactivation() 271
 - syncFromDataObject() 272
 - syncToDataObject() 272
 - uninitForPassivation() 271
- managed object framework
 - methods 268
 - overview 251
 - Primary Key Class 273, 275

K

- Key Class Construction 124
- Keyed Reference Collections 148

L

- languages, programming
 - C++ 13
 - Interface Definition Language 12
 - Java 13
 - naming conventions 14
 - other 13
- life insurance application example, personal
 - application model 33
 - design model 34
 - object model 32
- local only development process 351
- Local Only Development Process 127

M

- Managed and Non-Managed Objects 52
- Managed Object Framework (MOFW) 51
 - Introduction 51
 - overview 51
- meet in the middle application development 8, 33
- model, object
 - Agent 38
 - Beneficiary 39
 - Claim 45
 - Customer 40
 - object identity 37
 - PayoutFraction 41
 - Person 41
 - Policy 42
 - PolicyHolder 44
- model, programming
 - overview 1
 - submodels 1
- Modify Beneficiary 48
- ModifyPolicy 47
- Module Scoping 105
- MOFW (Managed Object Framework)
 - Introduction 51
 - overview 51
- MOFW Client Programming Model
 - Advanced Concepts 133
 - conventions and guidelines 133
 - create_object() method 133
 - expanding the client
 - programming interface 133
 - queries, iterations and specialized homes 133
 - session service 133
 - transactions 133
 - using keyed reference collections 133
- MOFW Objects, understanding 55
- MOFW Server Programming Model
 - Advanced Concepts 157
- Multiple Interfaces to Business Objects 80

N

- Name Scoping and Modules in the C++ Bindings 508
- name space 70
 - NamingStringSyntax 70
 - Navigating 70
 - using the naming service 70
- naming conventions 14

- Naming Service 70
 - NamingStringSyntax 70
 - navigating 70
 - using the naming service 70
- New Object — Create From Key, Creating a
 - Factory 92
 - factory finder 92

O

- Object, Extending a Business
 - choosing an inheritance pattern 160
 - essential state extensions 159
 - Extending a Business Object Interface 158
 - extension summary 168
 - implement the additional business logic 161
 - MOFW IManegeable requirements 163
 - MOFW requirements —
 - IManaged server 165
 - more copy helper classes 167
 - more key classes 167
 - other variations to consider 168
- Object Basics, Business
 - Artifacts Table for Business Objects 55, 129
 - summary 55
- Object Customization for Cardinality Relations, Data
 - Bottom-Up Customizations 398
 - Cardinality-1 Relationships 400
 - Cardinality-N Relationships 402
 - enabling additional clients 417
 - mapping helpers 403
 - packaging for client and server (VA C++) 406
 - create functions for dynamic DLL loading 408
 - DLL packaging 406
 - exposing interfaces to business object builders 409
 - exposing interfaces to clients 409
 - packaging the DLL for the ActiveX visual C++ client 409
 - packaging the Java client code 410
 - summarizing relationships implementations 403
 - Top-Down Customizations 396

- Object Customization for Cardinality Relations, Data *(continued)*
 - Top-Down Versus Bottom-Up Relations 396
 - Object Handle 25
 - object identity 37
 - object model
 - Agent 38
 - Beneficiary 39
 - Claim 45
 - Customer 40
 - object identity 37
 - PayoutFraction 41
 - Person 41
 - Policy 42
 - PolicyHolder 44
 - object-oriented application development
 - bottom-up 4, 34
 - combined approaches 10
 - meet in the middle 8, 33
 - top-down 3, 32
 - object references 100, 237
 - object relationships
 - cardinality-1 relationship 169
 - "uses a" and "has a" cardinality-1 relationship 181
 - implementing the relationship interface 187
 - making cardinality-1 relationships persistent 183
 - optional or required cardinality-1 relationships 172
 - cardinality-N relationship 184
 - Optional Copy Helper Class 125
 - other programming languages 13
 - overview 1
- P**
- Patterns for Handling State (Caching) 111
 - Patterns for Handling State (Delegating) 112
 - PayoutFraction 41
 - Person 41
 - personal life insurance application example
 - application model 33
 - design model 34
 - object model 32
 - Policy 42
 - PolicyHolder 44
 - Preparing Managed Objects for Remote Access 234
 - primary key class 122
 - Process Claim 48
 - Programming, C++ CORBA
 - C++ binding restrictions 523
 - C++ bindings for data types 483
 - C++ bindings for interfaces 508
 - C++ client bindings 520
 - C++ server bindings 521
 - CC++ bindings 481
 - CORBA types and business objects 482
 - exceptions 503
 - name scoping and modules in the C++ bindings 508
 - storage management and `_var` types 512
 - storage management responsibilities for arguments 516
 - Programming Interface, Client
 - QOS interfaces for non-transactional support 434, 437
 - Quality of Service Interface 434, 435
 - programming languages and conventions
 - C++ 13
 - Interface Definition Language 12
 - Java 13
 - naming conventions 14
 - other 13
 - programming model
 - overview 1
 - submodels 1
 - Programming Model: Basic Tasks, ActiveX Client
 - concurrency control 228
 - creating a managed object 221
 - Creating a New Object - Create from Copy 223
 - Creating a New Object — Create From Key 221
 - finding a managed object 217
 - bound in the naming service 218
 - by methods on held objects 221
 - using the PrimaryKey helper class 219
 - initializing the CB client environment 216
 - releasing and deleting objects 228
 - Programming Model: Basic Tasks, ActiveX Client *(continued)*
 - remembering your favorite objects 227
 - using a managed object 221
 - using sets of objects 225
 - When References Explode 228
 - Programming Model, Java Client
 - Basic Tasks 235
 - Component Broker Applications 231
 - creating a new object 241
 - Finding a Managed Object 237
 - Initializing the Component Broker Client Environment 235
 - Java exception handling 247
 - Preparing Managed Objects for Remote Access 234
 - releasing and deleting objects 246
 - remembering your favorite object 245
 - using a managed object 241
 - using sets of objects 243
 - when references explode 246
 - Programming Model, Java Server
 - business object methods 262
 - 'this' references in business objects 266
 - managing memory 265
 - reference scoping 266
 - CosStream:: Streamable::
 - externalize_to_stream 269
 - CosStream:: Streamable::
 - internalize_from_stream 269
 - developing an interface to the business object 253
 - IManageable::
 - getPrimaryKeyString 268
 - IManagedClient:: IManageable::
 - getHandleString 268
 - IManagedServer::
 - IManagedObject
 - initForCreation 271
 - uninitForDestruction 271
 - IManagedServer::
 - IManagedObjectWithDataObject
 - initForReactivation() 271
 - syncFromDataObject() 272
 - syncToDataObject() 272
 - uninitForPassivation() 271
 - managed object framework methods 268
 - overview 251

- Programming Model, Java Server
(*continued*)
 - Primary Key Class 273, 275
- Programming Model — Advanced Concepts, MOFW Client
 - conventions and guidelines 133
 - create_object() method 133
 - expanding the client
 - programming interface 133
 - queries, iterations and specialized homes 133
 - session service 133
 - transactions 133
 - using keyed reference collections 133
- programming roles and responsibilities 11
- pseudo-code, life insurance example
 - Create Customer 47
 - Modify Beneficiary 48
 - ModifyPolicy 47
 - Process Claim 48

Q

- Queries, Iterations and Specialized Homes 142
 - using iterated homes — specific function 142
 - using queryable homes — specific functions 142, 145

R

- Reference Collection 95
- Reference Collection Interfaces 98, 226, 244
- Releasing and Deleting Objects 65
- Remembering your Favorite Objects 98

S

- Sample Framework Flows 442
- Selecting a Pattern for Handling Essential State 107
- Service, Naming
 - NamingStringSyntax 70
 - navigating 70
 - using the naming service 70
- session service 138
- Sets of Objects 95
- space, name
 - NamingStringSyntax 70
 - Navigating 70
 - using the naming service 70

- Specialized Homes, Queries, Iterations and
 - using iterated homes — specific function 142
 - using queryable homes — specific functions 142, 145
- Storage Management and _var Types 512
- summary of interfaces 101, 228, 248

T

- top-down application development 3, 32
- Transactions 81
 - CosTransactions Module 81
 - transactions, exceptions, and timeouts 81
 - using iterated homes — specific function 142
 - using queryable homes — specific functions 142, 145
- Transactions, Exceptions, and Timeouts 136
- Transient Sets 98, 226, 244

U

- using a managed object 91
- Using C++ 'this' References in Business Objects 110



Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

SC09-4442-00



Spine information:



WebSphere

Programming Guide

Version 3.0

SC09-4442-00