



Sun Microsystems

Enterprise JavaBeansTM Specification, Version 2.1

This is the specification of the Enterprise JavaBeansTM architecture. The Enterprise JavaBeans architecture is a component architecture for the development and deployment of component-based distributed business applications. Applications written using the Enterprise JavaBeans architecture are scalable, transactional, and multi-user secure. These applications may be written once, and then deployed on any server platform that supports the Enterprise JavaBeans specification.

Specification Lead:

Linda G. DeMichiel

Please send technical comments on this specification draft to:

ejb21-comments@sun.com

Enterprise JavaBeans™ Specification ("Specification")**Version: 2.1****Status: FCS****Release: November 24, 2003****Copyright 2003 Sun Microsystems, Inc.****4150 Network Circle, Santa Clara, California 95054, U.S.A****All rights reserved.****NOTICE; LIMITED LICENSE GRANTS**

Sun Microsystems, Inc. ("Sun") hereby grants you a fully-paid, non-exclusive, non-transferable, world-wide, limited license (without the right to sublicense), under the Sun's applicable intellectual property rights to view, download, use and reproduce the Specification only for the purpose of internal evaluation, which shall be understood to include developing applications intended to run on an implementation of the Specification provided that such applications do not themselves implement any portion(s) of the Specification.

Sun also grants you a perpetual, non-exclusive, worldwide, fully paid-up, royalty free, limited license (without the right to sublicense) under any applicable copyrights or patent rights it may have in the Specification to create and/or distribute an Independent Implementation of the Specification that: (i) fully implements the Spec(s) including all its required interfaces and functionality; (ii) does not modify, subset, superset or otherwise extend the Licensor Name Space, or include any public or protected packages, classes, Java interfaces, fields or methods within the Licensor Name Space other than those required/authorized by the Specification or Specifications being implemented; and (iii) passes the TCK (including satisfying the requirements of the applicable TCK Users Guide) for such Specification. The foregoing license is expressly conditioned on your not acting outside its scope. No license is granted hereunder for any other purpose.

You need not include limitations (i)-(iii) from the previous paragraph or any other particular "pass through" requirements in any license You grant concerning the use of your Independent Implementation or products derived from it. However, except with respect to implementations of the Specification (and products derived from them) that satisfy limitations (i)-(iii) from the previous paragraph, You may neither: (a) grant or otherwise pass through to your licensees any licenses under Sun's applicable intellectual property rights; nor (b) authorize your licensees to make any claims concerning their implementation's compliance with the Spec in question.

For the purposes of this Agreement: "*Independent Implementation*" shall mean an implementation of the Specification that neither derives from any of Sun's source code or binary code materials nor, except with an appropriate and separate license from Sun, includes any of Sun's source code or binary code materials; and "*Licensor Name Space*" shall mean the public class or interface declarations whose names begin with "java", "javax", "com.sun" or their equivalents in any subsequent naming convention adopted by Sun through the Java Community Process, or any recognized successors or replacements thereof.

This Agreement will terminate immediately without notice from Sun if you fail to comply with any material provision of or act outside the scope of the licenses granted above.

TRADEMARKS

No right, title, or interest in or to any trademarks, service marks, or trade names of Sun or Sun's licensors is granted hereunder. Sun, Sun Microsystems, the Sun logo, Java, the Java Coffee Cup logo, EJB, En-

Enterprise JavaBeans, and JavaBeans are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

DISCLAIMER OF WARRANTIES

THE SPECIFICATION IS PROVIDED "AS IS". SUN MAKES NO REPRESENTATIONS OR WARRANTIES, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT, THAT THE CONTENTS OF THE SPECIFICATION ARE SUITABLE FOR ANY PURPOSE OR THAT ANY PRACTICE OR IMPLEMENTATION OF SUCH CONTENTS WILL NOT INFRINGE ANY THIRD PARTY PATENTS, COPYRIGHTS, TRADE SECRETS OR OTHER RIGHTS. This document does not represent any commitment to release or implement any portion of the Specification in any product.

THE SPECIFICATION COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION THEREIN; THESE CHANGES WILL BE INCORPORATED INTO NEW VERSIONS OF THE SPECIFICATION, IF ANY. SUN MAY MAKE IMPROVEMENTS AND/OR CHANGES TO THE PRODUCT(S) AND/OR THE PROGRAM(S) DESCRIBED IN THE SPECIFICATION AT ANY TIME. Any use of such changes in the Specification will be governed by the then-current license for the applicable version of the Specification.

LIMITATION OF LIABILITY

TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL SUN OR ITS LICENSORS BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION, LOST REVENUE, PROFITS OR DATA, OR FOR SPECIAL, INDIRECT, CONSEQUENTIAL, INCIDENTAL OR PUNITIVE DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF OR RELATED TO ANY FURNISHING, PRACTICING, MODIFYING OR ANY USE OF THE SPECIFICATION, EVEN IF SUN AND/OR ITS LICENSORS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

You will indemnify, hold harmless, and defend Sun and its licensors from any claims arising or resulting from: (i) your use of the Specification; (ii) the use or distribution of your Java application, applet and/or clean room implementation; and/or (iii) any claims that later versions or releases of any Specification furnished to you are incompatible with the Specification provided to you under this license.

RESTRICTED RIGHTS LEGEND

U.S. Government: If this Specification is being acquired by or on behalf of the U.S. Government or by a U.S. Government prime contractor or subcontractor (at any tier), then the Government's rights in the Specification and accompanying documentation shall be only as set forth in this license; this is in accordance with 48 C.F.R. 227.7201 through 227.7202-4 (for Department of Defense (DoD) acquisitions) and with 48 C.F.R. 2.101 and 12.212 (for non-DoD acquisitions).

REPORT

You may wish to report any ambiguities, inconsistencies or inaccuracies you may find in connection with your use of the Specification ("Feedback"). To the extent that you provide Sun with any Feedback, you hereby: (i) agree that such Feedback is provided on a non-proprietary and non-confidential basis, and (ii) grant Sun a perpetual, non-exclusive, worldwide, fully paid-up, irrevocable license, with the right to sublicense through multiple levels of sublicensees, to incorporate, disclose, and use without limitation the Feedback for any purpose related to the Specification and future versions, implementations, and test suites thereof.

(LFI#135809/Form ID#011801)

Table of Contents

Chapter 1	Introduction	27
	1.1 Target Audience	27
	1.2 What is New in EJB 2.1	27
	1.3 Acknowledgments	28
	1.4 Organization	28
	1.5 Document Conventions	29
Chapter 2	Goals	31
	2.1 Overall Goals	31
	2.2 EJB Releases 1.0 and 1.1	32
	2.3 EJB Release 2.0	32
	2.4 Goals for the EJB 2.1 Release	33
Chapter 3	EJB Roles and Scenarios	35
	3.1 EJB Roles	35
	3.1.1 Enterprise Bean Provider	36
	3.1.2 Application Assembler	36
	3.1.3 Deployer	36
	3.1.4 EJB Server Provider	37
	3.1.5 EJB Container Provider	37
	3.1.6 System Administrator	38
	3.2 Scenario: Development, Assembly, and Deployment	38
Chapter 4	Overview	43
	4.1 Enterprise Beans as Components	43
	4.1.1 Component Characteristics	43
	4.1.2 Flexible Component Model	44
	4.2 Enterprise JavaBeans Contracts	45
	4.2.1 Client-view Contracts for Session Beans and Entity Beans	45
	4.2.2 Component Contract	47
	4.2.3 Ejb-jar File	48
	4.2.4 Contracts Summary	49
	4.3 Session, Entity, and Message-Driven Objects	49
	4.3.1 Session Objects	50
	4.3.2 Entity Objects	50
	4.3.3 Message-Driven Objects	50
	4.4 Standard Mapping to CORBA Protocols	51
	4.5 Mapping to Web Service Protocols	51

Chapter 5	Local, Remote, and Web Service Client Views	53
	5.1 Overview.....	53
	5.2 Remote Clients.....	54
	5.3 Local Clients	54
	5.4 Choosing Between a Local or Remote Client View	55
	5.5 Web Service Clients.....	56
Chapter 6	Client View of a Session Bean	59
	6.1 Overview.....	59
	6.2 EJB Container	61
	6.2.1 Locating a Session Bean's Home Interface.....	61
	6.2.2 What a Container Provides.....	62
	6.3 Session Bean's Remote Home Interface.....	62
	6.3.1 Creating a Session Object	63
	6.3.2 Removing a Session Object.....	63
	6.4 Session Bean's Local Home Interface	64
	6.4.1 Creating a Session Object	64
	6.4.2 Removing a Session Object.....	64
	6.5 EJBObject and EJBLocalObject.....	65
	6.6 Session Object Identity	66
	6.7 Client view of Session Object's Life Cycle.....	66
	6.7.1 References to Session Object Remote Interfaces	67
	6.7.2 References to Session Object Local Interfaces	68
	6.8 Creating and Using a Session Object.....	68
	6.9 Object Identity	70
	6.9.1 Stateful Session Beans	70
	6.9.2 Stateless Session Beans	70
	6.9.3 getPrimaryKey()	71
	6.10 Type Narrowing	71
	6.11 The Web Service Client View of a Stateless Session Bean	71
Chapter 7	Session Bean Component Contract.....	75
	7.1 Overview.....	75
	7.2 Goals	76
	7.3 A Container's Management of its Working Set.....	76
	7.4 Conversational State of a Stateful Session Bean	77
	7.4.1 Instance Passivation and Conversational State.....	77
	7.4.2 The Effect of Transaction Rollback on Conversational State	79
	7.5 Protocol Between a Session Bean Instance and its Container.....	79
	7.5.1 The Required SessionBean Interface	79
	7.5.2 The SessionContext Interface.....	80
	7.5.3 Use of the MessageContext Interface by Stateless Session Beans....	81
	7.5.4 The Optional SessionSynchronization Interface for Stateful Session Beans	81

7.5.5	The Optional TimedObject Interface for Stateless Session Beans....	82
7.5.6	Business Method Delegation.....	82
7.5.7	Session Bean's ejbCreate<METHOD> Methods.....	82
7.5.8	Serializing Session Bean Methods	83
7.5.9	Transaction Context of Session Bean Methods.....	83
7.6	Stateful Session Bean State Diagram	84
7.6.1	Operations Allowed in the Methods of a Stateful Session Bean Class	86
7.6.2	Dealing with Exceptions	88
7.6.3	Missed ejbRemove Calls.....	89
7.6.4	Restrictions for Transactions.....	89
7.7	Object Interaction Diagrams for a Stateful Session Bean	90
7.7.1	Notes	90
7.7.2	Creating a Session Object	90
7.7.3	Starting a Transaction.....	91
7.7.4	Committing a Transaction.....	92
7.7.5	Passivating and Activating an Instance Between Transactions	93
7.7.6	Removing a Session Object	94
7.8	Stateless Session Beans	95
7.8.1	Stateless Session Bean State Diagram	96
7.8.2	Operations Allowed in the Methods of a Stateless Session Bean Class	98
7.8.3	Dealing with Exceptions	102
7.9	Object Interaction Diagrams for a Stateless Session Bean with a Home and Component Interface	102
7.9.1	Client-invoked create	103
7.9.2	Business Method Invocation	103
7.9.3	Client-invoked remove	104
7.9.4	Adding an Instance to the Pool	105
7.10	Object Interaction Diagrams for a Stateless Session Bean with a Web Service Endpoint Interface	106
7.10.1	Business Method Invocation	106
7.10.2	Adding an Instance to the Pool	107
7.11	The Responsibilities of the Bean Provider	109
7.11.1	Classes and Interfaces	109
7.11.2	Session Bean Class.....	109
7.11.3	ejbCreate<METHOD> Methods.....	110
7.11.4	Business Methods	111
7.11.5	Session Bean's Remote Interface	111
7.11.6	Session Bean's Remote Home Interface	112
7.11.7	Session Bean's Local Interface	112
7.11.8	Session Bean's Local Home Interface.....	113
7.11.9	Session Bean's Web Service Endpoint Interface.....	113
7.12	The Responsibilities of the Container Provider.....	114
7.12.1	Generation of Implementation Classes	115
7.12.2	Generation of WSDL	115
7.12.3	Session EJBHome Class	115
7.12.4	Session EJBObject Class	116
7.12.5	Session EJBLocalHome Class	116
7.12.6	Session EJBLocalObject Class	116

	7.12.7	Web Service Endpoint Implementation Class	116
	7.12.8	Handle Classes	116
	7.12.9	EJBMetaData Class	116
	7.12.10	Non-reentrant Instances	117
	7.12.11	Transaction Scoping, Security, Exceptions	117
	7.12.12	JAX-RPC Message Handlers for Web Service Endpoints	117
	7.12.13	SessionContext	117
Chapter 8		Example Session Scenario	119
	8.1	Overview	119
	8.2	Inheritance Relationship	119
	8.2.1	What the Session Bean Provider is Responsible For	121
	8.2.2	Classes Supplied by Container Provider	121
	8.2.3	What the Container Provider is Responsible For	121
Chapter 9		Client View of an Entity	123
	9.1	Overview	123
	9.2	Remote Clients	124
	9.3	Local Clients	124
	9.4	EJB Container	125
	9.4.1	Locating an Entity Bean's Home Interface	126
	9.4.2	What a Container Provides	126
	9.5	Entity Bean's Remote Home Interface	127
	9.5.1	Create Methods	128
	9.5.2	Finder Methods	129
	9.5.3	Remove Methods	130
	9.5.4	Home Methods	130
	9.6	Entity Bean's Local Home Interface	131
	9.6.1	Create Methods	131
	9.6.2	Finder Methods	132
	9.6.3	Remove Methods	133
	9.6.4	Home Methods	133
	9.7	Entity Object's Life Cycle	133
	9.7.1	References to Entity Object Remote Interfaces	135
	9.7.2	References to Entity Object Local Interfaces	135
	9.8	Primary Key and Object Identity	136
	9.9	Entity Bean's Remote Interface	137
	9.10	Entity Bean's Local Interface	138
	9.11	Entity Bean's Handle	139
	9.12	Entity Home Handles	140
	9.13	Type Narrowing of Object References	140
Chapter 10		Entity Bean Component Contract for Container-Managed Persistence	141
	10.1	Overview	142

10.2	Container-Managed Entity Persistence and Data Independence	142
10.3	The Entity Bean Provider's View of Container-Managed Persistence	144
10.3.1	The Entity Bean Provider's Programming Contract	145
10.3.2	The Entity Bean Provider's View of Persistent Relationships	147
10.3.3	Dependent Value Classes	147
10.3.4	Remove Protocols	148
10.3.4.1	Remove Methods	148
10.3.4.2	Cascade-delete	149
10.3.5	Identity of Entity Objects	149
10.3.6	Semantics of Assignment for Relationships	150
10.3.6.1	Use of the java.util.Collection API to Update Relationships	150
10.3.6.2	Use of Set Accessor Methods to Update Relationships	152
10.3.7	Assignment Rules for Relationships	153
10.3.7.1	One-to-one Bidirectional Relationships	154
10.3.7.2	One-to-one Unidirectional Relationships	155
10.3.7.3	One-to-many Bidirectional Relationships	156
10.3.7.4	One-to-many Unidirectional Relationships	160
10.3.7.5	Many-to-one Unidirectional Relationships	163
10.3.7.6	Many-to-many Bidirectional Relationships	165
10.3.7.7	Many-to-many Unidirectional Relationships	169
10.3.8	Collections Managed by the Container	172
10.3.9	Non-persistent State	172
10.3.10	The Relationship Between the Internal View and the Client View ...	173
10.3.10.1	Restrictions on Remote Interfaces	173
10.3.11	Mapping Data to a Persistent Store	173
10.3.12	Example	174
10.3.13	The Bean Provider's View of the Deployment Descriptor	177
10.4	The Entity Bean Component Contract	181
10.4.1	Runtime Execution Model of Entity Beans	181
10.4.2	Container Responsibilities	183
10.4.2.1	Container-Managed Fields	183
10.4.2.2	Container-Managed Relationships	183
10.5	Instance Life Cycle Contract Between the Bean and the Container	184
10.5.1	Instance Life Cycle	185
10.5.2	Bean Provider's Entity Bean Instance's View	187
10.5.3	Container's View	191
10.5.4	The EntityContext Interface	195
10.5.5	Operations Allowed in the Methods of the Entity Bean Class	196
10.5.6	Finder Methods	199
10.5.6.1	Single-Object Finder Methods	199
10.5.6.2	Multi-Object Finder Methods	200
10.5.7	Select Methods	201
10.5.7.1	Single-Object Select Methods	201
10.5.7.2	Multi-Object Select Methods	202
10.5.8	Timer Notifications	202
10.5.9	Standard Application Exceptions for Entities	202
10.5.9.1	CreateException	203
10.5.9.2	DuplicateKeyException	203
10.5.9.3	FinderException	204

10.5.9.4	ObjectNotFoundException	204
10.5.9.5	RemoveException	204
10.5.10	Commit Options	204
10.5.11	Concurrent Access from Multiple Transactions	206
10.5.12	Non-reentrant and Re-entrant Instances	208
10.6	Responsibilities of the Enterprise Bean Provider	209
10.6.1	Classes and Interfaces	209
10.6.2	Enterprise Bean Class.....	209
10.6.3	Dependent Value Classes	210
10.6.4	ejbCreate<METHOD> Methods	210
10.6.5	ejbPostCreate<METHOD> Methods	211
10.6.6	ejbHome<METHOD> Methods.....	212
10.6.7	ejbSelect<METHOD> Methods.....	212
10.6.8	Business Methods.....	212
10.6.9	Entity Bean's Remote Interface.....	213
10.6.10	Entity Bean's Remote Home Interface.....	213
10.6.11	Entity Bean's Local Interface.....	214
10.6.12	Entity Bean's Local Home Interface	215
10.6.13	Entity Bean's Primary Key Class	216
10.6.14	Entity Bean's Deployment Descriptor.....	216
10.7	The Responsibilities of the Container Provider.....	216
10.7.1	Generation of Implementation Classes	217
10.7.2	Enterprise Bean Class.....	217
10.7.3	ejbFind<METHOD> Methods	218
10.7.4	ejbSelect<METHOD> Methods.....	218
10.7.5	Entity EJBHome Class	219
10.7.6	Entity EJBObject Class	219
10.7.7	Entity EJBLocalHome Class	219
10.7.8	Entity EJBLocalObject Class	220
10.7.9	Handle Class.....	220
10.7.10	Home Handle Class	220
10.7.11	Metadata Class	221
10.7.12	Instance's Re-entrance.....	221
10.7.13	Transaction Scoping, Security, Exceptions	221
10.7.14	Implementation of Object References.....	221
10.7.15	EntityContext.....	221
10.8	Primary Keys	222
10.8.1	Primary Key That Maps to a Single Field in the Entity Bean Class	222
10.8.2	Primary Key That Maps to Multiple Fields in the Entity Bean Class.....	222
10.8.3	Special Case: Unknown Primary Key Class.....	222
10.9	Object Interaction Diagrams.....	223
10.9.1	Notes.....	223
10.9.2	Creating an Entity Object	224
10.9.3	Passivating and Activating an Instance in a Transaction.....	226
10.9.4	Committing a Transaction	228
10.9.5	Starting the Next Transaction.....	229
10.9.6	Removing an Entity Object	231
10.9.7	Finding an Entity Object	232
10.9.8	Adding and Removing an Instance from the Pool	232

Chapter 11	EJB QL: EJB Query Language for Container-Managed Persistence Query Methods ..	235
11.1	Overview	235
11.2	EJB QL Definition.....	236
11.2.1	Abstract Schema Types and Query Domains.....	237
11.2.2	Query Methods.....	238
11.2.3	Naming.....	238
11.2.4	Examples	238
11.2.5	The FROM Clause and Navigational Declarations	240
11.2.5.1	Identifiers	240
11.2.5.2	Identification Variables	241
11.2.5.3	Range Variable Declarations.....	242
11.2.5.4	Collection Member Declarations	242
11.2.5.5	Example	243
11.2.5.6	Path Expressions	243
11.2.6	WHERE Clause and Conditional Expressions.....	244
11.2.6.1	Literals	244
11.2.6.2	Identification Variables	245
11.2.6.3	Path Expressions	245
11.2.6.4	Input Parameters	245
11.2.6.5	Conditional Expression Composition	246
11.2.6.6	Operators and Operator Precedence	246
11.2.6.7	Between <i>Expressions</i>	246
11.2.6.8	In Expressions.....	247
11.2.6.9	Like <i>Expressions</i>	247
11.2.6.10	Null Comparison Expressions	248
11.2.6.11	Empty Collection Comparison Expressions	248
11.2.6.12	Collection Member Expressions.....	249
11.2.6.13	Functional Expressions.....	249
11.2.7	SELECT Clause	250
11.2.7.1	Null Values in the Query Result	251
11.2.7.2	Aggregate Functions in the SELECT Clause	251
11.2.7.3	Examples.....	252
11.2.8	ORDER BY Clause.....	253
11.2.9	Return Value Types	254
11.2.10	Null Values	255
11.2.11	Equality and Comparison Semantics	257
11.2.12	Restrictions.....	257
11.3	Examples	257
11.3.1	Simple Queries.....	258
11.3.2	Queries with Relationships	258
11.3.3	Queries Using Input Parameters.....	259
11.3.4	Queries for Select Methods.....	259
11.3.5	EJB QL and SQL	260
11.4	EJB QL BNF	261
Chapter 12	Entity Bean Component Contract for Bean-Managed Persistence	265
12.1	Overview of Bean-Managed Entity Persistence	265
12.1.1	Entity Bean Provider's View of Persistence.....	266

12.1.2	Runtime Execution Model.....	267
12.1.3	Instance Life Cycle.....	269
12.1.4	The Entity Bean Component Contract	271
12.1.4.1	Entity Bean Instance's View	271
12.1.4.2	Container's View.....	275
12.1.5	The EntityContext Interface	278
12.1.6	Operations Allowed in the Methods of the Entity Bean Class.....	278
12.1.7	Caching of Entity State and the ejbLoad and ejbStore Methods.....	281
12.1.7.1	ejbLoad and ejbStore with the NotSupported Transaction Attribute	282
12.1.8	Finder Method Return Type	282
12.1.8.1	Single-Object Finder.....	282
12.1.8.2	Multi-Object Finders.....	283
12.1.9	Timer Notifications.....	285
12.1.10	Standard Application Exceptions for Entities	285
12.1.10.1	CreateException.....	285
12.1.10.2	DuplicateKeyException	285
12.1.10.3	FinderException.....	286
12.1.10.4	ObjectNotFoundException	286
12.1.10.5	RemoveException	286
12.1.11	Commit Options	287
12.1.12	Concurrent Access from Multiple Transactions.....	288
12.1.13	Non-reentrant and Re-entrant Instances	289
12.2	Responsibilities of the Enterprise Bean Provider	290
12.2.1	Classes and Interfaces	290
12.2.2	Enterprise Bean Class.....	290
12.2.3	ejbCreate<METHOD> Methods	291
12.2.4	ejbPostCreate<METHOD> Methods	292
12.2.5	ejbFind Methods.....	292
12.2.6	ejbHome<METHOD> Methods.....	293
12.2.7	Business Methods.....	293
12.2.8	Entity Bean's Remote Interface.....	294
12.2.9	Entity Bean's Remote Home Interface.....	295
12.2.10	Entity Bean's Local Interface	296
12.2.11	Entity Bean's Local Home Interface	296
12.2.12	Entity Bean's Primary Key Class	297
12.3	The Responsibilities of the Container Provider.....	298
12.3.1	Generation of Implementation Classes	298
12.3.2	Entity EJBHome Class	299
12.3.3	Entity EJBObject Class	299
12.3.4	Entity EJBLocalHome Class	299
12.3.5	Entity EJBLocalObject Class	300
12.3.6	Handle Class.....	300
12.3.7	Home Handle Class.....	300
12.3.8	Metadata Class	301
12.3.9	Instance's Re-entrance.....	301
12.3.10	Transaction Scoping, Security, Exceptions	301
12.3.11	Implementation of Object References	301
12.3.12	EntityContext.....	301

	12.4	Object Interaction Diagrams.....	301
	12.4.1	Notes	302
	12.4.2	Creating an Entity Object.....	303
	12.4.3	Passivating and Activating an Instance in a Transaction.....	304
	12.4.4	Committing a Transaction.....	305
	12.4.5	Starting the Next Transaction.....	305
	12.4.6	Removing an Entity Object.....	307
	12.4.7	Finding an Entity Object.....	308
	12.4.8	Adding and Removing an Instance from the Pool	308
Chapter 13		Example Bean-Managed Persistence Entity Scenario	311
	13.1	Overview	311
	13.2	Inheritance Relationship.....	312
	13.2.1	What the Entity Bean Provider is Responsible For.....	313
	13.2.2	Classes Supplied by Container Provider.....	313
	13.2.3	What the Container Provider is Responsible For.....	313
Chapter 14		EJB 1.1 Entity Bean Component Contract for Container-Managed Persistence.....	315
	14.1	EJB 1.1 Entity Beans with Container-Managed Persistence.....	315
	14.1.1	Container-Managed Fields	316
	14.1.2	ejbCreate, ejbPostCreate	317
	14.1.3	ejbRemove.....	318
	14.1.4	ejbLoad.....	319
	14.1.5	ejbStore	319
	14.1.6	Finder Hethods	319
	14.1.7	Home Methods	319
	14.1.8	Create Methods	320
	14.1.9	Primary Key Type	320
	14.1.9.1	Primary Key that Maps to a Single Field in the Entity Bean Class 320	
	14.1.9.2	Primary Key that Maps to Multiple Fields in the Entity Bean Class 320	
	14.1.9.3	Special Case: Unknown Primary Key Class.....	321
	14.2	Object Interaction Diagrams.....	321
	14.2.1	Notes	321
	14.2.2	Creating an Entity Object.....	322
	14.2.3	Passivating and Activating an Instance in a Transaction.....	324
	14.2.4	Committing a Transaction.....	326
	14.2.5	Starting the Next Transaction.....	327
	14.2.6	Removing an Entity Object.....	329
	14.2.7	Finding an Entity Object.....	330
	14.2.8	Adding and Removing an Instance from the Pool	330
Chapter 15		Message-Driven Bean Component Contract.....	333
	15.1	Overview	333
	15.2	Goals.....	334

15.3	Client View of a Message-Driven Bean.....	334
15.4	Protocol Between a Message-Driven Bean Instance and its Container.....	336
15.4.1	The Required MessageDrivenBean Interface.....	336
15.4.2	The Required Message Listener Interface.....	336
15.4.3	The Optional TimedObject Interface.....	337
15.4.4	The MessageDrivenContext Interface.....	337
15.4.5	Message-Driven Bean's ejbCreate Method.....	338
15.4.6	Serializing Message-Driven Bean Methods.....	338
15.4.7	Concurrency of Message Processing.....	338
15.4.8	Transaction Context of Message-Driven Bean Methods.....	338
15.4.9	Activation Configuration Properties.....	339
15.4.10	Message Acknowledgment for JMS Message-Driven Beans.....	339
15.4.11	Message Selectors for JMS Message-Driven Beans.....	339
15.4.12	Association of a Message-Driven Bean with a Destination or Endpoint.....	340
15.4.12.1	JMS Message-Driven Beans.....	340
15.4.13	Dealing with Exceptions.....	340
15.4.14	Missed ejbRemove Calls.....	341
15.4.15	Replying to a JMS Message.....	341
15.5	Message-Driven Bean State Diagram.....	341
15.5.1	Operations Allowed in the Methods of a Message-Driven Bean Class.....	343
15.6	Object Interaction Diagrams for a Message-Driven Bean.....	345
15.6.1	Message Receipt: Message Listener Method Invocation.....	345
15.6.2	Adding an Instance to the Pool.....	345
15.6.3	Removing an Instance from the Pool.....	346
15.7	The Responsibilities of the Bean Provider.....	347
15.7.1	Classes and Interfaces.....	347
15.7.2	Message-Driven Bean Class.....	347
15.7.3	ejbCreate Method.....	348
15.7.4	Message Listener Method.....	348
15.7.5	ejbRemove Method.....	349
15.8	The Responsibilities of the Container Provider.....	349
15.8.1	Generation of Implementation Classes.....	349
15.8.2	Deployment of JMS Message-Driven Beans.....	349
15.8.3	Request/Response Messaging Types.....	350
15.8.4	Non-reentrant Instances.....	350
15.8.5	Transaction Scoping, Security, Exceptions.....	350
Chapter 16	Example Message-Driven Bean Scenario.....	351
16.1	Overview.....	351
16.2	Inheritance Relationship.....	352
16.2.1	What the Message-Driven Bean Provider is Responsible For.....	354
16.2.2	Classes Supplied by Container Provider.....	354
16.2.3	What the Container Provider is Responsible For.....	354
Chapter 17	Support for Transactions.....	355
17.1	Overview.....	355

17.1.1	Transactions	355
17.1.2	Transaction Model.....	356
17.1.3	Relationship to JTA and JTS	357
17.2	Sample Scenarios.....	357
17.2.1	Update of Multiple Databases	357
17.2.2	Messages Sent or Received Over JMS Sessions and Update of Multiple Databases.....	358
17.2.3	Update of Databases via Multiple EJB Servers	360
17.2.4	Client-Managed Demarcation	361
17.2.5	Container-Managed Demarcation	362
17.3	Bean Provider's Responsibilities	363
17.3.1	Bean-Managed Versus Container-Managed Transaction Demarcation.....	363
17.3.1.1	Non-Transactional Execution	363
17.3.2	Isolation Levels	364
17.3.3	Enterprise Beans Using Bean-Managed Transaction Demarcation ..	364
17.3.3.1	getRollbackOnly and setRollbackOnly Methods	371
17.3.4	Enterprise Beans Using Container-Managed Transaction Demarcation.....	372
17.3.4.1	javax.ejb.SessionSynchronization Interface	373
17.3.4.2	javax.ejb.EJBContext.setRollbackOnly Method	373
17.3.4.3	javax.ejb.EJBContext.getRollbackOnly method	374
17.3.5	Use of JMS APIs in Transactions	374
17.3.6	Declaration in Deployment Descriptor	374
17.4	Application Assembler's Responsibilities.....	375
17.4.1	Transaction Attributes	375
17.5	Deployer's Responsibilities	379
17.6	Container Provider Responsibilities	379
17.6.1	Bean-Managed Transaction Demarcation.....	380
17.6.2	Container-Managed Transaction Demarcation for Session and Entity Beans	383
17.6.2.1	NotSupported.....	383
17.6.2.2	Required.....	383
17.6.2.3	Supports	383
17.6.2.4	RequiresNew.....	384
17.6.2.5	Mandatory.....	384
17.6.2.6	Never.....	384
17.6.2.7	Transaction Attribute Summary.....	385
17.6.2.8	Handling of setRollbackOnly Method.....	386
17.6.2.9	Handling of getRollbackOnly Method	386
17.6.2.10	Handling of getUserTransaction Method	386
17.6.2.11	javax.ejb.SessionSynchronization Callbacks.....	386
17.6.3	Container-Managed Transaction Demarcation for Message-Driven Beans	387
17.6.3.1	NotSupported.....	387
17.6.3.2	Required.....	387
17.6.3.3	Handling of setRollbackOnly Method.....	388
17.6.3.4	Handling of getRollbackOnly Method	388
17.6.3.5	Handling of getUserTransaction Method	388
17.6.4	Local Transaction Optimization.....	388
17.6.5	Handling of Methods that Run with "an unspecified transaction context".....	389

17.7	Access from Multiple Clients in the Same Transaction Context.....	390
17.7.1	Transaction “Diamond” Scenario with an Entity Object	390
17.7.2	Container Provider’s Responsibilities	391
17.7.3	Bean Provider’s Responsibilities.....	393
17.7.4	Application Assembler and Deployer’s Responsibilities	393
17.7.5	Transaction Diamonds involving Session Objects	393
Chapter 18	Exception Handling.....	395
18.1	Overview and Concepts	395
18.1.1	Application Exceptions	395
18.1.2	Goals for Exception Handling	396
18.2	Bean Provider’s Responsibilities	396
18.2.1	Application Exceptions	396
18.2.2	System Exceptions	397
18.2.2.1	javax.ejb.NoSuchEntityException	398
18.3	Container Provider Responsibilities	399
18.3.1	Exceptions from a Session or Entity Bean’s Business Methods	399
18.3.2	Exceptions from ejbCreate and ejbRemove Methods of a Stateless Session Bean with Web Service Client View.....	402
18.3.3	Exceptions from Message-Driven Bean Message Listener Methods	403
18.3.4	Exceptions from ejbCreate and ejbRemove Methods of a Message-Driven Bean.....	404
18.3.5	Exceptions from an Enterprise Bean’s ejbTimeout Method	404
18.3.6	Exceptions from Container-invoked Callbacks	405
18.3.7	javax.ejb.NoSuchEntityException.....	406
18.3.8	Non-existing Stateful Session or Entity Object.....	406
18.3.9	Exceptions from the Management of Container-Managed Transactions.....	407
18.3.10	Release of Resources.....	407
18.3.11	Support for Deprecated Use of java.rmi.RemoteException	407
18.4	Client’s View of Exceptions	408
18.4.1	Application Exception.....	408
18.4.1.1	Local and Remote Clients.....	408
18.4.1.2	Web Service Clients	409
18.4.2	java.rmi.RemoteException and javax.ejb.EJBException	409
18.4.2.1	javax.transaction.TransactionRolledbackException and javax.ejb.TransactionRolledbackLocalException.....	410
18.4.2.2	javax.transaction.TransactionRequiredException and javax.ejb.TransactionRequiredLocalException.....	410
18.4.2.3	java.rmi.NoSuchObjectException and javax.ejb.NoSuchObjectLocalException.....	411
18.5	System Administrator’s Responsibilities	411
18.6	Differences from EJB 1.0	411
Chapter 19	Support for Distribution and Interoperability	413
19.1	Support for Distribution.....	413
19.1.1	Client-Side Objects in a Distributed Environment	414
19.2	Interoperability Overview	414

19.2.1	Interoperability Goals	415
19.3	Interoperability Scenarios.....	416
19.3.1	Interactions Between Web Containers and EJB Containers for E-Commerce Applications.....	416
19.3.2	Interactions Between Application Client Containers and EJB Containers Within an Enterprise's Intranet.....	417
19.3.3	Interactions Between Two EJB Containers in an Enterprise's Intranet.....	418
19.3.4	Intranet Application Interactions Between Web Containers and EJB Containers.....	419
19.4	Overview of Interoperability Requirements	419
19.5	Remote Invocation Interoperability	420
19.5.1	Mapping Java Remote Interfaces to IDL	421
19.5.2	Mapping Value Objects to IDL	421
19.5.3	Mapping of System Exceptions	421
19.5.4	Obtaining Stub and Client View Classes	422
19.5.5	System Value Classes.....	422
19.5.5.1	HandleDelegate SPI.....	423
19.6	Transaction Interoperability	424
19.6.1	Transaction Interoperability Requirements	424
19.6.1.1	Transaction Context Wire Format.....	424
19.6.1.2	Two-Phase Commit Protocol	424
19.6.1.3	Transactional Policies of Enterprise Bean References	426
19.6.1.4	Exception Handling Behavior.....	426
19.6.2	Interoperating with Containers that do not Implement Transaction Interoperability.....	426
19.6.2.1	Client Container Requirements.....	427
19.6.2.2	EJB container requirements.....	427
19.7	Naming Interoperability	429
19.8	Security Interoperability	430
19.8.1	Introduction	430
19.8.1.1	Trust Relationships Between Containers, Principal Propagation.....	431
19.8.1.2	Application Client Authentication	432
19.8.2	Securing EJB Invocations	432
19.8.2.1	Secure Transport Protocol	433
19.8.2.2	Security Information in IORs	434
19.8.2.3	Propagating Principals and Authentication Data in IIOP Messages	434
19.8.2.4	Security Configuration for Containers.....	436
19.8.2.5	Runtime Behavior	436
Chapter 20	Enterprise Bean Environment	439
20.1	Overview	439
20.2	Enterprise Bean's Environment as a JNDI Naming Context.....	440
20.2.1	Bean Provider's Responsibilities.....	441
20.2.1.1	Access to Enterprise Bean's Environment.....	441
20.2.1.2	Declaration of Environment Entries	442
20.2.2	Application Assembler's Responsibility	444
20.2.3	Deployer's Responsibility	445

20.2.4	Container Provider Responsibility	445
20.3	EJB References	445
20.3.1	Bean Provider's Responsibilities	445
20.3.1.1	EJB Reference Programming Interfaces	446
20.3.1.2	Declaration of EJB References in Deployment Descriptor	446
20.3.2	Application Assembler's Responsibilities	448
20.3.3	Deployer's Responsibility	450
20.3.4	Container Provider's Responsibility	450
20.4	Web Service References	451
20.4.1	Bean Provider's Responsibilities	451
20.4.1.1	Web Service Reference Programming Interfaces	451
20.4.1.2	Declaration of Web Service References in Deployment Descriptor	452
20.4.2	Deployer's Responsibility	452
20.4.3	Container Provider's Responsibility	453
20.5	Resource Manager Connection Factory References	453
20.5.1	Bean Provider's Responsibilities	453
20.5.1.1	Programming Interfaces for Resource Manager Connection Factory References	453
20.5.1.2	Declaration of Resource Manager Connection Factory References in Deployment Descriptor	455
20.5.1.3	Standard Resource Manager Connection Factory Types	457
20.5.2	Deployer's Responsibility	457
20.5.3	Container Provider Responsibility	458
20.5.4	System Administrator's Responsibility	459
20.6	Resource Environment References	459
20.6.1	Bean Provider's Responsibilities	459
20.6.1.1	Resource Environment Reference Programming Interfaces	460
20.6.1.2	Declaration of Resource Environment References in Deployment Descriptor	460
20.6.2	Deployer's Responsibility	460
20.6.3	Container Provider's Responsibility	461
20.7	Message Destination References	461
20.7.1	Bean Provider's Responsibilities	461
20.7.1.1	Message Destination Reference Programming Interfaces	461
20.7.1.2	Declaration of Message Destination References in Deployment Descriptor	462
20.7.2	Application Assembler's Responsibilities	463
20.7.3	Deployer's Responsibility	466
20.7.4	Container Provider's Responsibility	466
20.8	Deprecated EJBContext.getEnvironment Method	467
20.9	UserTransaction Interface	468

Chapter 21	Security Management	471
21.1	Overview	471
21.2	Bean Provider's Responsibilities	473
21.2.1	Invocation of Other Enterprise Beans	473
21.2.2	Resource Access	473

21.2.3	Access of Underlying OS Resources	473
21.2.4	Programming Style Recommendations	473
21.2.5	Programmatic Access to Caller's Security Context	474
21.2.5.1	Use of <code>getCallerPrincipal</code>	475
21.2.5.2	Use of <code>isCallerInRole</code>	476
21.2.5.3	Declaration of Security Roles Referenced from the Bean's Code 477	
21.3	Application Assembler's Responsibilities	478
21.3.1	Security Roles	479
21.3.2	Method Permissions	480
21.3.3	Linking Security Role References to Security Roles	484
21.3.4	Specification of Security Identities in the Deployment Descriptor... 484	
21.3.4.1	Run-as	485
21.4	Deployer's Responsibilities	486
21.4.1	Security Domain and Principal Realm Assignment.....	486
21.4.2	Assignment of Security Roles	486
21.4.3	Principal Delegation.....	487
21.4.4	Security Management of Resource Access	487
21.4.5	General Notes on Deployment Descriptor Processing.....	487
21.5	EJB Client Responsibilities	488
21.6	EJB Container Provider's Responsibilities.....	488
21.6.1	Deployment Tools	488
21.6.2	Security Domain(s)	488
21.6.3	Security Mechanisms	489
21.6.4	Passing Principals on EJB Calls.....	489
21.6.5	Security Methods in <code>javax.ejb.EJBContext</code>	489
21.6.6	Secure Access to Resource Managers.....	490
21.6.7	Principal Mapping.....	490
21.6.8	System Principal	490
21.6.9	Runtime Security Enforcement	490
21.6.10	Audit Trail	491
21.7	System Administrator's Responsibilities.....	491
21.7.1	Security Domain Administration	491
21.7.2	Principal Mapping.....	492
21.7.3	Audit Trail Review	492
Chapter 22	Timer Service.....	493
22.1	Overview	493
22.2	Bean Provider's View of the Timer Service	494
22.2.1	The Timer Service Interface.....	495
22.2.2	The <code>TimedObject</code> Interface	496
22.2.3	The Timer and <code>TimerHandle</code> Interfaces	496
22.2.4	Timer Identity.....	497
22.2.5	Transactions	497
22.3	Bean Provider's Responsibilities	497
22.3.1	Enterprise Bean Class	498
22.3.2	<code>TimerHandle</code>	498
22.4	Container's Responsibilities	498

	22.4.1	TimerService, Timer, and TimerHandle Interfaces	498
	22.4.2	Timer Expiration and ejbTimeout Method	498
	22.4.3	Timer Cancellation	499
	22.4.4	Entity Bean Removal	499
Chapter 23		Deployment Descriptor	501
	23.1	Overview	501
	23.2	Bean Provider's Responsibilities	502
	23.3	Application Assembler's Responsibility	504
	23.4	Container Provider's Responsibilities	506
	23.5	Deployment Descriptor XML Schema	507
Chapter 24		Ejb-jar File	555
	24.1	Overview	555
	24.2	Deployment Descriptor	556
	24.3	Ejb-jar File Requirements	556
	24.4	The Client View and the ejb-client JAR File	556
	24.5	Requirements for Clients	557
	24.6	Example	558
	24.7	Removed in EJB 1.1	558
	24.7.1	ejb-jar Manifest	558
	24.7.2	Serialized Deployment Descriptor JavaBeans™ Components	559
Chapter 25		Runtime Environment	561
	25.1	Bean Provider's Responsibilities	561
	25.1.1	APIs Provided by Container	561
	25.1.2	Programming Restrictions	562
	25.2	Container Provider's Responsibility	565
	25.2.1	Java 2 APIs Requirements	565
	25.2.2	EJB 2.1 Requirements	567
	25.2.3	JNDI Requirements	567
	25.2.4	JTA 1.0.1 Requirements	568
	25.2.5	JDBC™ 3.0 Extension Requirements	568
	25.2.6	JMS 1.1 Requirements	568
	25.2.7	Argument Passing Semantics	569
Chapter 26		Responsibilities of EJB Roles	571
	26.1	Bean Provider's Responsibilities	571
	26.1.1	API Requirements	571
	26.1.2	Packaging Requirements	571
	26.2	Application Assembler's Responsibilities	572
	26.3	EJB Container Provider's Responsibilities	572
	26.4	Deployer's Responsibilities	572

	26.5	System Administrator's Responsibilities	572
	26.6	Client Programmer's Responsibilities	572
Chapter 27		Enterprise JavaBeans™ API Reference	573
		package javax.ejb.....	574
		package javax.ejb.spi	574
		package javax.ejb.deployment	575
Chapter 28		Related Documents	577
Appendix A		Features Deferred to Future Releases	579
Appendix B		EJB 2.0 Deployment Descriptor	581
	B.1	Overview	581
	B.2	Bean Provider's Responsibilities	582
	B.3	Application Assembler's Responsibility	584
	B.4	Container Provider's Responsibilities	586
	B.5	Deployment Descriptor DTD	586
Appendix C		EJB 1.1 Deployment Descriptor	613
	C.1	Overview	613
	C.2	Bean Provider's Responsibilities	614
	C.3	Application Assembler's Responsibility	616
	C.4	Container Provider's Responsibilities	617
	C.5	Deployment Descriptor DTD	617
	C.6	Deployment Descriptor Example	632
Appendix D		Frequently Asked Questions	639
	D.1	Client-Demarcated Transactions	639
	D.2	Container-Managed Persistence	640
	D.3	Inheritance	640
	D.4	How to Obtain Database Connections	640
	D.5	Session Beans and Primary Key	641
	D.6	Copying of Parameters Required for Remote EJB Calls Within the Same JVM.....	641
Appendix E		Revision History	643
	E.1	Version 0.1	643
	E.2	Version 0.2.....	643
	E.3	Version 0.3.....	644
	E.4	Version 0.4.....	644

E.5	Version 0.5	644
E.6	Version 0.6: Community Review Draft	645
E.7	Version 0.7	645
E.8	Version 0.8: Public Draft	645
E.9	Version 0.9: Proposed Final Draft	645
E.10	Proposed Final Draft 2.....	646
E.11	Final Release.....	646

List of Figures

Figure 1	Enterprise JavaBeans Contracts	49
Figure 2	Client View of Session Beans Deployed in a Container	62
Figure 3	Life Cycle of a Session Object	67
Figure 4	Session Bean Example Objects	69
Figure 5	Web Service Client View of Stateless Session Beans Deployed in a Container	72
Figure 6	Life Cycle of a Stateful Session Bean Instance	84
Figure 7	OID for Creation of a Session Object	91
Figure 8	OID for Session Object at Start of a Transaction	92
Figure 9	OID for Session Object Transaction Synchronization	93
Figure 10	OID for Passivation and Activation of a Session Object	94
Figure 11	OID for the Removal of a Session Object	95
Figure 12	Life Cycle of a Stateless Session Bean	97
Figure 13	OID for Creation of a Stateless Session Object	103
Figure 14	OID for Invocation of Business Method on a Stateless Session Object	104
Figure 15	OID for Removal of a Stateless Session Object	104
Figure 16	OID for Container Adding Instance of a Stateless Session Bean to a Method-Ready Pool	105
Figure 17	OID for a Container Removing an Instance of a Stateless Session Bean from Ready Pool	106
Figure 18	OID for Invocation of Business Method on a Stateless Session Bean	107
Figure 19	OID for Container Adding Instance of a Stateless Session Bean to a Method-Ready Pool	108
Figure 20	OID for a Container Removing an Instance of a Stateless Session Bean from Ready Pool	108
Figure 21	Example of Inheritance Relationships Between EJB Classes	120
Figure 22	Client View of Entity Beans Deployed in a Container	127
Figure 23	Client View of Entity Object Life Cycle	134
Figure 24	View of Underlying Data Sources Accessed Through Entity Bean	144
Figure 25	Relationship Example	174
Figure 26	Overview of the Entity Bean Runtime Execution Model	182
Figure 27	Life Cycle of an Entity Bean Instance	185
Figure 28	Multiple Clients Can Access the Same Entity Object Using Multiple Instances	207
Figure 29	Multiple Clients Can Access the Same Entity Object Using Single Instance	208
Figure 30	OID of Creation of an Entity Object with Container-Managed Persistence	225
Figure 31	OID of Passivation and Reactivation of an Entity Bean Instance with Container-Managed Persistence ... 227	
Figure 32	OID of Transaction Commit Protocol for an Entity Bean Instance with Container-Managed Persistence 228	
Figure 33	OID of Start of Transaction for an Entity Bean Instance with Container-Managed Persistence	230
Figure 34	OID of Removal of an Entity Bean Object with Container-Managed Persistence	231
Figure 35	OID of Execution of a Finder Method on an Entity Bean Instance with Container-Managed Persistence 232	
Figure 36	OID of a Container Adding an Instance to the Pool	233
Figure 37	OID of a Container Removing an Instance from the Pool	233

Figure 38	Several Entity Beans with Abstract Persistence Schemas Defined in the Same Ejb-jar File.....	238
Figure 39	Client View of Underlying Data Sources Accessed Through Entity Bean	266
Figure 40	Overview of the Entity Bean Runtime Execution Model.....	268
Figure 41	Life Cycle of an Entity Bean Instance.....	269
Figure 42	Multiple Clients Can Access the Same Entity Object Using Multiple Instances.....	288
Figure 43	Multiple Clients Can Access the Same Entity Object Using Single Instance	289
Figure 44	OID of Creation of an Entity Object with Bean-Managed Persistence.....	303
Figure 45	OID of Passivation and Reactivation of an Entity Bean Instance with Bean-Managed Persistence	304
Figure 46	OID of Transaction Commit Protocol for an Entity Bean Instance with Bean-Managed Persistence	305
Figure 47	OID of Start of Transaction for an Entity Bean Instance with Bean-Managed Persistence.....	306
Figure 48	OID of Removal of an Entity Bean Object with Bean-Managed Persistence	307
Figure 49	OID of Execution of a Finder Method on an Entity Bean Instance with Bean-Managed Persistence.....	308
Figure 50	OID of a Container Adding an Instance to the Pool.....	309
Figure 51	OID of a Container Removing an Instance from the Pool	309
Figure 52	Example of the Inheritance Relationship Between the Remote Interfaces and Classes:.....	312
Figure 53	OID of Creation of an Entity Object with EJB 1.1 Container-Managed Persistence.....	323
Figure 54	OID of Passivation and Reactivation of an Entity Bean Instance with EJB 1.1 CMP	325
Figure 55	OID of Transaction Commit Protocol for an Entity Bean Instance with EJB 1.1 Container-Managed Persistence.....	326
Figure 56	OID of Start of Transaction for an Entity Bean Instance with EJB 1.1 Container-Managed Persistence.....	328
Figure 57	OID of Removal of an Entity Bean Object with EJB 1.1 Container-Managed Persistence	329
Figure 58	OID of Execution of a Finder Method on an Entity Bean Instance with EJB 1.1 Container-Managed Persistence.....	330
Figure 59	OID of a Container Adding an Instance to the Pool.....	331
Figure 60	OID of a Container Removing an Instance from the Pool	331
Figure 61	Client view of Message-Driven Beans Deployed in a Container	335
Figure 62	Life Cycle of a Message-Driven Bean.....	342
Figure 63	OID for Invocation of Message Listener Method on Message-Driven Bean Instance.....	345
Figure 64	OID for Container Adding Instance of a Message-Driven Bean to a Method-Ready Pool	346
Figure 65	OID for a Container Removing an Instance of Message-Driven Bean from Ready Pool.....	347
Figure 66	Example of Inheritance Relationships Between EJB Classes	353
Figure 67	Updates to Simultaneous Databases	358
Figure 68	Message Sent to JMS Queue and Updates to Multiple Databases	359
Figure 69	Message Sent to JMS Queue Serviced by Message-Driven Bean and Updates to Multiple Databases.....	360
Figure 70	Updates to Multiple Databases in Same Transaction	360
Figure 71	Updates on Multiple Databases on Multiple Servers	361
Figure 72	Update of Multiple Databases from Non-Transactional Client.....	362
Figure 73	Transaction Diamond Scenario with Entity Object.....	390
Figure 74	Handling of Diamonds by a Multi-Process Container	392
Figure 75	Transaction Diamond Scenario with a Session Bean	393
Figure 76	Location of EJB Client Stubs.	414
Figure 77	Heterogeneous EJB Environment	415

Figure 78 Transaction Context Propagation	425
---	-----

List of Tables

Table 1	EJB Roles in the Example Scenarios.....	41
Table 2	Operations Allowed in the Methods of a Stateful Session Bean	87
Table 3	Operations Allowed in the Methods of a Stateless Session Bean	100
Table 4	Operations Allowed in the Methods of an Entity Bean.....	197
Table 5	Comparison of Finder and Select Methods	201
Table 6	Summary of Commit-Time Options	206
Table 7	Definition of the AND Operator.....	256
Table 8	Definition of the OR Operator	256
Table 9	Definition of the NOT Operator	256
Table 10	Operations Allowed in the Methods of an Entity Bean.....	279
Table 11	Summary of Commit-Time Options	287
Table 12	Operations Allowed in the Methods of a Message-Driven Bean	344
Table 13	Container's Actions for Methods of Beans with Bean-Managed Transaction	381
Table 14	Transaction Attribute Summary	385
Table 15	Handling of Exceptions Thrown by a Business Method of a Bean with Container-Managed Transaction Demarcation.....	400
Table 16	Handling of Exceptions Thrown by a Business Method of a Session Bean with Bean-Managed Transaction Demarcation.....	402
Table 17	Handling of Exceptions Thrown by an ejbCreate or ejbRemove Method of a Stateless Session Bean with Web Service Client View.....	402
Table 18	Handling of Exceptions Thrown by a Message Listener Method of a Message-Driven Bean with Container-Managed Transaction Demarcation.....	403
Table 19	Handling of Exceptions Thrown by a Message Listener Method of a Message-Driven Bean with Bean-Managed Transaction Demarcation.....	404
Table 20	Handling of Exceptions Thrown by an ejbCreate or ejbRemove Method of a Message-Driven Bean.....	404
Table 21	Handling of Exceptions Thrown by the ejbTimeout Method of an Enterprise Bean with Container-Managed Transaction Demarcation.....	405
Table 22	Handling of Exceptions Thrown by the ejbTimeout Method of an Enterprise Bean with Bean-Managed Transaction Demarcation.....	405
Table 23	Java 2 Platform Security Policy for a Standard EJB Container.....	567

Introduction

This is the specification of the Enterprise JavaBeans™ architecture. The Enterprise JavaBeans architecture is a component architecture for the development and deployment of component-based distributed business applications. Applications written using the Enterprise JavaBeans architecture are scalable, transactional, and multi-user secure. These applications may be written once, and then deployed on any server platform that supports the Enterprise JavaBeans specification.

1.1 Target Audience

The target audiences for this specification are the vendors of transaction processing platforms, vendors of enterprise application tools, and other vendors who want to support the Enterprise JavaBeans (EJB) technology in their products.

Many concepts described in this document are system-level issues that are transparent to the Enterprise JavaBeans application programmer.

1.2 What is New in EJB 2.1

The Enterprise JavaBeans 2.1 architecture presented in this document extends Enterprise JavaBeans to include the following new functionality:

- *We have added support for web services. We have added the ability for a stateless session bean to implement a web service endpoint. We have added support to allow all enterprise beans to utilize external web services.*
- *We have added a container-managed timer service. The timer service provides coarse-grained, transactional, time-based event notifications to enable enterprise beans to model and manage higher-level business processes.*
- *We have generalized the architecture for message-driven beans to support the use of messaging types in addition to JMS.*
- *We have enhanced EJB QL to include support for aggregate functions, ordering of results, and additional scalar functions, and have clarified the rules for allowing null values to be returned from finder and select methods.*
- *We have added deployment descriptor elements to enable the specification of the flow of messages from producers to consumers by means of common message destinations.*

1.3 Acknowledgments

The Enterprise JavaBeans 2.1 specification builds upon and extends the EJB 1.1 specification by Vlada Matena and Mark Hapner and the EJB 2.0 specification by Linda DeMichiel, Umit Yalcinalp, and Sanjeev Krishnan.

The Enterprise JavaBeans architecture is a broad effort that includes contributions from numerous groups at Sun and partner companies. In particular, we would like to thank the members of the EJB 2.1 Expert Group for their contributions to this version of the specification. The participants in the EJB 2.1 Expert Group have included: Nathan Abramson, Art Technology Group; Seth White, BEA; Jishnu Mitra, Borland; Anton Vorsamer, Fujitsu-Siemens; Russell Gold, Hewlett-Packard; Jim Knutson, IBM; David Hayes, IONA; Larry Allen, Novell; Sastry Malladi, Oracle; James Pinpin, Persistence Software; Amit Khanna, Pramati; Alan Davies, SeeBeyond Technology; Linda DeMichiel, Sun Microsystems; Evan Ireland, Sybase; Jon Dart, Tibco; Dan MacKinnon, Webgain; and Richard Monson-Haefel.

This specification draft has also had the benefit of numerous technical suggestions from the members of the J2EE team at Sun Microsystems. In particular, we would like to thank Mark Hapner, Mahesh Kanan, Sanjeev Krishnan, Ken Saks, Bill Shannon, and Umit Yalcinalp for the technical input they have provided. The J2EE 1.4 deployment descriptors, including the EJB 2.1 deployment descriptor, were converted to the use of XML Schema by Umit Yalcinalp.

1.4 Organization

Chapter 2, “Goals” discusses the advantages of the Enterprise JavaBeans architecture.

Chapter 3, “Roles and Scenarios” discusses the responsibilities of the Bean Provider, Application Assembler, Deployer, EJB Container and Server Providers, and System Administrator with respect to the Enterprise JavaBeans architecture.

Chapter 4, “Fundamentals” defines the scope of the Enterprise JavaBeans specification.

Chapter 5 discusses the difference between remote, local, and web service client views.

Chapters 6 through 8 define Session Beans: Chapter 6 discusses the client view; Chapter 7 presents the Session Bean component contract; and Chapter 8 outlines an example Session Bean scenario.

Chapters 9 through 14 define Entity Beans: Chapter 9 discusses the client view; Chapter 10 presents the Entity Bean component contract for container-managed persistence; Chapter 11 presents EJB QL, the query language for Entity Beans with container-managed persistence; Chapter 12 presents the Entity Bean component contract for bean-managed persistence; Chapter 13 outlines an example Entity Bean scenario; and Chapter 14 specifies the EJB 1.1 Entity Bean component contract for container-managed persistence.

Chapters 15 through 16 define Message-Driven Beans: Chapter 15 presents the Message-Driven Bean component contract; and Chapter 16 outlines an example Message-Driven Bean scenario.

Chapters 17 through 21 discuss transactions, exceptions, distribution and interoperability, environment, and security.

Chapter 22 describes the container-managed timer service.

Chapters 23 and 24 describe the format of the ejb-jar file and its deployment descriptor.

Chapter 25 defines the runtime APIs that a compliant EJB container must provide to the enterprise bean instances at runtime. The chapter also specifies the programming restrictions for portable enterprise beans.

Chapter 26 summarizes the responsibilities of the individual EJB Roles.

Chapter 27 is the Enterprise JavaBeans API Reference.

Chapter 28 provides a list of related documents.

1.5 Document Conventions

The regular Times font is used for information that is prescriptive by the EJB specification.

The italic Times font is used for paragraphs that contain descriptive information, such as notes describing typical use, or notes clarifying the text with prescriptive specification.

The Courier font is used for code examples.

The Helvetica font is used to specify the BNF of EJB QL.

Goals

2.1 Overall Goals

The Enterprise JavaBeans (EJB) architecture has the following goals:

- *The Enterprise JavaBeans architecture will be the standard component architecture for building distributed object-oriented business applications in the Java™ programming language.*
- *The Enterprise JavaBeans architecture will support the development, deployment, and use of web services.*
- *The Enterprise JavaBeans architecture will make it easy to write applications: Application developers will not have to understand low-level transaction and state management details, multi-threading, connection pooling, or other complex low-level APIs.*
- *Enterprise JavaBeans applications will follow the Write Once, Run Anywhere™ philosophy of the Java programming language. An enterprise bean can be developed once, and then deployed on multiple platforms without recompilation or source code modification.*
- *The Enterprise JavaBeans architecture will address the development, deployment, and runtime aspects of an enterprise application's life cycle.*

- *The Enterprise JavaBeans architecture will define the contracts that enable tools from multiple vendors to develop and deploy components that can interoperate at runtime.*
- *The Enterprise JavaBeans architecture will make it possible to build applications by combining components developed using tools from different vendors.*
- *The Enterprise JavaBeans architecture will provide interoperability between enterprise beans and Java 2 Platform, Enterprise Edition (J2EE) components as well as non-Java programming language applications.*
- *The Enterprise JavaBeans architecture will be compatible with existing server platforms. Vendors will be able to extend their existing products to support Enterprise JavaBeans.*
- *The Enterprise JavaBeans architecture will be compatible with other Java programming language APIs.*
- *The Enterprise JavaBeans architecture will be compatible with the CORBA protocols.*

2.2 EJB Releases 1.0 and 1.1

Enterprise JavaBeans Release 1.0 focused on the following aspects:

- *Defined the distinct “EJB Roles” that are assumed by the component architecture.*
- *Defined the client view of enterprise beans.*
- *Defined the enterprise bean developer’s view.*
- *Defined the responsibilities of an EJB Container Provider and Server Provider; together these make up a system that supports the deployment and execution of enterprise beans.*
- *Defined the format of the ejb-jar file, EJB’s unit of deployment.*

Release 1.1 augmented this with a focus on the following:

- *Provided better support for application assembly and deployment.*
- *Specified in greater detail the responsibilities of the individual EJB roles.*

2.3 EJB Release 2.0

Enterprise JavaBeans Release 2.0 focused on the following aspects:

- *Defined the message-driven bean type and the integration of EJB with the Java Message Service.*

- *Provided a local client view and support for efficient, lightweight access to enterprise beans from local clients.*
- *Defined a new architecture for container-managed persistence.*
- *Provided support for the management of relationships among entity beans.*
- *Provided a declarative query syntax for finder and select methods for entity beans with container-managed persistence.*
- *Provided support for additional methods in the home interface.*
- *Provided support for a run-as security identity.*
- *Provided for network interoperability among EJB servers.*

2.4 Goals for the EJB 2.1 Release

Enterprise JavaBeans Release 2.1 focuses on the following aspects:

- *Enabling enterprise beans to implement and utilize web services.*
- *Providing a container-managed timer service.*
- *Enhancing EJB QL with the addition of ORDER BY and aggregate operators.*
- *Extending the message-driven bean component type to other messaging types.*

EJB Roles and Scenarios

3.1 EJB Roles

The Enterprise JavaBeans architecture defines six distinct roles in the application development and deployment life cycle. Each EJB Role may be performed by a different party. The EJB architecture specifies the contracts that ensure that the product of each EJB Role is compatible with the product of the other EJB Roles. The EJB specification focuses on those contracts that are required to support the development and deployment of ISV-written enterprise beans.

In some scenarios, a single party may perform several EJB Roles. For example, the Container Provider and the EJB Server Provider may be the same vendor. Or a single programmer may perform the two EJB Roles of the Enterprise Bean Provider and the Application Assembler.

The following sections define the six EJB Roles.

3.1.1 Enterprise Bean Provider

The Enterprise Bean Provider (Bean Provider for short) is the producer of enterprise beans. His or her output is an ejb-jar file that contains one or more enterprise beans. The Bean Provider is responsible for the Java classes that implement the enterprise bean's business methods; the definition of the bean's home, component, and/or web service endpoint interfaces; and the bean's deployment descriptor. The deployment descriptor includes the structural information (e.g. the name of the enterprise bean class) of the enterprise bean and declares all the enterprise bean's external dependencies (e.g. the names and types of resources that the enterprise bean uses).

The Enterprise Bean Provider is typically an application domain expert. The Bean Provider develops reusable enterprise beans that typically implement business tasks or business entities.

The Bean Provider is not required to be an expert at system-level programming. Therefore, the Bean Provider usually does not program transactions, concurrency, security, distribution, or other services into the enterprise beans. The Bean Provider relies on the EJB container for these services.

A Bean Provider of multiple enterprise beans often performs the EJB Role of the Application Assembler.

3.1.2 Application Assembler

The Application Assembler combines enterprise beans into larger deployable application units. The input to the Application Assembler is one or more ejb-jar files produced by the Bean Provider(s). The Application Assembler outputs one or more ejb-jar files that contain the enterprise beans along with their application assembly instructions. The Application Assembler inserts the application assembly instructions into the deployment descriptors.

The Application Assembler can also combine enterprise beans with other types of application components (e.g. JavaServer Pages™) when composing an application.

The EJB specification describes the case in which the application assembly step occurs *before* the deployment of the enterprise beans. However, the EJB architecture does not preclude the case that application assembly is performed *after* the deployment of all or some of the enterprise beans.

The Application Assembler is a domain expert who composes applications that use enterprise beans. The Application Assembler works with the enterprise bean's deployment descriptor and the enterprise bean's client-view contract. Although the Assembler must be familiar with the functionality provided by the enterprise bean's client-view interfaces, he or she does not need to have any knowledge of the enterprise bean's implementation.

3.1.3 Deployer

The Deployer takes one or more ejb-jar files produced by a Bean Provider or Application Assembler and deploys the enterprise beans contained in the ejb-jar files in a specific operational environment. The operational environment includes a specific EJB server and container.

The Deployer must resolve all the external dependencies declared by the Bean Provider (e.g. the Deployer must ensure that all resource manager connection factories used by the enterprise beans are present in the operational environment, and he or she must bind them to the resource manager connection factory references declared in the deployment descriptor), and must follow the application assembly instructions defined by the Application Assembler. To perform his or her role, the Deployer uses tools provided by the EJB Container Provider.

The Deployer's output is a set of enterprise beans (or an assembled application that includes enterprise beans) that have been customized for the target operational environment, and that are deployed in a specific EJB container.

The Deployer is an expert at a specific operational environment and is responsible for the deployment of enterprise beans. For example, the Deployer is responsible for mapping the security roles defined by the Application Assembler to the user groups and accounts that exist in the operational environment in which the enterprise beans are deployed.

The Deployer uses tools supplied by the EJB Container Provider to perform the deployment tasks. The deployment process is typically two-stage:

- *The Deployer first generates the additional classes and interfaces that enable the container to manage the enterprise beans at runtime. These classes are container-specific.*
- *The Deployer performs the actual installation of the enterprise beans and the additional classes and interfaces into the EJB container.*

In some cases, a qualified Deployer may customize the business logic of the enterprise beans at their deployment. Such a Deployer would typically use the Container Provider's tools to write relatively simple application code that wraps the enterprise bean's business methods.

3.1.4 EJB Server Provider

The EJB Server Provider is a specialist in the area of distributed transaction management, distributed objects, and other lower-level system-level services. A typical EJB Server Provider is an OS vendor, middleware vendor, or database vendor.

The current EJB architecture assumes that the EJB Server Provider and the EJB Container Provider roles are the same vendor. Therefore, it does not define any interface requirements for the EJB Server Provider.

3.1.5 EJB Container Provider

The EJB Container Provider (Container Provider for short) provides:

- The deployment tools necessary for the deployment of enterprise beans.
- The runtime support for the deployed enterprise bean instances.

From the perspective of the enterprise beans, the container is a part of the target operational environment. The container runtime provides the deployed enterprise beans with transaction and security management, network distribution of remote clients, scalable management of resources, and other services that are generally required as part of a manageable server platform.

The “EJB Container Provider’s responsibilities” defined by the EJB architecture are meant to be requirements for the implementation of the EJB container and server. Since the EJB specification does not architect the interface between the EJB container and server, it is left up to the vendor how to split the implementation of the required functionality between the EJB container and server.

The expertise of the Container Provider is system-level programming, possibly combined with some application-domain expertise. The focus of a Container Provider is on the development of a scalable, secure, transaction-enabled container that is integrated with an EJB server. The Container Provider insulates the enterprise bean from the specifics of an underlying EJB server by providing a simple, standard API between the enterprise bean and the container. This API is the Enterprise JavaBeans component contract.

The Container Provider typically provides support for versioning the installed enterprise bean components. For example, the Container Provider may allow enterprise bean classes to be upgraded without invalidating existing clients or losing existing enterprise bean objects.

The Container Provider typically provides tools that allow the System Administrator to monitor and manage the container and the beans running in the container at runtime.

3.1.6 System Administrator

The System Administrator is responsible for the configuration and administration of the enterprise’s computing and networking infrastructure that includes the EJB server and container. The System Administrator is also responsible for overseeing the well-being of the deployed enterprise beans applications at runtime.

The EJB architecture does not define the contracts for system management and administration. The System Administrator typically uses runtime monitoring and management tools provided by the EJB Server and Container Providers to accomplish these tasks.

3.2 Scenario: Development, Assembly, and Deployment

*Aardvark Inc. specializes in application integration. Aardvark developed the AardvarkPayroll enterprise bean, which is a generic payroll access component that allows Java™ applications to access the payroll modules of the leading ERP systems. Aardvark packages the AardvarkPayroll enterprise bean in a standard ejb-jar file and markets it as a customizable enterprise bean to application developers. In the terms of the EJB architecture, Aardvark is the **Bean Provider** of the AardvarkPayroll bean.*

*Wombat Inc. is a Web-application development company. Wombat developed an employee self-service application. The application allows a target enterprise's employees to access and update employee record information. The application includes the EmployeeService, EmployeeServiceAdmin, and EmployeeRecord enterprise beans. The EmployeeRecord bean is a container-managed entity that allows deployment-time integration with an enterprise's existing human resource applications. In terms of the EJB architecture, Wombat is the **Bean Provider** of the EmployeeService, EmployeeServiceAdmin, and EmployeeRecord enterprise beans.*

In addition to providing access to employee records, Wombat would like to provide employee access to the enterprise's payroll and pension plan systems. To provide payroll access, Wombat licenses the AardvarkPayroll enterprise bean from Aardvark, and includes it as part of the Wombat application. Because there is no available generic enterprise bean for pension plan access, Wombat decides that a suitable pension plan enterprise bean will have to be developed at deployment time. The pension plan bean will implement the necessary application integration logic, and it is likely that the pension plan bean will be specific to each Wombat customer.

In order to provide a complete solution, Wombat also develops the necessary non-EJB components of the employee self-service application, such as the JavaServer Pages™ (JSP) that invoke the enterprise beans and generate the HTML presentation to the clients. Both the JSP pages and enterprise beans are customizable at deployment time because they are intended to be sold to a number of target enterprises that are Wombat customers.


*The Wombat application is packaged as a collection of JAR files. A single ejb-jar file contains all the enterprise beans developed by Wombat and also the AardvarkPayroll enterprise bean developed by Aardvark; the other JAR files contain the non-EJB application components, such as the JSP components. The ejb-jar file contains the application assembly instructions describing how the enterprise beans are composed into an application. In terms of the EJB architecture, Wombat performs the role of the **Application Assembler**.*

*Acme Corporation is a server software vendor. Acme developed an EJB server and container. In terms of the EJB architecture, Acme performs the **EJB Container Provider** and **EJB Server Provider** roles.*

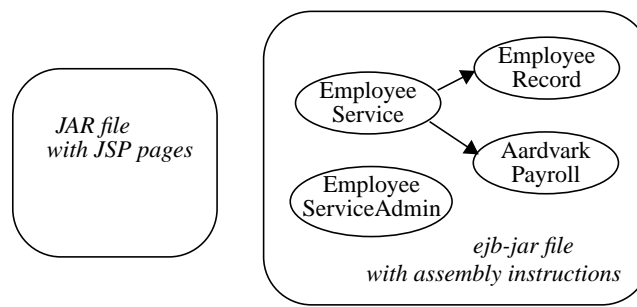
*The ABC Enterprise wants to enable its employees to access and update employee records, payroll information, and pension plan information over the Web. The information is stored in ABC's ERP system. ABC buys the employee self-service application from Wombat. To host the application, ABC buys the EJB container and server from Acme. ABC's Information Technology (IT) department, with the help of Wombat's consulting services, deploys the Wombat self-service application. In terms of the EJB architecture, ABC's IT department and Wombat consulting services perform the **Deployer** role. ABC's IT department also develops the ABCPensionPlan enterprise bean that provides the Wombat application with access to ABC's existing pension plan application.*

*ABC's IT staff is responsible for configuring the Acme product and integrating it with ABC's existing network infrastructure. The IT staff is responsible for the following tasks: security administration, such as adding and removing employee accounts; adding employees to user groups such as the payroll department; and mapping principals from digital certificates that identify employees on VPN connections from home computers to the Kerberos user accounts that are used on ABC's intranet. ABC's IT staff also monitors the well-being of the Wombat application at runtime, and is responsible for servicing any error conditions raised by the application. In terms of the EJB architecture, ABC's IT staff performs the role of the **System Administrator**.*

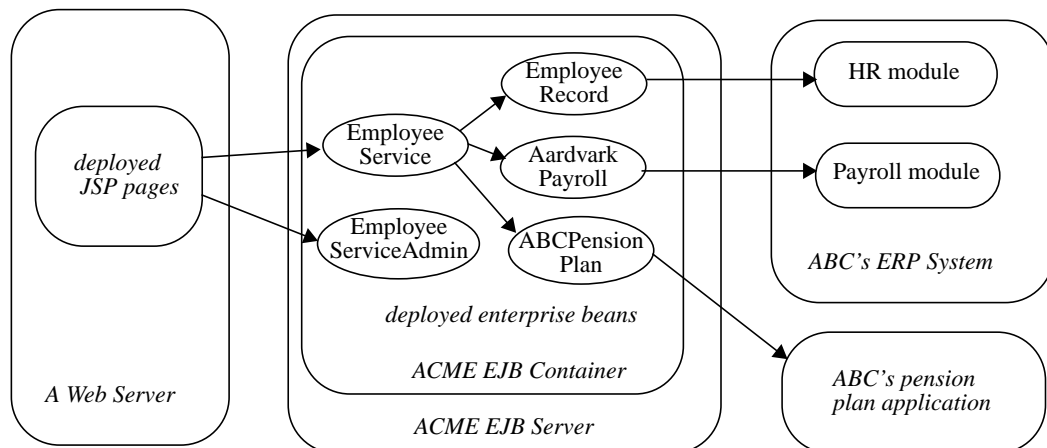
The following diagrams illustrates the products of the various EJB Roles.



(a) Aardvark's product is an ejb-jar file with an enterprise bean



(b) Wombat's product is an ejb-jar file with several enterprise beans assembled into an application. Wombat's product also includes non-EJB components.



(c) Wombat's application is deployed in ACME's EJB Container at the ABC enterprise.

The following table summarizes the EJB Roles of the organizations involved in the scenario.

Table 1 EJB Roles in the Example Scenarios

Organization	EJB Roles
Aardvark Inc.	Bean Provider
Wombat Inc.	Bean Provider Application Assembler
Acme Corporation	EJB Container Provider EJB Server Provider
ABC Enterprise's IT staff	Deployer Bean Provider (of <code>ABCPensionPlan</code>) System Administrator

Overview

This chapter provides an overview of the Enterprise JavaBeans specification.

4.1 Enterprise Beans as Components

Enterprise JavaBeans is an architecture for component-based computing. Enterprise beans are components of transaction-oriented enterprise applications.

4.1.1 Component Characteristics

The essential characteristics of an enterprise bean are:

- An enterprise bean typically contains business logic that operates on the enterprise's data.
- An enterprise bean's instances are created and managed at runtime by a container.
- An enterprise bean can be customized at deployment time by editing its environment entries.

- Various service information, such as transaction and security attributes, are separate from the enterprise bean class. This allows the service information to be managed by tools during application assembly and deployment.
- Client access is mediated by the container in which the enterprise bean is deployed.
- If an enterprise bean uses only the services defined by the EJB specification, the enterprise bean can be deployed in any compliant EJB container. Specialized containers can provide additional services beyond those defined by the EJB specification. An enterprise bean that depends on such a service can be deployed only in a container that supports that service.
- An enterprise bean can be included in an assembled application without requiring source code changes or recompilation of the enterprise bean.
- The Bean Provider defines a client view of an enterprise bean. The Bean Provider can manually define the client view or it can be generated automatically by application development tools. The client view is unaffected by the container and server in which the bean is deployed. This ensures that both the beans and their clients can be deployed in multiple execution environments without changes or recompilation.

4.1.2 Flexible Component Model

The enterprise bean architecture is flexible enough to implement components such as the following:

- An object that represents a stateless service.
- An object that represents a stateless service and that implements a web service endpoint.
- An object that represents a stateless service and whose invocation is asynchronous, driven by the arrival of messages.
- An object that represents a conversational session with a particular client. Such session objects automatically maintain their conversational state across multiple client-invoked methods.
- An entity object that represents a business object that can be shared among multiple clients.
- An entity object that represents a fine-grained persistent object that embodies the persistent state of a coarse-grained business object.

Enterprise beans that are remotely accessible components are intended to be relatively coarse-grained business objects (e.g. purchase order, employee record). Fine-grained objects (e.g. line items on a purchase order, employee's address) should not be modeled as remotely accessible enterprise bean components, but rather as locally accessible enterprise beans or as the dependent classes of enterprise beans.

Although the state management protocol defined by the Enterprise JavaBeans architecture is simple, it provides an enterprise bean developer great flexibility in managing a bean's state.

4.2 Enterprise JavaBeans Contracts

This section provides an overview of the Enterprise JavaBeans contracts. The contracts are described in detail in the following chapters of this document.

4.2.1 Client-view Contracts for Session Beans and Entity Beans

The client-view contracts are contracts between a client and a container. The client-view contracts provide a uniform development model for applications using enterprise beans as components. This uniform model enables the use of higher level development tools and allows greater reuse of components.

A client of an enterprise bean can be another enterprise bean deployed in the same or different container. Or it can be an arbitrary Java program, such as an application, applet, or servlet. The client of an enterprise bean can also be a non-Java client, such as a web service client or a CORBA client that is not written in the Java™ programming language.

The client of a session or entity bean may be a *remote client* or the client may be a *local client*. The client of a stateless session bean may be a *web service client*.

The enterprise bean remote client view is remotable—that is, both local and remote programs can access an enterprise bean using the same remote view of the enterprise bean. The objects that implement the interfaces for the remote client view are remote Java objects, and are accessible from a client through the standard Java™ APIs for remote object invocation [3].

The enterprise bean local client view is not remotable—only local components can access the enterprise bean through its local client view. Local components execute in the same Java Virtual Machine (JVM) as the locally accessed enterprise bean. The local client view is based on the use of standard Java™ programming language interfaces.

The enterprise bean web service client view is remotable. The web service endpoint that is implemented by a stateless session bean is accessed by Java clients through the JAX-RPC client APIs, which utilize an XML-based protocol (SOAP 1.1 [24]) over an HTTP transport. The service that includes the web service endpoint implemented by a stateless session bean is typically exported through a WSDL document that describes the external web service client view.

While it is possible for the Bean Provider to provide more than one client view for an enterprise bean, typically only one will be provided.

Both the local and the remote client view include:

- Home interface
- Component interface
- Object identity

The remote client view additionally includes:

- Metadata interface
- Handle

For the web service client, the existence of the enterprise bean is encapsulated behind the external web service client view as defined by the WSDL document that describes the web service endpoint that the bean implements. A bean that implements a web service endpoint does not have a home or component interface unless it also provides a local or remote client view. The Bean Provider defines the web service endpoint interface that the bean provides to the container as part of the implementation of the web service.

The enterprise bean's **home interface** defines the methods for the client to create, remove, and find EJB objects of the same type (i.e., they are implemented by the same enterprise bean) as well as home business methods for entity beans (business methods that are not specific to a particular bean instance). The home interface is specified by the Bean Provider; the container creates a class that implements the home interface. Enterprise beans that provide a remote client view provide a remote home interface; enterprise beans that provide a local client view provide a local home interface. A **remote home interface** extends the `javax.ejb.EJBHome` interface; a **local home interface** extends the `javax.ejb.EJBLocalHome` interface.

A client can locate an enterprise bean home interface through the standard Java Naming and Directory Interface™ (JNDI) API.

An EJB object is accessible via the enterprise bean's **component interface**. The component interface defines the business methods callable by the client. The component interface is specified by the Bean Provider; the container creates a class that implements the component interface. The component interface is either a **remote interface** or a **local interface**. A remote interface extends the `javax.ejb.EJBObject` interface. A local interface extends the `javax.ejb.EJBLocalObject` interface.

Each EJB object lives in a home, and has a unique identity within its home. For session beans, the container is responsible for generating a new unique identifier for each session object. The identifier is not exposed to the client. However, a client may test if two object references refer to the same session object. For entity beans, the Bean Provider is responsible for supplying a primary key at entity object creation time^[1]; the container uses the primary key to identify the entity object within its home. A client may obtain an entity object's primary key via the `javax.ejb.EJBObject` or the `javax.ejb.EJBLocalObject` interface. The client may also test if two object references refer to the same entity object.

A remote client may also obtain the enterprise bean's metadata interface. The metadata interface is typically used by clients who need to perform dynamic invocation of the enterprise bean. (Dynamic invocation is needed if the classes that provide the enterprise client view were not available at the time the client program was compiled.)

Message-driven beans have no home, component, or web service endpoint interface. A client can locate the destination or endpoint to which it should send messages that are to be delivered to a message-driven bean by means of the standard JNDI API.

[1] In some situations, the primary key type can be specified at deployment time (see subsection 10.8.3).

4.2.2 Component Contract

This subsection describes the contract between an enterprise bean and its container. The main requirements of the contract follow. (This is only a partial list of requirements defined by the specification.)

- The requirement for the Bean Provider to implement the business methods in the enterprise bean class. The requirement for the container to delegate the client method invocations to these methods.
- For message-driven beans, the requirement for the Bean Provider to implement the message listener methods in the enterprise bean class. The requirement for the container to invoke the bean's message listener method when a message has arrived for the bean to service.
- The requirement for the Bean Provider to implement the `ejbCreate<METHOD>`, `ejbPostCreate<METHOD>`, and `ejbRemove` methods, and to implement the `ejbFind<METHOD>` methods if the bean is an entity with bean-managed persistence. The requirement for the container to invoke these methods during EJB object creation, removal, and lookup.
- The requirement for the Bean Provider to define the enterprise bean's home and component interfaces, if the bean is an entity bean or a session bean that provides a local or remote client view. The requirement for the container to provide classes that implement these interfaces.
- The requirement for the Bean Provider to define the enterprise bean's web service endpoint interface, if the bean is a stateless session bean that provides a web service client view. The requirement for the container to provide the JAX-RPC runtime that receives and processes the HTTP request for the service endpoint, extracts and processes the SOAP message, and returns the result, if any. The requirement for the container to map the request to a method invocation of the appropriate method of the session bean class.
- For session beans, the requirement for the Bean Provider to implement the container callbacks defined in the `javax.ejb.SessionBean` interface, and optionally the `javax.ejb.SessionSynchronization` interfaces. The requirement for the container to invoke these callbacks at the appropriate times.
- For message-driven beans, the requirement for the Bean Provider to implement the container callbacks defined in the `javax.ejb.MessageDrivenBean` interface. The requirement for the container to invoke these callbacks at the appropriate times.
- For entity beans, the requirement for the Bean Provider to implement the container callbacks defined in the `javax.ejb.EntityBean` interface. The requirement for the container to invoke these callbacks at the appropriate times.
- The requirement for the container to implement persistence for entity beans with container-managed persistence.
- The requirement for the Bean Provider to provide abstract accessor methods for persistent fields and relationships for an entity with container-managed persistence. The requirement for the container to provide the implementation of these methods.

- The requirement for the Bean Provider to provide a description of the relationships for entity beans with container-managed persistence. The requirement for the container to implement the relationships described by the Bean Provider.
- The requirement for the Bean Provider to provide the EJB QL query strings for the finder and select methods of entity beans with container-managed persistence. The requirement for the container to implement the finder and select methods according to the query strings provided by the Bean Provider.
- The requirement for the container to provide the `javax.ejb.SessionContext` interface to session bean instances, the `javax.ejb.EntityContext` interface to entity bean instances, and the `javax.ejb.MessageDrivenContext` interface to message-driven bean instances. The context interface allows the instance to obtain information from the container.
- The requirement for the container to provide to the bean instances the JNDI context that contains the enterprise bean's environment.
- The requirement for the container to manage transactions, security, and exceptions on behalf of the enterprise bean instances.
- The requirement for the container to provide the timer service and to provide timed notifications to enterprise beans that are registered with the timer service.
- The requirement for the Bean Provider to avoid programming practices that would interfere with the container's runtime management of the enterprise bean instances.

4.2.3 Ejb-jar File

An **ejb-jar file** is a standard format used by EJB tools for packaging enterprise beans with their declarative information. The ejb-jar file is intended to be processed by application assembly and deployment tools.

The ejb-jar file is a contract used both between the Bean Provider and the Application Assembler, and between the Application Assembler and the Deployer.

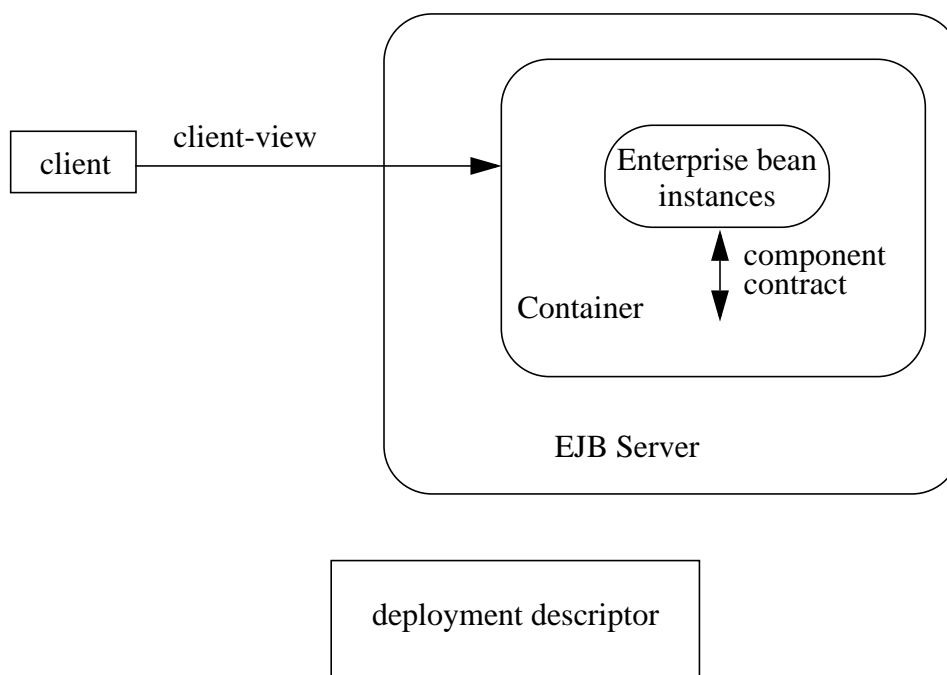
The ejb-jar file includes:

- Java class files for the enterprise beans and their home, component, and web service endpoint interfaces.
- An XML **deployment descriptor**. The deployment descriptor provides both the structural and application assembly information about the enterprise beans in the ejb-jar file. The application assembly information is optional. (Typically, only ejb-jar files with assembled applications include this information.)

4.2.4 Contracts Summary

The following figure illustrates the Enterprise JavaBeans contracts.

Figure 1 Enterprise JavaBeans Contracts



Note that while the figure illustrates only a client running outside of the container, the client-view APIs are also applicable to clients that are enterprise beans deployed in the same container.

4.3 Session, Entity, and Message-Driven Objects

The Enterprise JavaBeans architecture defines three types of enterprise bean objects:

- A session object.
- An entity object.
- A message-driven object.

4.3.1 Session Objects

A typical session object has the following characteristics:

- *Executes on behalf of a single client.*
- *Can be transaction-aware.*
- *Updates shared data in an underlying database.*
- *Does not represent directly shared data in the database, although it may access and update such data.*
- *Is relatively short-lived.*
- *Is removed when the EJB container crashes. The client has to re-establish a new session object to continue computation.*

A typical EJB container provides a scalable runtime environment to execute a large number of session objects concurrently.

*The EJB specification defines both **stateful** and **stateless session beans**. There are minor differences in the API between stateful session beans and stateless session beans.*

4.3.2 Entity Objects

A typical entity object has the following characteristics:

- *Provides an object view of data in the database.*
- *Allows shared access from multiple users.*
- *Can be long-lived (lives as long as the data in the database).*
- *The entity, its primary key, and its remote reference survive the crash of the EJB container. If the state of an entity was being updated by a transaction at the time the container crashed, the entity's state is automatically reset to the state of the last committed transaction. The crash is not fully transparent to the client—the client may receive an exception if it calls an entity in a container that has experienced a crash.*

A typical EJB container and server provide a scalable runtime environment for a large number of concurrently active entity objects.

4.3.3 Message-Driven Objects

A typical message-driven object has the following characteristics:

- *Executes upon receipt of a single client message.*

- *Is asynchronously invoked.*
- *Can be transaction-aware.*
- *May update shared data in an underlying database.*
- *Does not represent directly shared data in the database, although it may access and update such data.*
- *Is relatively short-lived.*
- *Is stateless.*
- *Is removed when the EJB container crashes. The container has to re-establish a new message-driven object to continue computation.*

A typical EJB container provides a scalable runtime environment to execute a large number of message-driven objects concurrently.

4.4 Standard Mapping to CORBA Protocols

To help interoperability for EJB environments that include systems from multiple vendors, the EJB specification requires compliant implementations to support the interoperability protocol based on CORBA/IIOP for remote invocations from J2EE clients on session beans and entity beans through their remote home and remote interfaces. Implementations may support other remote invocation protocols in addition to IIOP.

Chapter 19 summarizes the requirements for support for distribution and interoperability.

4.5 Mapping to Web Service Protocols

To support web service interoperability, the EJB 2.1 specification requires compliant implementations to support XML-based web service invocations using WSDL 1.1 and SOAP 1.1 over HTTP 1.1 in conformance with the requirements of the JAX-RPC [22] and Web Services for J2EE 1.1 [25] specifications.

Chapter 5

Local, Remote, and Web Service Client Views

This chapter describes some of the considerations the Bean Provider should take into account in determining the client view to provide for an enterprise bean.

5.1 Overview

A session bean is a component that implements some business logic running on the server. An entity bean is a component that represents an object-oriented view of some entities stored in a persistent storage, such as a database, or entities that are implemented by an existing enterprise application.

The client of a session bean or an entity bean may be a remote client or it may be a local client, depending on whether the client makes use of the enterprise bean's remote or local client view interfaces. The differences between remote and local clients are discussed in Sections 5.2 through 5.4.

The client of a stateless session bean may be a web service client. A web service client makes use of the enterprise bean's web service client view corresponding to the web service endpoint interface, as described by a WSDL document. Web service clients are discussed in Section 5.5.

5.2 Remote Clients

A remote client accesses a session bean or an entity bean through the bean's remote interface and remote home interface. The remote and remote home interfaces of the bean provide the remote client view.

The remote client view of an enterprise bean is location independent. A client running in the same JVM as a bean instance uses the same API to access the bean as a client running in a different JVM on the same or different machine.

The remote and remote home interfaces are Java™ RMI interfaces. The container provides classes that implement the bean's remote and remote home interfaces. The objects that implement the remote home and remote interfaces are remote Java objects, and are accessible from a client through the standard Java RMI APIs [3]. The arguments and results of the methods of the remote interface and remote home interface are passed by value.

A remote client can be another enterprise bean deployed in the same or different container; or a remote client can be an arbitrary Java program, such as an application, applet, or servlet. The remote client view can also be mapped to non-Java client environments, such as CORBA clients not written in the Java programming language.

5.3 Local Clients

Session and entity beans may have local clients. A local client is a client that is collocated in the same JVM with the session or entity bean that provides the local client view and which may be tightly coupled to the bean. A local client of a session bean or an entity bean may be another enterprise bean (a session bean, entity bean, or message-driven bean) or a web component.

Unlike the remote client view, the local client view of a bean is not location independent. Access to an enterprise bean through the local client view requires the collocation in the same JVM of both the local client and the enterprise bean that provides the local client view. The local client view therefore does not provide the location transparency provided by the remote client view.

A local client accesses a session or an entity bean through the bean's local interface and local home interface. The container provides classes that implement the bean's local and local home interfaces. The objects that implement the local home interface and local interface are local Java objects.

The arguments and results of the methods of the local interface and local home interface are passed "by reference"^[2]. Enterprise beans that provide a local client view should therefore be coded to assume that the state of any Java object that is passed as an argument or result is potentially shared by caller and callee.

[2] More literally, references are passed by value in the JVM: an argument variable of primitive type holds a value of that primitive type; an argument variable of a reference type holds a reference to the object. See [26].

In designing an application that makes use of the local programming model offered by local interfaces, the Bean Provider must be aware of the potential sharing of objects passed through the local interface. In particular, the Bean Provider must be careful that the state of one enterprise bean is not assigned as the state of another. In general, the references that are passed across the local interface cannot be used outside of the immediate call chain and must never be stored as part of the state of another enterprise bean.

The Bean Provider must also exercise caution in determining which objects to pass across the local interface. This caution applies particularly in the case where there is a change in transaction or security context.

5.4 Choosing Between a Local or Remote Client View

The following considerations should be taken into account in determining whether the local or remote programming model should be used for an enterprise bean.

- The remote programming model provides location independence and flexibility with regard to the distribution of components in the deployment environment. It provides a loose coupling between the client and the bean.
- Remote calls involve pass-by-value. This copy semantics provides a layer of isolation between caller and callee, and protects against the inadvertent modification of data. The client and the bean may be programmed to assume this parameter copying.
- Remote calls are potentially expensive. They involve network latency, overhead of the client and server software stacks, argument copying, etc. Remote calls are typically programmed in a coarse-grained manner with few interactions between the client and bean.
- The objects that are passed as parameters on remote calls must be serializable.
- Narrowing remote types requires the use of `javax.rmi.PortableRemoteObject.narrow` rather than Java language casts.
- Remote calls may involve error cases due to communication, resource usage on other servers, etc., which are not expected in local calls. The client has to explicitly program handlers for these remote exceptions.
- Because of the overhead of the remote programming model, it is typically used for relatively coarse-grained component access.
- Local calls involve pass-by-reference. The client and the bean may be programmed to rely on pass-by-reference semantics. For example, a client may have a large document which it wants to pass on to the bean to modify, and the bean further passes on. In the local programming model the sharing of state is possible. On the other hand, when the bean wants to return a data structure to the client but the bean does not want the client to modify it, the bean explicitly copies the data structure before returning it, while in the remote programming model the bean does not copy the data structure because it assumes that the system will do the copy.

- Because local calls involve pass-by-reference, the local client and the enterprise bean providing the local client view are collocated.
- The collocation entailed by the local programming model means that the enterprise bean cannot be deployed on a node different from that of its client—thus restricting the distribution of components.
- Because the local programming model provides more lightweight access to a component, it better supports more fine-grained component access.

Note that although collocation of the remote client and the enterprise bean may allow the container to reduce the overhead of calls through the remote interface, such calls are still likely to be less efficient than calls made using the local interface because any optimizations based on collocation must be done transparently.

The choice between the local and the remote programming model is a design decision that the Bean Provider makes when developing the enterprise bean.

In designing an enterprise bean with a local client view, the Bean Provider should be aware that full use of pass-by-reference semantics will potentially limit the ability to couple components more loosely and to replace the local client view with a remote client view in order to achieve distribution of components at a later date. Developers who want to preserve the ability to distribute components at a later date should rely on the efficiency benefits of pass by reference semantics, but not rely on the side-effects.

While it is possible to provide both a remote client view and a local client view for an enterprise bean, more typically only one or the other will be provided.

Entity beans that have container-managed relationships with other entity beans, as described in Chapter 10, “Entity Bean Component Contract for Container-Managed Persistence”, must be accessed in the same local scope as those related beans, and therefore typically provide a local client view to their related beans.

5.5 Web Service Clients

Stateless session beans may have web service clients.

A web service client accesses a stateless session bean through the web service client view. The web service client view is described by the WSDL document for the web service that the bean implements. WSDL is an XML format for describing a web service as a set of endpoints operating on messages. The abstract description of the service is bound to an XML based protocol (SOAP 1.1 [24]) and underlying transport (HTTP or HTTPS) by means of which the messages are conveyed between client and server. (See [22], [23], [25]).

The web service client view provided by a stateless session bean corresponds to the bean’s web service endpoint interface. The web service endpoint interface of the bean provides the basis of the web service client view of the bean that is exported through WSDL [23].

The Bean Provider chooses a web service endpoint interface for a stateless session bean when he or she wishes to expose the functionality of the bean as a web service endpoint through WSDL. A bean's web service client view may be initially defined by a WSDL document and then mapped to a web service endpoint interface that conforms to this, or an existing bean may be adapted to provide a web service client view. Reference [25] describes various design-time scenarios that may be used for EJB 2.1 web service endpoints.

The web service client view of an enterprise bean is location independent and remotable.

Web service clients may be Java clients and/or clients not written in the Java programming language. A web service client that is a Java client accesses the web service by means of the JAX-RPC client APIs [22], [25]. Access through web service clients occurs through SOAP 1.1 and HTTP(S).

Client View of a Session Bean

This chapter describes the client view of a session bean. The session bean itself implements the business logic. The bean's container provides functionality for remote access, security, concurrency, transactions, and so forth.

While classes implemented by the container provide the client view of the session bean, the container itself is transparent to the client.

6.1 Overview

For a client, a session object is a non-persistent object that implements some business logic running on the server. One way to think of a session object is as a logical extension of the client program that runs on the server. A session object is not shared among multiple clients.

A client accesses a session object through the session bean's component interface or web service client view. The session bean may provide a remote interface and/or it may provide a local interface. A stateless session bean may provide a web service endpoint interface instead of (or in addition to) a component and home interface. Only a *stateless* session bean may provide a web service endpoint interface.

The Java object that implements the remote interface is called a session **EJBObject**. A session EJBObject is a remote Java object accessible from a client through the standard Java™ APIs for remote object invocation [3].

The Java object that implements the local interface is called a session **EJBLocalObject**. A session EJBLocalObject is locally accessible, but not remotely accessible.

The bean's web service endpoint interface is a JAX-RPC interface [22]. It is described in a WSDL document in terms of a web service endpoint. The web service client view of a stateless session bean is described further in Section 6.11.

A remote client of an session bean can be another enterprise bean deployed in the same or different container; or it can be an arbitrary Java program, such as an application, applet, or servlet. The client view of a session bean can also be mapped to non-Java client environments, such as CORBA clients that are not written in the Java programming language.

The remote client view of a session bean is location-independent. A client running in the same JVM as the session object uses the same API as a client running in a different JVM on the same or different machine.

A local client of a session bean accesses the session bean through a local Java interface, which assumes collocation of the local client and the session. The local client view is not location-independent.

While it is possible to provide more than one client view for a session bean, typically only one will be provided.

The considerations that should be taken into account in determining the client view to be used for a session bean are further described in Chapter 5, "Local, Remote, and Web Service Client Views".

From its creation until destruction, a session object lives in a container. The container provides security, concurrency, transactions, swapping to secondary storage, and other services for the session object transparently to the client.

Each session object has an identity which, in general, *does not* survive a crash and restart of the container, although a high-end container implementation can mask container and server crashes to a remote or web service client.

Multiple enterprise beans can be installed in a container. The container allows the clients of session beans that provide local or remote client views to look up the home interfaces of the installed enterprise beans via JNDI. A session bean's home interface provides methods to create and remove the session objects of a particular session bean.

The client view of a session object is independent of the implementation of the session bean and the container.

6.2 EJB Container

An EJB container (container for short) is a system that functions as the “container” for enterprise beans. Multiple enterprise beans can be deployed in the same container. The container is responsible for making the home interfaces of its deployed enterprise beans available to the client through JNDI. Thus, the client can look up the home interface for a specific enterprise bean using JNDI.

6.2.1 Locating a Session Bean’s Home Interface

A client locates a session bean’s home interface using JNDI. For example, the remote home interface for the `Cart` session bean can be located using the following code segment:

```
Context initialContext = new InitialContext();
CartHome cartHome = (CartHome) javax.rmi.PortableRemoteObject.narrow(
    initialContext.lookup("java:comp/env/ejb/cart"),
    CartHome.class);
```

If the `Cart` session bean provides a local client view instead of a remote client view and `CartHome` is a local home interface, this lookup might be as follows:

```
Context initialContext = new InitialContext();
CartHome cartHome = (CartHome)
    initialContext.lookup("java:comp/env/ejb/cart");
```

A client’s JNDI name space may be configured to include the home interfaces of enterprise beans installed in multiple EJB containers located on multiple machines on a network. The actual locations of an enterprise bean and EJB container are, in general, transparent to the client using the remote interfaces of the enterprise bean.

The life cycle of the distributed object implementing the remote home interface (the `EJBHome` object) or the local Java object implementing the local home interface (the `EJBLocalHome` object) is container-specific. A client application should be able to obtain a home interface, and then use it multiple times, during the client application’s lifetime.

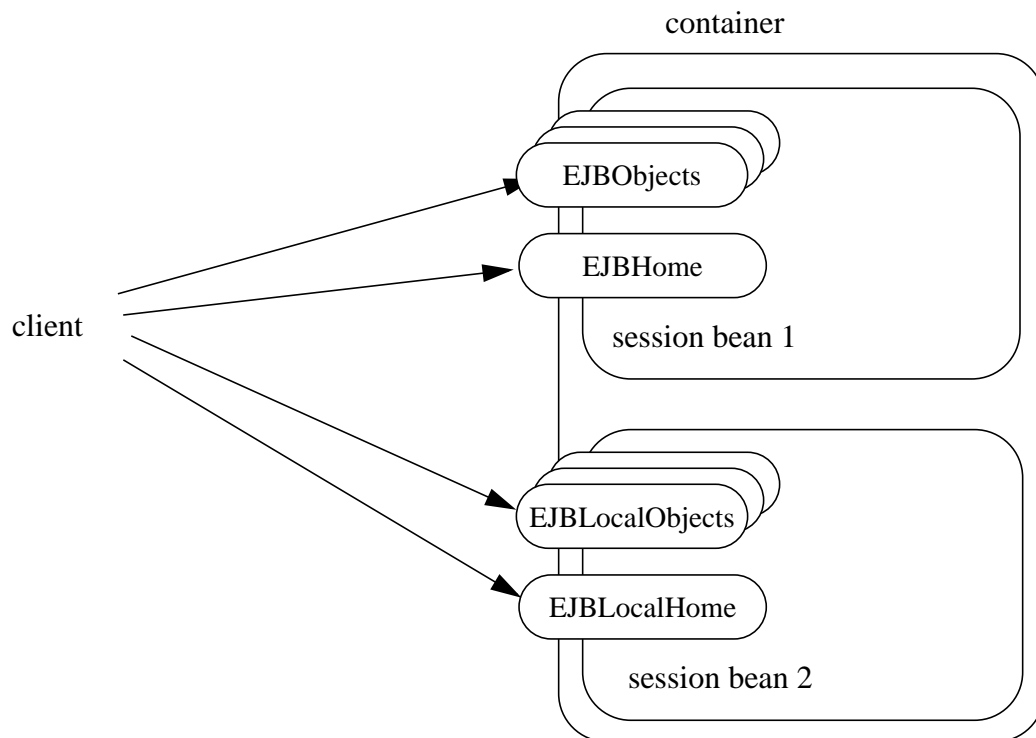
A client can pass a remote home object reference to another application. The receiving application can use the home interface in the same way that it would use a remote home object reference obtained via JNDI.

A client can pass a local home object reference to another application through its local interface. A local home object reference cannot be passed as an argument or result of a method on an enterprise bean’s remote home or remote interface.

6.2.2 What a Container Provides

The following diagram illustrates the view that a container provides to clients of session beans that provide local and/or remote client views. Note that a client may be a local client of some session beans and a remote client of others.

Figure 2 Client View of Session Beans Deployed in a Container



6.3 Session Bean's Remote Home Interface

This section is specific to session beans that provide a remote client view. Local home interfaces are described in Section 6.4.

The container provides the implementation of the remote home interface for each session bean that defines a remote home interface that is deployed in the container. The object that implements a session bean's remote home interface is called a session EJBHome object. The container makes the session bean's remote home interface available to the client through JNDI.

The remote home interface allows a client to do the following:

- Create a new session object.
- Remove a session object.
- Get the `javax.ejb.EJBMetaData` interface for the session bean. The `javax.ejb.EJBMetaData` interface is intended to allow application assembly tools to discover information about the session bean, and to allow loose client/server binding and client-side scripting.
- Obtain a handle for the remote home interface. The home handle can be serialized and written to stable storage. Later, possibly in a different JVM, the handle can be deserialized from stable storage and used to obtain back a reference of the remote home interface.

6.3.1 Creating a Session Object

A home interface defines one or more `create<METHOD>` methods, one for each way to create a session object. The arguments of the `create` methods are typically used to initialize the state of the created session object.

The return type of a `create<METHOD>` method on the remote home interface is the session bean's remote interface.

The following example illustrates a remote home interface that defines two `create<METHOD>` methods:

```
public interface CartHome extends javax.ejb.EJBHome {
    Cart create(String customerName, String account)
        throws RemoteException, BadAccountException,
        CreateException;
    Cart createLargeCart(String customerName, String account)
        throws RemoteException, BadAccountException,
        CreateException;
}
```

The following example illustrates how a client creates a new session object using a `create<METHOD>` method of the `CartHome` interface:

```
cartHome.create("John", "7506");
```

6.3.2 Removing a Session Object

A remote client may remove a session object using the `remove()` method of the `javax.ejb.EJBObject` interface, or the `remove(Handle handle)` method of the `javax.ejb.EJBHome` interface.

Because session objects do not have primary keys that are accessible to clients, invoking the `javax.ejb.EJBHome.remove(Object primaryKey)` method on a session results in a `javax.ejb.RemoveException`.

6.4 Session Bean's Local Home Interface

The container provides the implementation of the local home interface for each session bean that defines a local home interface that is deployed in the container. The object that implements a session bean's local home interface is called a session EJBLocalHome object. The container makes the session beans' local home interfaces available to the client through JNDI.

The local home interface allows a local client to do the following:

- Create a new session object.
- Remove a session object.

6.4.1 Creating a Session Object

A local home interface defines one or more `create<METHOD>` methods, one for each way to create a session object. The arguments of the `create` methods are typically used to initialize the state of the created session object.

The return type of a `create<METHOD>` method on the local home interface is the session bean's local interface.

The following example illustrates a local home interface that defines two `create<METHOD>` methods:

```
public interface CartHome extends javax.ejb.EJBLocalHome {  
    Cart create(String customerName, String account)  
        throws BadAccountException, CreateException;  
    Cart createLargeCart(String customerName, String account)  
        throws BadAccountException, CreateException;  
}
```

The following example illustrates how a client creates a new session object using a `create<METHOD>` method of the `CartHome` interface:

```
cartHome.create("John", "7506");
```

6.4.2 Removing a Session Object

A local client may remove a session object using the `remove()` method of the `javax.ejb.EJBLocalObject` interface.

Because session objects do not have primary keys that are accessible to clients, invoking the `javax.ejb.EJBLocalHome.remove(Object primaryKey)` method on a session results in a `javax.ejb.RemoveException`.

6.5 EJBObject and EJBLocalObject

A client never directly accesses instances of the session bean's class. A remote or local client always uses the session bean's component interface to access a session bean instance. The class that implements the session bean's component interface is provided by the container. Instances of a session bean's remote interface are called session **EJBObjects**. Instances of a session bean's local interface are called session **EJBLocalObjects**.

A session EJBObject supports:

- The business logic methods of the object. The session EJBObject delegates invocation of a business method to the session bean instance.
- The methods of the `javax.ejb.EJBObject` interface. These methods allow the client to:
 - Get the session object's remote home interface.
 - Get the session object's handle.
 - Test if the session object is identical with another session object.
 - Remove the session object.

A session EJBLocalObject supports:

- The business logic methods of the object. The session EJBLocalObject delegates invocation of a business method to the session bean instance.
- The methods of the `javax.ejb.EJBLocalObject` interface. These methods allow the client to:
 - Get the session object's local home interface.
 - Test if the session object is identical with another session object.
 - Remove the session object.

The implementation of the methods defined in the `javax.ejb.EJBObject` and `javax.ejb.EJBLocalObject` interfaces is provided by the container. They are not delegated to the instances of the session bean class.

6.6 Session Object Identity

Session objects are intended to be private resources used only by the client that created them. For this reason, session objects, from the client's perspective, appear anonymous. In contrast to entity objects, which expose their identity as a primary key, session objects hide their identity. As a result, the `EJBObject.getPrimaryKey()` method results in a `java.rmi.RemoteException` and the `EJBLocalObject.getPrimaryKey()` method results in a `javax.ejb.EJBException`, and the `EJBHome.remove(Object primaryKey)` and the `EJBLocalHome.remove(Object primaryKey)` methods result in a `javax.ejb.RemoveException` if called on a session bean. If the `EJBMetaData.getPrimaryKeyClass()` method is invoked on a `EJBMetaData` object for a session bean, the method throws the `java.lang.RuntimeException`.

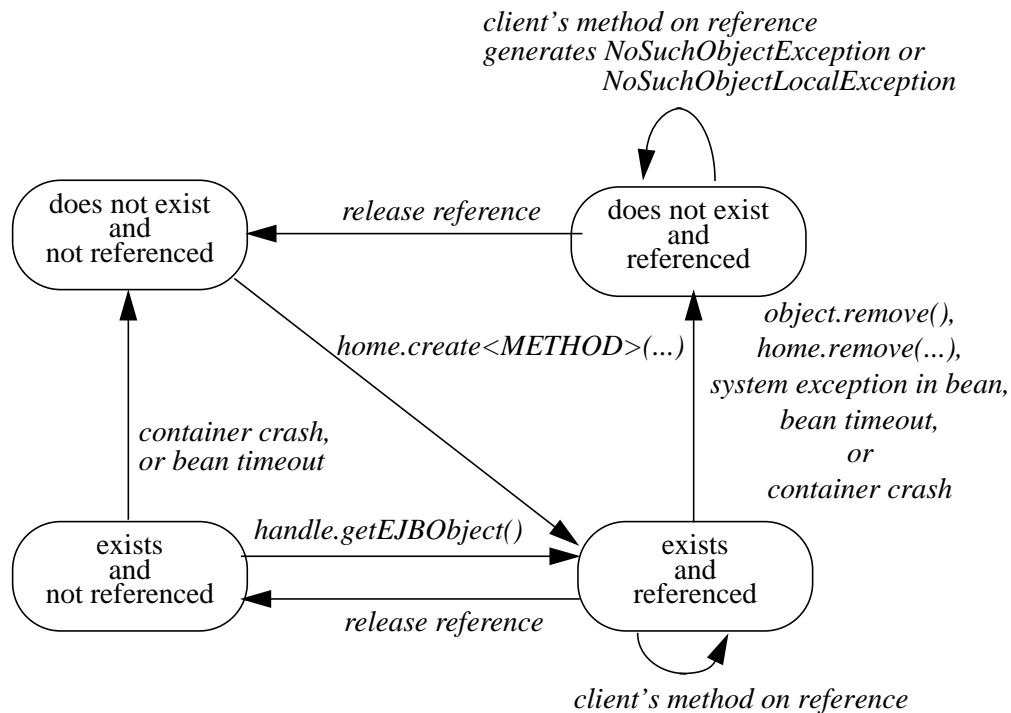
Since all session objects hide their identity, there is no need to provide a finder for them. The home interface of a session bean must not define any finder methods.

A session object handle can be held beyond the life of a client process by serializing the handle to persistent storage. When the handle is later deserialized, the session object it returns will work as long as the session object still exists on the server. (An earlier timeout or server crash may have destroyed the session object.)

A handle is not a capability, in the security sense, that would automatically grant its holder the right to invoke methods on the object. When a reference to a session object is obtained from a handle, and then a method on the session object is invoked, the container performs the usual access checks based on the caller's principal.

6.7 Client view of Session Object's Life Cycle

From the point of view of a local or remote client, the life cycle of a session object is illustrated below.

Figure 3 Life Cycle of a Session Object.

A session object does not exist until it is created. When a client creates a session object, the client has a reference to the newly created session object's component interface.

6.7.1 References to Session Object Remote Interfaces

A client that has a reference to a session object's remote interface can then do any of the following:

- Invoke business methods defined in the session object's remote interface.
- Get a reference to the session object's remote home interface.
- Get a handle for the session object.
- Pass the reference as a parameter or return value within the scope of the client.
- Remove the session object. A container may also remove the session object automatically when the session object's lifetime expires.

It is invalid to reference a session object that does not exist. Attempted remote invocations on a stateful session object that does not exist result in a `java.rmi.NoSuchObjectException`.^[3]

6.7.2 References to Session Object Local Interfaces

A client that has a reference to a session object's local interface can then do any of the following:

- Invoke business methods defined in the session object's local interface.
- Get a reference to the session object's local home interface.
- Pass the reference as a parameter or return value of a local interface method.
- Remove the session object. A container may also remove the session object automatically when the session object's lifetime expires.

It is invalid to reference a session object that does not exist. Attempted invocations on a stateful session object that does not exist result in `javax.ejb.NoSuchObjectLocalException`.^[4]


A client can pass a local object reference or local home object reference to another application through its local interface. A local object reference or local home object reference cannot be passed as an argument or result of a method on an enterprise bean's remote home or remote interface.

6.8 Creating and Using a Session Object

An example of the session bean runtime objects is illustrated by the following diagram:

[3] This may not apply to stateless session beans; see Section 7.8.

[4] This may not apply to stateless session beans; see Section 7.8.

Figure 4 Session Bean Example Objects

A client creates a remote `Cart` session object, which provides a shopping service, using a `create<METHOD>` method of the `Cart`'s remote home interface. The client then uses this session object to fill the cart with items and to purchase its contents.

Suppose that the end-user wishes to start the shopping session, suspend the shopping session temporarily for a day or two, and later complete the session. The client might implement this feature by getting the session object's handle, saving the serialized handle in persistent storage, and using it later to reestablish access to the original `Cart`.

For the following example, we start by looking up the `Cart`'s remote home interface in JNDI. We then use the remote home interface to create a `Cart` session object and add a few items to it:

```
CartHome cartHome = (CartHome)javax.rmi.PortableRemoteObject.narrow(
    initialContext.lookup(...), CartHome.class);
Cart cart = cartHome.createLargeCart(...);
cart.addItem(66);
cart.addItem(22);
```

Next we decide to complete this shopping session at a later time so we serialize a handle to this cart session object and store it in a file:

```
Handle cartHandle = cart.getHandle();
// serialize cartHandle, store in a file...
```

Finally we deserialize the handle at a later time, re-create the reference to the cart session object, and purchase the contents of the shopping cart:

```
Handle cartHandle = ...; // deserialize from a file...
Cart cart = (Cart)javax.rmi.PortableRemoteObject.narrow(
    cartHandle.getEJBObject(), Cart.class);
cart.purchase();
cart.remove();
```

6.9 Object Identity

6.9.1 Stateful Session Beans

A stateful session object has a unique identity that is assigned by the container at create time.

A remote client can determine if two remote object references refer to the same session object by invoking the `isIdentical(EJBObject otherEJBObject)` method on one of the references. A local client can determine if two local object references refer to the same session object by invoking the `isIdentical(EJBLocalObject otherEJBLocalObject)` method.

The following example illustrates the use of the `isIdentical` method for a stateful session object.

```
FooHome fooHome = ...; // obtain home of a stateful session bean
Foo foo1 = fooHome.create(...);
Foo foo2 = fooHome.create(...);

if (foo1.isIdentical(foo1)) { // this test must return true
    ...
}

if (foo1.isIdentical(foo2)) { // this test must return false
    ...
}
```

6.9.2 Stateless Session Beans

All session objects of the same stateless session bean within the same home have the same object identity, which is assigned by the container. If a stateless session bean is deployed multiple times (each deployment results in the creation of a distinct home), session objects from different homes will have a different identity.

The `isIdentical(EJBObject otherEJBObject)` and `isIdentical(EJBLocalObject otherEJBLocalObject)` methods always returns true when used to compare object references of two session objects of the same stateless session bean.

The following example illustrates the use of the `isIdentical` method for a stateless session object.

```
FooHome fooHome = ...; // obtain home of a stateless session bean
Foo foo1 = fooHome.create();
Foo foo2 = fooHome.create();

if (foo1.isIdentical(foo1)) { // this test returns true
    ...
}

if (foo1.isIdentical(foo2)) { // this test returns true
    ...
}
```

6.9.3 `getPrimaryKey()`

The object identifier of a session object is, in general, opaque to the client. The result of `getPrimaryKey()` on a session `EJBObject` reference results in `java.rmi.RemoteException`. The result of `getPrimaryKey()` on a session `EJBLocalObject` reference results in `javax.ejb.EJBException`.

6.10 Type Narrowing

A client program that is intended to be interoperable with all compliant EJB container implementations must use the `javax.rmi.PortableRemoteObject.narrow` method to perform type-narrowing of the client-side representations of the remote home and remote interfaces.

Note: Programs using the cast operator for narrowing the remote and remote home interfaces are likely to fail if the container implementation uses RMI-IIOP as the underlying communication transport.


6.11 The Web Service Client View of a Stateless Session Bean

From the perspective of the client, the existence of the stateless session bean is completely hidden behind the web service endpoint that the bean implements.

The web service client's access to the web service functionality provided by a stateless session bean occurs through a web service endpoint. In the case of Java clients, this endpoint is accessed as a JAX-RPC service endpoint using the JAX-RPC client view APIs, as described in [22].

The following diagram illustrates the view that is provided to J2EE web service clients of a stateless session bean through the JAX-RPC client view APIs.

Figure 5 Web Service Client View of Stateless Session Beans Deployed in a Container



The J2EE web service client uses JNDI to look up the service object that implements the `javax.xml.rpc.Service` interface. The service interface can be a generic `javax.xml.rpc.Service` interface or a generated service interface which extends the `javax.xml.rpc.Service` interface. The service interface is then used to obtain a stub or proxy that implements the stateless session bean's web service endpoint interface. The mechanisms and APIs for client web service access are described in the JAX-RPC specification [22] and in the Web Services for J2EE specification [25].

The following example illustrates how a J2EE client looks up a web service in JNDI using a logical name called a service reference (specified using the `service-ref` element), obtains a stub instance for a web service endpoint, and invokes a method on that endpoint.

```
Context ctx = new InitialContext();
com.example.StockQuoteService sqs = (com.example.StockQuoteService)
    ctx.lookup("java:comp/env/service/StockQuoteService");
com.example.StockQuoteProvider sqp =
    sqs.getStockQuoteProviderPort();
float quotePrice = sqp.getLastTradePrice("ACME");
...
```

The use of service references and the `service-ref` deployment descriptor element is described in further detail in [25]. Section 20.4 provides an overview of the use of the `service-ref` element.

Session Bean Component Contract

This chapter specifies the contract between a session bean and its container. It defines the life cycle of the session bean instances.

This chapter defines the developer's view of session bean state management and the container's responsibilities for managing session bean state.

7.1 Overview

A session bean instance is an instance of the session bean class. It holds the session object's state.

By definition, a session bean instance is an extension of the client that creates it:

- In the case of a stateful session bean, its fields contain a **conversational state** on behalf of the session object's client. This state describes the conversation represented by a specific client/session object pair.
- It typically reads and updates data in a database on behalf of the client. Within a transaction, some of this data may be cached in the instance.
- In the case of a stateful session bean, its lifetime is controlled by the client.

A container may also terminate a session bean instance's life after a deployer-specified time-out or as a result of the failure of the server on which the bean instance is running. For this reason, a client should be prepared to recreate a new session object if it loses the one it is using.

Typically, a session object's conversational state is not written to the database. A session bean developer simply stores it in the session bean instance's fields and assumes its value is retained for the lifetime of the instance.

A session bean must explicitly manage cached database data. A session bean instance must write any cached database updates prior to a transaction completion, and it must refresh its copy of any potentially stale database data at the beginning of the next transaction. A session bean must also refresh any `java.sql` Statement objects before they are used in a new transaction context. A session bean that caches the persistent state of an entity bean with container-managed persistence should refresh this stale state at the beginning of the next transaction. A session bean that caches references to the collection objects that represent container-managed relationships and iterators for container-managed collections must refresh this stale state before accessing these objects in a new transaction context.

7.2 Goals

The container manages the life cycle of the session bean instances. It notifies the instances when bean action may be necessary, and it provides a full range of services to ensure that the session bean implementation is scalable and can support a large number of clients.

The remainder of this section describes the session bean life cycle in detail and the protocol between the bean and its container.

7.3 A Container's Management of its Working Set

To help the container manage its state, a session bean is specified at deployment as having one of the following state management modes:

- *stateless*—the session bean instances contain no conversational state between methods; any instance can be used for any client.
- *stateful*—the session bean instances contain conversational state which must be retained across methods and transactions.

To efficiently manage the size of its working set, a session bean container may need to temporarily transfer the state of an idle *stateful* session bean instance to some form of secondary storage. The transfer from the working set to secondary storage is called instance *passivation*. The transfer back is called *activation*.

A container may only passivate a stateful session bean instance when the instance is *not* in a transaction.

A stateless session bean is never passivated.

7.4 Conversational State of a Stateful Session Bean

The conversational state of a *stateful* session object is defined as the session bean instance's field values, plus the transitive closure of the objects from the instance's fields reached by following Java object references.

In advanced cases, a session object's conversational state may contain open resources, such as open sockets and open database cursors. A container cannot retain such open resources when a session bean instance is passivated. A developer of a stateful session bean must close and open the resources in the `ejbPassivate` and `ejbActivate` methods.

7.4.1 Instance Passivation and Conversational State

The Bean Provider is required to ensure that the `ejbPassivate` method leaves the instance fields ready to be serialized by the container. The objects that are assigned to the instance's non-transient fields after the `ejbPassivate` method completes must be one of the following:

- A serializable object^[5].
- A `null`.
- An enterprise bean's remote interface reference, even if the stub class is not serializable.
- An enterprise bean's remote home interface reference, even if the stub class is not serializable.
- An entity bean's local interface reference, even if it is not serializable.
- An entity bean's local home interface reference, even if it is not serializable.
- A reference to the `SessionContext` object, even if it is not serializable.
- A reference to the environment naming context (that is, the `java:comp/env` JNDI context) or any of its subcontexts.
- A reference to the `UserTransaction` interface.
- A reference to a resource manager connection factory.
- A reference to a `javax.ejb.Timer` object.

[5] Note that the Java Serialization protocol dynamically determines whether or not an object is serializable. This means that it is possible to serialize an object of a serializable subclass of a non-serializable declared field type.

- An object that is not directly serializable, but becomes serializable by replacing the references to an enterprise bean's home and component interfaces, the references to the `SessionContext` object, the references to the `java:comp/env` JNDI context and its subcontexts, and the references to the `UserTransaction` interface by serializable objects during the object's serialization.

This means, for example, that the Bean Provider must close all JDBC™ connections in `ejbPassivate` and assign the instance's fields storing the connections to null.

The last bulleted item covers cases such as storing Collections of component interfaces in the conversational state.

The Bean Provider must assume that the content of transient fields may be lost between the `ejbPassivate` and `ejbActivate` notifications. Therefore, the Bean Provider should not store in a transient field a reference to any of the following objects: `SessionContext` object; environment JNDI naming context and any its subcontexts; home and component interfaces; and the `UserTransaction` interface.

The restrictions on the use of transient fields ensure that containers can use Java Serialization during passivation and activation.

The following are the requirements for the container.

The container performs the Java programming language Serialization (or its equivalent) of the instance's state after it invokes the `ejbPassivate` method on the instance.

The container must be able to properly save and restore the reference to the home and component interfaces of the enterprise beans stored in the instance's state even if the classes that implement the object references are not serializable.

The container must be able to properly save and restore references to timers stored in the instance's state even if the classes that implement the timers are not serializable.

The container may use, for example, the object replacement technique that is part of the `java.io.ObjectOutputStream` and `java.io.ObjectInputStream` protocol to externalize the home and component references.

If the session bean instance stores in its conversational state an object reference to the `javax.ejb.SessionContext` interface passed to the instance in the `setSessionContext` method, the container must be able to save and restore the reference across the instance's passivation. The container can replace the original `SessionContext` object with a different and functionally equivalent `SessionContext` object during activation.

If the session bean instance stores in its conversational state an object reference to the `java:comp/env` JNDI context or its subcontext, the container must be able to save and restore the object reference across the instance's passivation. The container can replace the original object with a different and functionally equivalent object during activation.

If the session bean instance stores in its conversational state an object reference to the `UserTransaction` interface, the container must be able to save and restore the object reference across the instance's passivation. The container can replace the original object with a different and functionally equivalent object during activation.

The container may destroy a session bean instance if the instance does not meet the requirements for serialization after `ejbPassivate`.

While the container is not required to use the Serialization protocol for the Java programming language to store the state of a passivated session instance, it must achieve the equivalent result. The one exception is that containers are not required to reset the value of `transient` fields during activation^[6]. Declaring the session bean's fields as `transient` is, in general, discouraged.

7.4.2 The Effect of Transaction Rollback on Conversational State

A session object's conversational state is not transactional. It is not automatically rolled back to its initial state if the transaction in which the object has participated rolls back.

If a rollback could result in an inconsistency between a session object's conversational state and the state of the underlying database, the bean developer (or the application development tools used by the developer) must use the `afterCompletion` notification to manually reset its state.

7.5 Protocol Between a Session Bean Instance and its Container

Containers themselves make no actual service demands on the session bean instances. The container makes calls on a bean instance to provide it with access to container services and to deliver notifications issued by the container.

7.5.1 The Required SessionBean Interface

All session beans must implement the `SessionBean` interface.

The bean's container calls the `setSessionContext` method to associate a session bean instance with its context maintained by the container. Typically, a session bean instance retains its session context as part of its conversational state.

The `ejbRemove` notification signals that the instance is in the process of being removed by the container. In the `ejbRemove` method, the instance typically releases the same resources that it releases in the `ejbPassivate` method.

[6] This is to allow the container to swap out an instance's state through techniques other than the Java Serialization protocol. For example, the container's Java Virtual Machine implementation may use a block of memory to keep the instance's variables, and the container swaps the whole memory block to the disk instead of performing Java Serialization on the instance.

The `ejbPassivate` notification signals the intent of the container to passivate the instance. The `ejbActivate` notification signals the instance it has just been reactivated. Because containers automatically maintain the conversational state of a stateful session bean instance when it is passivated, most session beans can ignore these notifications. Their purpose is to allow stateful session beans to maintain those open resources that need to be closed prior to an instance's passivation and then reopened during an instance's activation. The `ejbPassivate` and `ejbActivate` methods are only called on stateful session bean instances.

7.5.2 The SessionContext Interface

A container provides the session bean instances with a `SessionContext`, which gives the session bean instance access to the instance's context maintained by the container. The `SessionContext` interface has the following methods:

- The `getEJBObject` method returns the session bean's remote interface.
- The `getEJBHome` method returns the session bean's remote home interface.
- The `getEJBLocalObject` method returns the session bean's local interface.
- The `getEJBLocalHome` method returns the session bean's local home interface.
- The `getCallerPrincipal` method returns the `java.security.Principal` that identifies the invoker.
- The `isCallerInRole` method tests if the session bean instance's caller has a particular role.
- The `setRollbackOnly` method allows the instance to mark the current transaction such that the only outcome of the transaction is a rollback. Only instances of a session bean with container-managed transaction demarcation can use this method.
- The `getRollbackOnly` method allows the instance to test if the current transaction has been marked for rollback. Only instances of a session bean with container-managed transaction demarcation can use this method.
- The `getUserTransaction` method returns the `javax.transaction.UserTransaction` interface. The instance can use this interface to demarcate transactions and to obtain transaction status. Only instances of a session bean with bean-managed transaction demarcation can use this method.
- The `getTimerService` method returns the `javax.ejb.TimerService` interface. Only stateless session beans can use this method. Stateful session beans cannot be timed objects.
- The `getMessageContext` method returns the `javax.xml.rpc.handler.MessageContext` interface. Only stateless session beans with web service endpoint interfaces can use this method.

7.5.3 Use of the MessageContext Interface by Stateless Session Beans

A stateless session bean that implements a web service endpoint may access the JAX-RPC MessageContext interface by means of the `SessionContext.getMessageContext` method. The MessageContext interface allows the stateless session bean instance to see the SOAP message for the web service endpoint, as well as the properties set by the JAX-RPC SOAP message handlers, if any. The stateless session bean may use the MessageContext interface to set properties for the JAX-RPC message response handlers, if any.

7.5.4 The Optional SessionSynchronization Interface for Stateful Session Beans

A stateful session bean class can optionally implement the `javax.ejb.SessionSynchronization` interface. This interface provides the session bean instances with transaction synchronization notifications. The instances can use these notifications, for example, to manage database data they may cache within transactions.

The `afterBegin` notification signals a session bean instance that a new transaction has begun. The container invokes this method before the first business method within a transaction (which is not necessarily at the beginning of the transaction). The `afterBegin` notification is invoked with the transaction context. The instance may do any database work it requires within the scope of the transaction.

The `beforeCompletion` notification is issued when a session bean instance's client has completed work on its current transaction but prior to committing the resource managers used by the instance. At this time, the instance should write out any database updates it has cached. The instance can cause the transaction to roll back by invoking the `setRollbackOnly` method on its session context.

The `afterCompletion` notification signals that the current transaction has completed. A completion status of `true` indicates that the transaction has committed. A status of `false` indicates that a rollback has occurred. Since a session bean instance's conversational state is not transactional, it may need to manually reset its state if a rollback occurred.

All container providers must support SessionSynchronization. It is optional only for the bean implementor. If a bean class implements SessionSynchronization, the container must invoke the `afterBegin`, `beforeCompletion`, and `afterCompletion` notifications as required by the specification.

Only a stateful session bean with container-managed transaction demarcation may implement the SessionSynchronization interface. A stateless session bean must not implement the SessionSynchronization interface.

There is no need for a session bean with bean-managed transaction demarcation to rely on the synchronization call backs because the bean is in control of the commit—the bean knows when the transaction is about to be committed and it knows the outcome of the transaction commit.

7.5.5 The Optional TimedObject Interface for Stateless Session Beans

A stateless session bean can be registered with the EJB Timer Service for time-based event notifications if it implements the `javax.ejb.TimedObject` interface. The container invokes the bean instance's `ejbTimeout` method when a timer for the bean has expired. See Chapter 22, "Timer Service". Stateful session beans must not implement the `javax.ejb.TimedObject` interface.

7.5.6 Business Method Delegation

The session bean's component or web service endpoint interface defines the business methods callable by a client.

The session bean's remote interface is implemented by the session `EJBObject` class generated by the container tools; the session bean's local interface is implemented by the session `EJBLocalObject` class generated by the container tools. The class that handles requests to a web service endpoint is generated by the container tools. The session `EJBObject` or `EJBLocalObject` class delegates an invocation of a business method to the matching business method that is implemented in the session bean class. The class that handles requests to the web service endpoint invokes the stateless session bean method that matches the web service endpoint interface method corresponding to the SOAP request.

7.5.7 Session Bean's `ejbCreate<METHOD>` Methods

A local or remote client creates a session bean instance using one of the `create<METHOD>` methods defined in the session bean's home interface. The session bean's home interface is provided by the bean developer. Its implementation is generated by the deployment tools provided by the Container Provider. A stateless session bean that provides only a web service client view has no home interface.

The container creates an instance of a session bean in three steps. First, the container calls the bean class's `newInstance` method to create a new session bean instance. Second, the container calls the `setSessionContext` method to pass the context object to the instance. Third, the container calls the instance's `ejbCreate<METHOD>` method whose signature matches the signature of the `create<METHOD>` method invoked by the client. The input parameters sent from the client are passed to the `ejbCreate<METHOD>` method.

Each stateful session bean class must have at least one `ejbCreate<METHOD>` method. The number and signatures of a session bean's `create<METHOD>` methods are specific to each session bean class. Since a stateful session bean represents a specific, private conversation between the bean and its client, its `create` parameters typically contain the information the client uses to customize the bean instance for its use.

A stateless session bean that has a local or remote client view has a single `create` method. It and the corresponding `ejbCreate` method have no arguments. Since stateless session bean instances are typically pooled, the time of the client's invocation of the `create` method may not be directly related to the container's invocation of the `ejbCreate` method on the stateless session bean instance.

A stateless session bean that provides only a web service client view has a single `ejbCreate` method, and no corresponding `create` method. The `ejbCreate` method is invoked by the container when it needs to create a new session bean instance in order to service a client request.

A stateless session bean must always have an `ejbCreate` method.

7.5.8 Serializing Session Bean Methods

A container serializes calls to each session bean instance. Most containers will support many instances of a session bean executing concurrently; however, each instance sees only a serialized sequence of method calls. Therefore, a session bean does not have to be coded as reentrant.

The container must serialize all the container-invoked callbacks (that is, the methods `ejbPassivate`, `ejbTimeout`, `beforeCompletion`, and so on), and it must serialize these callbacks with the client-invoked business method calls.

Clients are not allowed to make concurrent calls to a stateful session object. If a client-invoked business method is in progress on an instance when another client-invoked call, from the same or different client, arrives at the same instance of a stateful session bean class, the container may throw the `java.rmi.RemoteException` to the second client^[7] if the client is a remote client, or the `javax.ejb.EJBException` if the client is a local client. This restriction does not apply to a stateless session bean because the container routes each request to a different instance of the session bean class.

7.5.9 Transaction Context of Session Bean Methods

The implementation of a business method defined in the component interface or web service endpoint interface or an `ejbTimeout` method is invoked in the scope of a transaction determined by the transaction attribute specified in the deployment descriptor.

A session bean's `afterBegin` and `beforeCompletion` methods are always called with the same transaction context as the business methods executed between the `afterBegin` and `beforeCompletion` methods.

A session bean's `newInstance`, `setSessionContext`, `ejbCreate`, `ejbRemove`, `ejbPassivate`, `ejbActivate`, and `afterCompletion` methods are called with an unspecified transaction context. Refer to Subsection 17.6.5 for how the container executes methods with an unspecified transaction context.

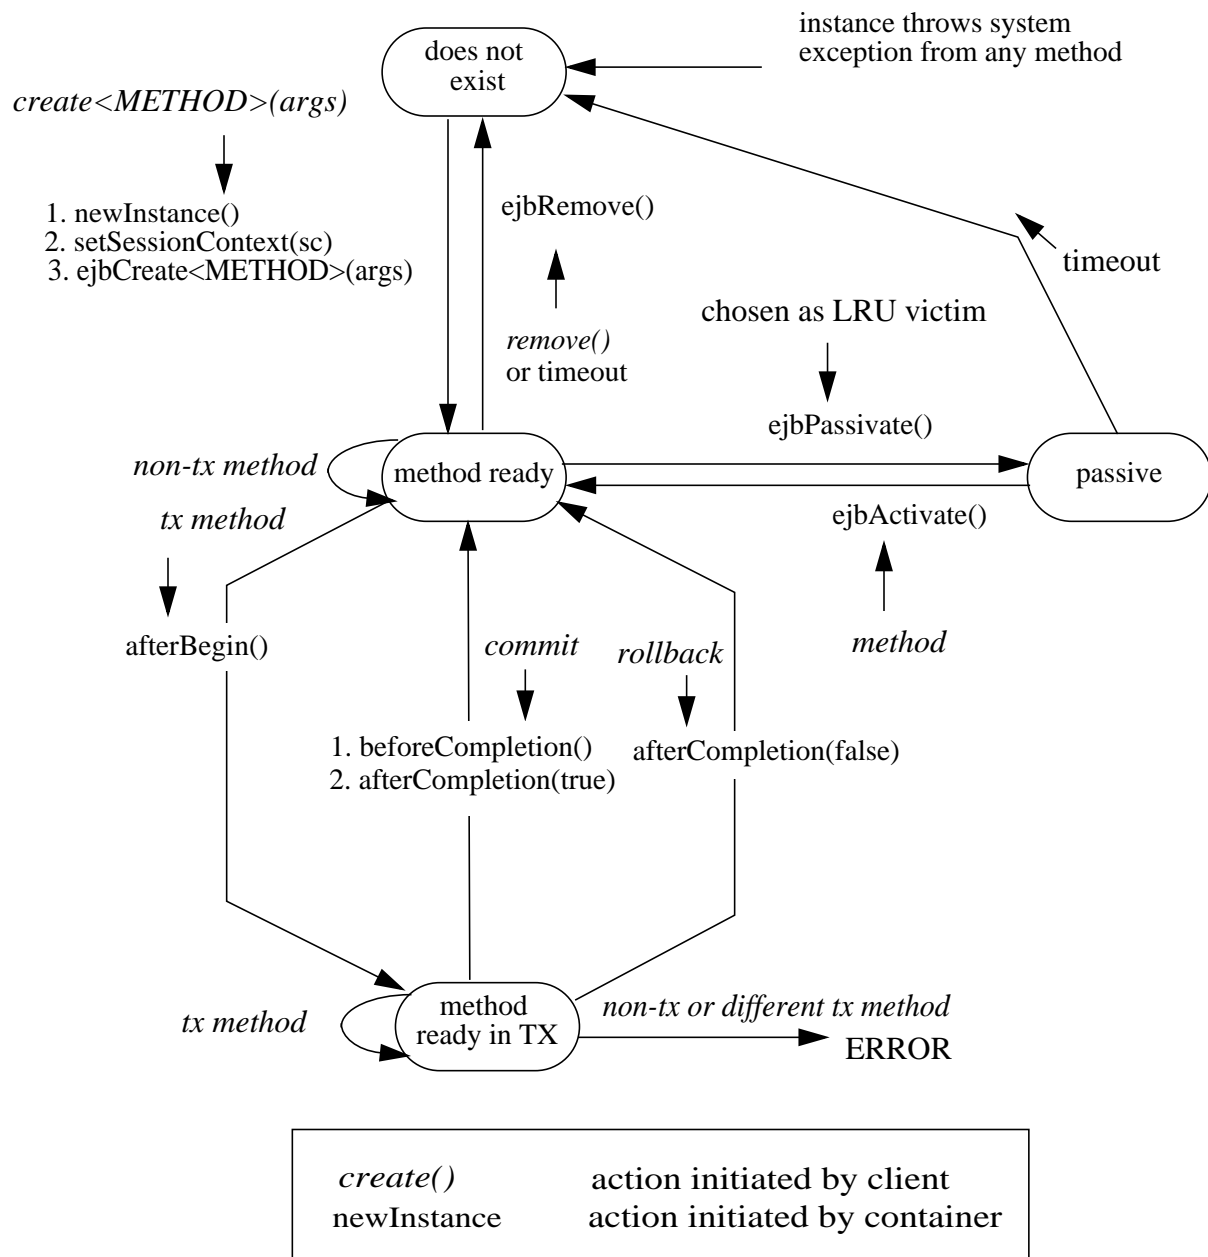
For example, it would be wrong to perform database operations within a session bean's `ejbCreate` or `ejbRemove` methods and to assume that the operations are part of the client's transaction. The `ejbCreate` and `ejbRemove` methods are not controlled by a transaction attribute because handling rollbacks in these methods would greatly complicate the session instance's state diagram.

[7] In certain special circumstances (e.g., to handle clustered web container architectures), the container may instead queue or serialize such concurrent requests. Clients, however, cannot rely on this behavior.

7.6 Stateful Session Bean State Diagram

The following figure illustrates the life cycle of a stateful session bean instance.

Figure 6 Life Cycle of a Stateful Session Bean Instance



The following steps describe the life cycle of a stateful session bean instance:

- A session bean instance's life starts when a client invokes a `create<METHOD>` method on the session bean's home interface. This causes the container to invoke `newInstance` on the session bean class to create a new session bean instance. Next, the container calls `setSessionContext` and `ejbCreate<METHOD>` on the instance and returns the session object reference to the client. The instance is now in the method ready state.
- The session bean instance is now ready for client's business methods. Based on the transaction attributes in the session bean's deployment descriptor and the transaction context associated with the client's invocation, a business method is executed either in a transaction context or with an unspecified transaction context (shown as "tx method" and "non-tx method" in the diagram). See Chapter 17 for how the container deals with transactions.
- A non-transactional method is executed while the instance is in the method ready state.
- An invocation of a transactional method causes the instance to be included in a transaction. When the session bean instance is included in a transaction, the container issues the `afterBegin` method on it. The `afterBegin` method is invoked on the instance before any business method is executed as part of the transaction. The instance becomes associated with the transaction and will remain associated with the transaction until the transaction completes.
- Session bean methods invoked by the client in this transaction can now be delegated to the bean instance. An error occurs if a client attempts to invoke a method on the session object and the deployment descriptor for the method requires that the container invoke the method in a different transaction context than the one with which the instance is currently associated or in an unspecified transaction context.
- If a transaction commit has been requested, the transaction service notifies the container of the commit request before actually committing the transaction, and the container issues a `beforeCompletion` on the instance. When `beforeCompletion` is invoked, the instance should write any cached updates to the database. If a transaction rollback had been requested instead, the rollback status is reached without the container issuing a `beforeCompletion`. The container may not call the `beforeCompletion` method if the transaction has been marked for rollback (nor does the instance write any cached updates to the database).
- The transaction service then attempts to commit the transaction, resulting in either a commit or rollback.
- When the transaction completes, the container issues `afterCompletion` on the instance, specifying the status of the completion (either commit or rollback). If a rollback occurred, the bean instance may need to reset its conversational state back to the value it had at the beginning of the transaction.
- The container's caching algorithm may decide that the bean instance should be evicted from memory. (This could be done at the end of each method, or by using an LRU policy). The container issues `ejbPassivate` on the instance. After this completes, the container saves the instance's state to secondary storage. A session bean can be passivated only between transactions, and not within a transaction.

- While the instance is in the passivated state, the container may remove the session object after the expiration of a timeout specified by the Deployer. All object references and handles for the session object become invalid. If a client attempts to invoke the session object, the container will throw the `java.rmi.NoSuchObjectException` if the client is a remote client, or the `javax.ejb.NoSuchObjectLocalException` if the client is a local client.
- If a client invokes a session object whose session bean instance has been passivated, the container will activate the instance. To activate the session bean instance, the container restores the instance's state from secondary storage and issues `ejbActivate` on it.
- The session bean instance is again ready for client methods.
- When the client calls `remove` on the home or component interface to remove the session object, the container invokes `ejbRemove` on the bean instance. This ends the life of the session bean instance and the associated session object. Any subsequent attempt by its client to invoke the session object causes the `java.rmi.NoSuchObjectException` to be thrown if the client is a remote client, or the `javax.ejb.NoSuchObjectLocalException` if the client is a local client. (The `java.rmi.NoSuchObjectException` is a subclass of the `java.rmi.RemoteException`; the `javax.ejb.NoSuchObjectLocalException` is a subclass of the `javax.ejb.EJBException`). The `ejbRemove` method cannot be called when the instance is participating in a transaction. An attempt to remove a session object while the object is in a transaction will cause the container to throw the `javax.ejb.RemoveException` to the client. Note that a container can also invoke the `ejbRemove` method on the instance without a client call to remove the session object after the lifetime of the EJB object has expired.

The container must call the `afterBegin`, `beforeCompletion`, and `afterCompletion` methods if the session bean class implements, directly or indirectly, the `SessionSynchronization` interface. The container does not call these methods if the session bean class does not implement the `SessionSynchronization` interface.

7.6.1 Operations Allowed in the Methods of a Stateful Session Bean Class

Table 2 defines the methods of a stateful session bean class from which the session bean instances can access the methods of the `javax.ejb.SessionContext` interface, the `java:comp/env` environment naming context, resource managers, Timer methods, and other enterprise beans.

If a session bean instance attempts to invoke a method of the `SessionContext` interface, and that access is not allowed in Table 2, the container must throw the `java.lang.IllegalStateException`.

If a session bean instance attempts to access a resource manager or an enterprise bean, and that access is not allowed in Table 2, the behavior is undefined by the EJB architecture.

If a session bean instance attempts to invoke a method of the `Timer` interface and the access is not allowed in Table 2, the container must throw the `java.lang.IllegalStateException`.

Table 2 Operations Allowed in the Methods of a Stateful Session Bean

Bean method	Bean method can perform the following operations	
	Container-managed transaction demarcation	Bean-managed transaction demarcation
constructor	-	-
setSessionContext	SessionContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> JNDI access to java:comp/env	SessionContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> JNDI access to java:comp/env
ejbCreate ejbRemove ejbActivate ejbPassivate	SessionContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getCallerPrincipal</i> , <i>isCallerInRole</i> , <i>getEJBObject</i> , <i>getEJBLocalObject</i> JNDI access to java:comp/env Resource manager access Enterprise bean access	SessionContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getCallerPrincipal</i> , <i>isCallerInRole</i> , <i>getEJBObject</i> , <i>getEJBLocalObject</i> , <i>getUserTransaction</i> UserTransaction methods JNDI access to java:comp/env Resource manager access Enterprise bean access
business method from component interface	SessionContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getCallerPrincipal</i> , <i>getRollbackOnly</i> , <i>isCallerInRole</i> , <i>setRollbackOnly</i> , <i>getEJBObject</i> , <i>getEJBLocalObject</i> JNDI access to java:comp/env Resource manager access Enterprise bean access Timer methods	SessionContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getCallerPrincipal</i> , <i>isCallerInRole</i> , <i>getEJBObject</i> , <i>getEJBLocalObject</i> , <i>getUserTransaction</i> UserTransaction methods JNDI access to java:comp/env Resource manager access Enterprise bean access Timer methods
afterBegin beforeCompletion	SessionContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getCallerPrincipal</i> , <i>getRollbackOnly</i> , <i>isCallerInRole</i> , <i>setRollbackOnly</i> , <i>getEJBObject</i> , <i>getEJBLocalObject</i> JNDI access to java:comp/env Resource manager access Enterprise bean access Timer methods	N/A (a bean with bean-managed transaction demarcation cannot implement the SessionSynchronization interface)
afterCompletion	SessionContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getCallerPrincipal</i> , <i>isCallerInRole</i> , <i>getEJBObject</i> , <i>getEJBLocalObject</i> JNDI access to java:comp/env	

Notes:

- The `ejbCreate<METHOD>`, `ejbRemove`, `ejbPassivate`, and `ejbActivate` methods of a session bean with container-managed transaction demarcation execute with an unspecified transaction context. Refer to Subsection 17.6.5 for how the container executes methods with an unspecified transaction context.

Additional restrictions:

- The `getRollbackOnly` and `setRollbackOnly` methods of the `SessionContext` interface should be used only in the session bean methods that execute in the context of a transaction. The container must throw the `java.lang.IllegalStateException` if the methods are invoked while the instance is not associated with a transaction.

The reasons for disallowing the operations in Table 2 follow:

- Invoking the `getEJBObject` and `getEJBLocalObject` methods is disallowed in the session bean methods in which there is no session object identity established for the instance.
- Invoking the `getEJBObject` and `getEJBHome` methods is disallowed if the session bean does not define a remote client view.
- Invoking the `getEJBLocalObject` and `getEJBLocalHome` methods is disallowed if the session bean does not define a local client view.
- Invoking the `getRollbackOnly` and `setRollbackOnly` methods is disallowed in the session bean methods for which the container does not have a meaningful transaction context, and to all session beans with bean-managed transaction demarcation.
- Accessing resource managers and enterprise beans is disallowed in the session bean methods for which the container does not have a meaningful transaction context or client security context.
- The `UserTransaction` interface is unavailable to enterprise beans with container-managed transaction demarcation.
- The `TimerService` interface is unavailable to stateful session beans.
- Invoking the `getMessageContext` method is disallowed for stateful session beans.

7.6.2 Dealing with Exceptions

A `RuntimeException` thrown from any method of the session bean class (including the business methods and the callbacks invoked by the container) results in the transition to the “does not exist” state. Exception handling is described in detail in Chapter 18.

From the client perspective, the corresponding session object does not exist any more. Subsequent invocations through the remote interface will result in `java.rmi.NoSuchObjectException`; subsequent invocations through the local interface will result in `javax.ejb.NoSuchObjectLocalException`.

7.6.3 Missed ejbRemove Calls

The Bean Provider cannot assume that the container will always invoke the `ejbRemove` method on a session bean instance. The following scenarios result in `ejbRemove` not being called on an instance:

- A crash of the EJB container.
- A system exception thrown from the instance's method to the container.
- A timeout of client inactivity while the instance is in the `passive` state. The timeout is specified by the Deployer in an EJB container implementation-specific way.

If the session bean instance allocates resources in the `ejbCreate<METHOD>` method and/or in the business methods, and normally releases the resources in the `ejbRemove` method, these resources will not be automatically released in the above scenarios. The application using the session bean should provide some clean up mechanism to periodically clean up the unreleased resources.

For example, if a shopping cart component is implemented as a session bean, and the session bean stores the shopping cart content in a database, the application should provide a program that runs periodically and removes "abandoned" shopping carts from the database.

7.6.4 Restrictions for Transactions

The state diagram implies the following restrictions on transaction scoping of the client invoked business methods. The restrictions are enforced by the container and must be observed by the client programmer.

- A session bean instance can participate in at most a single transaction at a time.
- If a session bean instance is participating in a transaction, it is an error for a client to invoke a method on the session object such that the transaction attribute in the deployment descriptor would cause the container to execute the method in a different transaction context or in an unspecified transaction context. In such a case, the container throws the `java.rmi.RemoteException` to the client if the client is a remote client, or the `javax.ejb.EJBException` if the client is a local client.
- If a session bean instance is participating in a transaction, it is an error for a client to invoke the `remove` method on the session object's home or component interface object. The container must detect such an attempt and throw the `javax.ejb.RemoveException` to the client. The container should not mark the client's transaction for rollback, thus allowing the client to recover.

7.7 Object Interaction Diagrams for a Stateful Session Bean

This section contains object interaction diagrams (OID) that illustrates the interaction of the classes.

7.7.1 Notes

The object interaction diagrams illustrate a box labeled “container-provided classes.” These are either classes that are part of the container, or classes that were generated by the container tools. These classes communicate with each other through protocols that are container implementation-specific. Therefore, the communication between these classes is not shown in the diagrams.

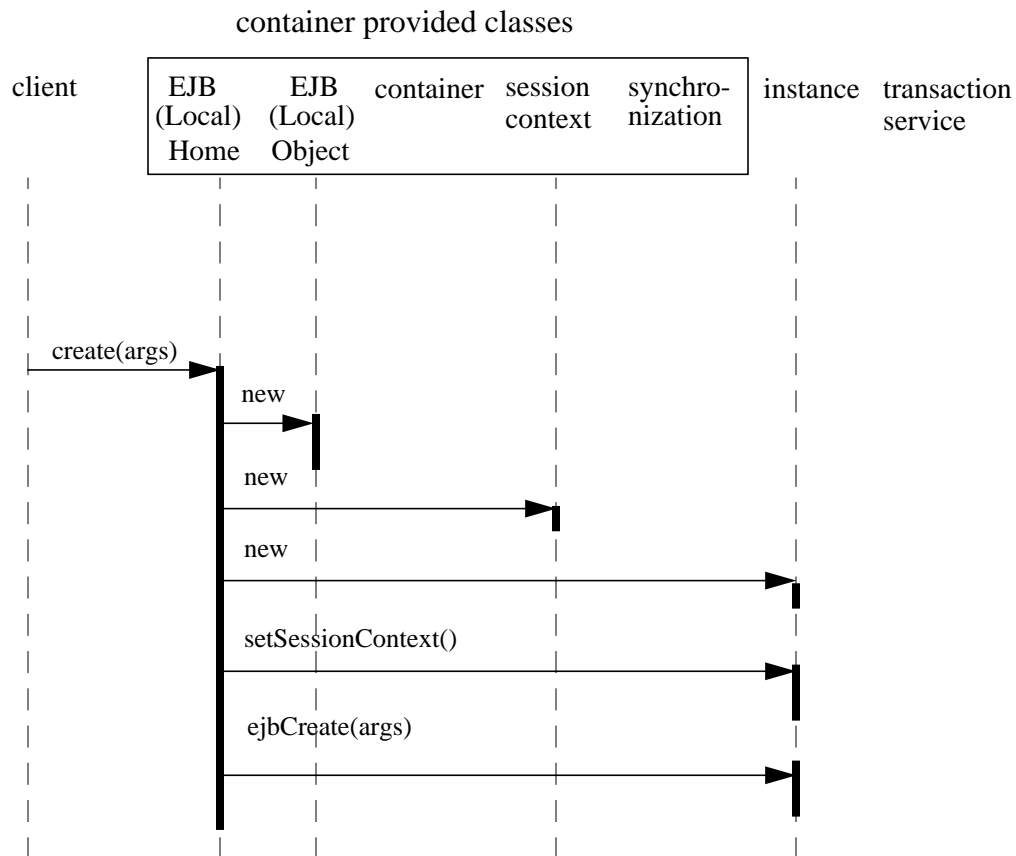
The classes shown in the diagrams should be considered as an illustrative implementation rather than as a prescriptive one.

7.7.2 Creating a Session Object

The following diagram illustrates the creation of a session object.

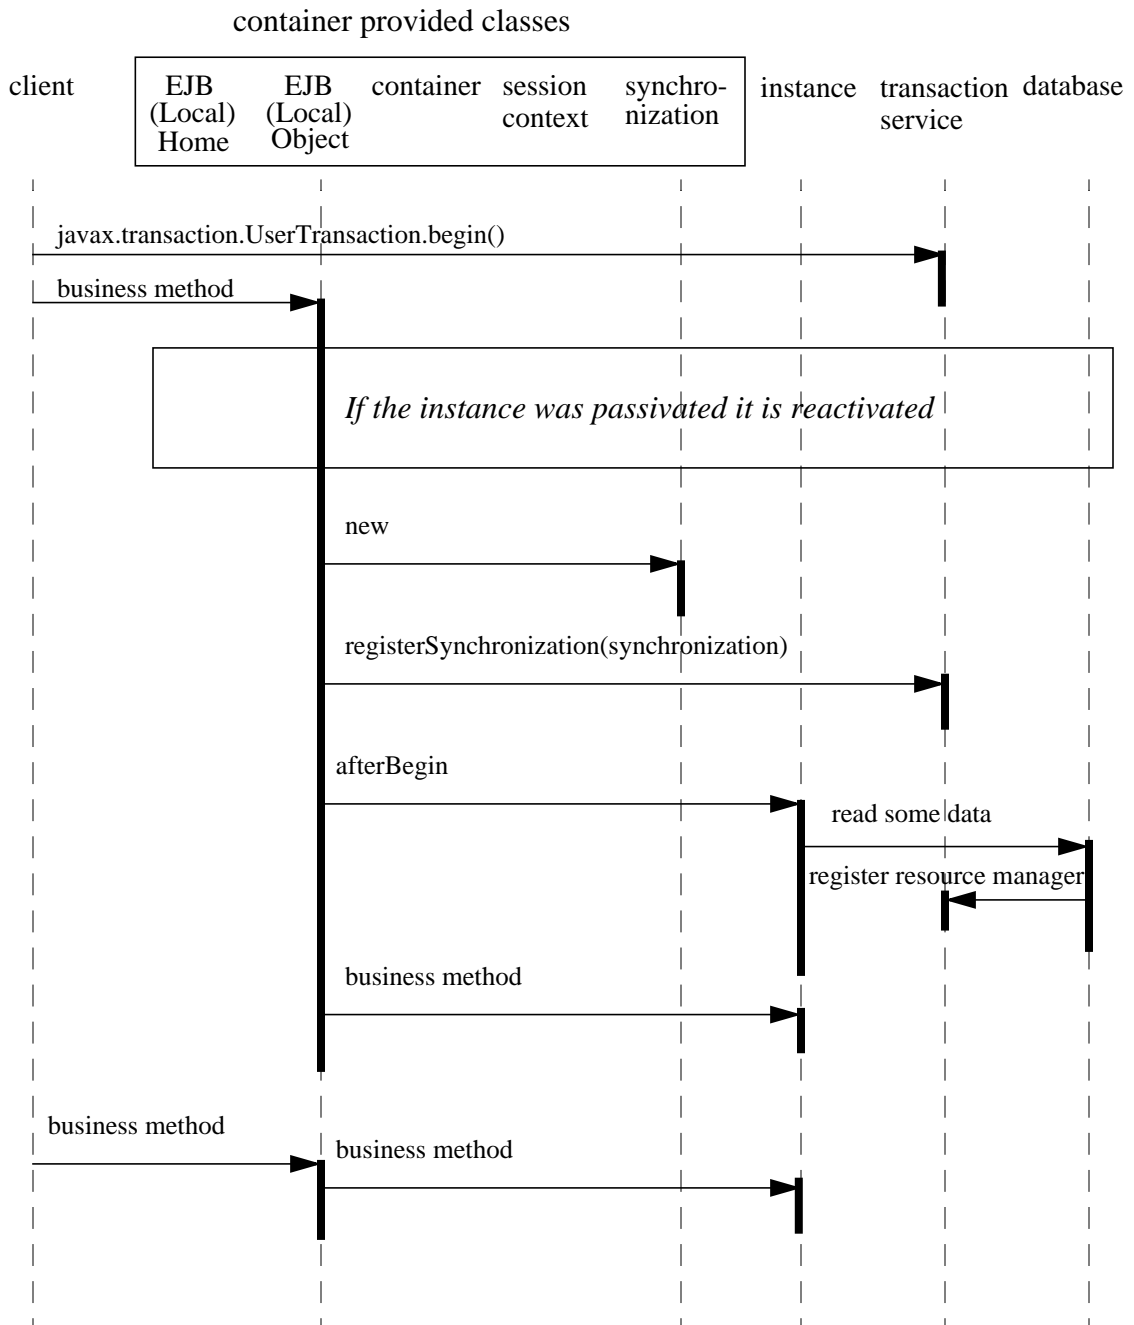
Figure 7

OID for Creation of a Session Object



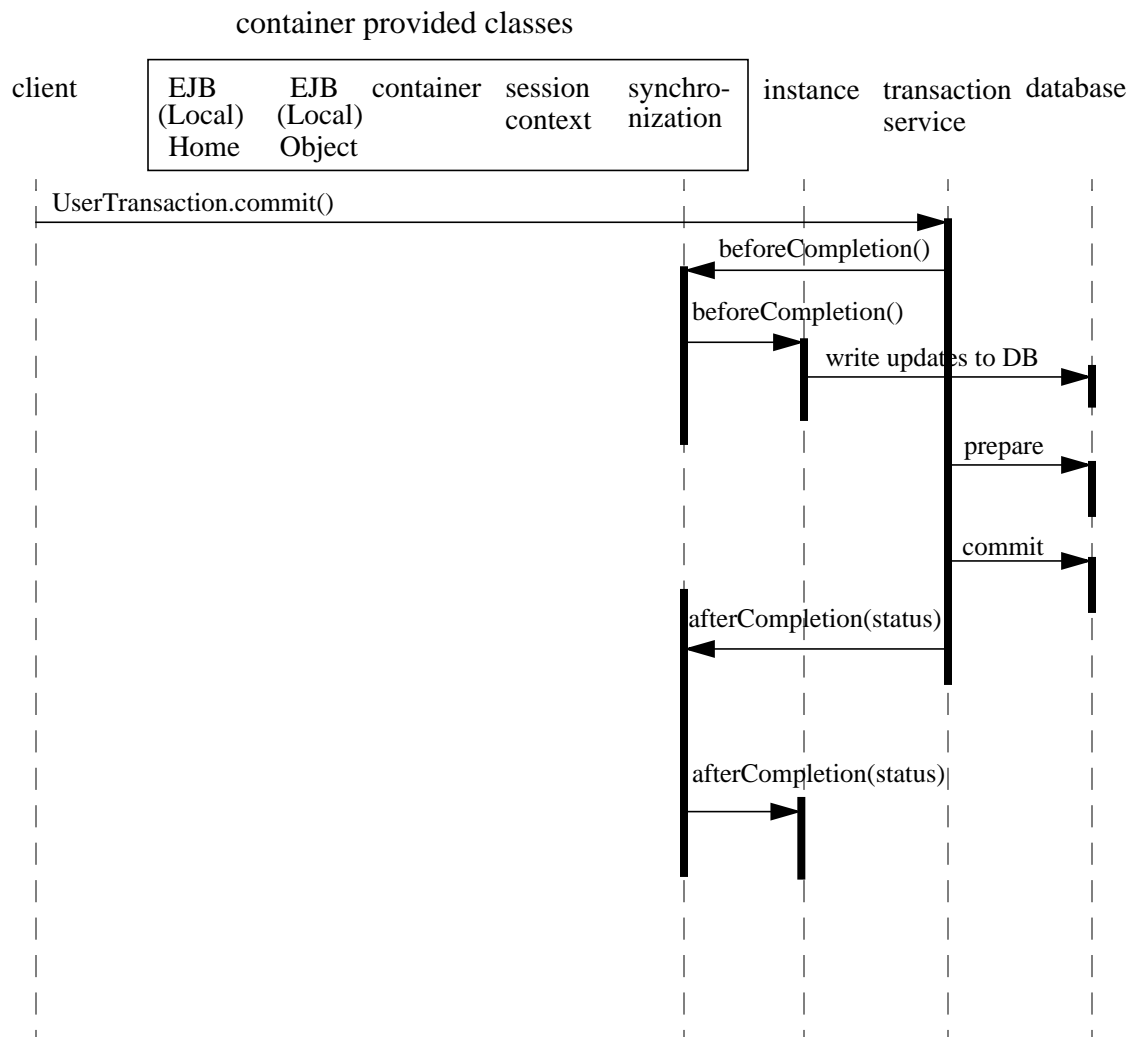
7.7.3 Starting a Transaction

The following diagram illustrates the protocol performed at the beginning of a transaction.

Figure 8 OID for Session Object at Start of a Transaction.

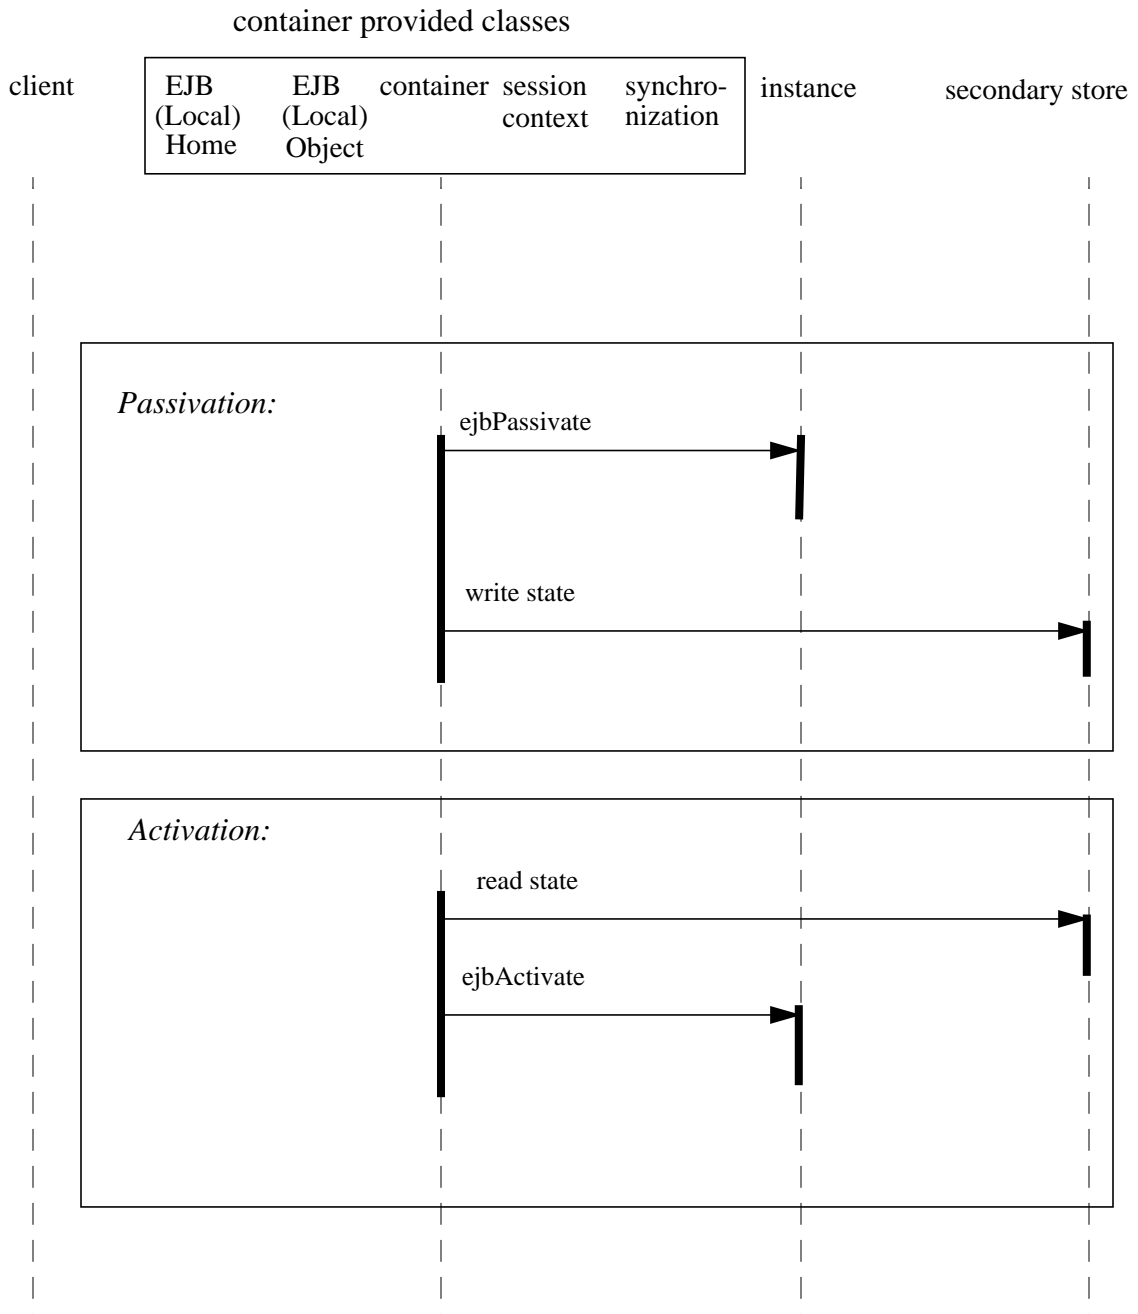
7.7.4 Committing a Transaction

The following diagram illustrates the transaction synchronization protocol for a session object.

Figure 9 OID for Session Object Transaction Synchronization

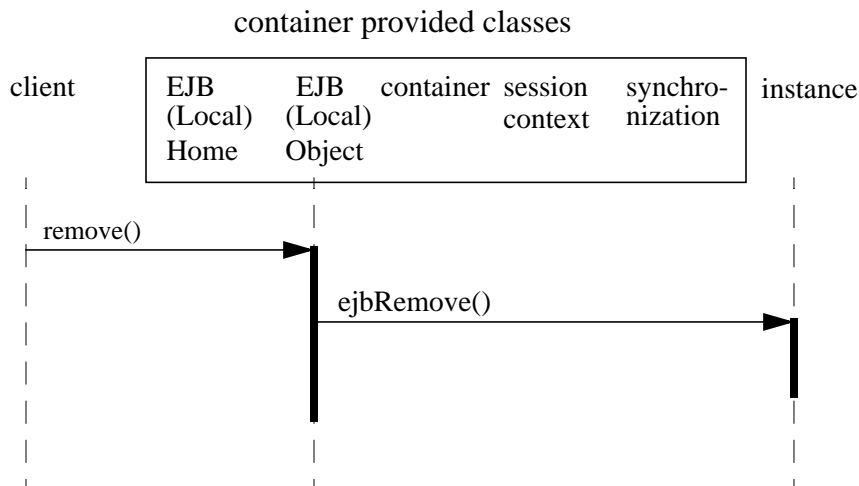
7.7.5 Passivating and Activating an Instance Between Transactions

The following diagram illustrates the passivation and reactivation of a session bean instance. Passivation typically happens spontaneously based on the needs of the container. Activation typically occurs when a client calls a method.

Figure 10 OID for Passivation and Activation of a Session Object

7.7.6 Removing a Session Object

The following diagram illustrates the removal of a session object.

Figure 11 OID for the Removal of a Session Object

7.8 Stateless Session Beans

Stateless session beans are session beans whose instances have no conversational state. This means that all bean instances are equivalent when they are not involved in servicing a client-invoked method.

The term “stateless” signifies that an instance has no state for a specific client. However, the instance variables of the instance can contain the state across client-invoked method calls. Examples of such states include an open database connection and an object reference to an EJB object.

The Bean Provider must exercise caution if retaining any application state across method calls. In particular, references to common bean state should not be returned through multiple local interface method calls.

The home interface of a stateless session bean must have one `create` method that takes no arguments. The `create` method of the remote home interface must return the session bean’s remote interface. The `create` method of the local home interface must return the session bean’s local interface. There can be no other `create` methods in the home interface. The session bean class must define a single `ejbCreate` method that takes no arguments, regardless of whether the session bean defines a home interface.

Because all instances of a stateless session bean are equivalent, the container can choose to delegate a client-invoked method to any available instance. This means, for example, that the container may delegate the requests from the same client within the same transaction to different instances, and that the container may interleave requests from multiple transactions to the same instance.

A container only needs to retain the number of instances required to service the current client load. Due to client “think time,” this number is typically much smaller than the number of active clients. Passivation is not needed for stateless sessions. The container creates another stateless session bean instance if one is needed to handle an increase in client work load. If a stateless session bean is not needed to handle the current client work load, the container can destroy it.

Because stateless session beans minimize the resources needed to support a large population of clients, depending on the implementation of the container, applications that use stateless session beans may scale somewhat better than those using stateful session beans. However, this benefit may be offset by the increased complexity of the client application that uses the stateless beans.

Local and remote clients use the `create` and `remove` methods on the home interface of a stateless session bean in the same way as on a stateful session bean. To the client, it appears as if the client controls the life cycle of the session object. However, the container handles the `create` and `remove` calls without necessarily creating and removing an EJB instance.

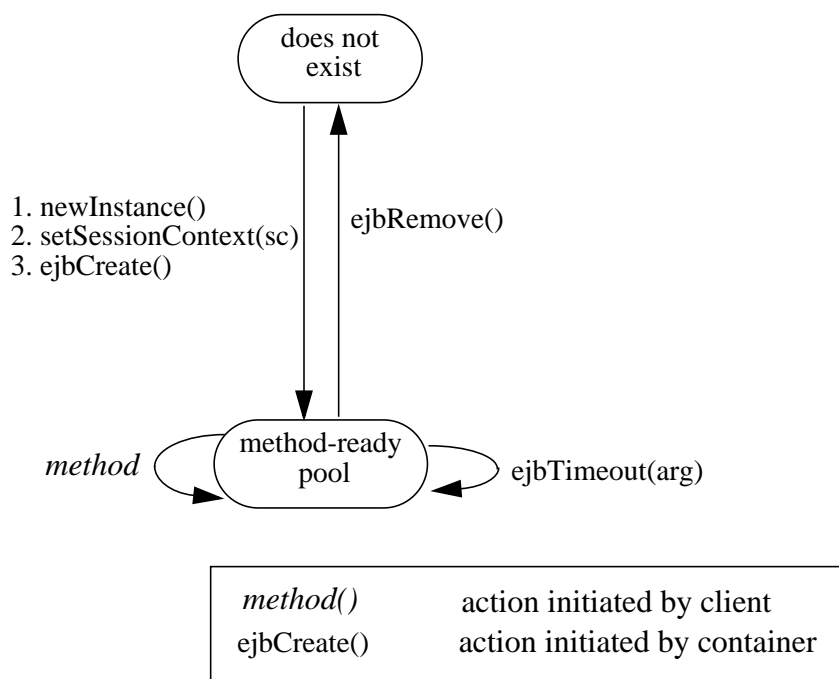
There is no fixed mapping between clients and stateless instances. The container simply delegates a client’s work to any available instance that is method-ready.

A stateless session bean must not implement the `javax.ejb.SessionSynchronization` interface.

7.8.1 Stateless Session Bean State Diagram

When a client calls a method on a stateless session object or invokes a method on a stateless session bean through its web service client view, the container selects one of its method-ready instances and delegates the method invocation to it.

The following figure illustrates the life cycle of a *stateless* session bean instance.

Figure 12 Life Cycle of a Stateless Session Bean

The following steps describe the life cycle of a session bean instance:

- A stateless session bean instance's life starts when the container invokes the `newInstance` method on the session bean class to create a new instance. Next, the container calls `setSessionContext` followed by `ejbCreate` on the instance. The container can perform the instance creation at any time—there is no relationship to a client's invocation of the `create` method.
- The session bean instance is now ready to be delegated a business method call from any client or a call from the container to the `ejbTimeout` method.
- When the container no longer needs the instance (usually when the container wants to reduce the number of instances in the method-ready pool), the container invokes `ejbRemove` on it. This ends the life of the stateless session bean instance.

7.8.2 Operations Allowed in the Methods of a Stateless Session Bean Class

Table 3 defines the methods of a stateless session bean class in which the session bean instances can access the methods of the `javax.ejb.SessionContext` interface, the `java:comp/env` environment naming context, resource managers, `TimerService` and `Timer` methods, and other enterprise beans.

If a session bean instance attempts to invoke a method of the `SessionContext` interface, and the access is not allowed in Table 3, the container must throw the `java.lang.IllegalStateException`.

If a session bean instance attempts to invoke a method of the `TimerService` or `Timer` interface and the access is not allowed in Table 3, the container must throw the `java.lang.IllegalStateException`.

If a session bean instance attempts to access a resource manager or an enterprise bean and the access is not allowed in Table 3, the behavior is undefined by the EJB architecture.

Table 3 Operations Allowed in the Methods of a Stateless Session Bean

Bean method	Bean method can perform the following operations	
	Container-managed transaction demarcation	Bean-managed transaction demarcation
constructor	-	-
setSessionContext	SessionContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> JNDI access to java:comp/env	SessionContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> JNDI access to java:comp/env
ejbCreate ejbRemove	SessionContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getEJBObject</i> , <i>getEJBLocalObject</i> , <i>getTimerService</i> JNDI access to java:comp/env	SessionContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getEJBObject</i> , <i>getEJBLocalObject</i> , <i>getUserTransaction</i> , <i>getTimerService</i> JNDI access to java:comp/env
business method from component interface	SessionContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getCallerPrincipal</i> , <i>isCallerInRole</i> , <i>getRollbackOnly</i> , <i>setRollbackOnly</i> , <i>getEJBObject</i> , <i>getEJBLocalObject</i> , <i>getTimerService</i> JNDI access to java:comp/env Resource manager access Enterprise bean access TimerService and Timer methods	SessionContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getCallerPrincipal</i> , <i>isCallerInRole</i> , <i>getEJBObject</i> , <i>getEJBLocalObject</i> , <i>getUserTransaction</i> , <i>getTimerService</i> UserTransaction methods JNDI access to java:comp/env Resource manager access Enterprise bean access TimerService and Timer methods
business method from web service endpoint interface	SessionContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getCallerPrincipal</i> , <i>isCallerInRole</i> , <i>getRollbackOnly</i> , <i>setRollbackOnly</i> , <i>getEJBObject</i> , <i>getEJBLocalObject</i> , <i>getTimerService</i> , <i>getMessageContext</i> Message context methods JNDI access to java:comp/env Resource manager access Enterprise bean access TimerService and Timer methods	SessionContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getCallerPrincipal</i> , <i>isCallerInRole</i> , <i>getEJBObject</i> , <i>getEJBLocalObject</i> , <i>getUserTransaction</i> , <i>getTimerService</i> , <i>getMessageContext</i> UserTransaction methods Message context methods JNDI access to java:comp/env Resource manager access Enterprise bean access TimerService and Timer methods

Table 3 Operations Allowed in the Methods of a Stateless Session Bean

Bean method	Bean method can perform the following operations	
	Container-managed transaction demarcation	Bean-managed transaction demarcation
ejbTimeout	SessionContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getCallerPrincipal</i> , <i>isCallerInRole</i> , <i>getRollbackOnly</i> , <i>setRollbackOnly</i> , <i>getEJBObject</i> , <i>getEJBLocalObject</i> , <i>getTimerService</i> JNDI access to java:comp/env Resource manager access Enterprise bean access TimerService and Timer methods	SessionContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getCallerPrincipal</i> , <i>isCallerInRole</i> , <i>getEJBObject</i> , <i>getEJBLocalObject</i> , <i>getUserTransaction</i> , <i>getTimerService</i> UserTransaction methods JNDI access to java:comp/env Resource manager access Enterprise bean access TimerService and Timer methods

Additional restrictions:

- The `getRollbackOnly` and `setRollbackOnly` methods of the `SessionContext` interface should be used only in the session bean methods that execute in the context of a transaction. The container must throw the `java.lang.IllegalStateException` if the methods are invoked while the instance is not associated with a transaction.

The reasons for disallowing operations in Table 3:

- Invoking the `getEJBObject` and `getEJBLocalObject` methods is disallowed in the session bean methods in which there is no session object identity associated with the instance. This applies to all methods of a stateless session bean that provides only a web service client view.
- Invoking the `getEJBObject` and `getEJBHome` methods is disallowed if the session bean does not define a remote client view.
- Invoking the `getEJBLocalObject` and `getEJBLocalHome` methods is disallowed if the session bean does not define a local client view.
- Invoking the `getRollbackOnly` and `setRollbackOnly` methods is disallowed in the session bean methods for which the container does not have a meaningful transaction context, and for all session beans with bean-managed transaction demarcation.
- Invoking the `getMessageContext` method is disallowed in session bean methods that were not invoked by the container through the JAX-RPC web service endpoint interface.
- Accessing resource managers and enterprise beans is disallowed in the session bean methods for which the container does not have a meaningful transaction context or client security context.
- The `UserTransaction` interface is unavailable to session beans with container-managed transaction demarcation.

7.8.3 Dealing with Exceptions

A `RuntimeException` thrown from any method of the enterprise bean class (including the business methods and the callbacks invoked by the container) results in the transition to the “does not exist” state. Exception handling is described in detail in Chapter 18.

From the client perspective, the session object continues to exist. The client can continue accessing the session object because the container can delegate the client’s requests to another instance.

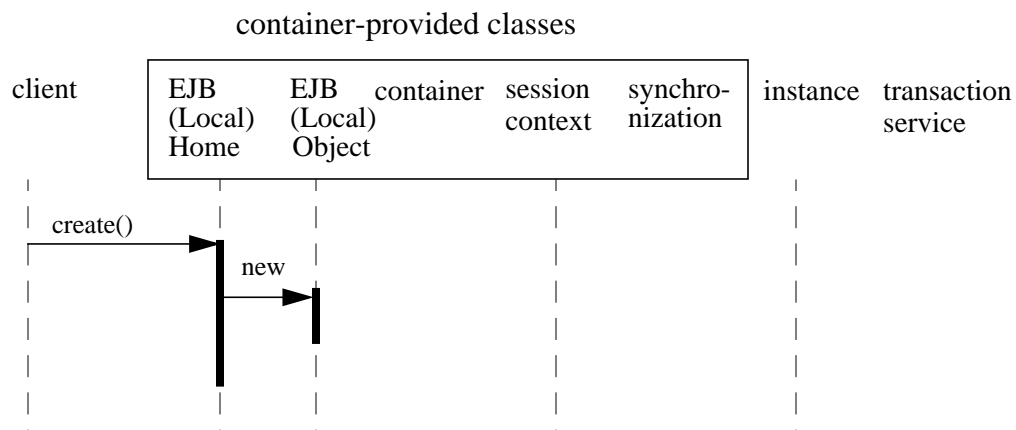
7.9 Object Interaction Diagrams for a Stateless Session Bean with a Home and Component Interface

This section contains object interaction diagrams that illustrates the interaction of the classes.

7.9.1 Client-invoked create

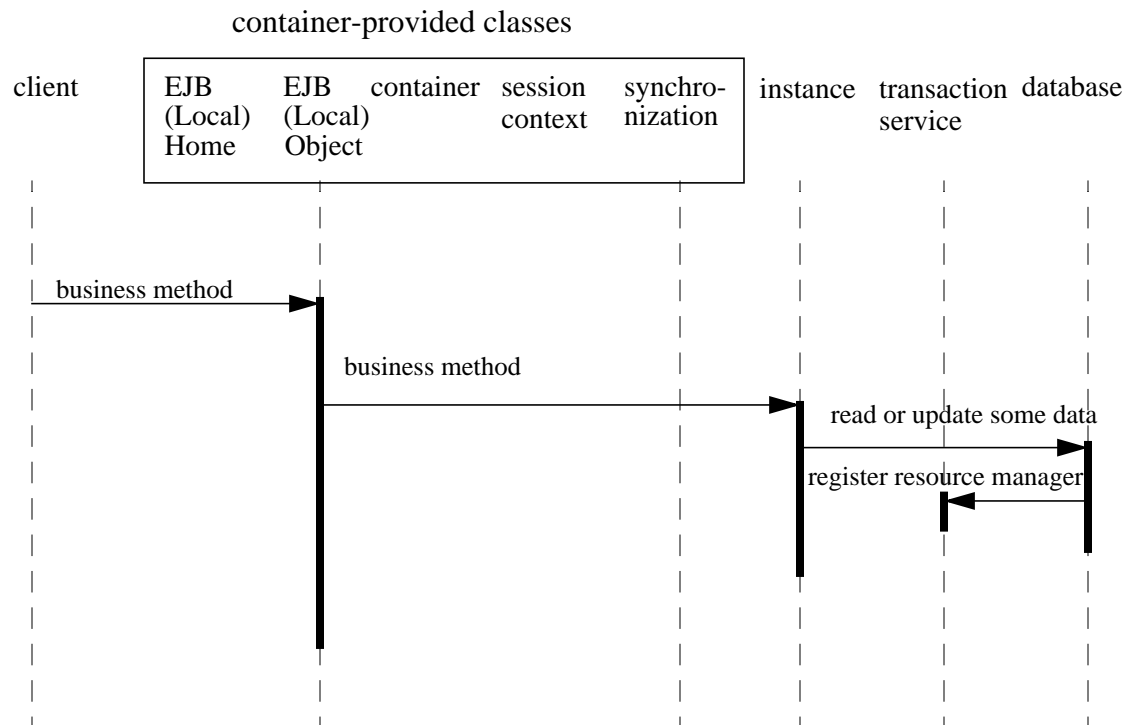
The following diagram illustrates the creation of a stateless session object.

Figure 13 OID for Creation of a Stateless Session Object



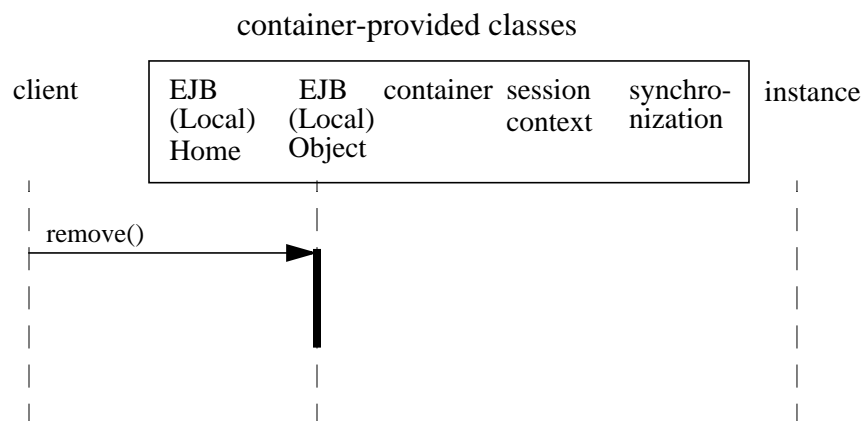
7.9.2 Business Method Invocation

The following diagram illustrates the invocation of a business method.

Figure 14 OID for Invocation of Business Method on a Stateless Session Object

7.9.3 Client-invoked remove

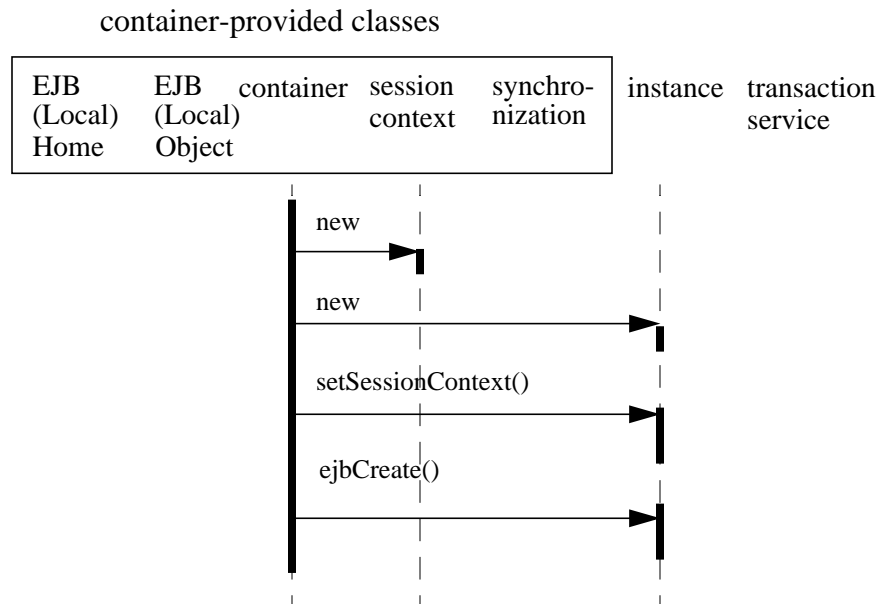
The following diagram illustrates the destruction of a stateless session object.

Figure 15 OID for Removal of a Stateless Session Object

7.9.4 Adding an Instance to the Pool

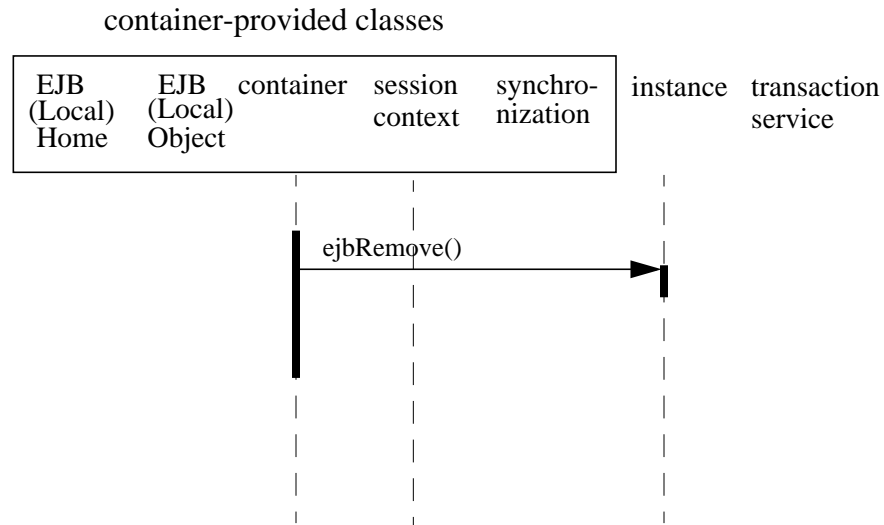
The following diagram illustrates the sequence for a container adding an instance to the method-ready pool.

Figure 16 OID for Container Adding Instance of a Stateless Session Bean to a Method-Ready Pool



The following diagram illustrates the sequence for a container removing an instance from the method-ready pool.

Figure 17 OID for a Container Removing an Instance of a Stateless Session Bean from Ready Pool

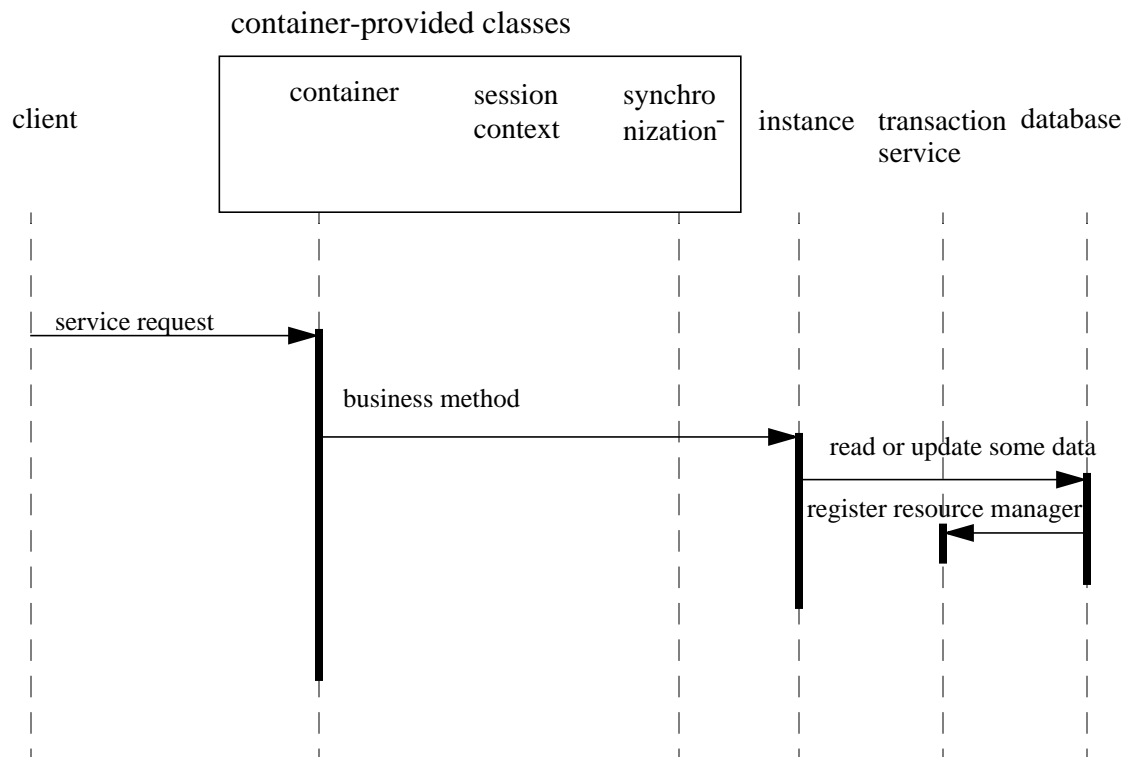


7.10 Object Interaction Diagrams for a Stateless Session Bean with a Web Service Endpoint Interface

This section contains object interaction diagrams that illustrate the interaction of the classes.

7.10.1 Business Method Invocation

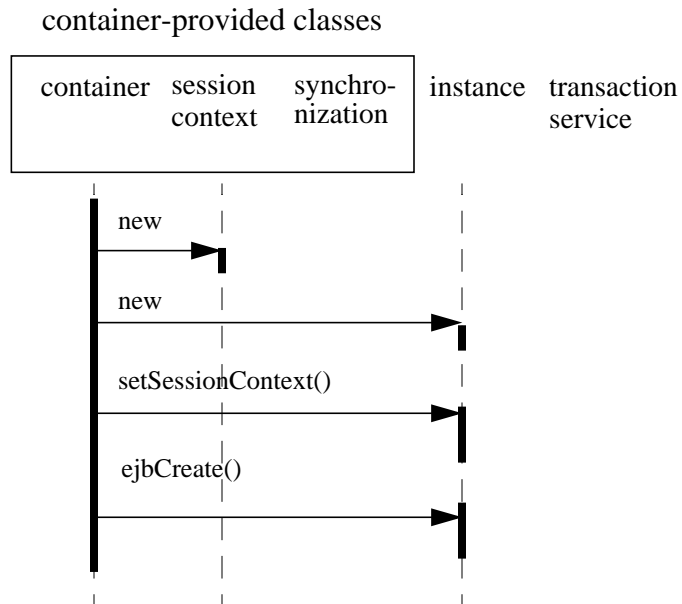
The following diagram illustrates the invocation of a business method.

Figure 18 OID for Invocation of Business Method on a Stateless Session Bean

7.10.2 Adding an Instance to the Pool

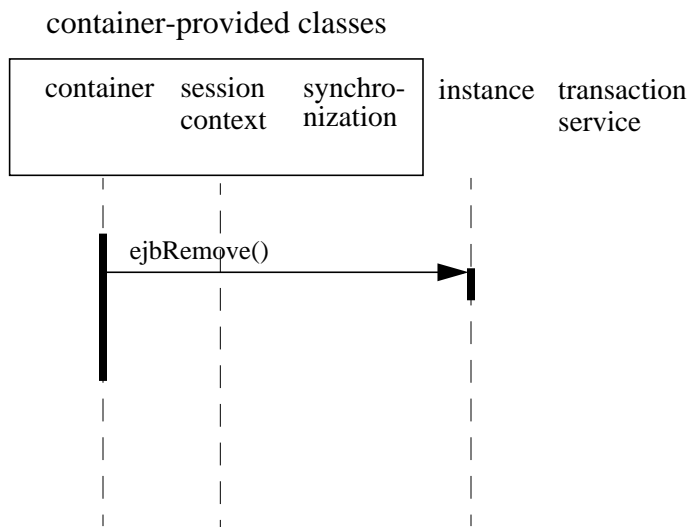
The following diagram illustrates the sequence for a container adding an instance to the method-ready pool.

Figure 19 OID for Container Adding Instance of a Stateless Session Bean to a Method-Ready Pool



The following diagram illustrates the sequence for a container removing an instance from the method-ready pool.

Figure 20 OID for a Container Removing an Instance of a Stateless Session Bean from Ready Pool



7.11 The Responsibilities of the Bean Provider

This section describes the responsibilities of the session Bean Provider to ensure that a session bean can be deployed in any EJB container.

7.11.1 Classes and Interfaces

The session Bean Provider is responsible for providing the following class files:

- Session bean class.
- Session bean's remote interface and remote home interface, if the session bean provides a remote client view.
- Session bean's local interface and local home interface, if the session bean provides a local client view.
- Session bean's web service endpoint interface, if the session bean provides a web service client view.

The Bean Provider may provide a remote interface and a remote home interface, and/or a local interface and a local home interface, and/or a web service endpoint interface for the bean. If a remote interface is provided, a remote home interface must be provided as well, and vice-versa. If a local interface is provided, a local home interface must be provided as well, and vice-versa.

The Bean Provider for a stateless session bean that provides a web service client view may also define JAX-RPC message handlers for the bean. The requirements for such message handlers are defined in [25].

7.11.2 Session Bean Class

The following are the requirements for the session bean class:

- The class must implement, directly or indirectly, the `javax.ejb.SessionBean` interface.
- The class must be defined as `public`, must not be `final`, and must not be `abstract`. The class must be a top level class.
- The class must have a `public` constructor that takes no parameters. The container uses this constructor to create instances of the session bean class.
- The class must not define the `finalize()` method.
- The class may, but is not required to, implement the session bean's web service endpoint or component interface^[8].
- The class must implement the business methods and the `ejbCreate` methods.

- If the class is a stateful session bean, it may optionally implement the `javax.ejb.SessionSynchronization` interface.
- If the class is a stateless session bean, it may optionally implement the `javax.ejb.TimeoutObject` interface.
- The session bean class may have superclasses and/or superinterfaces. If the session bean has superclasses, then the business methods, the `ejbCreate<METHOD>` methods, the methods of the `SessionBean` interface, the `ejbTimeout` method, and the methods of the optional `SessionSynchronization` interface may be defined in the session bean class, or in any of its superclasses.
- The session bean class is allowed to implement other methods (for example helper methods invoked internally by the business methods) in addition to the methods required by the EJB specification.

7.11.3 `ejbCreate<METHOD>` Methods

The session bean class must define one or more `ejbCreate<METHOD>` methods whose signatures must follow these rules:

- The method name must have `ejbCreate` as its prefix.
- The method must be declared as `public`.
- The method must not be declared as `final` or `static`.
- The return type must be `void`.
- The method arguments must be legal types for RMI/IIOP if there is a `create<METHOD>` corresponding to the `ejbCreate<METHOD>` method on the session bean's remote home interface.
- A stateless session bean must define a single `ejbCreate` method, with no arguments.
- The `throws` clause may define arbitrary application exceptions, possibly including the `javax.ejb.CreateException`.

Compatibility Note: EJB 1.0 allowed the `ejbCreate` method to throw the `java.rmi.RemoteException` to indicate a non-application exception. This practice was deprecated in EJB 1.1—an EJB 1.1 or EJB 2.0 or later compliant enterprise bean should throw the `javax.ejb.EJBException` or another `RuntimeException` to indicate non-application exceptions to the container (see Section 18.2.2). An EJB 2.0 or later compliant enterprise bean should not throw the `java.rmi.RemoteException` from the `ejbCreate` method.

[8] If the session bean class does implement the component interface, care must be taken to avoid the passing of `this` as a method argument or result. This potential error can be avoided by choosing not to implement the component interface in the session bean class.

7.11.4 Business Methods

The session bean class may define zero or more business methods whose signatures must follow these rules:

- The method names can be arbitrary, but they must not start with “ejb” to avoid conflicts with the callback methods used by the EJB architecture.
- The business method must be declared as `public`.
- The method must not be declared as `final` or `static`.
- The argument and return value types for a method must be legal types for RMI/IIOP if the method corresponds to a business method on the session bean’s remote interface.
- The argument and return value types for a method must be legal types for JAX-RPC if the method corresponds to a method on the session bean’s web service endpoint interface.
- The `throws` clause may define arbitrary application exceptions.

Compatibility Note: EJB 1.0 allowed the business methods to throw the `java.rmi.RemoteException` to indicate a non-application exception. This practice was deprecated in EJB 1.1—an EJB 1.1 or EJB 2.0 or later compliant enterprise bean should throw the `javax.ejb.EJBException` or another `RuntimeException` to indicate non-application exceptions to the container (see Section 18.2.2). An EJB 2.0 or later compliant enterprise bean should not throw the `java.rmi.RemoteException` from a business method.

7.11.5 Session Bean’s Remote Interface

The following are the requirements for the session bean’s remote interface:

- The interface must extend the `javax.ejb.EJBObject` interface.
- The methods defined in this interface must follow the rules for RMI/IIOP. This means that their argument and return values must be of valid types for RMI/IIOP, and their `throws` clauses must include the `java.rmi.RemoteException`.
- The remote interface is allowed to have superinterfaces. Use of interface inheritance is subject to the RMI/IIOP rules for the definition of remote interfaces.
- For each method defined in the remote interface, there must be a matching method in the session bean’s class. The matching method must have:
 - The same name.
 - The same number and types of arguments, and the same return type.
 - All the exceptions defined in the `throws` clause of the matching method of the session bean class must be defined in the `throws` clause of the method of the remote interface.

- The remote interface methods must not expose local interface types, local home interface types, timers or timer handles, or the managed collection classes that are used for entity beans with container-managed persistence as arguments or results.

7.11.6 Session Bean's Remote Home Interface

The following are the requirements for the session bean's remote home interface:

- The interface must extend the `javax.ejb.EJBHome` interface.
- The methods defined in this interface must follow the rules for RMI/IIOP. This means that their argument and return values must be of valid types for RMI/IIOP, and that their `throws` clauses must include the `java.rmi.RemoteException`.
- The remote home interface is allowed to have superinterfaces. Use of interface inheritance is subject to the RMI/IIOP rules for the definition of remote interfaces.
- A session bean's remote home interface must define one or more `create<METHOD>` methods. A stateless session bean must define exactly one `create` method with no arguments.
- Each `create` method must be named `create<METHOD>`, and it must match one of the `ejbCreate<METHOD>` methods defined in the session bean class. The matching `ejbCreate<METHOD>` method must have the same number and types of arguments. (Note that the return type is different.) The methods for a stateless session bean must be named "create" and "ejbCreate".
- The return type for a `create<METHOD>` method must be the session bean's remote interface type.
- All the exceptions defined in the `throws` clause of an `ejbCreate<METHOD>` method of the session bean class must be defined in the `throws` clause of the matching `create<METHOD>` method of the remote home interface.
- The `throws` clause must include `javax.ejb.CreateException`.

7.11.7 Session Bean's Local Interface

The following are the requirements for the session bean's local interface:

- The interface must extend the `javax.ejb.EJBLocalObject` interface.
- The `throws` clause of a method defined in the local interface must not include the `java.rmi.RemoteException`.
- The local interface is allowed to have superinterfaces.
- For each method defined in the local interface, there must be a matching method in the session bean's class. The matching method must have:

- The same name.
- The same number and types of arguments, and the same return type.
- All the exceptions defined in the `throws` clause of the matching method of the session bean class must be defined in the `throws` clause of the method of the local interface.

7.11.8 Session Bean's Local Home Interface

The following are the requirements for the session bean's local home interface:

- The interface must extend the `javax.ejb.EJBLocalHome` interface.
- The `throws` clause of a method in the local home interface must not include the `java.rmi.RemoteException`.
- The local home interface is allowed to have superinterfaces.
- A session bean's local home interface must define one or more `create<METHOD>` methods. A stateless session bean must define exactly one `create` method with no arguments.
- Each `create` method must be named "`create<METHOD>`", and it must match one of the `ejbCreate<METHOD>` methods defined in the session bean class. The matching `ejbCreate<METHOD>` method must have the same number and types of arguments. (Note that the return type is different.) The methods for a stateless session bean must be named "`create`" and "`ejbCreate`".
- The return type for a `create<METHOD>` method must be the session bean's local interface type.
- All the exceptions defined in the `throws` clause of an `ejbCreate<METHOD>` method of the session bean class must be defined in the `throws` clause of the matching `create<METHOD>` method of the local home interface.
- The `throws` clause must include `javax.ejb.CreateException`.

7.11.9 Session Bean's Web Service Endpoint Interface

The following are the requirements for a stateless session bean's web service endpoint interface. The web service endpoint interface must follow the rules for JAX-RPC service endpoint interfaces [22].

- The web service endpoint interface must extend the `java.rmi.Remote` interface.
- The methods defined in the interface must follow the rules for JAX-RPC service endpoint interfaces. This means that their argument and return values must be of valid types for JAX-RPC, and their `throws` clauses must include the `java.rmi.RemoteException`. The `throws` clause may additionally include application exceptions.

Note that JAX-RPC Holder classes may be used as method parameters. The JAX-RPC specification requires support for Holder classes as part of the standard Java mapping of WSDL operations in order to handle out and inout parameters. Holder classes implement the `javax.xml.rpc.holders.Holder` interface. See the JAX-RPC specification [22] for further details.

- For each method defined in the web service endpoint interface, there must be a matching method in the session bean's class. The matching method must have:
 - The same name.
 - The same number and types of arguments, and the same return type.
 - All the exceptions defined in the `throws` clause of the matching method of the session bean class must be defined in the `throws` clause of the method of the web service endpoint interface.
- The web service endpoint interface must not include an `EJBObject` or `EJBLocalObject` as either a parameter or return type. An array or JAX-RPC value type must not include an `EJBObject` or `EJBLocalObject` as a contained element. The web service endpoint interface methods must not expose local or remote interface types, local or remote home interface types, timers or timer handles, or the managed collection classes that are used for entity beans with container-managed persistence as arguments or results or as fields of value types.
- JAX-RPC serialization rules apply for any value types that are used by the web service endpoint interface. If it is important that Java serialization semantics apply, the Bean Provider should use the restricted set of JAX-RPC value types for which the semantics of Java serialization apply under JAX-RPC serialization. See the JAX-RPC specification [22] for details.
- The web service endpoint interface must not include constant (as `public final static`) declarations.
- The Bean Provider must designate the web service endpoint interface in the deployment descriptor by means of the `service-endpoint` element. The service endpoint itself is only exposed within a web service if it is referenced by a web service deployment descriptor as defined by [25].

7.12 The Responsibilities of the Container Provider

This section describes the responsibilities of the Container Provider to support a session bean. The Container Provider is responsible for providing the deployment tools and for managing the session bean instances at runtime.

Because the EJB specification does not define the API between deployment tools and the container, we assume that the deployment tools are provided by the Container Provider. Alternatively, the deployment tools may be provided by a different vendor who uses the container vendor's specific API.

7.12.1 Generation of Implementation Classes

The deployment tools provided by the container are responsible for the generation of additional classes when the session bean is deployed. The tools obtain the information that they need for generation of the additional classes by introspecting the classes and interfaces provided by the Bean Provider and by examining the session bean's deployment descriptor.

The deployment tools must generate the following classes:

- A class that implements the session bean's remote home interface (session `EJBHome` class).
- A class that implements the session bean's remote interface (session `EJBObject` class).
- A class that implements the session bean's local home interface (session `EJBLocalHome` class).
- A class that implements the session bean's local interface (session `EJBLocalObject` class).
- A class that implements the session bean's web service endpoint.

The deployment tools may also generate a class that mixes some container-specific code with the session bean class. This code may, for example, help the container to manage the bean instances at runtime. The tools can use subclassing, delegation, and code generation.

The deployment tools may also allow the generation of additional code that wraps the business methods and is used to customize the business logic to an existing operational environment. For example, a wrapper for a `debit` function on the `AccountManager` bean may check that the debited amount does not exceed a certain limit.

7.12.2 Generation of WSDL

Reference [25] describes the generation of a WSDL document for a web service endpoint interface. The Java to WSDL mapping must adhere to the requirements of JAX-RPC.

7.12.3 Session EJBHome Class

The session `EJBHome` class, which is generated by the deployment tools, implements the session bean's remote home interface. This class implements the methods of the `javax.ejb.EJBHome` interface and the `create<METHOD>` methods specific to the session bean.

The implementation of each `create<METHOD>` method invokes a matching `ejbCreate<METHOD>` method.

7.12.4 Session EJBObject Class

The session EJBObject class, which is generated by the deployment tools, implements the session bean's remote interface. It implements the methods of the `javax.ejb.EJBObject` interface and the business methods specific to the session bean.

The implementation of each business method must activate the instance (if the instance is in the passive state) and invoke the matching business method on the instance.

7.12.5 Session EJBLocalHome Class

The session EJBLocalHome class, which is generated by the deployment tools, implements the session bean's local home interface. This class implements the methods of the `javax.ejb.EJBLocalHome` interface and the `create<METHOD>` methods specific to the session bean.

The implementation of each `create<METHOD>` method invokes a matching `ejbCreate<METHOD>` method.

7.12.6 Session EJBLocalObject Class

The session EJBLocalObject class, which is generated by the deployment tools, implements the session bean's local interface. It implements the methods of the `javax.ejb.EJBLocalObject` interface and the business methods specific to the session bean.

The implementation of each business method must activate the instance (if the instance is in the passive state) and invoke the matching business method on the instance.

7.12.7 Web Service Endpoint Implementation Class

The implementation class for a stateless session bean's web service endpoint is generated by the container's deployment tools. This class must handle requests to the web service endpoint, unmarshall the SOAP request, and invoke the stateless session bean method that matches the web service endpoint interface method that corresponds to the request.

7.12.8 Handle Classes

The deployment tools are responsible for implementing the handle classes for the session bean's remote home and remote interfaces.

7.12.9 EJBMetaData Class

The deployment tools are responsible for implementing the class that provides metadata to the remote client view contract. The class must be a valid RMI Value class and must implement the `javax.ejb.EJBMetaData` interface.

7.12.10 Non-reentrant Instances

The container must ensure that only one thread can be executing an instance at any time. If a client request arrives for an instance while the instance is executing another request, the container may throw the `java.rmi.RemoteException` to the second request if the client is a remote client, or the `javax.ejb.EJBException` if the client is a local client.^[9]

Note that a session object is intended to support only a single client. Therefore, it would be an application error if two clients attempted to invoke the same session object.

One implication of this rule is that an application cannot make loopback calls to a session bean instance.

7.12.11 Transaction Scoping, Security, Exceptions

The container must follow the rules with respect to transaction scoping, security checking, and exception handling, as described in Chapters 17, 21, and 18, respectively.

7.12.12 JAX-RPC Message Handlers for Web Service Endpoints

The container must support the use of JAX-RPC message handlers for web service endpoints. Container requirements for support of message handlers are specified in [25].

7.12.13 SessionContext

The container must implement the `SessionContext.getEJBObject` method such that the bean instance can use the Java language cast to convert the returned value to the session bean's remote interface type. Specifically, the bean instance does not have to use the `PortableRemoteObject.narrow` method for the type conversion.

[9] In certain special circumstances (e.g., to handle clustered web container architectures), the container may instead queue or serialize such concurrent requests. Clients, however, cannot rely on this behavior.

Example Session Scenario

This chapter describes an example development and deployment scenario of a session bean. We use the scenario to explain the responsibilities of the Bean Provider and those of the Container Provider.

The classes generated by the Container Provider's tools in this scenario should be considered illustrative rather than prescriptive. Container providers are free to implement the contract between a session bean and its container in a different way, provided that it achieves an equivalent effect (from the perspectives of the Bean Provider and the client-side programmer).

8.1 Overview

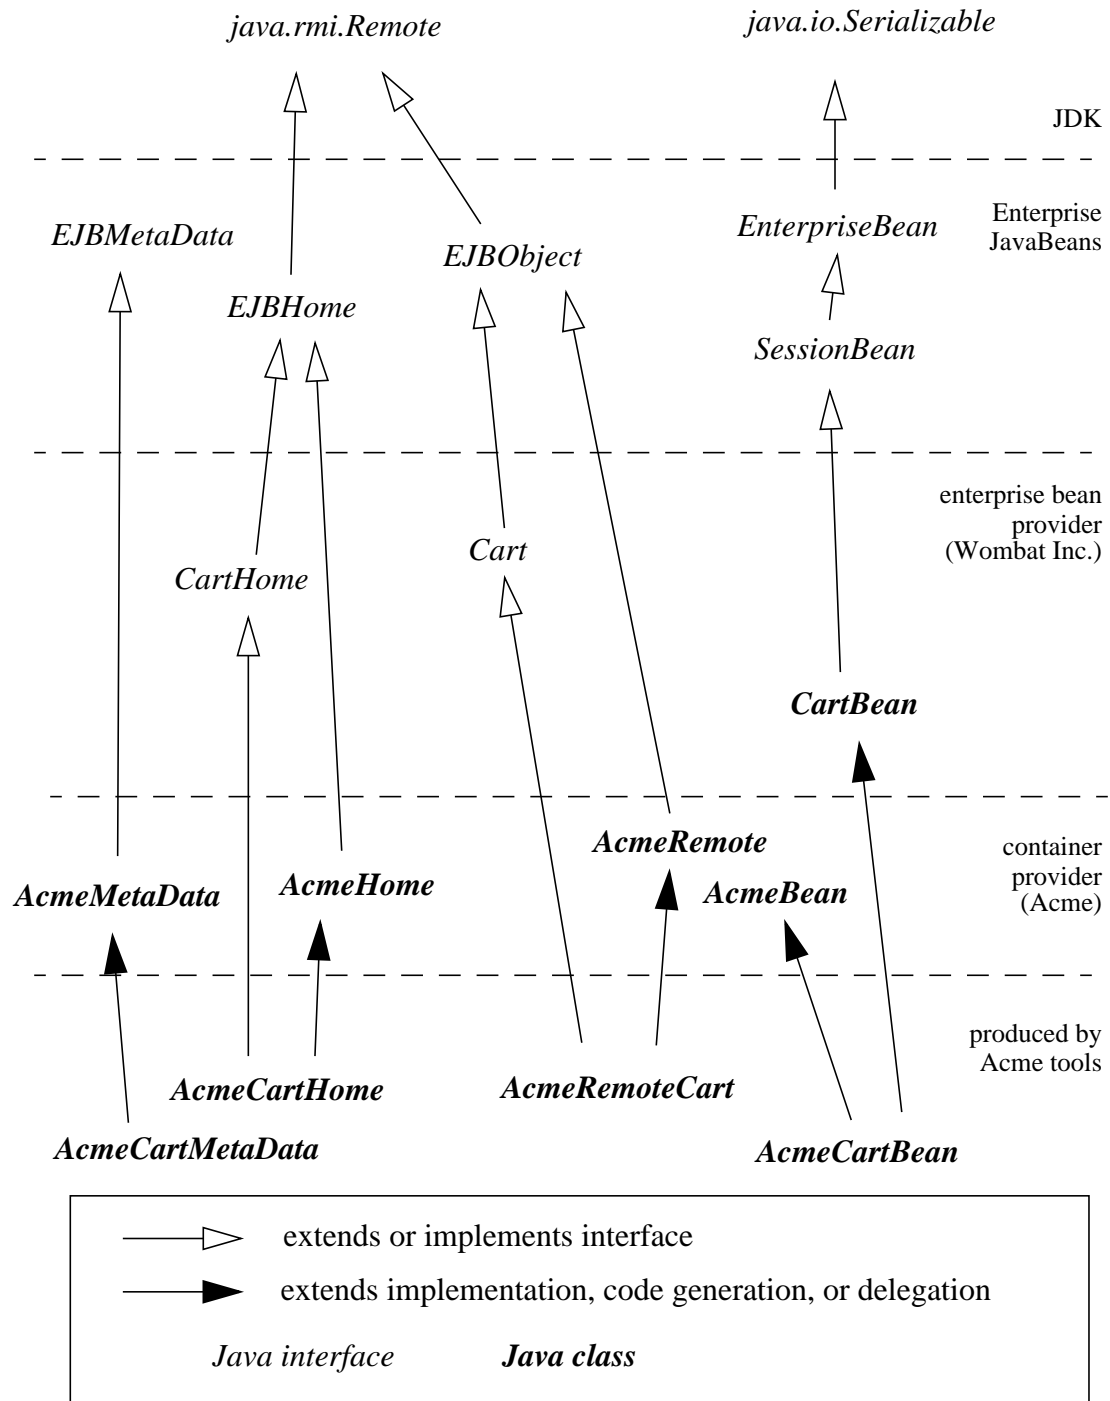
Wombat Inc. has developed the `CartBean` session bean. The `CartBean` is deployed in a container provided by the Acme Corporation. The `CartBean` provides a remote client view.

8.2 Inheritance Relationship

An example of the inheritance relationship between the remote interfaces and classes is illustrated in the following diagram:

Figure 21

Example of Inheritance Relationships Between EJB Classes



8.2.1 What the Session Bean Provider is Responsible For

Wombat Inc. is responsible for providing the following:

- *Define the session bean's remote interface (`Cart`). The remote interface defines the business methods callable by a client. The remote interface must extend the `javax.ejb.EJBObject` interface, and follow the standard rules for a RMI-IIOP remote interface. The remote interface must be defined as public.*
- *Write the business logic in the session bean class (`CartBean`). The enterprise bean class may, but is not required to, implement the enterprise bean's remote interface (`Cart`). The enterprise bean must implement the `javax.ejb.SessionBean` interface, and define the `ejbCreate<METHOD>` method(s) invoked at EJB object creation.*
- *Define a remote home interface (`CartHome`) for the enterprise bean. The remote home interface must be defined as public, extend the `javax.ejb.EJBHome` interface, and follow the standard rules for RMI-IIOP remote interfaces.*
- *Define a deployment descriptor specifying any declarative metadata that the session Bean Provider wishes to pass with the bean to the next stage of the development/deployment workflow.*

8.2.2 Classes Supplied by Container Provider

The following classes are supplied by the Container Provider Acme Corp:

The `AcmeHome` class provides the Acme implementation of the `javax.ejb.EJBHome` methods.

The `AcmeRemote` class provides the Acme implementation of the `javax.ejb.EJBObject` methods.

The `AcmeBean` class provides additional state and methods to allow Acme's container to manage its session bean instances. For example, if Acme's container uses an LRU algorithm, then `AcmeBean` may include the clock count and methods to use it.

The `AcmeMetaData` class provides the Acme implementation of the `javax.ejb.EJBMetaData` methods.

8.2.3 What the Container Provider is Responsible For

The tools provided by Acme Corporation are responsible for the following:

- *Generate the class (`AcmeRemoteCart`) that implements the session bean's remote interface. The tools also generate the classes that implement the communication protocol specific artifacts for the remote interface.*
- *Generate the implementation of the session bean class suitable for the Acme container (`AcmeCartBean`). `AcmeCartBean` includes the business logic from the `CartBean` class*

mixed with the services defined in the `AcmeBean` class. Acme tools can use inheritance, delegation, and code generation to achieve a mix-in of the two classes.

- *Generate the class (`AcmeCartHome`) that implements the session bean's remote home interface. The tools also generate the classes that implement the communication protocol specific artifacts for the remote home interface.*
- *Generate the class (`AcmeCartMetaData`) that implements the `javax.ejb.EJBMetaData` interface for the `Cart` bean.*

Many of the above classes and tools are container-specific (i.e., they reflect the way Acme Corp implemented them). Other container providers may use different mechanisms to produce their runtime classes, and these classes will likely be different from those generated by Acme's tools.

Client View of an Entity

This chapter describes the client view of an entity bean. It is actually a contract fulfilled by the container in which the entity bean is deployed. Only the business methods are supplied by the enterprise bean itself.

Although the client view of the deployed entity beans is provided by classes implemented by the container, the container itself is transparent to the client.

9.1 Overview

For a client, an entity bean is a component that represents an object-oriented view of some entities stored in a persistent storage, such as a database, or entities that are implemented by an existing enterprise application.

The client of an entity bean may be a local client or the client may be a remote client.

This section provides an overview of the entity bean client view that is independent of whether the client is a remote client or a local client. The differences between remote clients and local clients are discussed in the following sections.

From its creation until its destruction, an entity object lives in a container. Transparently to the client, the container provides security, concurrency, transactions, persistence, and other services for the entity objects that live in the container. The container is transparent to the client—there is no API that a client can use to manipulate the container.

Multiple clients can access an entity object concurrently. The container in which the entity bean is deployed properly synchronizes access to the entity object's state using transactions.

Each entity object has an identity which, in general, survives a crash and restart of the container in which the entity object has been created. The object identity is implemented by the container with the cooperation of the enterprise bean class.

Multiple enterprise beans can be deployed in a container. For each entity bean deployed in a container, the container provides a class that implements a home interface for the entity bean. This interface allows the client to create, find, and remove entity objects within the enterprise bean's home as well as to execute home business methods, which are not specific to a particular entity bean object. A client can look up the entity bean's home interface through JNDI. It is the responsibility of the container to make the entity bean's home interface available in the JNDI name space.

A client view of an entity bean is independent of the implementation of the entity bean and its container. This ensures that a client application is portable across all container implementations in which the entity bean might be deployed.

9.2 Remote Clients

A remote client accesses an entity bean through the entity bean's remote and remote home interfaces. The remote and remote home interfaces of the entity bean provide the remote client view.

The remote client view of an entity bean is location independent. A client running in the same JVM as an entity bean instance uses the same API to access the entity bean as a client running in a different JVM on the same or different machine.

The container provides classes that implement the entity bean's remote and remote home interfaces. The objects that implement the remote home and remote objects are remote Java objects, and are accessible from a client through the standard Java™ APIs for remote object invocation [3].

A remote client of an entity object can be another enterprise bean deployed in the same or different container or can be an arbitrary Java program, such as an application, applet, or servlet. The remote client view of an entity bean can also be mapped to non-Java client environments, such as CORBA clients not written in the Java programming language.

9.3 Local Clients

Entity beans may also have local clients.^[10] A local client is a client that is collocated with the entity bean and which may be tightly coupled to the bean.

Unlike the remote client view, the local client view of an entity bean is not location independent. The local client view requires the collocation in the same JVM of both the local client and the entity bean that provides the local client view. The local client view therefore does not provide the location transparency provided by the remote client view.

A local client accesses an entity bean through the entity bean's local and local home interfaces. The container provides classes that implement the entity bean's local and local home interfaces. The objects that implement the local home and local interfaces are local Java objects.

The arguments of the methods of the local interface and local home interface are passed by reference^[11]. Such entity beans and their clients must be coded to assume that the state of any Java object that is passed as an argument or result is potentially shared by caller and callee.

A local client of an entity bean may be a session bean, a message-driven bean, another entity bean, or a web-tier component.

The considerations that should be taken into account in determining the whether the local or remote programming model should be used for an entity bean are further described in Chapter 5, "Local, Remote, and Web Service Client Views".

The choice between the use of a local or remote programming model is a design decision that the Bean Provider makes when developing the entity bean application.

Entity beans that have container-managed relationships with other entity beans, as described in Chapter 10, "Entity Bean Component Contract for Container-Managed Persistence", must be accessed in the same local scope as those related beans, and therefore typically provide a local client view. In order to be the target of a container-managed relationship, an entity bean with container-managed persistence must provide a local interface.

While it is possible to provide both a client view and a local client view for an entity bean with container-managed persistence, it is more likely that the entity bean will be designed with the local view in mind.

9.4 EJB Container

An EJB container (container for short) is a system that functions as a runtime container for enterprise beans.

[10] Local clients are not supported for EJB 1.1 entity beans with container-managed persistence.

[11] More literally, references are passed by value in the JVM: an argument variable of primitive type holds a value of that primitive type; an argument variable of a reference type hold a reference to the object. See [26].

Multiple enterprise beans can be deployed in a single container. For each entity bean deployed in a container, the container provides a home interface that allows the client to create, find, and remove entity objects that belong to the entity bean. The home interface may also provide home business methods, which are not specific to a particular entity bean object. The container makes the entity bean's home interface (defined by the Bean Provider and implemented by the Container Provider) available in the JNDI name space for clients.

An EJB server may host one or multiple EJB containers. The containers are transparent to the client: there is no client-level API to manipulate the container.

9.4.1 Locating an Entity Bean's Home Interface

A client locates an entity bean's home interface using JNDI. A client's JNDI name space may be configured to include the home interfaces of enterprise beans deployed in multiple EJB containers located on multiple machines on a network. The actual location of an EJB container is, in general, transparent to the client.

For example, the remote home interface for the `Account` entity bean can be located using the following code segment:


```
Context initialContext = new InitialContext();
AccountHome accountHome = (AccountHome)
    javax.rmi.PortableRemoteObject.narrow(
        initialContext.lookup("java:comp/env/ejb/accounts"),
        AccountHome.class);
```

Locating an entity bean's local home interface using JNDI is done similarly. It does not, however, involve the APIs for remote access. For example, if the `Account` entity bean provided `AccountHome` as a local home interface rather than a remote home interface, a local client might use the following code segment:

```
Context initialContext = new InitialContext();
AccountHome accountHome = (AccountHome)
    initialContext.lookup("java:comp/env/ejb/accounts");
```

9.4.2 What a Container Provides

The following diagram illustrates the view that a container provides to the client of the entity beans deployed in the container. Note that a client may be a local client of some entity beans and a remote client of others.

Figure 22 Client View of Entity Beans Deployed in a Container

9.5 Entity Bean's Remote Home Interface

This section is specific to entity beans that provide a remote client view. Local home interfaces are described in Section 9.6.

The container provides the implementation of the remote home interface for each entity bean deployed in the container that defines a remote home interface. The container makes the remote home interface accessible to the clients through JNDI. An object that implements an entity bean's remote home interface is called an **EJBHome** object.

The entity bean's remote home interface allows a client to do the following:

- Create new entity objects within the home.
- Find existing entity objects within the home.

- Remove an entity object from the home.
- Execute a home business method.
- Get the `javax.ejb.EJBMetaData` interface for the entity bean. The `javax.ejb.EJBMetaData` interface is intended to allow application assembly tools to discover the metadata information about the entity bean. The metadata information allows loose client/server binding and scripting.
- Obtain a handle for the home interface. The home handle can be serialized and written to stable storage. Later, possibly in a different JVM, the handle can be deserialized from stable storage and used to obtain a reference to the home interface.

An entity bean's remote home interface must extend the `javax.ejb.EJBHome` interface and follow the standard rules for Java programming language remote interfaces.

9.5.1 Create Methods

An entity bean's remote home interface can define zero or more `create<METHOD>` methods, one for each way to create an entity object. The arguments of the `create` methods are typically used to initialize the state of the created entity object. The name of each `create` method starts with the prefix “create”.

The return type of a `create<METHOD>` method on the remote home interface is the entity bean's remote interface.

The `throws` clause of every `create<METHOD>` method on the remote home interface includes the `java.rmi.RemoteException` and the `javax.ejb.CreateException`. It may include additional application-level exceptions.

The following home interface illustrates three possible `create` methods:

```
public interface AccountHome extends javax.ejb.EJBHome {
    public Account create(String firstName, String lastName,
        double initialBalance)
        throws RemoteException, CreateException;
    public Account create(String accountNumber,
        double initialBalance)
        throws RemoteException, CreateException,
        LowInitialBalanceException;
    public Account createLargeAccount(String firstname,
        String lastname, double initialBalance)
        throws RemoteException, CreateException;
    ...
}
```


The following example illustrates how a client creates a new entity object:

```
AccountHome accountHome = ...;  
Account account = accountHome.create("John", "Smith", 500.00);
```

9.5.2 Finder Methods

An entity bean's remote home interface defines one or more finder methods^[12], one for each way to find an entity object or collection of entity objects within the home. The name of each finder method starts with the prefix "find", such as `findLargeAccounts`. The arguments of a finder method are used by the entity bean implementation to locate the requested entity objects. The return type of a finder method on the remote home interface must be the entity bean's remote interface, or a type representing a collection of objects that implement the entity bean's remote interface (see Subsections 10.5.6 and 12.1.8).

The `throws` clause of every finder method on the remote home interface includes the `java.rmi.RemoteException` and the `javax.ejb.FinderException` exceptions.

The remote home interface includes the `findByPrimaryKey(primaryKey)` method, which allows a client to locate an entity object using a primary key. The name of the method is always `findByPrimaryKey`; it has a single argument that is the same type as the entity bean's primary key type, and its return type is the entity bean's remote interface. There is a unique `findByPrimaryKey(primaryKey)` method for an entity bean on its remote home interface; this method must not be overloaded. The implementation of the `findByPrimaryKey(primaryKey)` method must ensure that the entity object exists.

The following example shows the `findByPrimaryKey` method:

```
public interface AccountHome extends javax.ejb.EJBHome {  
    ...  
    public Account findByPrimaryKey(String AccountNumber)  
        throws RemoteException, FinderException;  
}
```

The following example illustrates how a client uses the `findByPrimaryKey` method:

```
AccountHome = ...;  
Account account = accountHome.findByPrimaryKey("100-3450-3333");
```

[12] The `findByPrimaryKey` method is mandatory for the remote home interface of all entity beans.

9.5.3 Remove Methods

The `javax.ejb.EJBHome` interface defines several methods that allow the client to remove an entity object.

```
public interface EJBHome extends Remote {
    void remove(Handle handle) throws RemoteException,
        RemoveException;
    void remove(Object primaryKey) throws RemoteException,
        RemoveException;
}
```

After an entity object has been removed, subsequent attempts to access the entity object by a remote client result in the `java.rmi.NoSuchObjectException`.

9.5.4 Home Methods

An entity bean's remote home interface may define one or more home methods. Home methods are methods that the Bean Provider supplies for business logic that is not specific to an entity bean instance.

Home methods on the remote home interface can have arbitrary method names, but they must not start with "create", "find", or "remove". The arguments of a home method are used by the entity bean implementation in computations that do not depend on a specific entity bean instance. The method arguments and return value types of a home method on the remote home interface must be legal types for RMI-IIOP.

The `throws` clause of every home method on the remote home interface includes the `java.rmi.RemoteException`. It may also include additional application-level exceptions.

The following example shows two home methods:

```
public interface EmployeeHome extends javax.ejb.EJBHome {
    ...
    // this method returns a living index depending on
    // the state and the base salary of an employee:
    // the method is not specific to an instance
    public float livingIndex(String state, float Salary)
        throws RemoteException;

    // this method adds a bonus to all of the employees
    // based on a company profit-sharing index
    public void addBonus(float company_share_index)
        throws RemoteException, ShareIndexOutOfRangeException;

    ...
}
```

9.6 Entity Bean's Local Home Interface

The container provides the implementation of the local home interface for each entity bean deployed in the container that defines a local home interface. The container makes the local home interface accessible to local clients through JNDI. An object that implements an entity bean's local home interface is called an **EJBLocalHome** object.

The entity bean's local home interface allows a local client to do the following:

- Create new entity objects within the home.
- Find existing entity objects within the home.
- Remove an entity object from the home.
- Execute a home business method.

An entity bean's local home interface must extend the `javax.ejb.EJBLocalHome` interface.

9.6.1 Create Methods

An entity bean's local home interface can define zero or more `create<METHOD>` methods, one for each way to create an entity object. The arguments of the `create` methods are typically used to initialize the state of the created entity object. The name of each `create` method starts with the prefix “create”.

The return type of a `create<METHOD>` method on the local home interface is the entity bean's local interface.

The `throws` clause of every `create<METHOD>` method on the local home interface includes the `javax.ejb.CreateException`. It may include additional application-level exceptions. It must not include the `java.rmi.RemoteException`.

The following local home interface illustrates three possible `create` methods:

```
public interface AccountHome extends javax.ejb.EJBLocalHome {
    public Account create(String firstName, String lastName,
        double initialBalance)
        throws CreateException;
    public Account create(String accountNumber,
        double initialBalance)
        throws CreateException, LowInitialBalanceException;
    public Account createLargeAccount(String firstname,
        String lastname, double initialBalance)
        throws CreateException;
    ...
}
```

The following example illustrates how a client creates a new entity object:

```
AccountHome accountHome = ...;  
Account account = accountHome.create("John", "Smith", 500.00);
```

9.6.2 Finder Methods

An entity bean's local home interface defines one or more finder methods^[13], one for each way to find an entity object or collection of entity objects within the home. The name of each finder method starts with the prefix "find", such as `findLargeAccounts`. The arguments of a finder method are used by the entity bean implementation to locate the requested entity objects. The return type of a finder method on the local home interface must be the entity bean's local interface, or a type representing a collection of objects that implement the entity bean's local interface (see Subsections 10.5.6 and 12.1.8).

The `throws` clause of every finder method on the local home interface includes the `javax.ejb.FinderException`. The `throws` clause must not include the `java.rmi.RemoteException`.

The local home interface includes the `findByPrimaryKey(primaryKey)` method, which allows a client to locate an entity object using a primary key. The name of the method is always `findByPrimaryKey`; it has a single argument that is the same type as the entity bean's primary key type, and its return type is the entity bean's local interface. There is a unique `findByPrimaryKey(primaryKey)` method for an entity bean on its local home interface; this method must not be overloaded. The implementation of the `findByPrimaryKey` method must ensure that the entity object exists.

The following example shows the `findByPrimaryKey` method:

```
public interface AccountHome extends javax.ejb.EJBLocalHome {  
    ...  
    public Account findByPrimaryKey(String AccountNumber)  
        throws FinderException;  
}
```

The following example illustrates how a client uses the `findByPrimaryKey` method:

```
AccountHome = ...;  
Account account = accountHome.findByPrimaryKey("100-3450-3333");
```

[13] The `findByPrimaryKey` method is mandatory for the local home interface of all Entity Beans.

9.6.3 Remove Methods

The `javax.ejb.EJBLocalHome` interface defines the `remove` method to allow the client to remove an entity object.

```
public interface EJBLocalHome {  
    void remove(Object primaryKey) throws RemoveException,  
        EJBException;  
}
```

After an entity object has been removed, subsequent attempts to access the local entity object by the local client result in the `javax.ejb.NoSuchObjectLocalException`.

9.6.4 Home Methods

An entity bean's local home interface may define one or more home methods. Home methods are methods that the Bean Provider supplies for business logic that is not specific to an entity bean instance.

Home methods can have arbitrary method names, but they must not start with "create", "find", or "remove". The arguments of a home method are used by the entity bean implementation in computations that do not depend on a specific entity bean instance.

The `throws` clause of a home method on the local home interface may include additional application-level exceptions. It must not include the `java.rmi.RemoteException`.

The following example shows two home methods:

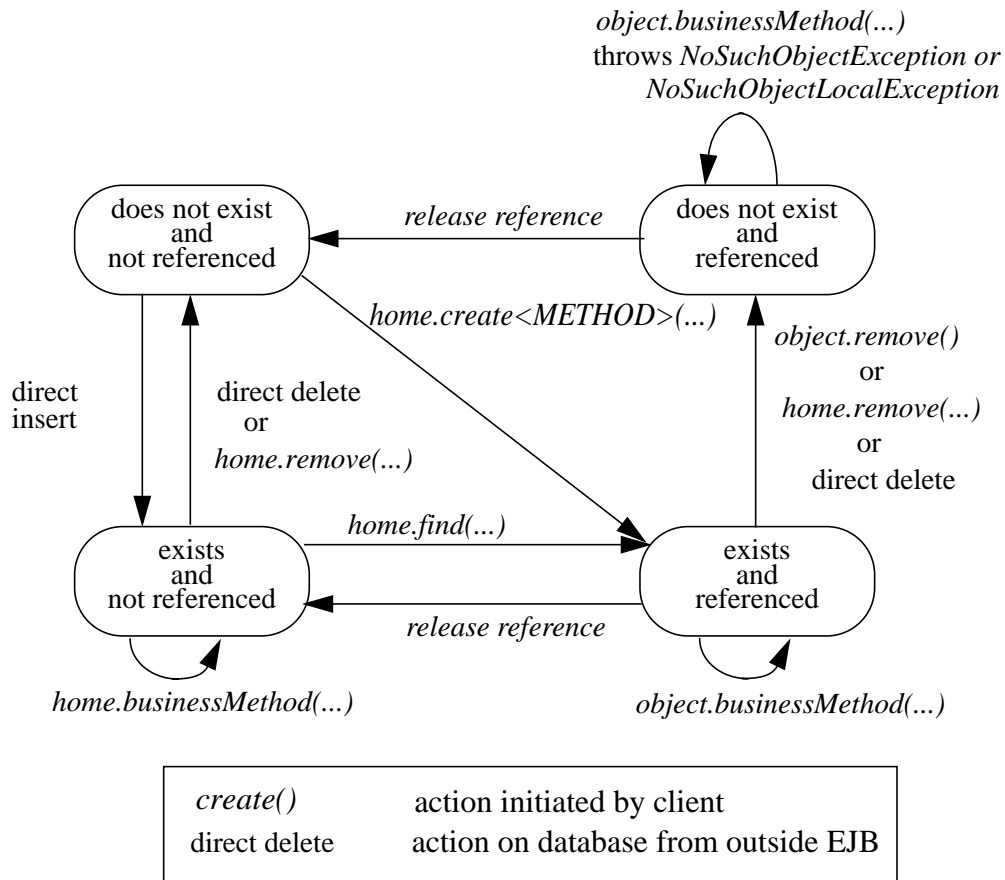
```
public interface EmployeeHome extends javax.ejb.EJBLocalHome {  
    ...  
    // this method returns a living index depending on  
    // the state and the base salary of an employee:  
    // the method is not specific to an instance  
    public float livingIndex(String state, float Salary);  
  
    // this method adds a bonus to all of the employees  
    // based on a company profit sharing index  
    public void addBonus(float company_share_index)  
        throws ShareIndexOutOfRangeException;  
    ...  
}
```

9.7 Entity Object's Life Cycle

This section describes the life cycle of an entity object from the perspective of a client.

The following diagram illustrates a client's point of view of an entity object life cycle. (The term “referenced” in the diagram means that the client program has a reference to the entity object's remote or local interface.)

Figure 23 Client View of Entity Object Life Cycle



An entity object does not exist until it is created. Until it is created, it has no identity. After it is created, it has identity. A client creates an entity object using the entity bean's home interface, whose class is implemented by the container. When a client creates an entity object, the client obtains a reference to the newly created entity object.

In an environment with legacy data, entity objects may “exist” before the container and entity bean are deployed. In addition, an entity object may be “created” in the environment via a mechanism other than by invoking a `create<METHOD>` method of the home interface (e.g. by inserting a database record), but still may be accessible via the finder methods. Also, an entity object may be deleted directly using other means than the `remove` operation (e.g. by deletion of a database record). The “direct insert” and “direct delete” transitions in the diagram represent such direct database manipulation.

All entity objects are considered persistent objects. The lifetime of an entity object is not limited by the lifetime of the Java Virtual Machine process in which the entity bean instance executes. While a crash of the Java Virtual Machine may result in a rollback of current transactions, it does not destroy previously created entity objects nor invalidate the references to the home and component interfaces held by clients.

Multiple clients can access the same entity object concurrently. Transactions are used to isolate the clients' work from each other.

9.7.1 References to Entity Object Remote Interfaces

A client can get a reference to an existing entity object's remote interface in any of the following ways:

- Receive the reference as a parameter in a method call (input parameter or result).
- Find the entity object using a finder method defined in the entity bean's remote home interface.
- Obtain the reference from the entity object's handle. (See Section 9.11).

A client that has a reference to an entity object's remote interface can do any of the following:

- Invoke business methods on the entity object through the remote interface.
- Obtain a reference to the enterprise bean's remote home interface.
- Pass the reference as a parameter or return value of a method call.
- Obtain the entity object's primary key.
- Obtain the entity object's handle.
- Remove the entity object.

All references to an entity object that does not exist are invalid. All attempted invocations on an entity object that does not exist result in an `java.rmi.NoSuchObjectException` being thrown.

9.7.2 References to Entity Object Local Interfaces

A local client can get a reference to an existing entity object's local interface in any of the following ways:

- Receive the reference as a result of a method call.
- Find the entity object using a finder method defined in the entity bean's local home interface.

A local client that has a reference to an entity object's local interface can do any of the following:

- Invoke business methods on the entity object through the local interface.

- Obtain a reference to the enterprise bean's local home interface.
- Pass the reference as a parameter or return value of a local method call.
- Obtain the entity object's primary key.
- Remove the entity object.

All local references to an entity object that does not exist are invalid. All attempted invocations on an entity object that does not exist result in a `javax.ejb.NoSuchObjectLocalException` being thrown.

A local interface type must not be passed as an argument or result of a remote interface method.

9.8 Primary Key and Object Identity

Every entity object has a unique identity within its home. If two entity objects have the same home and the same primary key, they are considered identical.

The Enterprise JavaBeans architecture allows a primary key class to be any class that is a legal Value Type in RMI-IIOP, subject to the restrictions defined in Subsections 10.6.13 and 12.2.12. The primary key class may be specific to an entity bean class (i.e., each entity bean class may define a different class for its primary key, but it is possible that multiple entity beans use the same primary key class).

A client that holds a reference to an entity object's component interface can determine the entity object's identity within its home by invoking the `getPrimaryKey` method on the reference.

The object identity associated with a reference does not change over the lifetime of the reference. (That is, `getPrimaryKey` always returns the same value when called on the same entity object reference). If an entity object has both a remote home interface and a local home interface, the result of invoking the `getPrimaryKey` method on a reference to the entity object's remote interface and on a reference to the entity object's local interface is the same.

A client can test whether two entity object references refer to the same entity object by using the `isIdentical` method. Alternatively, if a client obtains two entity object references from the same home, it can determine if they refer to the same entity by comparing their primary keys using the `equals` method.

The following code illustrates using the `isIdentical` method to test if two object references refer to the same entity object:

```
Account acc1 = ...;
Account acc2 = ...;

if (acc1.isIdentical(acc2)) {
    // acc1 and acc2 are the same entity object
} else {
    // acc1 and acc2 are different entity objects
}
```

A client that knows the primary key of an entity object can obtain a reference to the entity object by invoking the `findByPrimaryKey(key)` method on the entity bean's home interface.

Note that the Enterprise JavaBeans architecture does not specify "object equality" (i.e. use of the `==` operator) for entity object references. The result of comparing two object references using the Java programming language `Object.equals(Object obj)` method is unspecified. Performing the `Object.hashCode()` method on two object references that represent the entity object is not guaranteed to yield the same result. Therefore, a client should always use the `isIdentical` method to determine if two entity object references refer to the same entity object.

Note that the use of `isIdentical` for the comparison of object references applies to the implementation of the methods of the `java.util.Collection` API as well.

9.9 Entity Bean's Remote Interface

A client can access an entity object through the entity bean's remote interface. An entity bean's remote interface must extend the `javax.ejb.EJBObject` interface. A remote interface defines the business methods that are callable by remote clients.

The following example illustrates the definition of an entity bean's remote interface:

```
public interface Account extends javax.ejb.EJBObject {
    void debit(double amount)
        throws java.rmi.RemoteException,
               InsufficientBalanceException;
    void credit(double amount)
        throws java.rmi.RemoteException;
    double getBalance()
        throws java.rmi.RemoteException;
}
```

The `javax.ejb.EJBObject` interface defines the methods that allow the client to perform the following operations on an entity object's reference:

- Obtain the remote home interface for the entity object.
- Remove the entity object.

- Obtain the entity object's handle.
- Obtain the entity object's primary key.

The container provides the implementation of the methods defined in the `javax.ejb.EJBObject` interface. Only the business methods are delegated to the instances of the enterprise bean class.

Note that the entity object does not expose the methods of the `javax.ejb.EnterpriseBean` interface to the client. These methods are not intended for the client—they are used by the container to manage the enterprise bean instances.

9.10 Entity Bean's Local Interface

A local client can access an entity object through the entity bean's local interface. An entity bean's local interface must extend the `javax.ejb.EJBLocalObject` interface. A local interface defines the business methods that are callable by local clients.

The following example illustrates the definition of an entity bean's local interface:

```
public interface Account extends javax.ejb.EJBLocalObject {  
    void debit(double amount)  
        throws InsufficientBalanceException;  
    void credit(double amount);  
    double getBalance();  
}
```

Note that the methods of the entity bean's local interface must not throw the `java.rmi.RemoteException`.

The `javax.ejb.EJBLocalObject` interface defines the methods that allow the local client to perform the following operations on an entity object's local reference:

- Obtain the local home interface for the entity object.
- Remove the entity object.
- Obtain the entity object's primary key.

The container provides the implementation of the methods defined in the `javax.ejb.EJBLocalObject` interface. Only the business methods are delegated to the instances of the enterprise bean class.

Note that the entity object does not expose the methods of the `javax.ejb.EntityBean` or the optional `javax.ejb.TimedObject` interface to the local client. These methods are not intended for the local client—they are used by the container to manage the enterprise bean instances.

9.11 Entity Bean's Handle

An entity object's handle is an object that identifies the entity object on a network. A client that has a reference to an entity object's remote interface can obtain the entity object's handle by invoking the `getHandle` method on the remote interface. The `getHandle` method is only available on the remote interface.

Since a handle class extends `java.io.Serializable`, a client may serialize the handle. The client may use the serialized handle later, possibly in a different process or even system, to re-obtain a reference to the entity object identified by the handle.

The client code must use the `javax.rmi.PortableRemoteObject.narrow` method to convert the result of the `getEJBObject` method invoked on a handle to the entity bean's remote interface type.

The lifetime and scope of a handle is specific to the handle implementation. At the minimum, a program running in one JVM must be able to obtain and serialize the handle, and another program running in a different JVM must be able to deserialize it and re-create an object reference. An entity handle is typically implemented to be usable over a long period of time—it must be usable at least across a server restart.

Containers that store long-lived entities will typically provide handle implementations that allow clients to store a handle for a long time (possibly many years). Such a handle will be usable even if parts of the technology used by the container (e.g. ORB, DBMS, server) have been upgraded or replaced while the client has stored the handle. Support for this "quality of service" is not required by the EJB specification.

An EJB container is not required to accept a handle that was generated by another vendor's EJB container.

The use of a handle is illustrated by the following example:

```
// A client obtains a handle of an account entity object and
// stores the handle in stable storage.
//
OutputStream stream = ...;
Account account = ...;
Handle handle = account.getHandle();
stream.writeObject(handle);

// A client can read the handle from stable storage, and use the
// handle to resurrect an object reference to the
// account entity object.
//
InputStream stream = ...;
Handle handle = (Handle) stream.readObject(handle);
Account account = (Account) javax.rmi.PortableRemoteObject.narrow(
    handle.getEJBObject(), Account.class);
account.debit(100.00);
```

A handle is not a capability, in the security sense, that would automatically grant its holder the right to invoke methods on the object. When a reference to an object is obtained from a handle, and then a method on the object is invoked, the container performs the usual access checks based on the caller's principal.

9.12 Entity Home Handles

The EJB specification allows a client to obtain a handle for the remote home interface. The client can use the home handle to store a reference to an entity bean's remote home interface in stable storage, and re-create the reference later. This handle functionality may be useful to a client that needs to use the remote home interface in the future, but does not know the JNDI name of the remote home interface.

A handle to a remote home interface must implement the `javax.ejb.HomeHandle` interface.

The client code must use the `javax.rmi.PortableRemoteObject.narrow` method to convert the result of the `getEJBHome` method invoked on a handle to the home interface type.

The lifetime and scope of a handle is specific to the handle implementation. At a minimum, a program running in one JVM must be able to serialize the handle, and another program running in a different JVM must be able to deserialize it and re-create an object reference. An entity handle is typically implemented to be usable over a long period of time—it must be usable at least across a server restart.

9.13 Type Narrowing of Object References

A client program that is intended to be interoperable with all compliant EJB container implementations must use the `javax.rmi.PortableRemoteObject.narrow` method to perform type-narrowing of the client-side representations of the remote home and remote interfaces.

Note: Programs that use the cast operator to narrow the remote and remote home interfaces are likely to fail if the container implementation uses RMI-IIOP as the underlying communication transport.

Chapter 10

Entity Bean Component Contract for Container-Managed Persistence

The entity bean component contract for container-managed persistence is the contract between an entity bean and its container. It defines the life cycle of the entity bean instances, the model for method delegation of the business methods invoked by the client, and the model for the management of the entity bean's persistent state and relationships. The main goal of this contract is to ensure that an entity bean component using container-managed persistence is portable across all compliant EJB containers.

This chapter defines the Enterprise Bean Provider's view of this contract and responsibilities of the Container Provider for managing the life cycle of the enterprise bean instances and their persistent state and relationships.

10.1 Overview

In accordance with the architecture for container-managed persistence, the Bean Provider develops a set of entity beans for an application, and determines the relationships among them. The Bean Provider designs an abstract persistence schema for each entity bean, which defines its container-managed fields and relationships, and determines the methods for accessing them. The entity bean instance accesses its container-managed fields and relationships at runtime by means of the methods defined for its abstract persistence schema.

The abstract persistence schema is specified in the deployment descriptor that is produced by the Bean Provider. The Deployer, using the Container Provider's tools, determines how the persistent fields and relationships defined by the abstract persistence schema are mapped to a database or other persistent store, and generates the necessary additional classes and interfaces that enable the container to manage the persistent fields and relationships of the entity bean instances at runtime.

This chapter describes the component contract for an entity bean with container-managed persistence, and how data independence is maintained between the entity bean instance and its representation in the persistent store. It describes this contract from the viewpoints of both the Bean Provider and the container.

The entity bean component contract for container-managed persistence has been substantially changed in the EJB 2.0 specification. Entity beans that use the EJB 1.1 component contract for container-managed persistence must still be supported by EJB 2.0 and later containers. However, the contracts are separate, and the Bean Provider must choose one or the other. The EJB 1.1 entity bean contract for container-managed persistence is defined in Chapter 14, "EJB 1.1 Entity Bean Component Contract for Container-Managed Persistence".

10.2 Container-Managed Entity Persistence and Data Independence

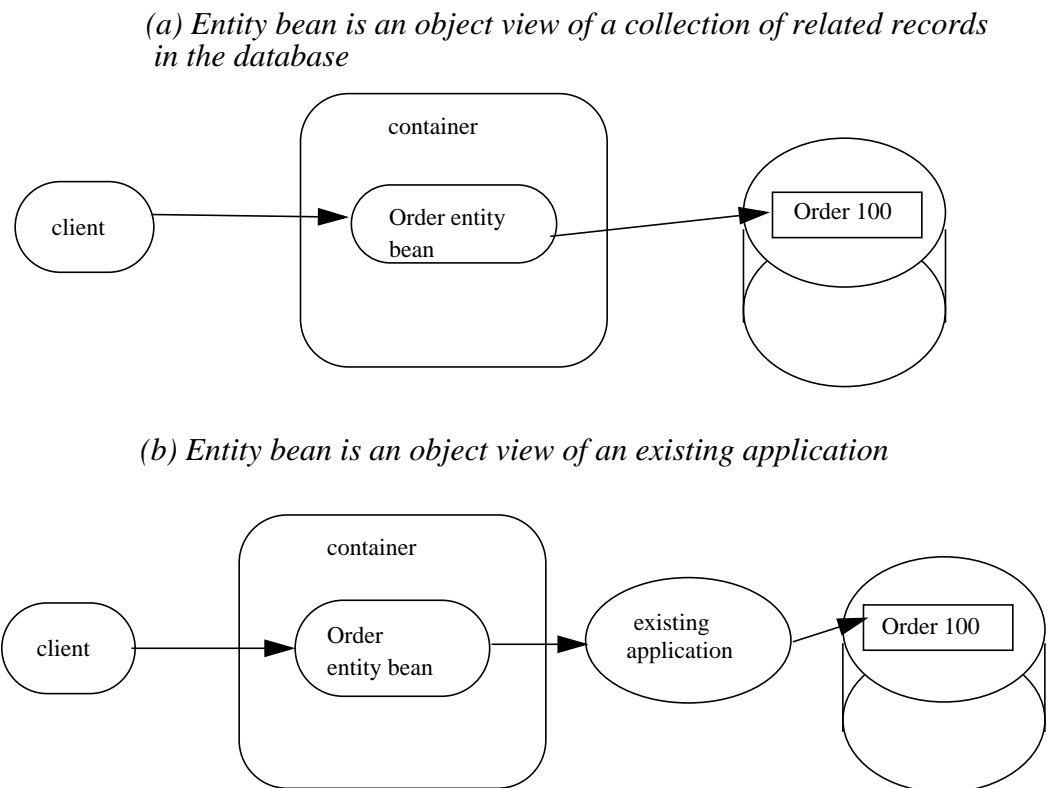
The EJB component model provides a separation between the client view of a bean (as presented by its home and component interfaces) and the entity bean class (which provides the implementation of the client view). The EJB architecture for container-managed persistence adds to this a separation between the entity bean class (as defined by the Bean Provider) and its persistent representation. The container-managed persistence architecture thus provides not only a layer of data independence between the client view of a bean as an *entity object* and the Bean Provider's internal view of the bean in terms of the entity bean instance, but also between the entity bean instance and its persistent representation. This allows an entity bean to be evolved independently from its clients, without requiring the redefinition or recompilation of those clients, and it allows an entity bean to be redeployed across different containers and different persistent data stores, without requiring the redefinition or recompilation of the entity bean class.

In container-managed persistence, unlike in bean-managed persistence, the Bean Provider does not write database access calls in the methods of the entity bean class. Instead, persistence is handled by the container at runtime. The entity Bean Provider must specify in the deployment descriptor those persistent fields and relationships for which the container must handle data access. The Bean Provider codes all persistent data access by using the accessor methods that are defined for the abstract persistence schema. The implementation of the persistent fields and relationships, as well as all data access, is deferred to the container.

It is the responsibility of the Deployer to map the abstract persistence schema of a set of interrelated entity bean classes and into the physical schema used by the underlying data store (e.g., into a relational schema) by using the Container Provider's tools. The Deployer uses the deployment descriptor as input to the Container Provider's tools to perform this mapping. The Container Provider's tools are also used to generate the concrete implementation of the entity bean classes, including the code that delegates calls to the accessor methods of the entity bean class to the runtime persistent data access layer of the container.

The EJB deployment descriptor describes *logical* relationships among entity beans. It does not provide a mechanism for specifying how the abstract persistence schema of an entity bean or of a set of interrelated entity beans is to be mapped to an underlying database. This is the responsibility of the Deployer, who, using the Container Provider's tools, uses the logical relationships that are specified in the deployment descriptor to map to the physical relationships that are specific to the underlying resource. It is the responsibility of the container to manage the mapping between the logical and physical relationships at runtime and to manage the referential integrity of the relationships.

The advantage of using container-managed persistence is that the entity bean can be logically independent of the data source in which the entity is stored. The Container Provider's tools can, for example, generate classes that use JDBC or SQLJ to access the entity state in a relational database; classes that implement access to a non-relational data source, such as an IMS database; or classes that implement function calls to existing enterprise applications. These tools are typically specific to each data source.

Figure 24 View of Underlying Data Sources Accessed Through Entity Bean

10.3 The Entity Bean Provider's View of Container-Managed Persistence

An entity bean implements an object view of a business entity or set of business entities stored in an underlying database or implemented by an existing enterprise application (for example, by a mainframe program or by an ERP application).

An entity bean with container-managed persistence typically consists of its entity bean class; a component interface which defines its client view business methods; a home interface which defines the create, remove, home, and finder methods of its client view; and its abstract persistence schema as specified in the deployment descriptor.

A client of an entity bean can control the life cycle of a bean by using the bean's home interface and can manipulate the bean as a business entity by using the methods defined by its component interface. The home and component interfaces of a bean define its client view.

An entity bean with container-managed persistence typically has container-managed relationships with other container-managed persistence entity beans, as defined by the `relationships` element of the deployment descriptor. The architecture for container-managed persistence thus allows the Bean Provider to implement complex applications by defining a complex abstract persistence schema encompassing multiple entity bean classes related by means of container-managed relationships.

An entity bean accesses related entity beans by means of the accessor methods for its container-managed relationship fields, which are specified by the `cmr-field` elements of its abstract persistence schema defined in the deployment descriptor. Entity bean relationships are defined in terms of the local interfaces of the related beans, and the view an entity bean presents to its related beans is defined by its local home and local interfaces. Thus, an entity bean can be the target of a relationship from another entity bean only if it has a local interface.

The Bean Provider programming an application that uses container-managed persistence typically avoids calls to the methods of the remote home and remote interfaces in favor of invoking related beans by means of the methods of their local interfaces. Unlike remote method calls, such internal method invocations are made using call-by-reference and commonly do not involve the checking of method permissions.

The Enterprise JavaBeans architecture for container-managed persistence provides great flexibility to the Bean Provider in designing an application.

For example, a group of related entity beans—`Order`, `LineItem`, and `Customer`—might all be defined as having only local interfaces, with a remotable session bean containing the business logic that drives their invocation. The individual entity beans form a coordinated whole that provides an interrelated set of services that are exposed by their several home and component interfaces. The services provided by the local network of persistent entities is exposed to the remote client view through the home and remote interfaces of the session bean, which offers a coarser grained remote service.

Alternatively, a single entity bean might represent an independent, remotable business object that forms a unit of distribution that is designed to be referenced remotely by multiple enterprise beans and/or other remote clients. Such a remotable entity bean might make use of other entity beans within its local scope to further model its complex internal state. For example, an `Order` entity bean might make use of a `LineItem` entity bean internally, not exposing it to remote clients. In this case, the `Order` entity bean might define both a remote and a local component interface, where the local interface is presented only to the related entity beans, such as `LineItem`, and the remote interface is presented to session beans and/or web-tier clients.

10.3.1 The Entity Bean Provider's Programming Contract

The Bean Provider must observe the following programming contract when defining an entity bean class that uses container-managed persistence:

- The Bean Provider must define the entity bean class as an abstract class. The container provides the implementation class that is used at runtime.
- The container-managed persistent fields and container-managed relationship fields must *not* be defined in the entity bean class. From the perspective of the Bean Provider, the container-man-

aged persistent fields and container-managed relationship fields are virtual fields only, and are accessed through get and set accessor methods. The implementation of the container-managed persistent fields and container-managed relationship fields is supplied by the container.

- The container-managed persistent fields and container-managed relationship fields must be specified in the deployment descriptor using the `cmp-field` and `cmr-field` elements respectively. The names of these fields must be valid Java identifiers and must begin with a lowercase letter, as determined by `java.lang.Character.isLowerCase`.
- The Bean Provider must define the accessor methods for the container-managed persistent fields and container-managed relationship fields as get and set methods, using the JavaBeans conventions. The implementation of the accessor methods is supplied by the container.
- The accessor methods must be public, must be abstract, and must bear the name of the container-managed persistent field (`cmp-field`) or container-managed relationship field (`cmr-field`) that is specified in the deployment descriptor, and in which the first letter of the name of the `cmp-field` or `cmr-field` has been uppercased and prefixed by “get” or “set”.
- The accessor methods for a container-managed relationship field must be defined in terms of the local interface of the related entity bean, as described in Section 10.3.2.
- The accessor methods for container-managed relationship fields for one-to-many or many-to-many relationships must utilize one of the following Collection interfaces: `java.util.Collection` or `java.util.Set`^[14]. The Collection interfaces used in relationships are specified in the deployment descriptor. The implementation of the collection classes used for the container-managed relationship fields is supplied by the container.
- An entity bean local interface type (or a collection of such) can be the type of a `cmr-field`. An entity bean local interface type (or a collection of such) cannot be the type of a `cmp-field`.
- The accessor methods for the container-managed relationship fields must not be exposed in the remote interface of an entity bean.
- The local interface types of the entity bean and of related entity beans must not be exposed through the remote interface of the entity bean.
- The collection classes that are used for container-managed relationships must not be exposed through the remote interface of the entity bean.
- Once the primary key for an entity bean has been set, the Bean Provider must not attempt to change it by use of set accessor methods on the primary key `cmp-fields`. The Bean Provider should therefore not expose the set accessor methods for the primary key `cmp-fields` in the component interface of the entity bean.
- The Bean Provider must ensure that the Java types assigned to the `cmp-fields` are restricted to the following: Java primitive types and Java serializable types.

[14] We expect to include `java.util.List` and `java.util.Map` in a later version of this specification.

10.3.2 The Entity Bean Provider's View of Persistent Relationships

An entity bean may have relationships with other entity beans with container-managed persistence.

Relationships may be one-to-one, one-to-many, or many-to-many relationships.

Container-managed relationships can exist only among entity beans within the same local relationship scope, as defined by the `relationships` element in the deployment descriptor. Container-managed relationships are defined in terms of the local interfaces of the related beans.

Relationships may be either bidirectional or unidirectional. If a relationship is bidirectional, it can be navigated in both directions, whereas a unidirectional relationship can be navigated in one direction only.

A unidirectional relationship is implemented with a `cmr-field` on the entity bean instance from which navigation can take place, and no related `cmr-field` on the entity bean instance that is the target of the relationship. Unidirectional relationships are typically used when the Bean Provider wishes to restrict the visibility of a relationship.

An entity bean that does not have a local interface can have only unidirectional relationships from itself to other entity beans. The lack of a local interface prevents other entity beans from having a relationship to it.

The bean developer navigates or manipulates relationships by using the get and set accessor methods for the container-managed relationship fields and the `java.util.Collection` API for collection-valued container-managed relationship fields.

The Bean Provider must consider the type and cardinality of relationships when the entity bean classes are programmed. The get method for a `cmr-field` must return either the local interface of the entity bean or a collection (either `java.util.Collection` or `java.util.Set`) of the same. The set method for the relationship must take as an argument the entity bean's local interface or a collection of the same.

10.3.3 Dependent Value Classes

A dependent value class is a concrete class that is the value of a `cmp-field`. A dependent value class may be a class that the Bean Provider wishes to use internally within an entity bean with container-managed persistence, and/or it may be a class that the Bean Provider chooses to expose through the remote (or local) interface of the entity bean.

A dependent value class can be the value of a `cmp-field`; it cannot be the value of a `cmr-field`.

The get accessor method for a `cmp-field` that corresponds to a dependent value class returns a *copy* of the dependent value class instance. The assignment of a dependent value class value to a `cmp-field` using the set accessor method causes the value to be copied to the target `cmp-field`.

A dependent value class must be serializable. The internal structure of a dependent value class is not described in the EJB deployment descriptor.

10.3.4 Remove Protocols

The Bean Provider can specify the removal of an entity object in two ways:

- By the use of a `remove` method on the entity bean's component interface or home interface.
- By the use of a `cascade-delete` specification in the deployment descriptor.

10.3.4.1 Remove Methods

When the `remove` method is invoked on an entity object, the container must invoke the entity Bean Provider's `ejbRemove` method as described in Section 10.5.3. After the Bean Provider's `ejbRemove` method returns (and prior to returning to the client), the container must remove the entity object from all relationships in which it participates, and then remove its persistent representation. ^[15]

- Once an entity has been removed from a relationship, the accessor methods for any relationships to the entity will reflect this removal. An accessor method for a one-to-one or many-to-one relationship to the entity will return null; and an accessor method for a many-to-many relationship to the entity will return a collection from which the entity object has been removed.
- The container must detect any subsequent attempt to invoke an accessor method on the removed entity object and throw the `java.rmi.NoSuchObjectException` if the client is a remote client or the `javax.ejb.NoSuchObjectLocalException` if the client is a local client. The container must detect an attempt to assign a removed entity object as the value of a cmr-field of another object (whether as an argument to a set accessor method or as an argument to a method of the `java.util.Collection` API) and throw the `java.lang.IllegalArgumentException`.

After removing the entity object from all relationships and removing its persistent representation, the container must then cascade the removal to all entity beans with which the entity had been previously in container-managed relationships for which the `cascade-delete` option was specified.

More than one relationship may be affected by the removal of an entity object, as in the following example. Once the shipping address object used by the `Order` bean has been removed, the billing address accessor method will also return null.

```
public void changeAddress()
    Address a = createAddress();
    setShippingAddress(a);
    setBillingAddress(a);
    //both relationships now reference the same entity object
    getShippingAddress().remove();
    if (getBillingAddress() == null) // it must be
        ...
    else ...
        // this is impossible....
```

[15] At this point it must appear to the application that the entity has been removed from the persistent store. If the container employs an optimistic caching strategy and defers the removal of the entity from the database (e.g., to the end of transaction), this must be invisible to the application.

The `remove` method, alone, causes only the entity on which it is invoked to be removed. It does not cause the deletion to be cascaded to other entity objects. In order for the deletion of one entity object to be automatically cascaded to another, the `cascade-delete` mechanism should be used.

10.3.4.2 Cascade-delete

The `cascade-delete` deployment descriptor element is used within a particular relationship to specify that the lifetime of one or more entity objects is dependent upon the lifetime of another entity object.

The `cascade-delete` deployment descriptor element is contained within the `ejb-relationship-role` element. The `cascade-delete` element can only be specified for an `ejb-relationship-role` element contained in an `ejb-relation` element if the *other* `ejb-relationship-role` element in the same `ejb-relation` element specifies a multiplicity of One. The `cascade-delete` option cannot be specified for a many-to-many relationship. The deletion of one entity object can only be cascaded to cause the deletion of other entity objects if the first entity object is in a one-to-one or one-to-many relationship with those other entity objects.

If an entity is deleted, and the `cascade-delete` deployment descriptor element is specified for a related entity bean, then the removal is cascaded to cause the removal of the related entity object or objects. As with the `remove` operation, the removal triggered by the `cascade-delete` option causes the container to invoke the `ejbRemove` method on the entity bean instance that is to be removed before the persistent representation of that entity object is removed. Once an entity has been removed from a relationship because of a cascaded delete, the accessor methods for any relationships to the entity will reflect this removal. An accessor method for a one-to-one or many-to-one relationship to the entity will return null; and an accessor method for a many-to-many relationship to the entity will return a collection from which the entity object has been removed. After removing the entity object from all relationships and removing its persistent representation, the container must then cascade the removal to all entity beans with which the entity had been previously been in container-managed relationships for which the `cascade-delete` option was specified.

The use of `cascade-delete` causes only the entity object or objects in the relationship for which it is specified to be deleted. It does not cause the deletion to be further cascaded to other entity objects, unless they are participants in relationship roles for which `cascade-delete` has also been specified.

10.3.5 Identity of Entity Objects

From the viewpoint of the Bean Provider, entity objects have a runtime object identity that is maintained by the container.

The container maintains the persistent identity of an entity object on the basis of its primary key.

The primary key of an entity bean may or may not be visible as one or more `cmp-fields` of the instance, depending on the way in which it is specified. The Bean Provider specifies the primary key as described in Section 10.8. Once it has been set, the Bean Provider must not attempt to change the value of a primary key field by means of a set method on its `cmp-fields`.

When a new instance of an entity bean whose primary key fields are visible in the entity bean class is created, the Bean Provider must use the `ejbCreate<METHOD>` method to set all the primary key fields of the entity bean instance before the instance can participate in a relationship, e.g. be used in a set accessor method for a cmr-field. The Bean Provider must not reset a primary key value by means of a set method on any of its cmp-fields after it has been set in the `ejbCreate<METHOD>` method. If the Bean Provider attempts to reset a primary key value, the container must throw the `java.lang.IllegalStateException`.

Note that the container's implementation of the referential integrity semantics for container-managed relationships must not cause the value of the primary key to change.

The Bean Provider should not use untrimmed or blank-padded string-valued primary key fields. Use of untrimmed primary key fields may cause comparison operations based on primary keys to fail, and may result in non-portable behavior. If untrimmed strings are used in primary key fields or other cmp-fields, the container or database system may trim them.

10.3.6 Semantics of Assignment for Relationships

The assignment operations for container-managed relationships have a special semantics that is determined by the referential integrity semantics for the relationship multiplicity.

In the case of a one-to-one relationship, when the Bean Provider uses a set accessor method to assign an object from a cmr-field in one instance to a cmr-field of the *same relationship type* (i.e., as defined by the `ejb-relation` and `ejb-relationship-role` deployment descriptor elements) in another instance, the object is effectively *moved* and the value of the source cmr-field is set to null in the same transaction context. If the argument to the set accessor method is not of the same type as the cmr-field, the container must throw the `java.lang.IllegalArgumentException`.

In the case of a one-to-many or many-to-many relationship, either the `java.util.Collection` API or a set accessor method may be used to manipulate the contents of a collection-valued cmr-field. These two approaches are discussed below.

10.3.6.1 Use of the `java.util.Collection` API to Update Relationships

The methods of the `java.util.Collection` API for the container-managed collections used for collection-valued cmr-fields have the usual semantics, with the following exception: the `add` and `addAll` methods applied to container-managed collections in one-to-many relationships have a special semantics that is determined by the referential integrity of one-to-many relationships.

- If the argument to the `add` method is already an element of a collection-valued relationship field of the *same relationship type* as the target collection (as defined by the `ejb-relation` and `ejb-relationship-role` deployment descriptor elements), it is removed from this first relationship and added, in the same transaction context, to the target relationship (i.e., it is effectively moved from one collection of the relationship type to the other). For example, if there is a one-to-many relationship between field offices and sales representatives, adding a sales representative to a new field office will have the effect of removing him or her from his or her current field office. If the argument to the `add` method is not an element of a collection-valued relationship of the *same relationship type*, it is simply added to the target collection and not removed from its current collection, if any.

- The `addAll` method, when applied to a target collection in a one-to-many relationship, has similar semantics, applied to the members of its collection argument individually.

Note that in the case of many-to-many relationships, adding an element or elements to the contents of a collection-valued cmr-field has no effect on the source collection, if any. For example, if there is a many-to-many relationship between customers and sales representatives, a customer can be added to the set of customers handled by a particular sales representative without affecting the set of customers handled by any other sales representative.

When the `java.util.Collection` API is used to manipulate the contents of container-managed relationship fields, the argument to any `Collection` method defined with a single `Object` parameter must be of the element type of the collection defined for the target cmr-field. The argument for any collection-valued parameter must be a `java.util.Collection` (or `java.util.Set`), all of whose elements are of the element type of the collection defined for the target cmr-field. If an argument is not of the correct type for the relationship, the container must throw the `java.lang.IllegalArgumentException`.

The Bean Provider should exercise caution when using an `Iterator` over a collection in a container-managed relationship. In particular, the Bean Provider should not modify the container-managed collection while the iteration is in progress in any way that causes elements to be added or removed, other than by the `java.util.Iterator.remove()` method. If elements are added or removed from the underlying container-managed collection used by an iterator other than by the `java.util.Iterator.remove()` method, the container should throw the `java.lang.IllegalStateException` on the next operation on the iterator.

The following example illustrates how operations on container-managed relationships that affect the contents of a collection-valued cmr-field viewed through an iterator can be avoided. Because there is a one-to-many relationship between field offices and sales representatives, adding a sales representative to a new field office causes the sales representative to be removed from the current field office.

```
Collection nySalesreps = nyOffice.getSalesreps();
Collection sfSalesreps = sfOffice.getSalesreps();

Iterator i = nySalesreps.iterator();
Salesrep salesrep;

// a wrong way to transfer the salesrep
while (i.hasNext()) {
    salesrep = (Salesrep)i.next();
    sfSalesreps.add(salesrep); // removes salesrep from nyOffice
}

// this is a correct and safe way to transfer the salesrep
while (i.hasNext()) {
    salesrep = (Salesrep)i.next();
    i.remove();
    sfSalesreps.add(salesrep);
}
```

10.3.6.2 Use of Set Accessor Methods to Update Relationships

The semantics of a set accessor method, when applied to a collection-valued cmr-field, is also determined by the referential integrity semantics associated with the multiplicity of the relationship. The identity of the collection object referenced by a cmr-field does not change when a set accessor method is executed.

In the case of a one-to-many relationship, if a collection of entity objects is assigned from a cmr-field of in one instance to a cmr-field of the same relationship type in another instance, the objects in the collection are effectively moved. The contents of the collection of the target instance are replaced with the contents of the collection of the source instance, but the *identity* of the collection object containing the instances in the relationship does not change. The source cmr-field references the same collection object as before (i.e., the identity of the collection object is preserved), but the collection is empty.

The Bean Provider can thus use the set method to move objects between the collections referenced by cmr-fields of the same relationship type in different instances. The set accessor method, when applied to a cmr-field in a one-to-many relationship thus has the semantics of the `java.util.Collection` methods `clear`, followed by `addAll`, applied to the target collection; and `clear`, applied to the source collection. It is the responsibility of the container to transfer the contents of the collection instances in the same transaction context.

Note that if the collection that is passed to the cmr setter method is an unmanaged collection (i.e., not itself the value of a collection-valued cmr-field), the same requirements apply in the case that the collection contains entity objects that already participate in a one-to-many relationship of the same relationship type as the target cmr-field.

In the following example, the telephone numbers associated with the billing address of an `Order` bean instance are transferred to the shipping address. Billing address and shipping address are different instances of the same local interface type, `Address`. `Address` is related to `TelephoneNumber` in a one-to-many relationship. The example illustrates how a Bean Provider uses the set method to move a set of instances.

```
public void changeTelephoneNumber() {
    Address a = getShippingAddress();
    Address b = getBillingAddress();
    Collection c = b.getTelephoneNumber();
    a.setTelephoneNumber(b.getTelephoneNumber());
    if (c.isEmpty()) { // must be true...
        ...
    }
}
```

In the case of a many-to-many relationship, if the value of a cmr-field is assigned to a cmr-field of the same relationship type in another instance, the objects in the collection of the first instance are assigned as the contents of the cmr-field of the second instance. The identities of the collection objects referenced by the cmr-fields do not change. The contents of the collections are shared, but not the collections themselves. The set accessor method, when applied to a cmr-field in a many-to-many relationship thus has the semantics of the `java.util.Collection` methods `clear`, followed by `addAll`, applied to the target collection.

For example, if there is a many-to-many relationship between customers and sales representatives, assigning the set of customers of one sales representative to the another sales representative will result in both sales representatives handling the same customers. If the second sales representative originally handled a different group of customers, those customers will no longer be handled by that sales representative.

```
public void shareCustomers(SalesRep rep) {  
    setCustomers(rep.getCustomers());  
    // the customers are shared among the sales reps  
}
```

The following section, 10.3.7, “Assignment Rules for Relationships”, defines the semantics of assignment for relationships in further detail.

10.3.7 Assignment Rules for Relationships


This section defines the semantics of assignment and collection manipulation in one-to-one, one-to-many, and many-to-many container-managed relationships.

The figures make use of two entity beans, with local interface types A and B. Instances with local interface type A are typically designated as a_1, \dots, a_n ; instances with local interface type B are typically designated as b_1, \dots, b_m . Interface A exposes accessor methods `getB` and `setB` for navigable relationships with B: `getB` returns an instance of B or a collection of instances of B, depending on the multiplicity of the relationship. Similarly, B exposes accessor methods `getA` and `setA` for navigable relationships with A.


All changes in each subsection are assumed to be applied to the figure labeled “Before change” at the beginning of the subsection (i.e., changes are not cumulative). The results of changes are designated graphically as well as in conditional expressions expressed in the JavaTM programming language.

10.3.7.1 One-to-one Bidirectional Relationships

A and B are in a one-to-one bidirectional relationship:



Before change:



Before change:

```

B b1 = a1.getB();
B b2 = a2.getB();
  
```

Change:

```


a1.setB(a2.getB());
  
```

Expected result:

```


(b2.isIdentical(a1.getB())) &&
(a2.getB() == null) &&
(b1.getA() == null) &&
(a1.isIdentical(b2.getA()))
  
```

After change:




10.3.7.2 One-to-one Unidirectional Relationships

A and B are in a one-to-one unidirectional relationship:



Before change:



Before change:

```

B b1 = a1.getB();
B b2 = a2.getB();
  
```

Change:

```


a1.setB(a2.getB());
  
```

Expected result:

```


(b2.isIdentical(a1.getB())) && (a2.getB() == null)
  
```

After change:




10.3.7.3 One-to-many Bidirectional Relationships

A and B are in a one-to-many bidirectional relationship:



Before change:



Before change:

```

Collection b1 = a1.getB();
Collection b2 = a2.getB();
B b11, b12, ... , b1n; // members of b1
B b21, b22, ... , b2m; // members of b2
  
```

Change:


```

a1.setB(a2.getB());
  
```

Expected result:

```
(a2.getB().isEmpty()) &&
(b2.isEmpty()) &&
(b1 == a1.getB()) &&
(b2 == a2.getB()) &&
(a1.getB().contains(b21)) &&
(a1.getB().contains(b22)) && ... &&
(a1.getB().contains(b2m)) &&
(b11.getA() == null) &&
(b12.getA() == null) && ... &&
(b1n.getA() == null) &&
(a1.isIdentical(b21.getA())) &&
(a1.isIdentical(b22.getA())) && ... &&
(a1.isIdentical(b2m.getA()))
```

After change:




Change:

```
b2m.setA(b1n.getA());
```

Expected result:

```
(b1.contains(b11)) &&
(b1.contains(b12)) && ... &&
(b1.contains(b1n)) &&
(b1.contains(b2m)) &&
(b2.contains(b21)) &&
(b2.contains(b22)) && ... &&
(b2.contains(b2m_1)) &&
(a1.isIdentical(b11.getA())) &&
(a1.isIdentical(b12.getA())) && ... &&
(a1.isIdentical(b1n.getA())) &&
(a2.isIdentical(b21.getA())) &&
(a2.isIdentical(b22.getA())) && ... &&
(a2.isIdentical(b2m_1.getA())) &&
(a1.isIdentical(b2m.getA()))
```

After change:




Change:

```
a1.getB().add(b2m);
```

Expected result:

```
(b1.contains(b11)) &&
(b1.contains(b12)) && ... &&
(b1.contains(b1n)) &&
(b1.contains(b2m)) &&
(b2.contains(b21)) &&
(b2.contains(b22)) && ... &&
(b2.contains(b2m_1)) &&
(a1.isIdentical(b11.getA())) &&
(a1.isIdentical(b12.getA())) && ... &&
(a1.isIdentical(b1n.getA())) &&
(a2.isIdentical(b21.getA())) &&
(a2.isIdentical(b22.getA())) && ... &&
(a2.isIdentical(b2m_1.getA())) &&
(a1.isIdentical(b2m.getA()))
```

After change:




Change:

```
a1.getB().remove(b1n);
```

Expected result:


```
(b1n.getA() == null) &&
(b1 == a1.getB()) &&
(b1.contains(b11)) &&
(b1.contains(b12)) && ... &&
(b1.contains(b1n_1)) &&
!(b1.contains(b1n))
```

After change:




10.3.7.4 One-to-many Unidirectional Relationships

A and B are in a one-to-many unidirectional relationship:



Before change:



Before change:

```
Collection b1 = a1.getB();
Collection b2 = a2.getB();
B b11, b12, ... , b1n; // members of b1
B b21, b22, ... , b2m; // members of b2
```


Change:

```
a1.setB(a2.getB());
```

Expected result:

```
(a2.getB().isEmpty()) &&
(b2.isEmpty()) &&
(b1 == a1.getB()) &&
(b2 == a2.getB()) &&
(a1.getB().contains(b21)) &&
(a1.getB().contains(b22)) && ... &&
(a1.getB().contains(b2m))
```

After change:




Change:

```
a1.getB().add(b2m);
```

Expected result:

```
(b1 == a1.getB()) &&
(b1.contains(b2m))
```

After change:




Change:

```
a1.getB().remove(b1n);
```

Expected result:


```
(a1.getB().contains(b11)) &&  
(a1.getB().contains(b12)) && ... &&  
(a1.getB().contains(b1n_1)) &&  
!(a1.getB().contains(b1n)) &&
```

After change:




10.3.7.5 Many-to-one Unidirectional Relationships

A and B are in a many-to-one unidirectional relationship:



Before change:



Before change:

```

B b11, b12, ... , b1n;
B b21, b22, ... , b2m;
// the following is true
// (a1.isIdentical(b11.getA())) && ... && (a1.isIdentical(b1n.getA()
)) &&
// (a2.isIdentical(b21.getA())) && ... && (a2.isIdentical(b2m.getA()
))
  
```

Change:


```

b1j.setA(b2k.getA());
  
```

Expected result:


```
(a1.isIdentical(b11.getA())) &&
(a1.isIdentical(b12.getA())) &&
...
(a2.isIdentical(b1j.getA())) &&
...
(a1.isIdentical(b1n.getA())) &&
(a2.isIdentical(b21.getA())) &&
(a2.isIdentical(b22.getA())) &&
...
(a2.isIdentical(b2k.getA())) &&
...
(a2.isIdentical(b2m.getA()))
```

After change:




10.3.7.6 Many-to-many Bidirectional Relationships

A and B are in a many-to-many bidirectional relationship:



Before change:



Before change the following holds:

```
(a1.getB().contains(b1)) &&
(a1.getB().contains(b2)) &&
(a2.getB().contains(b1)) &&
(a2.getB().contains(b2)) &&
(a2.getB().contains(b3)) &&
(a3.getB().contains(b2)) &&
(a3.getB().contains(b3)) &&
(a3.getB().contains(b4)) &&
(a4.getB().contains(b3)) &&
(a4.getB().contains(b4)) &&
(a4.getB().contains(b5)) &&
(a5.getB().contains(b4)) &&
(a5.getB().contains(b5)) &&
(b1.getA().contains(a1)) &&
(b1.getA().contains(a2)) &&
(b2.getA().contains(a1)) &&
(b2.getA().contains(a2)) &&
(b2.getA().contains(a3)) &&
(b3.getA().contains(a2)) &&
(b3.getA().contains(a3)) &&
(b3.getA().contains(a4)) &&
(b4.getA().contains(a3)) &&
(b4.getA().contains(a4)) &&
(b4.getA().contains(a5)) &&
(b5.getA().contains(a4)) &&
(b5.getA().contains(a5)) &&
```


Change:

```
a1.setB(a3.getB());
```

Expected result:

```
(a1.getB().contains(b2)) &&
(a1.getB().contains(b3)) &&
(a1.getB().contains(b4)) &&
(a3.getB().contains(b2)) &&
(a3.getB().contains(b3)) &&
(a3.getB().contains(b4)) &&
(b1.getA().contains(a2)) &&
(b2.getA().contains(a1)) &&
(b2.getA().contains(a2)) &&
(b2.getA().contains(a3)) &&
(b3.getA().contains(a1)) &&
(b3.getA().contains(a2)) &&
(b3.getA().contains(a3)) &&
(b3.getA().contains(a4)) &&
(b4.getA().contains(a1)) &&
(b4.getA().contains(a3)) &&
(b4.getA().contains(a4)) &&
(b4.getA().contains(a5))
```

After change:




Change:

```
a1.getB().add(b3);
```

Expected result:

```
(a1.getB().contains(b1)) &&  
(a1.getB().contains(b2)) &&  
(a1.getB().contains(b3)) &&  
(b3.getA().contains(a1)) &&  
(b3.getA().contains(a2)) &&  
(b3.getA().contains(a3)) &&  
(b3.getA().contains(a4)) &&
```


After change:

Change:

```
a2.getB().remove(b2);
```


Expected result:

```
(a2.getB().contains(b1)) &&
(a2.getB().contains(b3)) &&
(b2.getA().contains(a1)) &&
(b2.getA().contains(a3))
```


After change:

10.3.7.7 Many-to-many Unidirectional Relationships

A and B are in a many-to-many unidirectional relationship:



Before change:



Before change the following holds:

```

(a1.getB().contains(b1)) &&
(a1.getB().contains(b2)) &&
(a2.getB().contains(b1)) &&
(a2.getB().contains(b2)) &&
(a2.getB().contains(b3)) &&
(a3.getB().contains(b2)) &&
(a3.getB().contains(b3)) &&
(a3.getB().contains(b4)) &&
(a4.getB().contains(b3)) &&
(a4.getB().contains(b4)) &&
(a4.getB().contains(b5)) &&
(a5.getB().contains(b4)) &&
(a5.getB().contains(b5)) &&

```

Change:

```


a1.setB(a3.getB());

```

Expected Result:

```
(a1.getB().contains(b2)) &&  
(a1.getB().contains(b3)) &&  
(a1.getB().contains(b4)) &&  
(a3.getB().contains(b2)) &&  
(a3.getB().contains(b3)) &&  
(a3.getB().contains(b4)) &&
```

After change:




Change:

```
a1.getB().add(b3);
```

Expected result:

```
(a1.getB().contains(b1)) &&  
(a1.getB().contains(b2)) &&  
(a1.getB().contains(b3))
```

After change:




Change:

```
a2.getB().remove(b2);
```

Expected result:

```
(a2.getB().contains(b1)) &&  
(a2.getB().contains(b3))
```

After change:



10.3.8 Collections Managed by the Container

The collections that are used in the representation of one-to-many and many-to-many container-managed relationships are implemented and managed by the container. The following semantics apply to these collections:

- It is the responsibility of the container to preserve the runtime identity of the collection objects used in container-managed relationships.
- There is no constructor available to the Bean Provider for the container-managed collections.
- If there are no related values for a given container-managed relationship, the `get` accessor method for that cmr-field returns an empty collection (and not `null`).
- It is the responsibility of the container to raise the `java.lang.IllegalArgumentException` if the Bean Provider attempts to assign `null` as the value of a collection-valued cmr-field by means of the `set` accessor method.
- It is the responsibility of the container to ensure that when the `java.util.Collection` API is used to manipulate the contents of container-managed relationship fields, the argument to any `Collection` method defined with a single `Object` parameter must be of the element type of the collection defined for the target cmr-field. The argument for any collection-valued parameter must be a `java.util.Collection` (or `java.util.Set`), all of whose elements are of the element type of the collection defined for the target cmr-field. If an argument is not of the correct type for the relationship, the container must throw the `java.lang.IllegalArgumentException`.
- It is the responsibility of the container to throw the `java.lang.IllegalStateException` if an attempt is made to modify a container-managed collection corresponding to a multivalued cmr-field using the `java.util.Collection` API outside of the transaction context in which the collection object was initially materialized.
- It is the responsibility of the container to throw the `java.lang.IllegalStateException` if an attempt is made to use a `java.util.Iterator` for a container-managed collection in a transaction context other than that in which the iterator was obtained.

10.3.9 Non-persistent State

The Bean Provider may use instance variables in the entity bean instance to maintain non-persistent state, e.g. a JMS connection.

The Bean Provider can use instance variables to store values that depend on the persistent state of the entity bean instance, although this use is not encouraged. The Bean Provider should use the `ejbLoad` method to resynchronize the values of any instance variables that depend on the entity bean's persistent state. In general, any non-persistent state that depends on the persistent state of an entity bean should be recomputed during the `ejbLoad` method.

The Bean Provider should exercise care in passing the contents of instance variables as the arguments or results of method invocations when local interfaces are used. In general, the Bean Provider should avoid passing state that is maintained in instance variables as the argument or result of a local method invocation.

10.3.10 The Relationship Between the Internal View and the Client View

In designing the entity bean, the Bean Provider should keep in mind the following:

- The classes that are exposed by the remote interface are decoupled from the persistence layer. Instances of these classes are passed to and from the client by value.
- The classes that are exposed by the local interface of the bean may be tightly coupled to the bean's internal state. Instances of these classes are passed to and from the client by reference and may therefore be modified by the client. The Bean Provider should exercise care in determining what is exposed through the local interface of the bean.

10.3.10.1 Restrictions on Remote Interfaces

The following restrictions apply to the remote interface of an entity bean with container-managed persistence.

- The Bean Provider must not expose the get and set methods for container-managed relationship fields or the persistent `Collection` classes that are used in container-managed relationships through the remote interface of the bean.
- The Bean Provider must not expose local interface types or local home interface types through the remote interface or remote home interface of the bean.
- The Bean Provider must not expose the container-managed collection classes that are used for relationships through the remote interface of the bean.
- The Bean Provider must not expose timers or timer handles through the remote interface of the bean.

Dependent value classes can be exposed in the remote interface or remote home interface and can be included in the client `ejb-jar` file.

The Bean Provider is free to expose get and set methods that correspond to `cmp`-fields of the entity bean through the bean's remote interface.

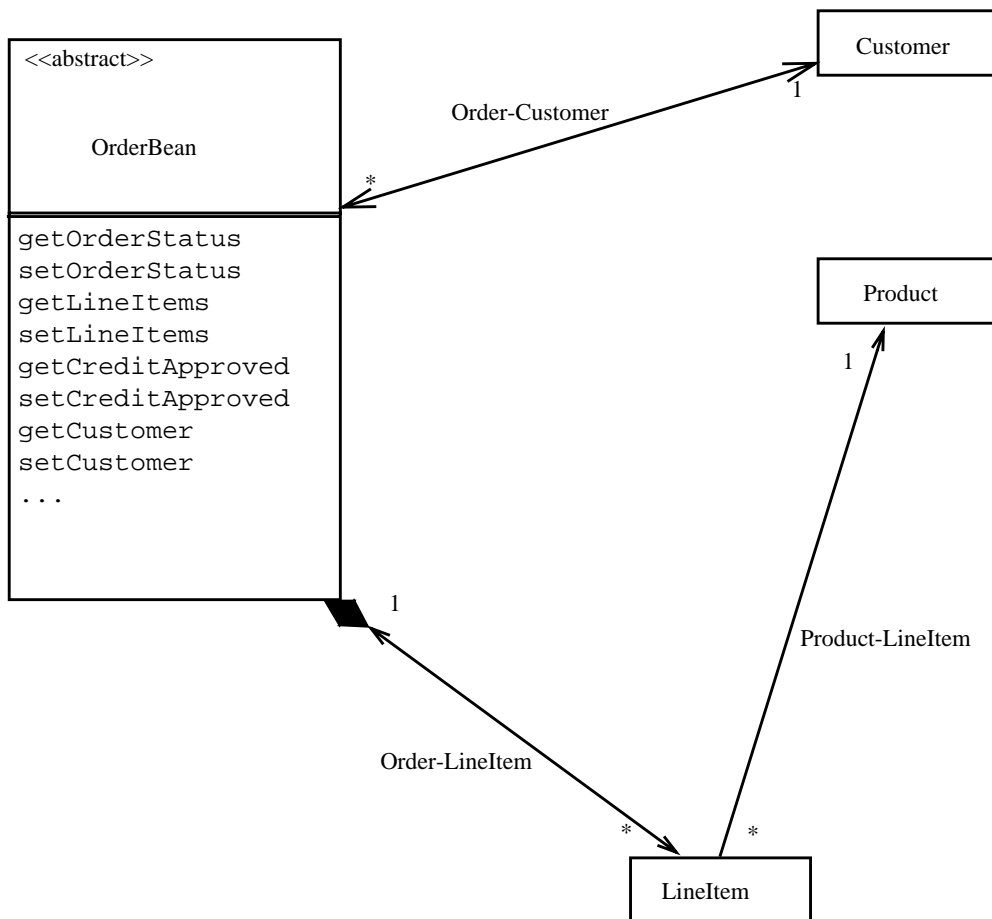
10.3.11 Mapping Data to a Persistent Store

This specification does not prescribe how the abstract persistence schema of an entity bean should be mapped to a relational (or other) schema of a persistent store, or define how such a mapping is described.

10.3.12 Example

Figure 25 illustrates an `Order` entity bean with relationships to line items and customers, which are other entity beans within the same local scope. `Product` is indirectly related to `Order` by means of the relationship between `LineItem` and `Product`. Sample code for the `OrderBean` class follows the figure.

Figure 25 Relationship Example



```
package com.acme.order;

// This example shows the implementation of OrderBean, the
// entity bean class for the OrderEJB entity bean. OrderEJB has
// container-managed relationships with the entity beans
// CustomerEJB and LineItemEJB.
// This example illustrates the use of local interfaces.

import java.util.Collection;
import java.util.Vector;
import java.util.Date;

import javax.naming.*;

public abstract class OrderBean implements javax.ejb.EntityBean {

    private javax.ejb.EntityContext context;

    // define status codes for processing

    static final int BACKORDER = 1;
    static final int SHIPPED = 2;
    static final int UNSHIPPED = 3;

    // get and set methods for the cmp fields

    public abstract int getOrderStatus();
    public abstract void setOrderStatus(int orderStatus);

    public abstract boolean getCreditApproved();
    public abstract void setCreditApproved(boolean creditapproved);

    public abstract Date getOrderDate();
    public abstract void setOrderDate(Date orderDate);

    // get and set methods for the relationship fields

    public abstract Collection getLineItems();
    public abstract void setLineItems(Collection lineitems);

    public abstract Customer getCustomer();
    public abstract void setCustomer(Customer customer);

    // business methods.

    // addLineItem:
    // This method is used to add a line item.
    // It creates the lineitem object and adds it to the
    // persistent managed relationship.

    public void addLineItem(Product product,
                           int quantity,
                           Address address)
        throws InsufficientInfoException
    {
        // create a new line item
    }
}
```

```

if (validAddress(address)) {
    // Address is a legacy class. It is a dependent value
    // class that is available both in the client and in
    // the entity bean, and is serializable.
    // We will use the address as the value of a cmp field
    // of lineItem.

    try {
        Context ic = new InitialContext();
        LineItemLocalHome litemLocalHome =
            (LineItemLocalHome)ic.lookup("LineItemEJB");
        LineItem litem = litemLocalHome.create();

        litem.setProduct(product);
        litem.setQuantity(quantity);
        litem.setTax(calculateTax(product.getPrice(),
                                   quantity,
                                   address));
        litem.setStatus(UNSHIPPED);
        // set the address for the line item to be shipped
        litem.setAddress(address);
        // The lineItem entity bean uses a dependent value
        // class to represent the dates for the order status.
        // This class holds shipment date, expected shipment
        // date, credit approval date, and inventory
        // dates which are internal to the order fulfillment
        // process. Not all this information will be available
        // to the client.

        Dates dates = new Dates();
        litem.setDates(dates);
        getLineItems().add(litem);
    } catch (Exception someexception) {}
} else {
    throw new InsufficientInfoException();
}
}

// getOrderLineItems:
// This method makes a view of the lineitems that are in this
// order available in the client. It makes only the relevant
// information visible to the client and hides the internal
// details of the representation of the lineitem
public Collection getOrderLineItems() {
    Vector clientlineitems = new Vector();
    Collection lineitems = getLineItems();
    java.util.Iterator iterator = lineitems.iterator();
    // ClientLineItem is a value class that is used in
    // the client view.
    // The entity bean provider abstracts from the persistent
    // representation of the line item to construct the client
    // view.
    ClientLineItem clitem;
    while (iterator.hasNext()) {
        LineItem litem = (LineItem)iterator.next();
        clitem = new ClientLineItem();
        // only the name of the product is available in the
        // client view
    }
}

```



```

        clitem.setProductName(litem.getProduct().getName());
        clitem.setQuantity(litem.getQuantity());
        // the client view gets a specific descriptive message
        // depending on the line item status.
        clitem.setCurrentStatus(
            statusCodeToString(litem.getStatus()));
        // address is not copied to the client view.
        // as this class includes other information with
        // respect to the order handling that should not be
        // available to the client. Only the relevant info
        // is copied.
        int lineitemStatus = litem.getStatus();
        if ( lineitemStatus == BACKORDER) {
            clitem.setShipDate(
                litem.getDates().getExpectedShipDate());
        } else if (lineitemStatus == SHIPPED) {
            clitem.setShipDate(
                litem.getDates().getShippedDate());
        }
        //add the new line item
        clientlineitems.add(clitem);
    }
    // return the value objects to the client
    return clientlineitems;
}

// other methods internal to the entity bean class
...

// other javax.ejb.EntityBean methods
...
}

```

10.3.13 The Bean Provider's View of the Deployment Descriptor

The persistent fields (cmp-fields) and relationships (cmr-fields) of an entity bean must be declared in the deployment descriptor.

The deployment descriptor provides the following information about the abstract persistence schemas of entity beans and their container-managed relationships:

- An `ejb-name` element for each entity bean. The `ejb-name` must be a valid Java identifier and must be unique within the `ejb-name` elements of the `ejb-jar` file.
- An `abstract-schema-name` element for each entity bean. The `abstract-schema-name` must be a valid Java identifier and must be unique within the `abstract-schema-name` elements of the `ejb-jar` file. The `abstract-schema-name` element is used in the specification of EJB QL queries.
- A set of `ejb-relation` elements, each of which contains a pair of `ejb-relation-ship-role` elements to describe the two roles in the relationship.^[16]

- Each `ejb-relationship-role` element describes a relationship role: its name, its multiplicity within a relation, and its navigability. It specifies the name of the `cmr-field` that is used from the perspective of the relationship participant. The `cmr-field-type` element must be specified if the type of the `cmr-field` is `java.util.Collection` or `java.util.Set`. Each relationship role refers to an entity bean by means of an `ejb-name` element contained in the `relationship-role-source` element.

[16] The relation names and the relationship role names are not used in the code provided by the Bean Provider.

The following example shows a deployment descriptor segment that defines the abstract persistence schema for a set of related entity beans. The deployment descriptor elements for container-managed persistence and relationships are described further in Chapter 23.

```
<ejb-jar>

...

<enterprise-beans>
...
</enterprise-beans>

<relationships>

<!--
ONE-TO-MANY: Order LineItem
-->

    <ejb-relation>
        <ejb-relation-name>Order-LineItem</ejb-relation-name>
        <ejb-relationship-role>
            <ejb-relationship-role-name>
                order-has-lineitems
            </ejb-relationship-role-name>
            <multiplicity>One</multiplicity>
            <relationship-role-source>
                <ejb-name>OrderEJB</ejb-name>
            </relationship-role-source>
            <cmr-field>
                <cmr-field-name>lineItems</cmr-field-name>
                <cmr-field-type>java.util.Collection
            </cmr-field-type>
            </cmr-field>
        </ejb-relationship-role>

        <ejb-relationship-role>
            <ejb-relationship-role-name>lineitem-belongsto-order
            </ejb-relationship-role-name>
            <multiplicity>Many</multiplicity>
            <cascade-delete/>
            <relationship-role-source>
                <ejb-name>LineItemEJB</ejb-name>
            </relationship-role-source>
            <cmr-field>
                <cmr-field-name>order</cmr-field-name>
            </cmr-field>
        </ejb-relationship-role>
    </ejb-relation>

<!--
ONE-TO-MANY unidirectional relationship:
Product is not aware of its relationship with LineItem
-->

    <ejb-relation>
        <ejb-relation-name>Product-LineItem</ejb-relation-name>

        <ejb-relationship-role>
            <ejb-relationship-role-name>
```

```

        product-has-lineitems
      </ejb-relationship-role-name>
      <multiplicity>One</multiplicity>
      <relationship-role-source>
        <ejb-name>ProductEJB</ejb-name>
      </relationship-role-source>
      <!-- since Product does not know about LineItem
           there is no cmr field in Product for accessing
           Lineitem
      -->
    </ejb-relationship-role>

    <ejb-relationship-role>
      <ejb-relationship-role-name>
        lineitem-for-product
      </ejb-relationship-role-name>
      <multiplicity>Many</multiplicity>
      <relationship-role-source>
        <ejb-name>LineItemEJB</ejb-name>
      </relationship-role-source>
      <cmr-field>
        <cmr-field-name>product</cmr-field-name>
      </cmr-field>
    </ejb-relationship-role>
  </ejb-relation>

  <!--
  ONE-TO-MANY: Order Customer:
  -->

  <ejb-relation>
    <ejb-relation-name>Order-Customer</ejb-relation-name>

    <ejb-relationship-role>
      <ejb-relationship-role-name>
        customer-has-orders
      </ejb-relationship-role-name>
      <multiplicity>One</multiplicity>
      <relationship-role-source>
        <ejb-name>CustomerEJB</ejb-name>
      </relationship-role-source>
      <cmr-field>
        <cmr-field-name>orders</cmr-field-name>
        <cmr-field-type>java.util.Collection
      </cmr-field-type>
      </cmr-field>
    </ejb-relationship-role>

    <ejb-relationship-role>
      <ejb-relationship-role-name>
        order-belongsto-customer
      </ejb-relationship-role-name>
      <multiplicity>Many</multiplicity>
      <relationship-role-source>
        <ejb-name>OrderEJB</ejb-name>
      </relationship-role-source>
      <cmr-field>

```

```
        <cmr-field-name>customer</cmr-field-name>
      </cmr-field>
    </ejb-relationship-role>
  </ejb-relation>

</relationships>

...

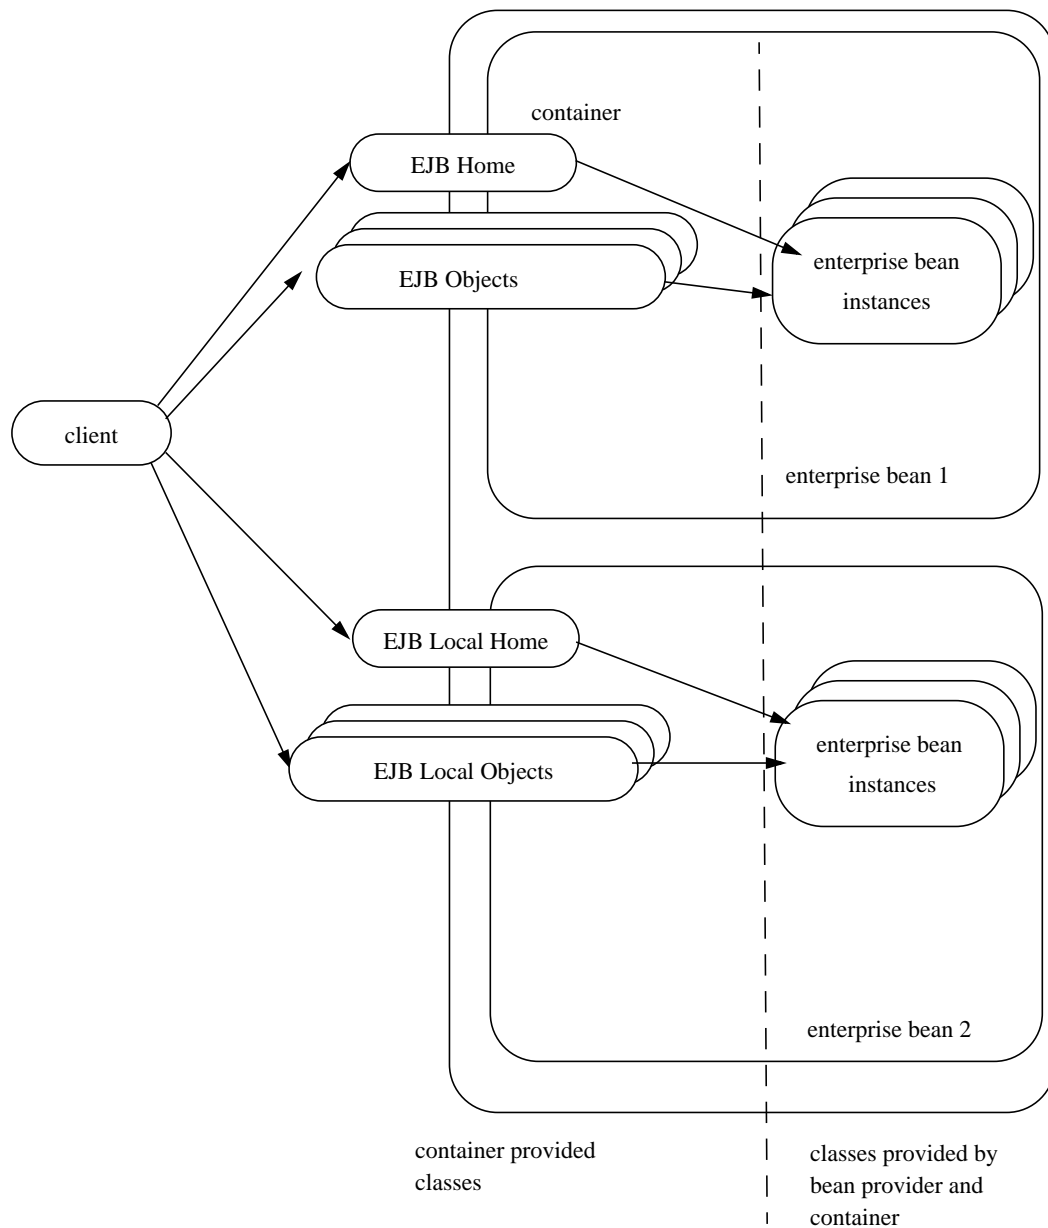
</ejb-jar>
```

10.4 The Entity Bean Component Contract

This section specifies the container-managed persistence contract between an entity bean and its container.

10.4.1 Runtime Execution Model of Entity Beans

This subsection describes the runtime model and the classes used in the description of the contract between an entity bean and its container. Figure 26 shows an overview of the runtime model. The client of an entity bean may be a local client or it may be a remote client.

Figure 26 Overview of the Entity Bean Runtime Execution Model

An enterprise bean is an object whose class is provided by the Bean Provider. The class of an entity bean with container-managed persistence is abstract. The concrete bean class is generated by the Container Provider's tools at deployment time. The container is also responsible for providing the implementation of the `java.util.Collection` classes that are used in maintaining the container-managed relationships of the entity bean.

An entity **EJBObject** or **EJBLocalObject** is an object whose class was generated at deployment time by the Container Provider's tools. A client never references an entity bean instance directly—a client always references an entity **EJBObject** or **EJBLocalObject** whose class is generated by the Container Provider's tools. The entity **EJBObject** class implements an entity bean's remote interface. The entity **EJBLocalObject** class implements an entity bean's local interface. A related entity bean never references another entity bean instance directly—a related entity bean, like any other local client of an entity bean, always references an entity **EJBLocalObject** whose class is generated by the Container Provider's tools.

An entity **EJBHome** or **EJBLocalHome** object provides life cycle operations (create, find, remove) for its entity objects as well as home business methods, which are business methods that are not specific to an entity bean instance. The class for the entity **EJBHome** or **EJBLocalHome** object is generated by the Container Provider's tools at deployment time. The entity **EJBHome** or **EJBLocalHome** object implements the entity bean's remote or local home interface that was defined by the Bean Provider.

10.4.2 Container Responsibilities

The following are the container responsibilities for the management of persistent state.

10.4.2.1 Container-Managed Fields

An entity bean with container-managed persistence relies on the container to perform persistent data access on behalf of the entity bean instances. The container transfers data between an entity bean instance and the underlying resource manager. The container also implements the creation, removal, and lookup of the entity object in the underlying database.

The container transfers data between the entity bean and the underlying data source as a result of the execution of the entity bean's methods. Because of the requirement that all data access occur through the accessor methods, the container can implement both eager and lazy loading and storing schemes.

The container is responsible for implementing the entity bean class by providing the implementation of the get and set accessor methods for its abstract persistence schema. The container is allowed to use Java serialization to store the container-managed persistent fields (cmp-fields).

The container must also manage the mapping between primary keys and **EJBLocalObjects** or **EJBObjects**. If both a remote and a local interface are specified for the entity bean, the container must manage the mapping between **EJBObjects** and **EJBLocalObjects**.

Because the container is free to optimize the delivery of persistent data to the bean instance (for example, by the use of lazy loading strategies), the contents of the entity bean instance and the contents of container-managed collections may not be fully materialized.

10.4.2.2 Container-Managed Relationships

The container maintains the relationships among entity beans.

- It is the responsibility of the container to maintain the referential integrity of the container-managed relationships, as described in Section 10.3.6, in accordance with the semantics of the relationship type as specified in the deployment descriptor. For example, if an entity

bean is added to a collection corresponding to the container-managed relationship field of another entity bean, the container-managed relationship field of the first entity bean must also be updated by the container in the same transaction context.

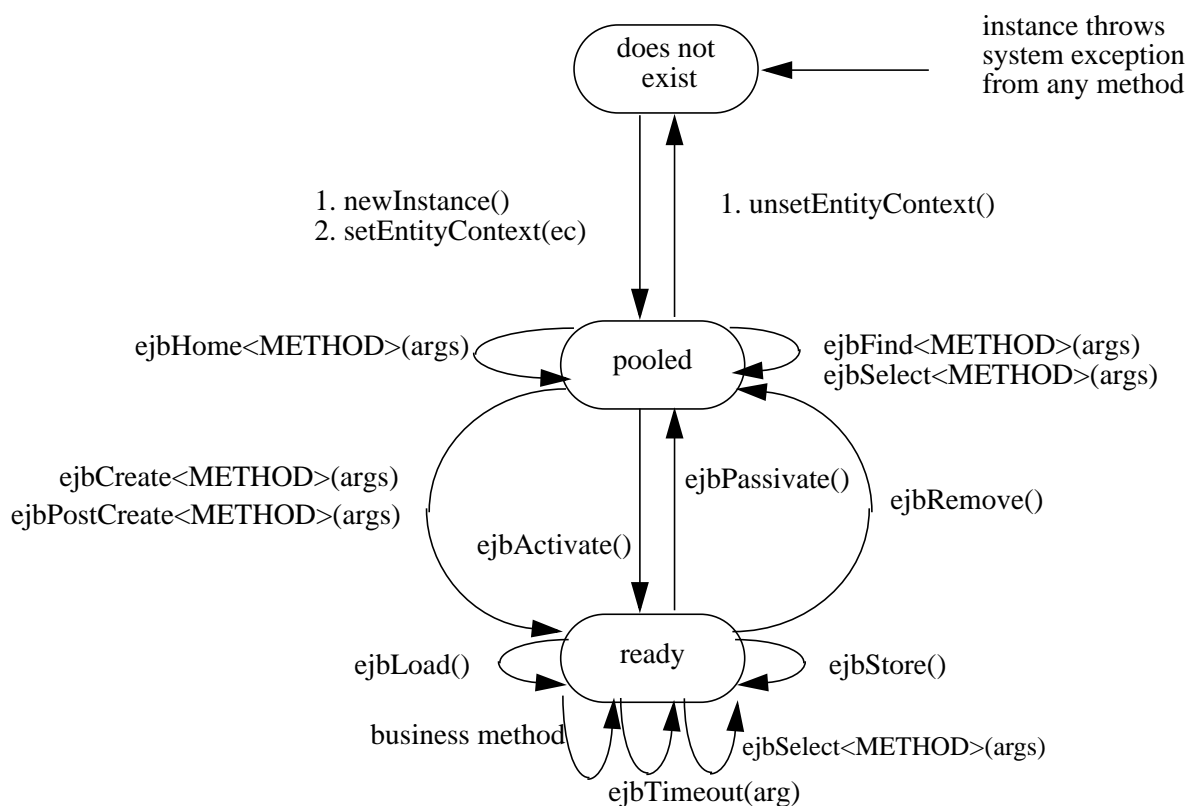
- It is the responsibility of the container to throw the `java.lang.IllegalArgumentException` when the argument to a set method in a relationship is an instance of the wrong relationship type or a collection containing instances of the wrong type, or when an argument to a method of the `java.util.Collection` API used to manipulate a collection-valued container-managed relationship field is an instance of the wrong type or a collection that contains instances of the wrong type (see Section 10.3.6).
- It is the responsibility of the container to throw the `java.lang.IllegalStateException` when a method of the `java.util.Collection` API is used to access a collection-valued cmr-field within a transaction context other than the transaction context in which the cmr-field was initially materialized. For example, if the container-managed collection is returned as the result of a local interface method with transaction attribute `RequiresNew`, and the client attempts to access the collection, the container must throw the `IllegalStateException`.
- It is the responsibility of the container to throw the `java.lang.IllegalStateException` when a `java.util.Iterator` is used to access a collection-valued cmr-field within a transaction context other than the transaction context in which the iterator was initially obtained.

10.5 Instance Life Cycle Contract Between the Bean and the Container

This section describes the part of the component contract between the entity bean and the container that relates to the management of the entity bean instance's life cycle.

10.5.1 Instance Life Cycle

Figure 27 Life Cycle of an Entity Bean Instance.



An entity bean instance is in one of the following three states:

- It does not exist.
- Pooled state. An instance in the pooled state is not associated with any particular entity object identity.
- Ready state. An instance in the ready state is assigned an entity object identity.

The following steps describe the life cycle of an entity bean instance:

- An entity bean instance's life starts when the container creates the instance using `newInstance()`. The container then invokes the `setEntityContext` method to pass the instance a reference to the `EntityContext` interface. The `EntityContext` interface allows the

instance to invoke services provided by the container and to obtain the information about the caller of a client-invoked method.

- The instance enters the pool of available instances. Each entity bean has its own pool. While the instance is in the available pool, the instance is not associated with any particular entity object identity. All instances in the pool are considered equivalent, and therefore any instance can be assigned by the container to any entity object identity at the transition to the ready state. While the instance is in the pooled state, the container may use the instance to execute any of the entity bean's finder methods (shown as `ejbFind<METHOD>` in the diagram) or any of the entity bean's home methods (shown `ejbHome<METHOD>` in the diagram). The instance does *not* move to the ready state during the execution of a finder or a home method. An `ejbSelect<METHOD>` method may be called by an entity bean's home method while the instance is in the pooled state.
- An instance transitions from the pooled state to the ready state when the container selects that instance to service a client call to an entity object or an `ejbTimeout` method. There are two possible transitions from the pooled to the ready state: through the `ejbCreate<METHOD>` and `ejbPostCreate<METHOD>` methods, or through the `ejbActivate` method. The container invokes the `ejbCreate<METHOD>` and `ejbPostCreate<METHOD>` methods when the instance is assigned to an entity object during entity object creation (i.e., when the client invokes a create method on the entity bean's home object). The container invokes the `ejbActivate` method on an instance when an instance needs to be activated to service an invocation on an existing entity object—this occurs because there is no suitable instance in the ready state to service the client's call or the `ejbTimeout` method.
- When an entity bean instance is in the ready state, the instance is associated with a specific entity object identity. While the instance is in the ready state, the container can synchronize the state of the instance with the state of the entity in the underlying data source whenever it determines the need to, in the process invoking the `ejbLoad` and `ejbStore` methods zero or more times. A business method can be invoked on the instance zero or more times. The `ejbTimeout` method can be invoked on the instance zero or more times. Invocations of the `ejbLoad` and `ejbStore` methods can be arbitrarily mixed with invocations of business methods and `ejbTimeout` method invocations. An `ejbSelect<METHOD>` method can be called by a business method (or `ejbLoad` or `ejbStore` method or `ejbTimeout` method) while the instance is in the ready state.
- The container can choose to passivate an entity bean instance within a transaction. To passivate an instance, the container first invokes the `ejbStore` method to allow the instance to prepare itself for the synchronization of the database state with the instance's state, and then the container invokes the `ejbPassivate` method to return the instance to the pooled state.
- Eventually, the container will transition the instance to the pooled state. There are three possible transitions from the ready to the pooled state: through the `ejbPassivate` method, through the `ejbRemove` method, and because of a transaction rollback for `ejbCreate`, `ejbPostCreate`, or `ejbRemove` (not shown in Figure 27). The container invokes the `ejbPassivate` method when the container wants to disassociate the instance from the entity object identity without removing the entity object. The container invokes the `ejbRemove` method when the container is removing the entity object (i.e., when the client invoked the `remove` method on the entity object's component interface or a `remove` method on the entity bean's home interface). If `ejbCreate`, `ejbPostCreate`, or `ejbRemove` is called

and the transaction rolls back, the container will transition the bean instance to the pooled state.

- When the instance is put back into the pool, it is no longer associated with an entity object identity. The container can assign the instance to any entity object within the same entity bean home.
- The container can remove an instance in the pool by calling the `unsetEntityContext` method on the instance.

Notes:

1. The `EntityContext` interface passed by the container to the instance in the `setEntityContext` method is an interface, not a class that contains static information. For example, the result of the `EntityContext.getPrimaryKey` method might be different each time an instance moves from the pooled state to the ready state, and the result of the `getCallerPrincipal` and `isCallerInRole` methods may be different in each business method.
2. A `RuntimeException` thrown from any method of an entity bean class (including the business methods and the callbacks invoked by the container) results in the transition to the “does not exist” state. The container must not invoke any method on the instance after a `RuntimeException` has been caught. From the caller’s perspective, the corresponding entity object continues to exist. The client can continue accessing the entity object through its component interface because the container can use a different entity bean instance to delegate the client’s requests. Exception handling is described further in Chapter 18.
3. The container is not required to maintain a pool of instances in the pooled state. The pooling approach is an example of a possible implementation, but it is not the required implementation. Whether the container uses a pool or not has no bearing on the entity bean coding style.

10.5.2 Bean Provider’s Entity Bean Instance’s View

The following describes the entity bean instance’s view of the contract as seen by the Bean Provider:

The entity Bean Provider is responsible for implementing the following methods in the abstract entity bean class:

- A public constructor that takes no arguments.
- `public void setEntityContext(EntityContext ic);`

A container uses this method to pass a reference to the `EntityContext` interface to the entity bean instance. If the entity bean instance needs to use the `EntityContext` interface during its lifetime, it must remember the `EntityContext` interface in an instance variable.

This method executes with an unspecified transaction context (Refer to Subsection 17.6.5 for how the container executes methods with an unspecified transaction context). An identity of an entity object is not available during this method. The entity bean must not attempt to access its persistent state and relationships using the accessor methods during this method.

The instance can take advantage of the `setEntityContext()` method to allocate any resources that are to be held by the instance for its lifetime. Such resources cannot be specific to an entity object identity because the instance might be reused during its lifetime to serve multiple entity object identities.

- `public void unsetEntityContext();`

A container invokes this method before terminating the life of the instance.

This method executes with an unspecified transaction context. An identity of an entity object is not available during this method. The entity bean must not attempt to access its persistent state and relationships using the accessor methods during this method.

The instance can take advantage of the `unsetEntityContext` method to free any resources that are held by the instance. (These resources typically had been allocated by the `setEntityContext` method.)

- `public PrimaryKeyClass ejbCreate<METHOD>(...);`

There are zero^[17] or more `ejbCreate<METHOD>` methods, whose signatures match the signatures of the `create<METHOD>` methods of the entity bean's home interface. The container invokes an `ejbCreate<METHOD>` method on an entity bean instance when a client invokes a matching `create<METHOD>` method on the entity bean's home interface.

The entity Bean Provider's responsibility is to initialize the instance in the `ejbCreate<METHOD>` methods from the input arguments, using the get and set accessor methods, such that when the `ejbCreate<METHOD>` method returns, the persistent representation of the instance can be created. The entity Bean Provider is guaranteed that the values that will be initially returned by the instance's get methods for container-managed fields will be the Java language defaults (e.g. 0 for integer, null for pointers), except for collection-valued cmr-fields, which will have the empty collection (or set) as their value. The entity Bean Provider must not attempt to modify the values of cmr-fields in an `ejbCreate<METHOD>` method. This should be done in the `ejbPostCreate<METHOD>` method instead.

The entity object created by the `ejbCreate<METHOD>` method must have a unique primary key. This means that the primary key must be different from the primary keys of all the existing entity objects within the same home. However, it is legal to reuse the primary key of a previously removed entity object. The implementation of the Bean Provider's `ejbCreate<METHOD>` methods should be coded to return a null.^[18]

An `ejbCreate<METHOD>` method executes in the transaction context determined by the transaction attribute of the matching `create<METHOD>` method. The database insert operations are performed by the container within the same transaction context after the Bean Provider's `ejbCreate<METHOD>` method completes.

- `public void ejbPostCreate<METHOD>(...);`

For each `ejbCreate<METHOD>` method, there is a matching `ejbPostCreate<METHOD>` method that has the same input parameters but whose return type is `void`. The container invokes the matching `ejbPostCreate<METHOD>` method on an instance

[17] An entity bean has no `ejbCreate<METHOD>` and `ejbPostCreate<METHOD>` methods if it does not define any create methods in its home interface. Such an entity bean does not allow its clients to create new EJB objects. The entity bean restricts the clients to accessing entities that were created through direct database inserts.

[18] The above requirement is to allow the creation of an entity bean with bean-managed persistence by subclassing an entity bean with container-managed persistence.

after it invokes the `ejbCreate<METHOD>` method with the same arguments. The instance can discover the primary key by calling `getPrimaryKey` on its entity context object.

The entity object identity is available during the `ejbPostCreate<METHOD>` method. The instance may, for example, obtain the component interface of the associated entity object and pass it to another enterprise bean as a method argument.

The entity Bean Provider may use the `ejbPostCreate<METHOD>` to set the values of cmr-fields to complete the initialization of the entity bean instance.

An `ejbPostCreate<METHOD>` method executes in the same transaction context as the previous `ejbCreate<METHOD>` method.

- `public void ejbActivate();`

The container invokes this method on the instance when the container picks the instance from the pool and assigns it to a specific entity object identity. The `ejbActivate` method gives the entity bean instance the chance to acquire additional resources that it needs while it is in the ready state.

This method executes with an unspecified transaction context. The entity bean must not attempt to access its persistent state or relationships using the accessor methods during this method.

The instance can obtain the identity of the entity object via the `getPrimaryKey`, `getEJBLocalObject`, or `getEJBObject` method on the entity context. The instance can rely on the fact that the primary key and entity object identity will remain associated with the instance until the completion of `ejbPassivate` or `ejbRemove`.

- `public void ejbPassivate();`

The container invokes this method on an instance when the container decides to disassociate the instance from an entity object identity, and to put the instance back into the pool of available instances. The `ejbPassivate` method gives the instance the chance to release any resources that should not be held while the instance is in the pool. (These resources typically had been allocated during the `ejbActivate` method.)

This method executes with an unspecified transaction context. The entity bean must not attempt to access its persistent state or relationships using the accessor methods during this method.

The instance can still obtain the identity of the entity object via the `getPrimaryKey`, `getEJBLocalObject`, or `getEJBObject` method of the `EntityContext` interface.

- `public void ejbRemove();`

The container invokes the `ejbRemove` method on an entity bean instance in response to a client-invoked `remove` operation on the entity bean's home or component interface or as the result of a cascade-delete operation. The instance is in the ready state when `ejbRemove` is invoked and it will be entered into the pool when the method completes.

The entity Bean Provider can use the `ejbRemove` method to implement any actions that must be done before the entity object's persistent representation is removed.

The container synchronizes the instance's state before it invokes the `ejbRemove` method. This means that the state of the instance at the beginning of the `ejbRemove` method is the same as it would be at the beginning of a business method.

This method and the database delete operation(s) execute in the transaction context determined by the transaction attribute of the `remove` method that triggered the `ejbRemove` method.

The instance can still obtain the identity of the entity object via the `getPrimaryKey`, `getEJBLocalObject`, or `getEJBObject` method of the `EntityContext` interface.

After the entity Bean Provider's `ejbRemove` returns, and in the same transaction context, the container removes the entity bean from all relationships in which it participates before removing the entity object's persistent representation.

Since the instance will be entered into the pool, the state of the instance at the end of this method must be equivalent to the state of a passivated instance. This means that the instance must release any resource that it would normally release in the `ejbPassivate` method.

- `public void ejbLoad();`

When the container needs to synchronize the state of an enterprise bean instance with the entity object's persistent state, the container calls the `ejbLoad` method.

The entity Bean Provider can assume that the instance's persistent state has been loaded just before the `ejbLoad` method is invoked. It is the responsibility of the Bean Provider to use the `ejbLoad` method to recompute or initialize the values of any instance variables that depend on the entity bean's persistent state. In general, any transient state that depends on the persistent state of an entity bean should be recalculated using the `ejbLoad` method. The entity bean can use the `ejbLoad` method, for instance, to perform some computation on the values returned by the accessor methods (for example, uncompressing text fields).

This method executes in the transaction context determined by the transaction attribute of the business method or `ejbTimeout` method that triggered the `ejbLoad` method.

- `public void ejbStore();`

When the container needs to synchronize the state of the entity object's persistent state with the state of the enterprise bean instance, the container first calls the `ejbStore` method on the instance.

The entity Bean Provider should use the `ejbStore` method to update the instance using the accessor methods before its persistent state is synchronized. For example, the `ejbStore` method may perform compression of text before the text is stored in the database.

The Bean Provider can assume that after the `ejbStore` method returns, the persistent state of the instance is synchronized.

This method executes in the same transaction context as the previous `ejbLoad` or `ejbCreate` method invoked on the instance. All business methods or the `ejbTimeout` method invoked between the previous `ejbLoad` or `ejbCreate`<METHOD> method and this `ejbStore` method are also invoked in the same transaction context.

- `public <primary key type or collection> ejbFind<METHOD>(...);`

The Bean Provider of an entity bean with container-managed persistence does not write the finder (`ejbFind`<METHOD>) methods.

The finder methods are generated at the entity bean deployment time using the Container Provider's tools. The syntax for the Bean Provider's specification of finder methods is described in Chapter 11, "EJB QL: EJB Query Language for Container-Managed Persistence Query Methods".

- `public <type> ejbHome<METHOD>(...);`

The container invokes this method on the instance when the container selects the instance to execute a matching client-invoked <METHOD> home method. The instance is in the pooled

state (i.e., it is not assigned to any particular entity object identity) when the container selects the instance to execute the `ejbHome<METHOD>` method on it, and it is returned to the pooled state when the execution of the `ejbHome<METHOD>` method completes.

The `ejbHome<METHOD>` method executes in the transaction context determined by the transaction attribute of the matching `<METHOD>` home method, as described in Section 17.6.2.

The entity Bean Provider provides the implementation of the `ejbHome<METHOD>` method. The entity bean must not attempt to access its persistent state or relationships using the accessor methods during this method because a home method is not specific to a particular bean instance.

- `public abstract <type> ejbSelect<METHOD>(...);`

The Bean Provider may provide zero or more select methods. A select method is a query method that is not directly exposed to the client in the home or component interface. The Bean Provider typically calls a select method within a business method.

The Bean Provider defines the select methods as abstract methods.

The select methods are generated at the entity bean deployment time using the Container Provider's tools.

The syntax for the specification of select methods is described in Chapter 11, "EJB QL: EJB Query Language for Container-Managed Persistence Query Methods".

The `ejbSelect<METHOD>` method executes in the transaction context determined by the transaction attribute of the invoking business method.

- `public void ejbTimeout(...);`

The container invokes the `ejbTimeout` method on an instance when a timer for the instance has expired. The `ejbTimeout` method notifies the instance of the time-based event and allows the instance to execute the business logic to handle it.

The `ejbTimeout` method executes in the transaction context determined by its transaction attribute.

10.5.3 Container's View

This subsection describes the container's view of the state management contract. The container must call the following methods:

- `public void setEntityContext(ec);`

The container invokes this method to pass a reference to the `EntityContext` interface to the entity bean instance. The container must invoke this method after it creates the instance, and before it puts the instance into the pool of available instances.

The container invokes this method with an unspecified transaction context. At this point, the `EntityContext` is not associated with any entity object identity.

- `public void unsetEntityContext();`

The container invokes this method when the container wants to reduce the number of instances in the pool. After this method completes, the container must not reuse this instance.

The container invokes this method with an unspecified transaction context.

- `public PrimaryKeyClass ejbCreate<METHOD>(...);`
`public void ejbPostCreate<METHOD>(...);`

The container invokes these two methods during the creation of an entity object as a result of a client invoking a `create<METHOD>` method on the entity bean's home interface.

The container invokes the `ejbCreate<METHOD>` method whose signature matches the `create<METHOD>` method invoked by the client.

Prior to invoking the `ejbCreate<METHOD>` method provided by the Bean Provider, the container must ensure that the values that will be initially returned by the instance's get methods for container-managed fields will be the Java language defaults (e.g. 0 for integer, null for pointers), except for collection-valued cmr-fields, which must have the empty collection (or set) as their value.

The container is responsible for calling the `ejbCreate<METHOD>` method, for obtaining the primary key fields of the newly created entity object persistent representation, and for creating an entity `EJBObject` reference and/or `EJBLocalObject` reference for the newly created entity object. The container must establish the primary key before it invokes the `ejbPostCreate<METHOD>` method.

The entity object created by the `ejbCreate<METHOD>` method must have a unique primary key. This means that the primary key must be different from the primary keys of all the existing entity objects within the same home. However, it is legal to reuse the primary key of a previously removed entity object. The container may, but is not required to, throw the `DuplicateKeyException` on the Bean Provider's attempt to create an entity object with a duplicate primary key^[19].

The container may create the representation of the entity in the database immediately, or it can defer it to a later time (for example to the time after the matching `ejbPostCreate<METHOD>` has been called, or to the end of the transaction), depending on the caching strategy that it uses.

The container then invokes the matching `ejbPostCreate<METHOD>` method with the same arguments on the instance to allow the instance to fully initialize itself. The instance can discover the primary key by calling the `getPrimaryKey` method on its entity context object. Finally, the container returns the entity object's remote interface (i.e., a reference to the entity `EJBObject`) to the client if the client is a remote client or the entity object's local interface (i.e., a reference to the entity `EJBLocalObject`) if the client is a local client.

The container must invoke the `ejbCreate<METHOD>` and `ejbPostCreate<METHOD>` methods and create the representation of the persistent instance in the database in the transaction context determined by the transaction attribute of the matching `create<METHOD>` method, as described in subsection 17.6.2.

- `public void ejbActivate();`

The container invokes this method on an entity bean instance at activation time (i.e., when the instance is taken from the pool and assigned to an entity object identity). The container must ensure that the primary key of the associated entity object is available to the instance if the instance invokes the `getPrimaryKey`, `getEJBLocalObject`, or `getEJBObject` method on its `EntityContext` interface.

The container invokes this method with an unspecified transaction context.

[19] Containers using optimistic caching strategies, for example, may rollback the transaction at a later point.

Note that instance is not yet ready for the delivery of a business method. The container must still invoke the `ejbLoad` method prior to a business method.

- `public void ejbPassivate();`

The container invokes this method on an entity bean instance at passivation time (i.e., when the instance is being disassociated from an entity object identity and moved into the pool). The container must ensure that the identity of the associated entity object is still available to the instance if the instance invokes the `getPrimaryKey`, `getEJBLocalObject`, or `getEJBObject` method on its entity context.

The container invokes this method with an unspecified transaction context.

Note that if the instance state has been updated by a transaction, the container must first invoke the `ejbStore` method on the instance before it invokes `ejbPassivate` on it.

- `public void ejbRemove();`

The container invokes the `ejbRemove` method in response to a client-invoked `remove` operation on the entity bean's home or component interface or as the result of a cascade-delete operation. The instance is in the ready state when `ejbRemove` is invoked and it will be entered into the pool when the method completes.

The container synchronizes the instance's state before it invokes the `ejbRemove` method. This means that the persistent state of the instance at the beginning of the `ejbRemove` method is the same as it would be at the beginning of a business method (i.e., if the instance is not already synchronized from the state in the database, the container must invoke `ejbLoad` before it invokes `ejbRemove`).

The container must ensure that the identity of the associated entity object is still available to the instance in the `ejbRemove` method (i.e., the instance can invoke the `getPrimaryKey`, `getEJBLocalObject`, or `getEJBObject` method on its `EntityContext` in the `ejbRemove` method).

After the entity Bean Provider's `ejbRemove` method returns, and in the same transaction context, the container removes the entity bean instance from all relationships in which it participates and then removes the entity object's persistent representation.

The container may delete the representation of the entity in the database immediately, or it can defer it to a later time (for example to the end of the transaction), depending on the caching strategy that it uses.

The container must ensure that the `ejbRemove` method and database delete operations are performed in the transaction context determined by the transaction attribute of the invoked `remove` method, as described in subsection 17.6.2.

- `public void ejbLoad();`

When the container needs to synchronize the state of an enterprise bean instance with the entity object's state in the database, the container calls the `ejbLoad` method. Depending on its caching strategy, the container may first read the entity object's state from the database, before invoking the `ejbLoad` method, or it may use a lazy loading strategy in making this state visible to the instance.

The exact times that the container invokes `ejbLoad` depend on the configuration of the component and the container, and are not defined by the EJB architecture. Typically, the container will call `ejbLoad` before the first business method within a transaction or before invoking the `ejbTimeout` method on an instance.

The container must invoke this method in the transaction context determined by the transaction attribute of the business method or `ejbTimeout` method that triggered the `ejbLoad` method.

- `public void ejbStore();`

When the container needs to synchronize the state of the entity object in the database with the state of the enterprise bean instance, the container calls the `ejbStore` method on the instance. This synchronization always happens at the end of a transaction. However, the container may also invoke this method when it passivates the instance in the middle of a transaction, or when it needs to transfer the most recent state of the entity object to another instance for the same entity object in the same transaction.

The container must invoke this method in the same transaction context as the previous `ejbLoad`, `ejbCreate<METHOD>`, or `ejbTimeout` method invoked on the instance. All business methods or the `ejbTimeout` method invoked between the previous `ejbLoad` or `ejbCreate <METHOD>` method and this `ejbStore` method are also invoked in the same transaction context.

After the `ejbStore` method returns, the container may store the persistent state of the instance to the database, depending on its caching strategy. If the container uses a lazy storing caching strategy, it is the container's responsibility to write the representation of the persistent object to the database in the same transaction context as that of the `ejbStore` method.

- `public <primary key type or collection> ejbFind<METHOD>(...);`

The implementation of the `ejbFind<METHOD>` method is supplied by the container.

The container invokes the `ejbFind<METHOD>` method on an instance when a client invokes a matching `find<METHOD>` method on the entity bean's home interface. The container must pick an instance that is in the pooled state (i.e., the instance is not associated with any entity object identity) for the execution of the `ejbFind<METHOD>` method. If there is no instance in the pooled state, the container creates one and calls the `setEntityContext` method on the instance before dispatching the finder method.

The container must invoke the `ejbFind<METHOD>` method in the transaction context determined by the transaction attribute of the matching `find` method, as described in subsection 17.6.2.

The container is responsible for ensuring that updates to the states of all entity beans in the same transaction context as the `ejbFind<METHOD>` method and whose abstract schema types are accessed in the method's EJB QL query are visible in the results of the `ejbFind<METHOD>` method. Before invoking the `ejbFind<METHOD>` method, the container must first synchronize the state of those entity bean instances by invoking the `ejbStore` method on them. This requirement does not apply to the `ejbFindByPrimaryKey` method. The results of the `ejbFindByPrimaryKey` method, however, must reflect the entities that have been created or removed within the same transaction context.

After the `ejbFind<METHOD>` method completes, the instance remains in the pooled state. The container may, but is not required to, immediately activate the objects that were located by the finder using the transition through the `ejbActivate` method.

If the `ejbFind<METHOD>` method is declared to return a single primary key, the container creates an entity `EJBObject` (`EJBLocalObject`) reference for the primary key and returns it to the client (local client). If the `ejbFind<METHOD>` method is declared to return a collection of primary keys, the container creates a collection of entity `EJBObject` (`EJBLocalObject`) ref-

ferences for the primary keys returned from the `ejbFind<METHOD>` method, and returns the collection to the client (local client).

The implementations of the finder methods are generated at the entity bean deployment time using the Container Provider's tools.

- `public <type> ejbSelect<METHOD>(...);`

A select method is a query method that is not directly exposed to the client in the home or component interface. The Bean Provider typically calls a select method within a business method or home method.

A select method executes in the transaction context determined by the transaction attribute of the invoking business method.

The container is responsible for ensuring that all updates to the states of all entity beans in the same transaction context as the `ejbSelect<METHOD>` method and whose abstract schema types are accessed in the EJB QL query for the `ejbSelect<METHOD>` method are visible in the results of the `ejbSelect<METHOD>` method by invoking the `ejbStore` method on those entity bean instances.

The implementations of the select methods are generated at the entity bean deployment time using the Container Provider's tools.

- `public <type> ejbHome<METHOD>(...);`

The container invokes the `ejbHome<METHOD>` method on an instance when a client invokes a matching `<METHOD>` home method on the entity bean's home interface. The container must pick an instance that is in the pooled state (i.e., the instance is not associated with any entity object identity) for the execution of the `ejbHome<METHOD>` method. If there is no instance in the pooled state, the container creates one and calls the `setEntityContext` method on the instance before dispatching the home method.

After the `ejbHome<METHOD>` method completes, the instance remains in the pooled state.

The container must invoke the `ejbHome<METHOD>` method in the transaction context determined by the transaction attribute of the matching `<METHOD>` home method, as described in subsection 17.6.2.

- `public void ejbTimeout(...);`

The container invokes the `ejbTimeout` method on the instance when a timer with which the entity has been registered expires. If there is no suitable instance in the ready state, the container must activate an instance, invoking the `ejbActivate` method and transitioning it to the ready state.

The container invokes the `ejbTimeout` method in the context of a transaction determined by its transaction attribute.

10.5.4 The EntityContext Interface

A container provides the entity bean instances with an `EntityContext`, which gives the entity bean instance access to the instance's context maintained by the container. The `EntityContext` interface has the following methods:

- The `getEJBObject` method returns the entity bean's remote interface.

- The `getEJBHome` method returns the entity bean's remote home interface.
- The `getEJBLocalObject` method returns the entity bean's local interface.
- The `getEJBLocalHome` method returns the entity bean's local home interface.
- The `getCallerPrincipal` method returns the `java.security.Principal` that identifies the invoker.
- The `isCallerInRole` method tests if the entity bean instance's caller has a particular role.
- The `setRollbackOnly` method allows the instance to mark the current transaction such that the only outcome of the transaction is a rollback.
- The `getRollbackOnly` method allows the instance to test if the current transaction has been marked for rollback.
- The `getPrimaryKey` method returns the entity bean's primary key.
- The `getTimerService` method returns the `javax.ejb.TimerService` interface.
- The `getUserTransaction` method returns the `javax.transaction.UserTransaction` interface. Entity bean instances must not call this method.

10.5.5 Operations Allowed in the Methods of the Entity Bean Class

Table 4 defines the methods of an entity bean class in which the enterprise bean instances can access the methods of the `javax.ejb.EntityContext` interface, the `java:comp/env` environment naming context, resource managers, `TimerService` and `Timer` methods, and other enterprise beans.

If an entity bean instance attempts to invoke a method of the `EntityContext` interface, and the access is not allowed in Table 4, the container must throw the `java.lang.IllegalStateException`.

If a entity bean instance attempts to invoke a method of the `TimerService` or `Timer` interface and the access is not allowed in Table 4, the container must throw the `java.lang.IllegalStateException`.

If an entity bean instance attempts to access a resource manager or an enterprise bean, and the access is not allowed in Table 4, the behavior is undefined by the EJB architecture.

Table 4 Operations Allowed in the Methods of an Entity Bean

Bean method	Bean method can perform the following operations
constructor	-
setEntityContext unsetEntityContext	EntityContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> JNDI access to <code>java:comp/env</code>
ejbCreate	EntityContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getCallerPrincipal</i> , <i>getRollbackOnly</i> , <i>isCallerInRole</i> , <i>setRollbackOnly</i> , <i>getTimerService</i> JNDI access to <code>java:comp/env</code> Resource manager access Enterprise bean access
ejbPostCreate	EntityContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getCallerPrincipal</i> , <i>getRollbackOnly</i> , <i>isCallerInRole</i> , <i>setRollbackOnly</i> , <i>getEJBObject</i> , <i>getEJBLocalObject</i> , <i>getPrimaryKey</i> , <i>getTimerService</i> JNDI access to <code>java:comp/env</code> Resource manager access Enterprise bean access Timer service or Timer methods
ejbRemove	EntityContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getCallerPrincipal</i> , <i>getRollbackOnly</i> , <i>isCallerInRole</i> , <i>setRollbackOnly</i> , <i>getEJBObject</i> , <i>getEJBLocalObject</i> , <i>getPrimaryKey</i> , <i>getTimerService</i> JNDI access to <code>java:comp/env</code> Resource manager access Enterprise bean access Timer service or Timer methods
ejbHome	EntityContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getCallerPrincipal</i> , <i>getRollbackOnly</i> , <i>isCallerInRole</i> , <i>setRollbackOnly</i> , <i>getTimerService</i> JNDI access to <code>java:comp/env</code> Resource manager access Enterprise bean access
ejbActivate ejbPassivate	EntityContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getEJBObject</i> , <i>getEJBLocalObject</i> , <i>getPrimaryKey</i> , <i>getTimerService</i> JNDI access to <code>java:comp/env</code>

Table 4 Operations Allowed in the Methods of an Entity Bean

Bean method	Bean method can perform the following operations
ejbLoad ejbStore	EntityContext methods: <i>getEJBHome, getEJBLocalHome, getCallerPrincipal, getRollbackOnly, isCallerInRole, setRollbackOnly, getEJBObject, getEJBLocalObject, getPrimaryKey, getTimerService</i> JNDI access to java:comp/env Resource manager access Enterprise bean access Timer service or Timer methods
business method from component interface	EntityContext methods: <i>getEJBHome, getEJBLocalHome, getCallerPrincipal, getRollbackOnly, isCallerInRole, setRollbackOnly, getEJBObject, getEJBLocalObject, getPrimaryKey, getTimerService</i> JNDI access to java:comp/env Resource manager access Enterprise bean access Timer service or Timer methods
ejbTimeout	EntityContext methods: <i>getEJBHome, getEJBLocalHome, getRollbackOnly, setRollbackOnly, getCallerPrincipal, isCallerInRole, getEJBObject, getEJBLocalObject, getPrimaryKey, getTimerService</i> JNDI access to java:comp/env Resource manager access Enterprise bean access Timer service or Timer methods

Additional restrictions:

- The `getRollbackOnly` and `setRollbackOnly` methods of the `EntityContext` interface should be used only in the enterprise bean methods that execute in the context of a transaction. The container must throw the `java.lang.IllegalStateException` if the methods are invoked while the instance is not associated with a transaction.

Reasons for disallowing operations:

- Invoking the `getEJBObject`, `getEJBLocalObject`, and `getPrimaryKey` methods is disallowed in the entity bean methods in which there is no entity object identity associated with the instance.
- Invoking the `getEJBObject` and `getEJBHome` methods is disallowed if the entity bean does not define a remote client view.
- Invoking the `getEJBLocalObject` and `getEJBLocalHome` methods is disallowed if the entity bean does not define a local client view.
- Invoking the `getRollbackOnly` and `setRollbackOnly` methods is disallowed in the entity bean methods for which the container does not have a meaningful transaction context.

- Accessing resource managers and enterprise beans, including accessing the persistent state of an entity bean instance, is disallowed in the entity bean methods for which the container does not have a meaningful transaction context or client security context.

10.5.6 Finder Methods

An entity bean's home interface defines one or more `finder` methods^[20], one for each way to find an entity object or collection of entity objects within the home. The name of each finder method starts with the prefix “find”, such as `findLargeAccounts`. The arguments of a finder method are used in the implementation of the query for the finder method to locate the requested entity objects.

Every finder method except `findByPrimaryKey(key)` must be associated with a `query` element in the deployment descriptor. The entity Bean Provider declaratively specifies the EJB QL finder query and associates it with the finder method in the deployment descriptor. A finder method is normally characterized by an EJB QL query string specified in the `query` element. EJB QL is described in Chapter 11, “EJB QL: EJB Query Language for Container-Managed Persistence Query Methods”.

In the case that both the remote home interface and local home interface define a finder method with the same name and argument types, the EJB QL query string specified by the `query` element defines the semantics of both methods.

10.5.6.1 Single-Object Finder Methods

Some finder methods (such as `findByPrimaryKey`) are designed to return at most one entity object. For single-object finders, the result type of a `find<METHOD>` method defined in the entity bean's remote home interface is the entity bean's remote interface, and the result type of the `find<METHOD>` method defined in the entity bean's local home interface is the entity bean's local interface.

The following code illustrates the definition of a single-object finder defined on the remote home interface.

```
// Entity's home interface
public interface AccountHome extends javax.ejb.EJBHome {
    ...
    Account findByPrimaryKey(AccountPrimaryKey primkey)
        throws FinderException, RemoteException;
    ...
}
```

Note that a finder method defined on the local home interface must not throw the `RemoteException`.

[20] The `findByPrimaryKey` method is mandatory for all entity beans.

In general, when defining a single-object finder method other than `findByPrimaryKey`, the entity Bean Provider should be sure that the finder method will always return only a single entity object. This may occur, for example, if the EJB QL query string that is used to specify the finder query includes an equality test on the entity bean's primary key fields. If the entity Bean Provider uses an unknown primary key class (see Section 10.8.3), the Bean Provider will typically define the finder method as a multi-object finder.

Note that a single-object finder method may return a null value. If the result set of the query consists of a single null value, the container must return the null value as the result of the method. If the result set of a query for a single-object finder method contains more than one value (whether non-null, null, or a combination), the container must throw the `FinderException` from the finder method. If the result set of the query contains no values, the container must throw the `ObjectNotFoundException`. See Section 11.2.7.1.

10.5.6.2 Multi-Object Finder Methods

Some finder methods are designed to return multiple entity objects. For multi-object finders defined on the entity bean's local home interface, the result type of the `find<METHOD>` method is a collection of objects implementing the entity bean's local interface. For multi-object finders defined on the entity bean's remote home interface, the result type of the `find<METHOD>` method is a collection of objects implementing the entity bean's remote interface.

The Bean Provider uses the Java™ 2 `java.util.Collection` interface to define a collection type for the result type of a finder method for an entity bean with container-managed persistence.

The collection of values returned by the container may contain duplicates if `DISTINCT` is not specified in the `SELECT` clause of the query for the finder method.

The collection of values returned by the container may contain null values if the finder method returns the values of a cmr-field and null values are not eliminated by the query.

A portable client program must use the `PortableRemoteObject.narrow` method to convert the objects contained in the collections returned by a finder method on the entity bean's remote home interface to the entity bean's remote interface type.

The following is an example of a multi-object finder method defined on the remote home interface:

```
// Entity's home interface
public interface AccountHome extends javax.ejb.EJBHome {
    ...
    java.util.Collection findLargeAccounts(double limit)
        throws FinderException, RemoteException;
    ...
}
```

Note that if this finder method were defined on the local home interface, it would not throw the `RemoteException`.

10.5.7 Select Methods

Select methods are query methods for use by the Bean Provider within an entity bean instance. Unlike finder methods, select methods are not specified in the entity bean's home interface. A select method is an abstract method defined by the Bean Provider on an entity bean class. A select method must not be exposed in the home or component interface of an entity bean.

The semantics of a select method, like those of a finder method, are defined by an EJB QL query string. A select method is similar to a finder method, but unlike a finder method, but it can return values that correspond to any cmp- or cmr-field type.

Every select method must be associated with a `query` element in the deployment descriptor. The entity Bean Provider declaratively specifies the EJB QL query and associates it with the select method in the deployment descriptor. A select method is normally characterized by an EJB QL query string specified in the `query` element. EJB QL is described in Chapter 11, "EJB QL: EJB Query Language for Container-Managed Persistence Query Methods"

Typically an `ejbSelect<METHOD>` method that returns entity objects returns these as `EJBLocalObjects`. If the `ejbSelect<METHOD>` method returns an `EJBObject` or collection of `EJBObjects`, the Bean Provider must specify the value of the `result-type-mapping` element in the `query` deployment descriptor element for the select method as `Remote`.

An `ejbSelect<METHOD>` is not based on the identity of the entity bean instance on which it is invoked. However, the Bean Provider can use the primary key of an entity bean as an argument to an `ejbSelect<METHOD>` to define a query that is logically scoped to a particular entity bean instance.

The following table illustrates the semantics of finder and select methods.

Table 5 Comparison of Finder and Select Methods

	Finder methods	Select methods
method	<code>find<METHOD></code>	<code>ejbSelect<METHOD></code>
visibility	exposed to client	internal to entity bean class
instance	arbitrary bean instance in pooled state	instance: current instance (could be bean instance in pooled state or ready state)
return value	<code>EJBObjects</code> or <code>EJBLocalObjects</code> of the same type as the entity bean	<code>EJBObjects</code> , <code>EJBLocalObjects</code> , or cmp-field types

10.5.7.1 Single-Object Select Methods

Some select methods are designed to return at most one value. In general, when defining a single-object select method, the entity Bean Provider must be sure that the select method will always return only a single object or value. If the query specified by the select method returns multiple values of the designated type, the container must throw the `FinderException`.

Note that a single-object select method may return a null value. If the result set of the query consists of a single null value, the container must return the null value as the result of the method. If the result set of a query for a single-object select method contains more than one value (whether non-null, null, or a combination), the container must throw the `FinderException` from the select method. If the result set of the query contains no values, the container must throw the `ObjectNotFoundException`. See Section 11.2.7.1.

The Bean Provider will typically define a select method as a multi-object select method.

10.5.7.2 Multi-Object Select Methods

Some select methods are designed to return multiple values. For these multi-object select methods, the result type of the `ejbSelect<METHOD>method` is a collection of objects.

The Bean Provider uses the Java™ 2 `java.util.Collection` interface or `java.util.Set` interface to define a collection type for the result type of a select method. The type of the elements of the collection is determined by the type of the `SELECT` clause of the corresponding EJB QL query. If the Bean Provider uses the `java.util.Collection` interface, the collection of values returned by the container may contain duplicates if `DISTINCT` is not specified in the `SELECT` clause of the query. If a query for a select method whose result type is `java.util.Set` does not specify `DISTINCT` in its `SELECT` clause, the container must interpret the query as if `SELECT DISTINCT` had been specified.

The collection of values returned by the container may contain null values if the select method returns the values of a `cmr-field` or `cmp-field` and null values are not eliminated by the query. See Section 11.2.7.1.

The following is an example of a multi-object select method definition in the `OrderBean` class:

```
// OrderBean implementation class
public abstract class OrderBean implements javax.ejb.EntityBean{
    ...
    public abstract java.util.Collection
        ejbSelectAllOrderedProducts(Customer customer)
        throws FinderException;
    // internal finder method to find all products ordered
```

10.5.8 Timer Notifications

An entity bean can be registered with the EJB timer service for time-based event notifications if it implements the `javax.ejb.TimerObject` interface. The container invokes the bean instance's `ejbTimeout` method when a timer for the bean has expired. See Chapter 22, “Timer Service”.

10.5.9 Standard Application Exceptions for Entities

The EJB specification defines the following standard application exceptions:

- `javax.ejb.CreateException`

- `javax.ejb.DuplicateKeyException`
- `javax.ejb.FinderException`
- `javax.ejb.ObjectNotFoundException`
- `javax.ejb.RemoveException`

This section describes the use of these exceptions by entity beans with container-managed persistence.

10.5.9.1 `CreateException`

From the client's perspective, a `CreateException` (or a subclass of `CreateException`) indicates that an application level error occurred during a `create<METHOD>` operation. If a client receives this exception, the client does not know, in general, whether the entity object was created but not fully initialized, or not created at all. Also, the client does not know whether or not the transaction has been marked for rollback. (However, the client may determine the transaction status using the `UserTransaction` interface or the `setRollbackOnly` method of the `EJBContext` interface.)

Both the container and the Bean Provider may throw the `CreateException` (or subclass of `CreateException`) from the `create<METHOD>`, `ejbCreate<METHOD>` and `ejbPostCreate<METHOD>` methods to indicate an application-level error from the create or initialization operation. Optionally, the container or Bean Provider may mark the transaction for rollback before throwing this exception.

The container or Bean Provider is encouraged to mark the transaction for rollback only if data integrity would be lost if the transaction were committed by the client. Typically, when a `CreateException` is thrown, it leaves the database in a consistent state, allowing the client to recover. For example, the `ejbCreate<METHOD>` method may throw the `CreateException` to indicate that the some of the arguments to the `create<METHOD>` method are invalid.

The container treats the `CreateException` as any other application exception. See Section 18.3.

10.5.9.2 `DuplicateKeyException`

The `DuplicateKeyException` is a subclass of `CreateException`. It may be thrown by the container to indicate to the client or local client that the entity object cannot be created because an entity object with the same key already exists. The unique key causing the violation may be the primary key, or another key defined in the underlying database.

Normally, the container should not mark the transaction for rollback before throwing the exception.

When the client or local client receives a `DuplicateKeyException`, the client knows that the entity was not created, and that the transaction has not typically been marked for rollback.

10.5.9.3 FinderException

From the client's perspective, a `FinderException` (or a subclass of `FinderException`) indicates that an application level error occurred during the `find` operation. Typically, the transaction has not been marked for rollback because of the `FinderException`.

The container throws the `FinderException` (or subclass of `FinderException`) from the implementation of a `finder` or `select` method to indicate an application-level error in the `finder` or `select` method. The container should not, typically, mark the transaction for rollback before throwing the `FinderException`.

The container treats the `FinderException` as any other application exception. See Section 18.3.

10.5.9.4 ObjectNotFoundException

The `ObjectNotFoundException` is a subclass of `FinderException`. The container throws the `ObjectNotFoundException` from the implementation of a `finder` or `select` method to indicate that the requested object does not exist.

Only single-object `finder` or `select` methods (see Subsections 10.5.6 and 10.5.7) should throw this exception. Multi-object `finder` or `select` methods must not throw this exception. Multi-object `finder` or `select` methods should return an empty collection as an indication that no matching objects were found.

10.5.9.5 RemoveException

From the client's perspective, a `RemoveException` (or a subclass of `RemoveException`) indicates that an application level error occurred during a `remove` operation. If a client receives this exception, the client does not know, in general, whether the entity object was removed or not. The client also does not know if the transaction has been marked for rollback. (However, the client may determine the transaction status using the `UserTransaction` interface.)

The container or Bean Provider throws the `RemoveException` (or subclass of `RemoveException`) from a `remove` method to indicate an application-level error from the entity object removal operation. Optionally, the container or Bean Provider may mark the transaction for rollback before throwing this exception.

The container or Bean Provider is encouraged to mark the transaction for rollback only if data integrity would be lost if the transaction were committed by the client. Typically, when a `RemoveException` is thrown, it leaves the database in a consistent state, allowing the client to recover.

The container treats the `RemoveException` as any other application exception. See Section 18.3.

10.5.10 Commit Options

The Entity Bean protocol is designed to give the container the flexibility to select the disposition of the instance state at transaction commit time. This flexibility allows the container to optimally manage the association of an entity object identity with the enterprise bean instances.

The container can select from the following commit-time options:

- **Option A:** The container caches a “ready” instance between transactions. The container knows that the bean instance has exclusive access to the state of the object in the persistent storage. Therefore, the container does not have to synchronize the instance’s state from the persistent storage at the beginning of the next transaction or have to verify that the instance’s state is in sync with the persistent storage at the beginning of the next transaction.
- **Option B:** The container caches a “ready” instance between transactions. In contrast to Option A, in this option the instance may not have exclusive access to the state of the object in the persistent storage. Therefore, the container must synchronize the instance’s state from the persistent storage at the beginning of the next transaction if the instance’s state in the persistent storage has changed. Containers using optimistic concurrency control strategies may instead choose to rollback the transaction if this invariant has not been met: The container must ensure that in order for a transaction to be successfully committed, the transaction must only operate on instance data that is in sync with the persistent storage at the beginning of the transaction.
- **Option C:** The container does not cache a “ready” instance between transactions. The container returns the instance to the pool of available instances after a transaction has completed.

Variants of these strategies that capture the same semantics from the Bean Provider’s viewpoint may be employed, e.g., to optimize data access.

The following illustrative lazy loading strategies are consistent with the intent of these requirements:

- If `ejbLoad` is called at the beginning of the transaction without the instance’s persistent state having been loaded from the persistent storage, the persistent state must be faulted in when `ejbLoad` causes the bean’s getter accessor methods to be invoked. If the `ejbLoad` method is empty, data may be faulted in as needed in the course of executing the business methods of the bean.
- If the instance’s persistent state is cached between transactions, `ejbLoad` need not be called and persistent data need not be faulted in from the persistent storage (unless it has not previously been accessed). In this case, because `ejbLoad` has been previously called when the instance was entered into the ready state for the first time, and because the bean instance’s state is consistent with its persistent state, there is no need to call `ejbLoad` unless the instance’s state in the persistent storage has changed. In this case, the container must ensure that in order for the transaction to be successfully committed, the instance’s persistent state was in sync with the persistent storage at the beginning of the transaction.

The following table provides a summary of the commit-time options.

Table 6 Summary of Commit-Time Options

	Write instance state to database	Instance stays ready	Instance state remains valid
Option A	Yes	Yes	Yes
Option B	Yes	Yes	No
Option C	Yes	No	No

Note that the container synchronizes the instance's state with the persistent storage at transaction commit for all three options.

The selection of the commit option is transparent to the entity bean implementation—the entity bean will work correctly regardless of the commit-time option chosen by the container. The Bean Provider writes the entity bean in the same way.

The object interaction diagrams in Section 10.9 illustrate the three alternative commit options in detail.

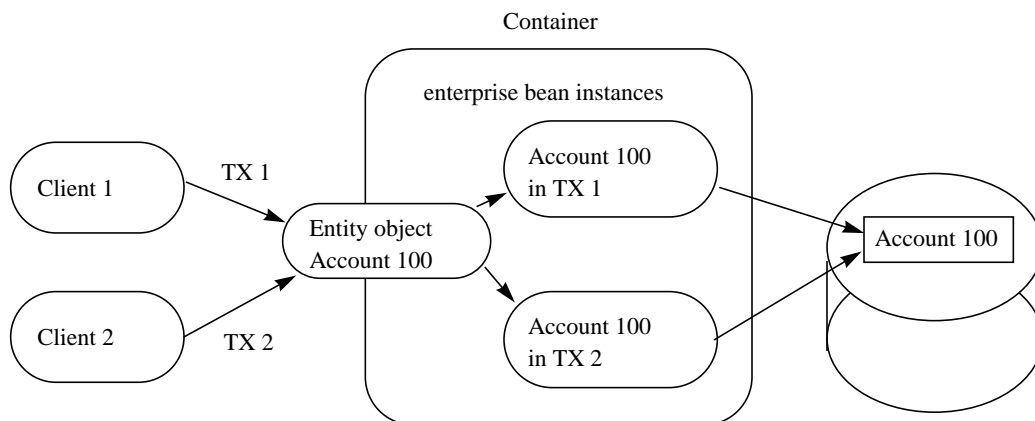
Note: The Bean Provider relies on the `ejbLoad` method to be invoked in order to resynchronize the bean's transient state with its persistent state. It is the responsibility of the container to call the `ejbLoad` method at the beginning of a new transaction if the bean instance's persistent data has changed..

10.5.11 Concurrent Access from Multiple Transactions

When writing the entity bean business methods, the Bean Provider does not have to worry about concurrent access from multiple transactions. The Bean Provider may assume that the container will ensure appropriate synchronization for entity objects that are accessed concurrently from multiple transactions.

The container typically uses one of the following implementation strategies to achieve proper synchronization. (These strategies are illustrative, not prescriptive.)

- The container activates multiple instances of the entity bean, one for each transaction in which the entity object is being accessed. The transaction synchronization is performed by the underlying database during the accessor method calls performed by the business methods, the `ejbTimeout` method, and by the `ejbLoad`, `ejbCreate<METHOD>`, `ejbStore`, and `ejbRemove` methods. The commit-time options B and C in Subsection 10.5.10 apply to this type of container.

Figure 28 Multiple Clients Can Access the Same Entity Object Using Multiple Instances

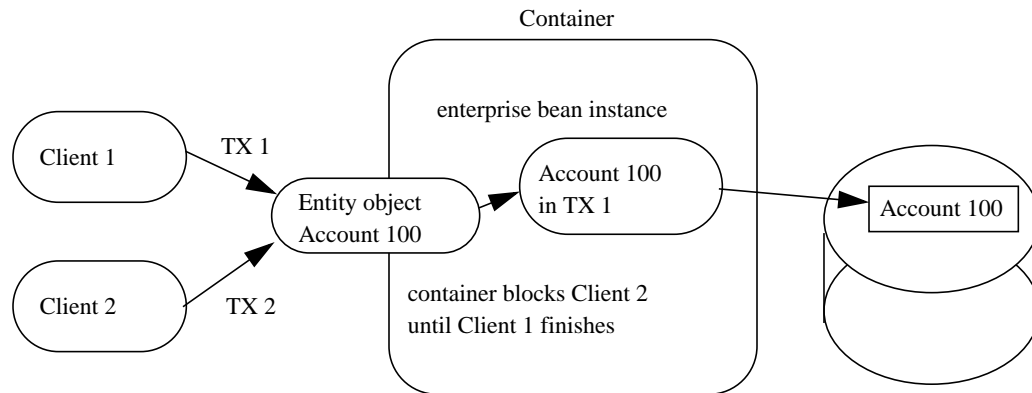
With this strategy, the type of lock acquired by `ejbLoad` or `get accessor method` (if a lazy loading cache management strategy is used) leads to a trade-off. If `ejbLoad` or the accessor method acquires an exclusive lock on the instance's state in the database, the throughput of read-only transactions could be impacted. If `ejbLoad` or the accessor method acquires a shared lock and the instance is updated, then either `ejbStore` or a `set accessor method` will need to promote the lock to an exclusive lock (which may cause a deadlock if it happens concurrently under multiple transactions), or, if the container uses an optimistic cache concurrency control strategy, the container will need to validate the state of the cache against the database at transaction commit (which may result in a rollback of the transaction).

It is expected that containers will provide deployment-time configuration options that will allow control to be exercised over the logical transaction isolation levels that their caching strategies provide.

- The container acquires exclusive access to the entity object's state in the database. The container activates a single instance and serializes the access from multiple transactions to this instance. The commit-time option A in Subsection 10.5.10 applies to this type of container.

Figure 29

Multiple Clients Can Access the Same Entity Object Using Single Instance



10.5.12 Non-reentrant and Re-entrant Instances

An entity Bean Provider can specify that an entity bean is non-reentrant. If an instance of a non-reentrant entity bean executes a client request in a given transaction context, and another request with the same transaction context arrives for the same entity object, the container will throw an exception to the second request. This rule allows the Bean Provider to program the entity bean as single-threaded, non-reentrant code.

The functionality of entity beans with container-managed persistence may require loopbacks in the same transaction context. An example of a loopback is when the client calls entity object A, A calls entity object B, and B calls back A in the same transaction context. The entity bean's method invoked by the loopback shares the current execution context (which includes the transaction and security contexts) with the Bean's method invoked by the client.

If the entity bean is specified as non-reentrant in the deployment descriptor, the container must reject an attempt to re-enter the instance via the entity bean's component interface while the instance is executing a business method. (This can happen, for example, if the instance has invoked another enterprise bean, and the other enterprise bean tries to make a loopback call.) If the attempt is made to reenter the instance through the remote interface, the container must throw the `java.rmi.RemoteException` to the caller. If the attempt is made to reenter the instance through the local interface, the container must throw the `javax.ejb.EJBException` to the caller. The container must allow the call if the Bean's deployment descriptor specifies that the entity bean is re-entrant.

Re-entrant entity beans must be programmed and used with caution. First, the Bean Provider must code the entity bean with the anticipation of a loopback call. Second, since the container cannot, in general, tell a loopback from a concurrent call from a different client, the client programmer must be careful to avoid code that could lead to a concurrent call in the same transaction context.

Concurrent calls in the same transaction context targeted at the same entity object are illegal and may lead to unpredictable results. Since the container cannot, in general, distinguish between an illegal concurrent call and a legal loopback, application programmers are encouraged to avoid using loopbacks. Entity beans that do not need callbacks should be marked as non-reentrant in the deployment descriptor, allowing the container to detect and prevent illegal concurrent calls from clients.

10.6 Responsibilities of the Enterprise Bean Provider

This section describes the responsibilities of an entity Bean Provider to ensure that an entity bean with container-managed persistence can be deployed in any EJB container.

10.6.1 Classes and Interfaces

The entity Bean Provider is responsible for providing the following class files:

- Entity bean class and any dependent classes
- Primary key class
- Entity bean's remote interface and entity bean's remote home interface, if the entity bean provides a remote client view
- Entity bean's local interface and local home interface, if the entity bean provides a local client view

The Bean Provider must provide a remote interface and a remote home interface or a local interface and a local home interface for the bean. The Bean Provider may provide a remote interface, remote home interface, local interface, and local home interface for the bean. Other combinations are not allowed.

10.6.2 Enterprise Bean Class

The following are the requirements for an entity bean class:

The class must implement, directly or indirectly, the `javax.ejb.EntityBean` interface.

The class may implement, directly or indirectly, the `javax.ejb.TimerObject` interface.

The class must be defined as `public` and must be `abstract`. The class must be a top level class.

The class must define a public constructor that takes no arguments.

The class must not define the `finalize()` method.

The class may, but is not required to, implement the entity bean's component interface^[21]. If the class implements the entity bean's component interface, the class must provide no-op implementations of the methods defined by that interface. The container will never invoke these methods on the bean instances at runtime.

The entity bean class must implement the business methods, and the `ejbCreate<METHOD>` and `ejbPostCreate<METHOD>` methods as described later in this section.

The entity bean class must implement the `ejbHome<METHOD>` methods that correspond to the home business methods specified in the bean's home interface. These methods are executed on an instance in the pooled state; hence they must not access state that is particular to a specific bean instance (e.g., the accessor methods for the bean's abstract persistence schema must not be used by these methods).

The entity bean class must implement the get and set accessor methods of the bean's abstract persistence schema as abstract methods.

The entity bean class may have superclasses and/or superinterfaces. If the entity bean has superclasses, the business methods, the `ejbCreate<METHOD>` and `ejbPostCreate<METHOD>` methods, and the methods of the `EntityBean` interface and/or the `TimedObject` interface may be implemented in the enterprise bean class or in any of its superclasses.

The entity bean class is allowed to implement other methods (for example helper methods invoked internally by the business methods) in addition to the methods required by the EJB specification.

The entity bean class does not implement the finder methods. The implementations of the finder methods are provided by the container.

The entity bean class must implement any `ejbSelect<METHOD>` methods as abstract methods.

10.6.3 Dependent Value Classes

The following are the requirements for a dependent value class:

The class must be defined as `public` and must not be `abstract`.

The class must be serializable.

10.6.4 ejbCreate<METHOD> Methods

The entity bean class must implement the `ejbCreate<METHOD>` methods that correspond to the `create<METHOD>` methods specified in the entity bean's home interface or local home interface.

The entity bean class may define zero or more `ejbCreate<METHOD>` methods whose signatures must follow these rules:

[21] If the entity bean class does implement the component interface, care must be taken to avoid passing of `this` as a method argument or result. This potential error can be avoided by choosing not to implement the component interface in the entity bean class.

The method name must have `ejbCreate` as its prefix.

The method must be declared as `public`.

The method must not be declared as `final` or `static`.

The return type must be the entity bean's primary key type.

If the `ejbCreate<METHOD>` method corresponds to a `create<METHOD>` on the entity bean's remote home interface, the method arguments and return value types must be legal types for RMI-IIOP.

The `throws` clause must define the `javax.ejb.CreateException`. The `throws` clause may define arbitrary application specific exceptions.

Compatibility Note: EJB 1.0 allowed the `ejbCreate` method to throw the `java.rmi.RemoteException` to indicate a non-application exception. This practice was deprecated in EJB 1.1—an EJB 1.1 or EJB 2.0 or later compliant enterprise bean should throw the `javax.ejb.EJBException` or another `java.lang.RuntimeException` to indicate non-application exceptions to the container (see Section 18.2.2). The `ejbCreate` method of an entity bean with `cmp-version 2.x` must not throw the `java.rmi.RemoteException`.

10.6.5 `ejbPostCreate<METHOD>` Methods

For each `ejbCreate<METHOD>` method, the entity bean class must define a matching `ejbPostCreate<METHOD>` method, using the following rules:

The method name must have `ejbPostCreate` as its prefix.

The method must be declared as `public`.

The method must not be declared as `final` or `static`.

The return type must be `void`.

The method arguments must be the same as the arguments of the matching `ejbCreate<METHOD>` method.

The `throws` clause may define arbitrary application specific exceptions, including the `javax.ejb.CreateException`.

Compatibility Note: EJB 1.0 allowed the `ejbPostCreate` method to throw the `java.rmi.RemoteException` to indicate a non-application exception. This practice was deprecated in EJB 1.1—an EJB 1.1 or EJB 2.0 or later compliant enterprise bean should throw the `javax.ejb.EJBException` or another `java.lang.RuntimeException` to indicate non-application exceptions to the container (see Section 18.2.2). The `ejbPostCreate` method of an entity bean with `cmp-version 2.x` must not throw the `java.rmi.RemoteException`.

10.6.6 ejbHome<METHOD> Methods

The entity bean class may define zero or more home methods whose signatures must follow the following rules:

An `ejbHome<METHOD>` method must exist for every home `<METHOD>` method on the entity bean's remote home or local home interface. The method name must have `ejbHome` as its prefix followed by the name of the `<METHOD>` method in which the first character has been uppercased.

The method must be declared as `public`.

The method must not be declared as `static`.

If the `ejbHome<METHOD>` method corresponds to a home `<METHOD>` on the entity bean's remote home interface, the method argument and return value types must be legal types for RMI-IIOP.

The `throws` clause may define arbitrary application specific exceptions. The `throws` clause must not throw the `java.rmi.RemoteException`.

10.6.7 ejbSelect<METHOD> Methods

The entity bean class may define one or more select methods whose signatures must follow the following rules:

The method name must have `ejbSelect` as its prefix.

The method must be declared as `public`.

The method must be declared as `abstract`.

The `throws` clause must define the `javax.ejb.FinderException`. The `throws` clause may define arbitrary application specific exceptions.

10.6.8 Business Methods

The entity bean class may define zero or more business methods whose signatures must follow these rules:

The method names can be arbitrary, but they must not start with `'ejb'` to avoid conflicts with the callback methods used by the EJB architecture.

The business method must be declared as `public`.

The method must not be declared as `final` or `static`.

If the business method corresponds to a method of the entity bean's remote interface, the method argument and return value types must be legal types for RMI-IIOP.

The `throws` clause may define arbitrary application specific exceptions.

Compatibility Note: EJB 1.0 allowed the business methods to throw the `java.rmi.RemoteException` to indicate a non-application exception. This practice was deprecated in EJB 1.1—an EJB 1.1 or EJB 2.0 or later compliant enterprise bean should throw the `javax.ejb.EJBException` or another `java.lang.RuntimeException` to indicate non-application exceptions to the container (see Section 18.2.2). The business methods of an entity bean with `cmp-version 2.x` must not throw the `java.rmi.RemoteException`.

10.6.9 Entity Bean's Remote Interface

The following are the requirements for the entity bean's remote interface:

The interface must extend the `javax.ejb.EJBObject` interface.

The methods defined in the remote interface must follow the rules for RMI-IIOP. This means that their argument and return value types must be valid types for RMI-IIOP, and their `throws` clauses must include the `java.rmi.RemoteException`.

The remote interface is allowed to have superinterfaces. Use of interface inheritance is subject to the RMI-IIOP rules for the definition of remote interfaces.

For each method defined in the remote interface, there must be a matching method in the entity bean's class. The matching method must have:

- The same name.
- The same number and types of its arguments, and the same return type.
- All the exceptions defined in the `throws` clause of the matching method of the enterprise Bean class must be defined in the `throws` clause of the method of the remote interface.

The remote interface methods must not expose local interface types, local home interface types, timer handles, or the managed collection classes that are used for entity beans with container-managed persistence as arguments or results.

10.6.10 Entity Bean's Remote Home Interface

The following are the requirements for the entity bean's home interface:

The interface must extend the `javax.ejb.EJBHome` interface.

The methods defined in this interface must follow the rules for RMI-IIOP. This means that their argument and return types must be of valid types for RMI-IIOP, and their `throws` clauses must include the `java.rmi.RemoteException`.

The remote home interface is allowed to have superinterfaces. Use of interface inheritance is subject to the RMI-IIOP rules for the definition of remote interfaces.

Each method defined in the remote home interface must be one of the following:

- A create method.
- A finder method.
- A home method.

Each create method must be named “create<METHOD>”, e.g. createLargeAccounts. Each create method name must match one of the ejbCreate<METHOD> methods defined in the enterprise bean class. The matching ejbCreate<METHOD> method must have the same number and types of its arguments. (Note that the return type is different.)

The return type for a create<METHOD> method must be the entity bean’s remote interface type.

All the exceptions defined in the throws clause of the matching ejbCreate<METHOD> and ejbPostCreate<METHOD> methods of the enterprise bean class must be included in the throws clause of the matching create method of the home interface (i.e., the set of exceptions defined for the create method must be a superset of the union of exceptions defined for the ejbCreate<METHOD> and ejbPostCreate<METHOD> methods).

The throws clause of a create<METHOD> method must include the javax.ejb.CreateException.

Each finder method must be named “find<METHOD>” (e.g. findLargeAccounts).

The return type for a find<METHOD> method must be the entity bean’s remote interface type (for a single-object finder), or a collection thereof (for a multi-object finder).

The remote home interface must always include the findByPrimaryKey method, which is always a single-object finder. The method must declare the primary key class as the method argument.

The throws clause of a finder method must include the javax.ejb.FinderException.

Home methods can have arbitrary names, but they must not start with “create”, “find”, or “remove”. Their argument and return types must be of valid types for RMI-IIOP, and their throws clauses must include the java.rmi.RemoteException. The matching ejbHome method specified in the entity bean class must have the same number and types of arguments and must return the same type as the home method as specified in the remote home interface of the bean.

The remote home interface methods must not expose local interface types, local home interface types, timers or timer handles, or the managed collection classes that are used for entity beans with container-managed persistence as arguments or results.

10.6.11 Entity Bean’s Local Interface

The following are the requirements for the entity bean’s local interface:

The interface must extend the javax.ejb.EJBLocalObject interface.

For each method defined in the local interface, there must be a matching method in the entity bean's class. The matching method must have:

- The same name.
- The same number and types of its arguments, and the same return type.
- All the exceptions defined in the `throws` clause of the matching method of the enterprise Bean class must be defined in the `throws` clause of the method of the local interface.

10.6.12 Entity Bean's Local Home Interface

The following are the requirements for the entity bean's local home interface:

The interface must extend the `javax.ejb.EJBLocalHome` interface.

Each method defined in the home interface must be one of the following:

- A create method.
- A finder method.
- A home method.

Each `create` method must be named "`create<METHOD>`", e.g. `createLargeAccounts`. Each `create` method name must match one of the `ejbCreate<METHOD>` methods defined in the enterprise bean class. The matching `ejbCreate<METHOD>` method must have the same number and types of its arguments. (Note that the return type is different.)

The return type for a `create<METHOD>` method on the local home interface must be the entity bean's local interface type.

All the exceptions defined in the `throws` clause of the matching `ejbCreate<METHOD>` and `ejbPostCreate<METHOD>` methods of the enterprise bean class must be included in the `throws` clause of the matching `create` method of the local home interface (i.e., the set of exceptions defined for the `create` method must be a superset of the union of exceptions defined for the `ejbCreate<METHOD>` and `ejbPostCreate<METHOD>` methods).

The `throws` clause of a `create<METHOD>` method must include the `javax.ejb.CreateException`.

Each `finder` method must be named "`find<METHOD>`" (e.g. `findLargeAccounts`).

The return type for a `find<METHOD>` method defined on the local home interface must be the entity bean's local interface type (for a single-object finder), or a collection thereof (for a multi-object finder).

The local home interface must always include the `findByPrimaryKey` method, which is always a single-object finder. The method must declare the primary key class as the method argument.

The `throws` clause of a `finder` method must include the `javax.ejb.FinderException`.

Home methods can have arbitrary names, but they must not start with “create”, “find”, or “remove”. The matching `ejbHome` method specified in the entity bean class must have the same number and types of arguments and must return the same type as the home method as specified in the home interface of the bean. The `throws` clause of a home method defined on the local home interface must not include the `java.rmi.RemoteException`.

10.6.13 Entity Bean’s Primary Key Class

The Bean Provider must specify a primary key class in the deployment descriptor.

The primary key type must be a legal Value Type in RMI-IIOP.

The class must provide suitable implementation of the `hashCode()` and `equals(Object other)` methods to simplify the management of the primary keys by the container.

10.6.14 Entity Bean’s Deployment Descriptor

The Bean Provider must specify the relationships in which the entity beans participate in the `relationships` element.

The Bean Provider must provide unique names to designate entity beans as follows, and as described in Section 10.3.13.

- The Bean Provider must specify unique names for entity beans which are defined in the `ejb-jar` file by using the `ejb-name` element.
- The Bean Provider must specify a unique abstract schema name for an entity bean using the `abstract-schema-name` deployment descriptor element.

The Bean Provider must define a query for each finder or select method except `findByPrimaryKey(key)`. Typically this will be provided as the content of the `ejb-ql` element contained in the query element for the entity bean. The syntax of EJB QL is defined in Chapter 11.

Since EJB QL query strings are embedded in the deployment descriptor, which is an XML document, it may be necessary to encode the following characters in the query string: “>”, “<”.

10.7 The Responsibilities of the Container Provider

This section describes the responsibilities of the Container Provider to support entity beans. The Container Provider is responsible for providing the deployment tools, and for managing the entity beans at runtime, including their persistent state and relationships.

Because the EJB specification does not define the API between deployment tools and the container, we assume that the deployment tools described in this section are provided by the Container Provider. Alternatively, the deployment tools may be provided by a different vendor who uses the container vendor's specific API.

10.7.1 Generation of Implementation Classes

The deployment tools provided by the Container Provider are responsible for the generation of additional classes when the entity bean is deployed. The tools obtain the information that they need for generation of the additional classes by introspecting the classes and interfaces provided by the Bean Provider and by examining the entity bean's deployment descriptor.

The deployment tools must generate the following classes:

- A class that implements the entity bean's remote home interface (i.e., the entity EJBHome class).
- A class that implements the entity bean's remote interface (i.e., the entity EJBObject class).
- A class that implements the entity bean's local home interface (i.e., the entity EJBLocalHome class).
- A class that implements the entity bean's local interface (i.e., the EJBLocalObject class).
- A class that implements the entity bean class (i.e., a concrete class corresponding to the abstract entity bean class that was provided by the Bean Provider).

The deployment tools may also generate a class that mixes some container-specific code with the entity bean class. The code may, for example, help the container to manage the entity bean instances at runtime. Tools can use subclassing, delegation, and code generation.

The deployment tools may also allow generation of additional code that wraps the business methods and that is used to customize the business logic for an existing operational environment. For example, a wrapper for a `debit` function on the `Account` bean may check that the debited amount does not exceed a certain limit, or perform security checking that is specific to the operational environment.

10.7.2 Enterprise Bean Class

The following are the requirements for a concrete entity bean class:

The class must extend the abstract entity bean class provided by the Bean Provider.

The class must be defined as `public` and must not be `abstract`.

The class must define a public constructor that takes no arguments.

The class must implement the `get` and `set` accessor methods of the bean's abstract persistence schema.

The class must not define the `finalize` method.

The entity bean class must implement the `ejbFind<METHOD>` methods.

The entity bean class must implement the `ejbSelect<METHOD>` methods.

The entity bean class is allowed to implement other methods in addition to the methods required by the EJB specification.

10.7.3 `ejbFind<METHOD>` Methods

For each `find<METHOD>` method in the remote home interface or local home interface of the entity bean, there must be a corresponding `ejbFind<METHOD>` method with the same argument types in the concrete entity bean class.

The method name must have `ejbFind` as its prefix.

The method must be declared as `public`.

If the `ejbFind<METHOD>` method corresponds to a `find<METHOD>` on the entity bean's remote home interface, the method argument and return value types must be legal types for RMI-IIOP.

The return type of an `ejbFind<METHOD>` method must be the entity bean's primary key type, or a collection of primary keys.

The `throws` clause must define the `javax.ejb.FinderException`. The `throws` clause may define arbitrary application specific exceptions.

Every finder method except `ejbFindByPrimaryKey(key)` is specified in the query deployment descriptor element for the entity. The container must use the EJB QL query string that is the content of the `ejb-ql` element or the descriptive query specification contained in the `description` element as the definition of the query of the corresponding `ejbFind<METHOD>` method.

10.7.4 `ejbSelect<METHOD>` Methods

For each `ejbSelect<METHOD>` method in the abstract entity bean class, there must be a method with the same argument and result types in the concrete entity bean class.

Every select method is specified in a query deployment descriptor element for the entity. The container must use the EJB QL query string that is the content of the `ejb-ql` element or the descriptive query specification that is contained in the `description` element as the definition of the query of the corresponding `ejbSelect<METHOD>` method.

The container must use the contents of the query element, the corresponding EJB QL string and the type of the values selected as specified by the `SELECT` clause to determine the type of the values returned by a select method.

The container must ensure that there are no duplicates returned by a select method if the return type is `java.util.Set`.

10.7.5 Entity EJBHome Class

The entity EJBHome class, which is generated by deployment tools, implements the entity bean's remote home interface. This class implements the methods of the `javax.ejb.EJBHome` interface, and the type-specific `create` and `finder` methods specific to the entity bean.

The implementation of each `create<METHOD>` method invokes a matching `ejbCreate<METHOD>` method, followed by the matching `ejbPostCreate<METHOD>` method, passing the `create<METHOD>` parameters to these matching methods.

The implementation of the `remove` methods defined in the `javax.ejb.EJBHome` interface must activate an instance (if an instance is not already in the ready state) and invoke the `ejbRemove` method on the instance.

The implementation of each `find<METHOD>` method invokes a matching `ejbFind<METHOD>` method. The implementation of the `find<METHOD>` method must create an entity object reference for the primary key returned from the `ejbFind<METHOD>` and return the entity object reference to the client. If the `ejbFind<METHOD>` method returns a collection of primary keys, the implementation of the `find<METHOD>` method must create a collection of entity object references for the primary keys and return the collection to the client.

The implementation of each `<METHOD>` home method invokes a matching `ejbHome<METHOD>` method (in which the first character of `<METHOD>` is uppercased in the name of the `ejbHome<METHOD>` method), passing the parameters of the `<METHOD>` method to the matching `ejbHome<METHOD>` method.

10.7.6 Entity EJBObject Class

The entity EJBObject class, which is generated by deployment tools, implements the entity bean's remote interface. It implements the methods of the `javax.ejb.EJBObject` interface and the remote business methods specific to the entity bean.

The implementation of the `remove` method (defined in the `javax.ejb.EJBObject` interface) must activate an instance (if an instance is not already in the ready state) and invoke the `ejbRemove` method on the instance.

The implementation of each remote business method must activate an instance (if an instance is not already in the ready state) and invoke the matching business method on the instance.

10.7.7 Entity EJBLocalHome Class

The entity EJBLocalHome class, which is generated by deployment tools, implements the entity bean's local home interface. This class implements the methods of the `javax.ejb.EJBLocalHome` interface, and the type-specific `create` and `finder` methods specific to the entity bean.

The implementation of each `create<METHOD>` method invokes a matching `ejbCreate<METHOD>` method, followed by the matching `ejbPostCreate<METHOD>` method, passing the `create<METHOD>` parameters to these matching methods.

The implementation of the `remove` method defined in the `javax.ejb.EJBLocalHome` interface must activate an instance (if an instance is not already in the ready state) and invoke the `ejbRemove` method on the instance.

The implementation of each `find<METHOD>` method invokes a matching `ejbFind<METHOD>` method. The implementation of the `find<METHOD>` method must create a local entity object reference for the primary key returned from the `ejbFind<METHOD>` and return the local entity object reference to the local client. If the `ejbFind<METHOD>` method returns a collection of primary keys, the implementation of the `find<METHOD>` method must create a collection of local entity object references for the primary keys and return the collection to the local client.

The implementation of each `<METHOD>` home method invokes a matching `ejbHome<METHOD>` method (in which the first character of `<METHOD>` is uppercased in the name of the `ejbHome<METHOD>` method), passing the parameters of the `<METHOD>` method to the matching `ejbHome<METHOD>` method.

10.7.8 Entity EJBLocalObject Class

The entity `EJBLocalObject` class, which is generated by deployment tools, implements the entity bean's local interface. It implements the methods of the `javax.ejb.EJBLocalObject` interface and the local business methods specific to the entity bean.

The implementation of the `remove` method (defined in the `javax.ejb.EJBLocalObject` interface) must activate an instance (if an instance is not already in the ready state) and invoke the `ejbRemove` method on the instance.

The implementation of each local business method must activate an instance (if an instance is not already in the ready state) and invoke the matching business method on the instance.

10.7.9 Handle Class

The deployment tools are responsible for implementing the handle class for the entity bean. The handle class must be serializable by the Java Serialization protocol.

As the handle class is not entity bean specific, the container may, but is not required to, use a single class for all deployed entity beans.

10.7.10 Home Handle Class

The deployment tools responsible for implementing the home handle class for the entity bean. The handle class must be serializable by the Java Serialization protocol.

Because the home handle class is not entity bean specific, the container may, but is not required to, use a single class for the home handles of all deployed entity beans.

10.7.11 Metadata Class

The deployment tools are responsible for implementing the class that provides metadata information to the remote client view contract. The class must be a valid RMI-IIOP Value Type, and must implement the `javax.ejb.EJBMetaData` interface.

Because the metadata class is not entity bean specific, the container may, but is not required to, use a single class for all deployed enterprise beans.

10.7.12 Instance's Re-entrance

The container runtime must enforce the rules defined in Section 10.5.12.

10.7.13 Transaction Scoping, Security, Exceptions

The container runtime must follow the rules on transaction scoping, security checking, and exception handling described in Chapters 17, 21, and 18.

10.7.14 Implementation of Object References

The container should implement the distribution protocol between the remote client and the container such that the object references of the remote home and remote interfaces used by entity bean clients are usable for a long period of time. Ideally, a remote client should be able to use an object reference across a server crash and restart. An object reference should become invalid only when the entity object has been removed, or after a reconfiguration of the server environment (for example, when the entity bean is moved to a different EJB server or container).

The motivation for this is to simplify the programming model for the entity bean client. While the client code needs to have a recovery handler for the system exceptions thrown from the individual method invocations on the remote home and remote interface, the client should not be forced to re-obtain the object references.

10.7.15 EntityContext

The container must implement the `EntityContext.getEJBObject` method such that the bean instance can use the Java language cast to convert the returned value to the entity bean's remote interface type. Specifically, the bean instance does not have to use the `PortableRemoteObject.narrow` method for the type conversion.

10.8 Primary Keys

The container must be able to manipulate the primary key type of an entity bean. Therefore, the primary key type for an entity bean with container-managed persistence must follow the rules in this subsection, in addition to those specified in Subsection 10.6.13.

There are two ways to specify a primary key class for an entity bean with container-managed persistence:

- Primary key that maps to a single field in the entity bean class.
- Primary key that maps to multiple fields in the entity bean class.

The second method is necessary for implementing compound keys, and the first method is convenient for single-field keys. Without the first method, simple types such as `String` would have to be wrapped in a user-defined class.

10.8.1 Primary Key That Maps to a Single Field in the Entity Bean Class

The Bean Provider uses the `primkey-field` element of the deployment descriptor to specify the container-managed field of the entity bean class that contains the primary key. The field's type must be the primary key type.

10.8.2 Primary Key That Maps to Multiple Fields in the Entity Bean Class

The primary key class must be `public`, and must have a `public` constructor with no parameters.

All fields in the primary key class must be declared as `public`.

The names of the fields in the primary key class must be a subset of the names of the container-managed fields. (This allows the container to extract the primary key fields from an instance's container-managed fields, and vice versa.)

10.8.3 Special Case: Unknown Primary Key Class

In special situations, the entity Bean Provider may choose not to specify the primary key class or the primary key fields for an entity bean with container-managed persistence. This case usually happens when the entity bean does not have a natural primary key, and/or the Bean Provider wants to allow the Deployer using the Container Provider's tools to select the primary key fields at deployment time. The entity bean's primary key type will usually be derived from the primary key type used by the underlying database system that stores the entity objects. The primary key used by the database system may not be known to the Bean Provider.

In this special case, the type of the argument of the `findByPrimaryKey` method must be declared as `java.lang.Object`. The Bean Provider must specify the primary key class in the deployment descriptor as of the type `java.lang.Object`.

When defining the primary key for the enterprise bean, the Deployer using the Container Provider's tools will typically add additional container-managed fields to the concrete subclass of the entity bean class (this typically happens for entity beans that do not have a natural primary key, and the primary keys are system-generated by the underlying database system that stores the entity objects). In this case, the container must generate the primary key value when the entity bean instance is created (and before `ejbPostCreate` is invoked on the instance.)

The primary key class is specified at deployment time in the situations when the Bean Provider develops an entity bean that is intended to be used with multiple back-ends that provide persistence, and when these multiple back-ends require different primary key structures.

Use of entity beans with a deferred primary key type specification limits the client application programming model, because the clients written prior to deployment of the entity bean may not use, in general, the methods that rely on the knowledge of the primary key type.

The implementation of the enterprise bean class methods must be done carefully. For example, the methods should not depend on the type of the object returned from `EntityContext.getPrimaryKey`, because the return type is determined by the Deployer after the EJB class has been written.

10.9 Object Interaction Diagrams

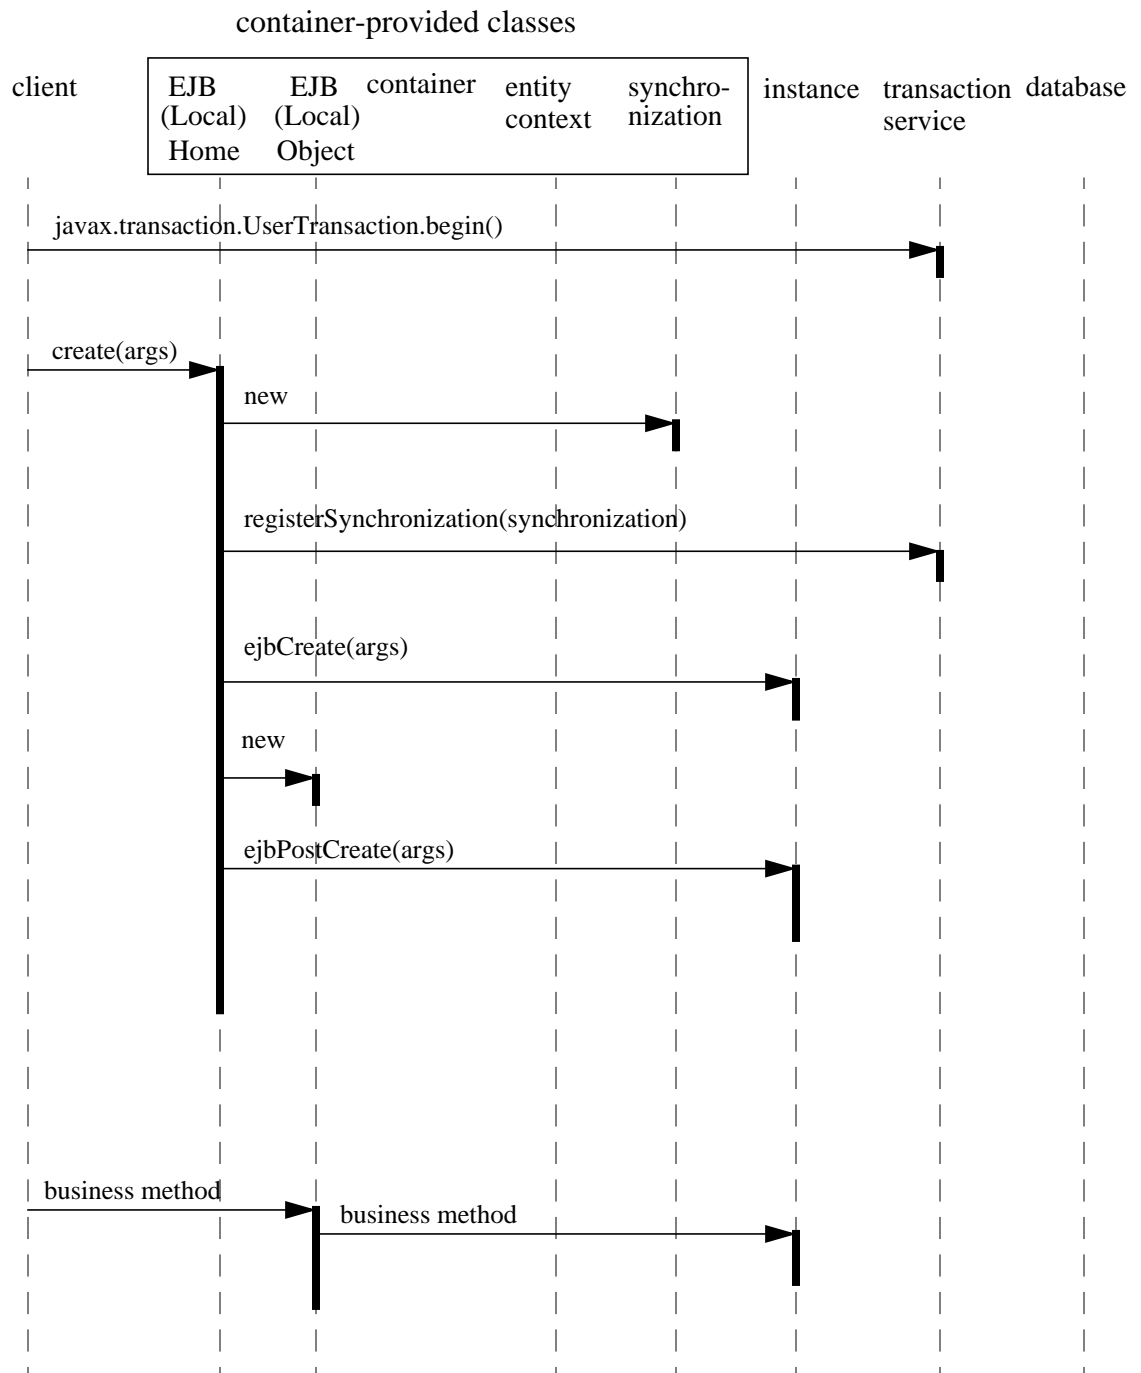
This section uses object interaction diagrams to illustrate the interactions between an entity bean instance and its container.

10.9.1 Notes

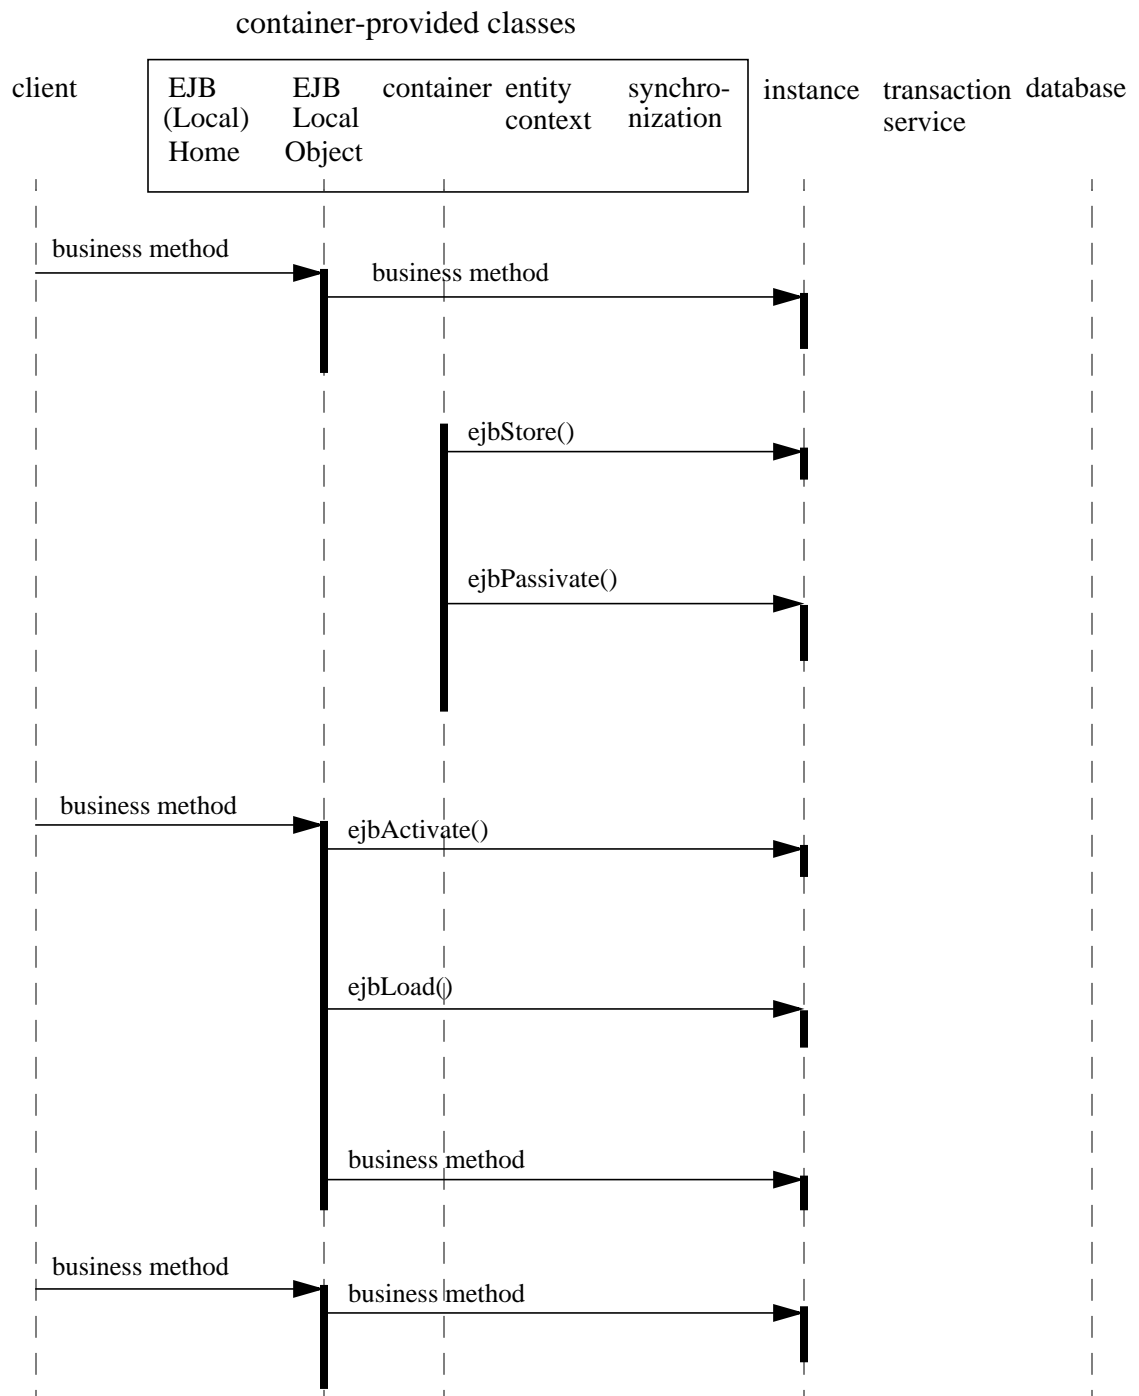
The object interaction diagrams illustrate a box labeled "container-provided classes." These classes are either part of the container or are generated by the container tools. These classes communicate with each other through protocols that are container implementation specific. Therefore, the communication between these classes is not shown in the diagrams.

The classes and actions shown in the diagrams should be considered as an illustrative implementation rather than as a prescriptive one.

10.9.2 Creating an Entity Object

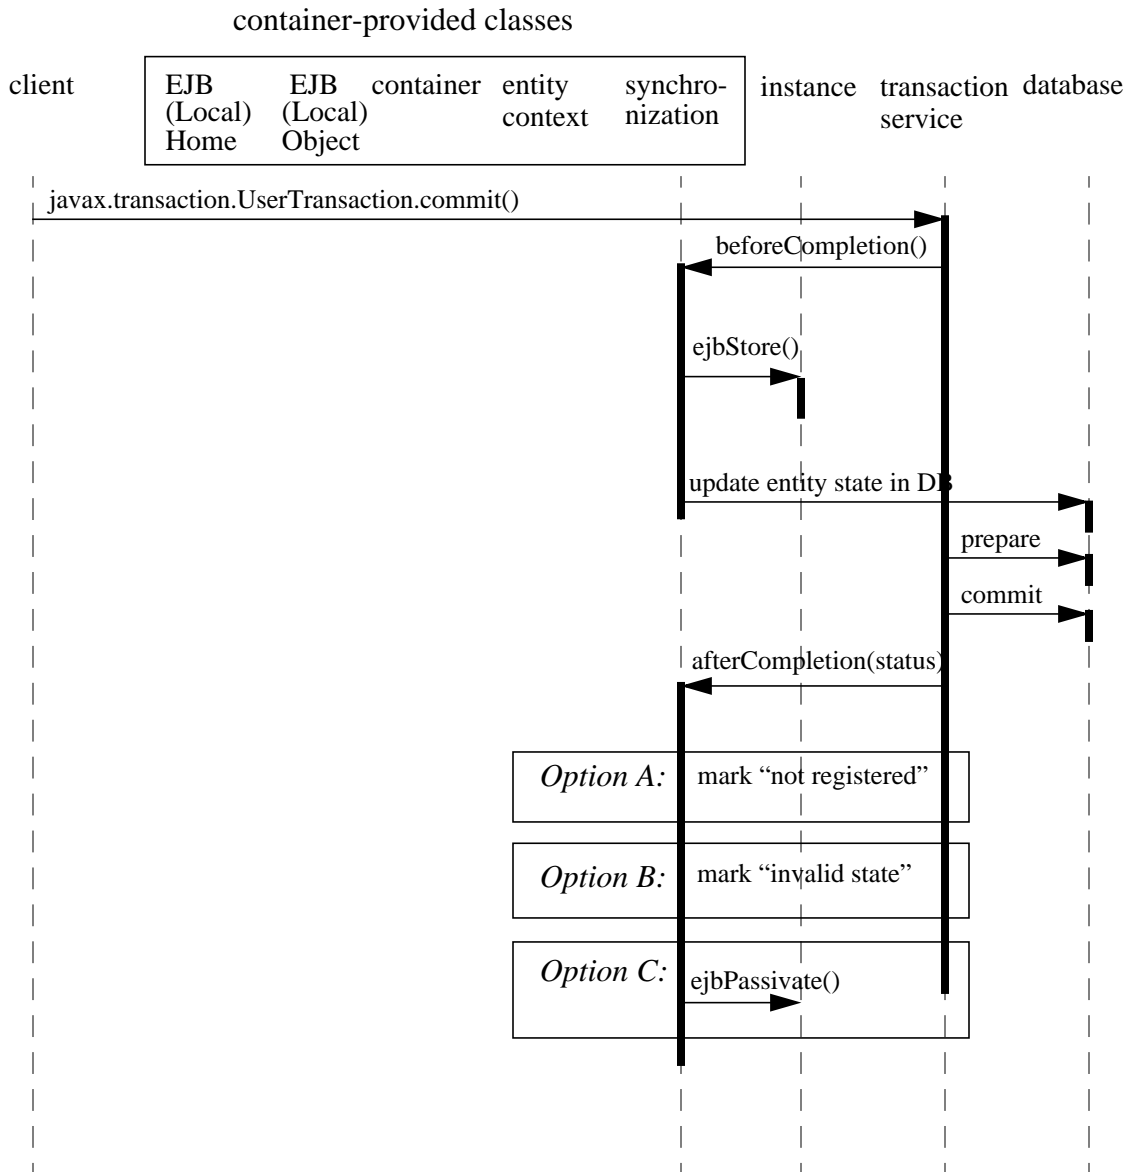
Figure 30 OID of Creation of an Entity Object with Container-Managed Persistence

10.9.3 Passivating and Activating an Instance in a Transaction

Figure 31 OID of Passivation and Reactivation of an Entity Bean Instance with Container-Managed Persistence

10.9.4 Committing a Transaction

Figure 32 OID of Transaction Commit Protocol for an Entity Bean Instance with Container-Managed Persistence

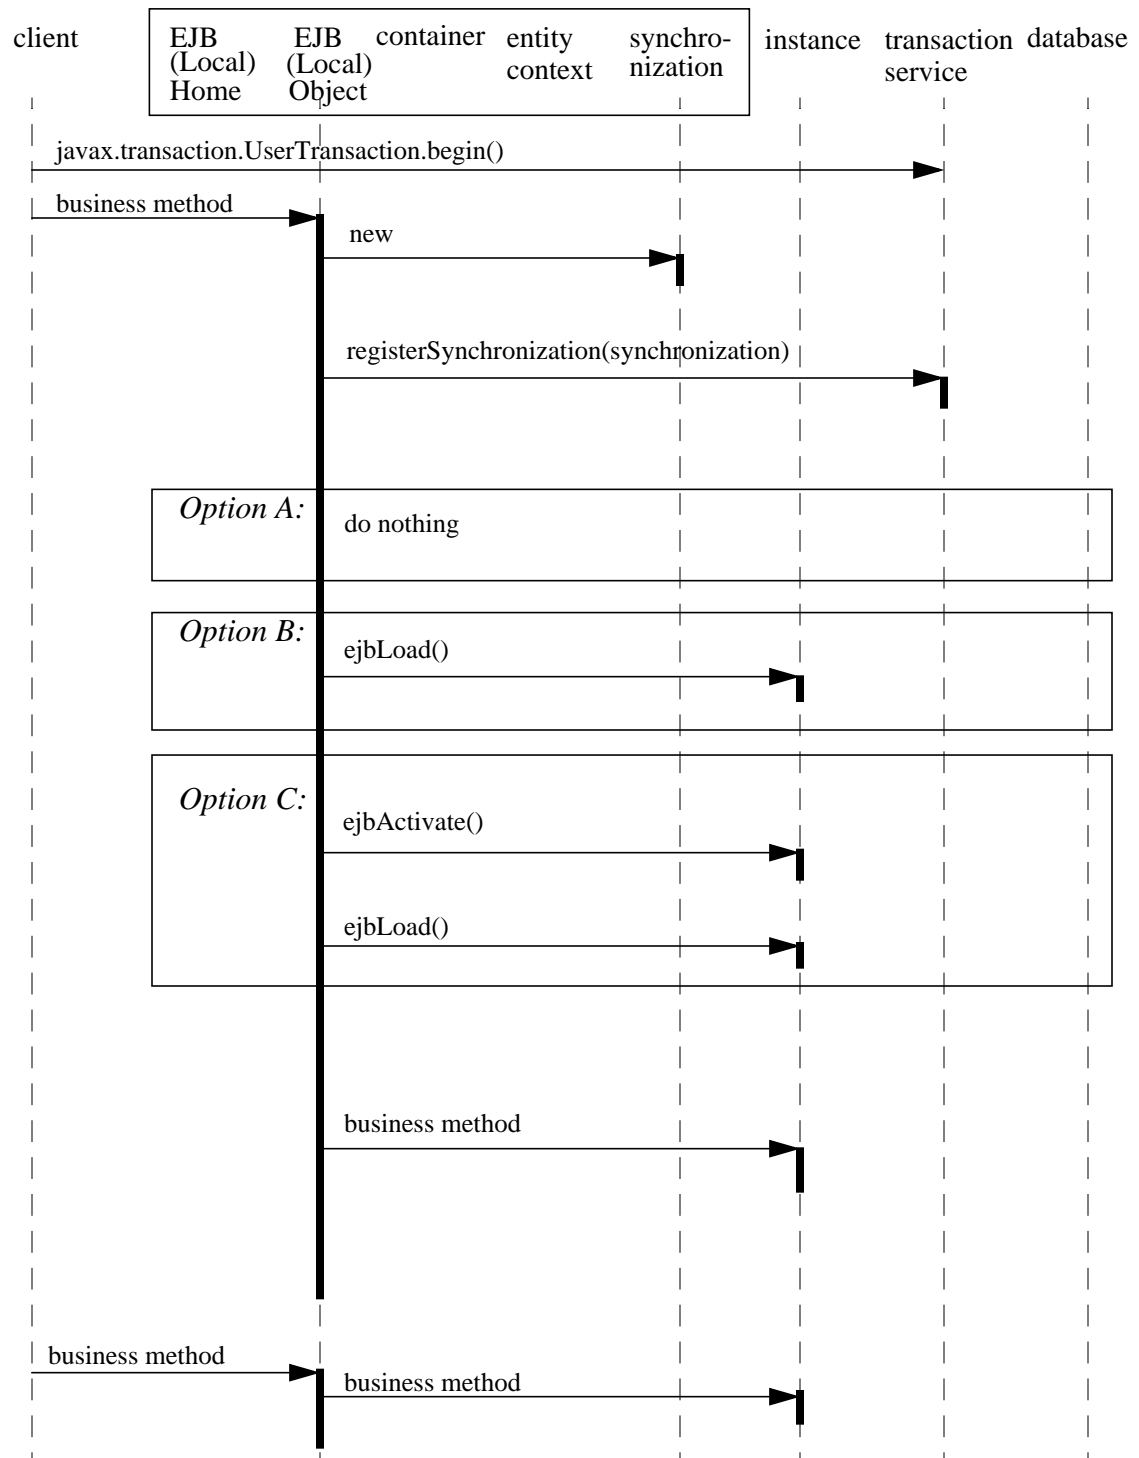


10.9.5 Starting the Next Transaction

The following diagram illustrates the protocol performed for an entity bean instance with container-managed persistence at the beginning of a new transaction. The three options illustrated in the diagram correspond to the three commit options in the previous subsection.

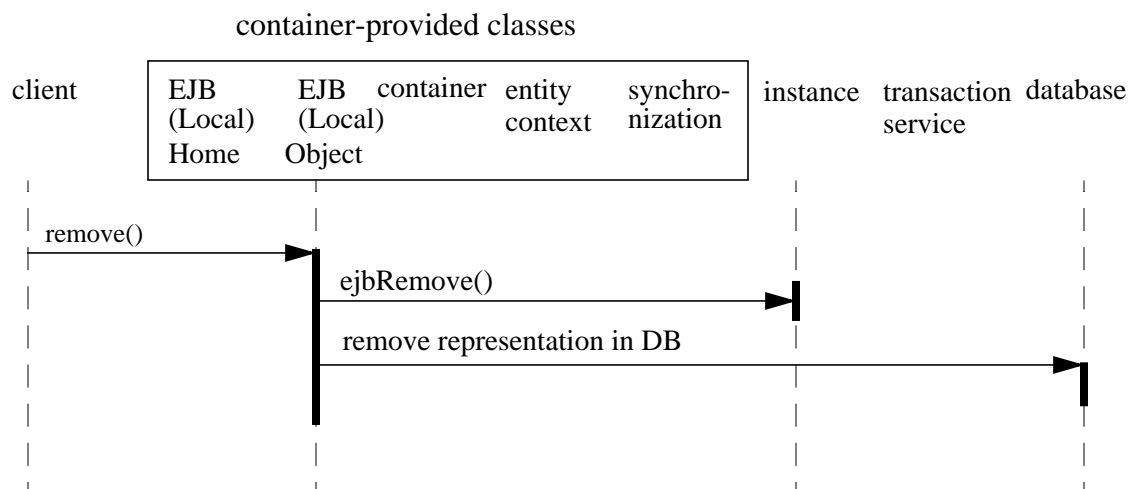
Figure 33 OID of Start of Transaction for an Entity Bean Instance with Container-Managed Persistence

container-provided classes




10.9.6 Removing an Entity Object

Figure 34 OID of Removal of an Entity Bean Object with Container-Managed Persistence



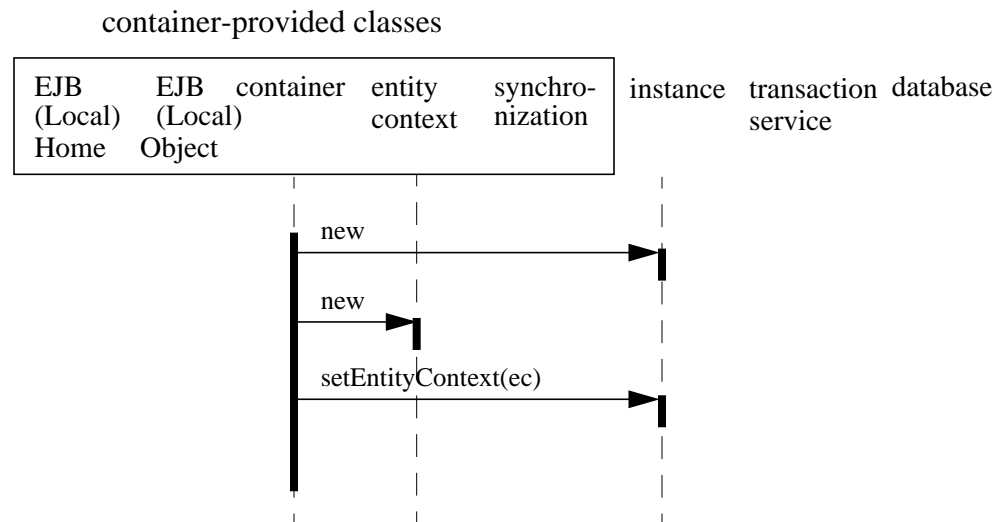
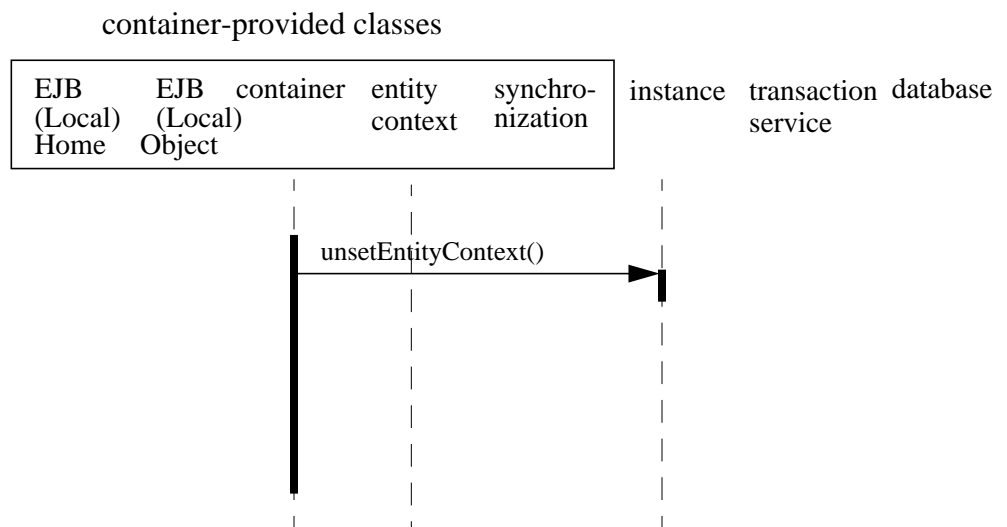
10.9.7 Finding an Entity Object

Figure 35 OID of Execution of a Finder Method on an Entity Bean Instance with Container-Managed Persistence



10.9.8 Adding and Removing an Instance from the Pool

The diagrams in Subsections 10.9.2 through 10.9.7 did not show the sequences between the “does not exist” and “pooled” state (see the diagram in Section 10.5.1).

Figure 36 OID of a Container Adding an Instance to the Pool**Figure 37** OID of a Container Removing an Instance from the Pool

EJB QL: EJB Query Language for Container-Managed Persistence Query Methods

The Enterprise JavaBeans query language, EJB QL, is used to define queries for entity beans with container-managed persistence. EJB QL enables the Bean Provider to specify the semantics of query methods in a portable way.

11.1 Overview

EJB QL is a query specification language for the finder and select methods of entity beans with container-managed persistence. EJB QL can be compiled to a target language, such as SQL, of a database or other persistent store. This allows the execution of queries to be shifted to the native language facilities provided by the persistent store, instead of requiring queries to be executed on the runtime representation of the entity beans' state. As a result, query methods can be optimizable as well as portable.

The Enterprise JavaBeans query language uses the abstract persistence schemas of entity beans, including their relationships, for its data model. It defines operators and expressions based on this data model.

The Bean Provider uses EJB QL to write queries based on the abstract persistence schemas and the relationships defined in the deployment descriptor. EJB QL depends on navigation and selection based on the cmp-fields and cmr-fields of the related entity beans. The Bean Provider can navigate from an entity bean to other entity beans by using the names of cmr-fields in EJB QL queries.

EJB QL allows the Bean Provider to use the abstract schema types of entity beans in a query if the abstract persistence schemas of the beans are defined in the same deployment descriptor as the query.

It is possible to parse and validate EJB QL queries before entity beans are deployed because EJB QL is based on the abstract schema types of entity beans.

EJB QL queries can be used in two different ways:

- as queries for selecting entity objects through finder methods defined in the home interface. Finder methods allow the results of an EJB QL query to be used by the clients of the entity bean.
- as queries for selecting entity objects or other values derived from an entity bean's abstract schema type through select methods defined on the entity bean class. Select methods allow the Bean Provider to use EJB QL to find objects or values related to the state of an entity bean without directly exposing the results to the client.

11.2 EJB QL Definition

EJB QL uses a SQL-like syntax to select objects or values based on the abstract schema types and relationships of entity beans. The path expressions of EJB QL allow the Bean Provider to navigate over relationships defined by the cmr-fields of the abstract schema types of entity beans.

This chapter provides the full definition of the language.

An EJB QL query is a string which consists of the following clauses:

- a SELECT clause, which determines the type of the objects or values to be selected.
- a FROM clause, which provides declarations that designate the domain to which the expressions specified in the SELECT clause and WHERE clause of the query apply.
- an optional WHERE clause, which may be used to restrict the results that are returned by the query.
- an optional ORDER BY clause, which may be used to order the results that are returned by the query.

In BNF syntax, an EJB QL query is defined as:

EJB QL ::= select_clause from_clause [where_clause] [orderby_clause]

An EJB QL query must always have a SELECT and a FROM clause. The square brackets [] indicate that the WHERE and ORDER BY clauses are optional.

An EJB QL query may have parameters that correspond to the parameters of the finder or select method for which it is defined.

An EJB QL query is statically defined in the `ejb-ql` deployment descriptor element.

11.2.1 Abstract Schema Types and Query Domains

EJB QL is a typed language whose design is based on the type model of EJB 2.0 container-managed persistence. Every expression in EJB QL has a type. The type of the expression is derived from the structure of the expression; the abstract schema types of the identification variable declarations; the types to which the `cmp-fields` and `cmr-fields` evaluate; and the types of literals. The allowable types in EJB QL are the abstract schema types of entity beans and `cmp-fields`.

The abstract schema type of an entity bean is derived from its entity bean class and the information provided in the deployment descriptor. There is a one-to-one mapping between entity bean abstract schema types and entity bean homes. Abstract schema names, as specified by the `abstract-schema-name` elements in the deployment descriptor, are used to denote entity bean abstract schema types in EJB QL.

Informally, the abstract schema type of an entity bean can be characterized as follows:

- *For every get accessor method of the entity bean class that corresponds to a `cmp-field` element in the deployment descriptor, there is a field (“`cmp-field`”) whose abstract schema type corresponds to the result type of the accessor method.*
- *For every get accessor method of the entity bean that corresponds to a `cmr-field` element in the deployment descriptor, there is a field (“`cmr-field`”) whose type is the abstract schema type of the entity bean denoted by the `ejb-name` element contained in the corresponding `ejb-relationship-role` element (or, if the role has a multiplicity of `Many`, a collection of such).*

Abstract schema types are specific to the EJB QL data model. The container is not required to implement or otherwise materialize an abstract schema type.

The domain of an EJB QL query consists of the abstract schema types of all entity beans with container-managed persistence that are defined in the same deployment descriptor.

The Bean Provider creates an `ejb-jar` file which contains a deployment descriptor describing several entity beans and their relationships. EJB QL assumes that a single deployment descriptor in an `ejb-jar` file constitutes a nondecomposable unit for the container responsible for implementing the abstract persistence schemas of the entity beans and the relationships defined in the deployment descriptor and the `ejb-jar` file. Queries can be written by utilizing navigation over the `cmr-fields` of related beans supplied in the same `ejb-jar` by the Bean Provider because they are implemented and managed by the same container.

The domain of a query may be restricted by the *navigability* of the relationships of the entity bean on which it is based. The cmr-fields of an entity bean's abstract schema type determine navigability. Using the cmr-fields and their values, a query can select related entity beans and use their abstract schema types in the query.

11.2.2 Query Methods

EJB QL is used for two types of query methods:

- **Finder methods**—Finder methods are defined in the home interface(s) of an entity bean and return entity objects or local entity objects. A finder method that is defined on the remote home interface must return either an EJBObject or a collection of EJBObjects; a finder method that is defined on the local home interface must return either an EJBLocalObject or a collection of EJBLocalObjects. The result type of a finder method defined on the remote home interface of an entity bean is the entity bean's remote interface (or a collection of objects implementing the entity bean's remote interface). The result type of a finder method defined on the local home interface of an entity bean is the entity bean's local interface (or a collection of objects implementing the entity bean's local interface).
- **Select methods**—Select methods are a special type of query method not directly exposed through the client view. The Bean Provider typically uses select methods to select the persistent state of an entity object or to select entities that are related to the entity bean for which the query is defined. The result type of a select method can be an EJBLocalObject (or a collection of EJBLocalObjects), an EJBObject (or a collection of EJBObjects), a cmp-field value (or a collection of such), or the result of an aggregate function.

11.2.3 Naming

Entity beans are designated in EJB QL query strings by their abstract schema names. The Bean Provider assigns unique abstract schema names to entity beans as part of the development process so that they can be used within queries. These unique names are scoped within the deployment descriptor file.


11.2.4 Examples

The following convention refers to the names used for entity beans in subsequent examples: An entity bean as a whole is designated by <name>EJB, and its entity bean class and abstract schema type are designated by <name>, following the convention used to name the local interface of an entity bean.

The first example assumes that the Bean Provider provides several entity beans, OrderEJB, ProductEJB, LineItemEJB, ShippingAddressEJB, and BillingAddressEJB. The abstract schema types for these entity beans are Order, Product, LineItem, ShippingAddress, and BillingAddress respectively. These beans are logically in the same ejb-jar file, as shown in Figure 38. Only two of the entity beans, OrderEJB and ProductEJB, have remote interfaces and remote home interfaces.

Figure 38

Several Entity Beans with Abstract Persistence Schemas Defined in the Same Ejb-jar File.



The entity beans `ShippingAddress` and `BillingAddress` each have one-to-many relationships with `Order`. There is also a one-to-many relationship between `Order` and `LineItem`. The entity bean `LineItem` is related to `Product` in a many-to-one relationship.

EJB QL allows the Bean Provider to specify finder queries for `OrderEJB` by navigating over the cmr-fields and cmp-fields defined by `Order` and `LineItem`. A finder method query to find all orders with pending line items might be written as follows:

```
SELECT DISTINCT OBJECT(o)
FROM Order AS o, IN(o.lineItems) AS l
WHERE l.shipped = FALSE
```

This query navigates over the cmr-field `lineItems` of the abstract schema type `Order` to find line items, and uses the cmp-field `shipped` of `LineItem` to select those orders that have at least one line item that has not yet shipped. (Note that this query does not select orders that have no line items.)

Although predefined reserved identifiers, such as `DISTINCT`, `OBJECT`, `FROM`, `AS`, `IN`, `WHERE`, and `FALSE` appear in upper case in this example, predefined reserved identifiers are case insensitive.

The `SELECT` clause of this example designates the return type of this query to be of type `Order`. If this query is defined for a finder method on the entity bean's remote home interface, the finder method will return objects of the entity bean's remote interface type corresponding to the abstract schema type instances selected by the query. If this same query is defined for a finder method on the entity bean's local home interface, the finder method will return objects of the entity bean's local interface type corresponding to these same abstract schema type instances. Finder methods must always return `EJBObjects` or `EJBLocalObjects` of the bean type for which the query method is defined.

Because the same deployment descriptor defines the abstract persistence schemas of the related entity beans, the Bean Provider can also specify a query for `OrderEJB` that utilizes the abstract schema type of `ProductEJB`, and hence the `cmp`-fields and `cmr`-fields of both the abstract schema types `Order` and `Product`. For example, if the abstract schema type `Product` has a `cmp`-field named `product_type`, a finder query for `OrderEJB` can be specified using this `cmp`-field. Such a finder query might be: “Find all orders for products with product type *office supplies*”. An EJB QL query string for this might be as follows.

```
SELECT DISTINCT OBJECT(o)
FROM Order o, IN(o.lineItems) l
WHERE l.product.product_type = 'office_supplies'
```

Because `Order` is related to `Product` by means of the relationships between `Order` and `LineItem` and between `LineItem` and `Product`, navigation using the `cmr`-fields `lineItems` and `product` is needed to express the query. This query is specified by using the `abstract-schema-name` for `OrderEJB`, namely `Order`, which designates the abstract schema type over which the query ranges. The basis for the navigation is provided by the `cmr`-fields `lineItems` and `product` of the abstract schema types `Order` and `LineItem` respectively.

11.2.5 The FROM Clause and Navigational Declarations

The `FROM` clause of an EJB QL query defines the domain of the query by declaring identification variables. The domain of the query may be constrained by path expressions.

Identification variables designate instances of a particular entity bean abstract schema type. The `FROM` clause can contain multiple identification variable declarations separated by a comma (,).

```
from_clause ::= FROM identification_variable_declaration
                [, identification_variable_declaration]*
identification_variable_declaration ::= collection_member_declaration /
                range_variable_declaration
collection_member_declaration ::= IN (collection_valued_path_expression) [AS] identifier
range_variable_declaration ::= abstract_schema_name [AS] identifier
```

The following subsections discuss the constructs used in the `FROM` clause.

11.2.5.1 Identifiers

An identifier is a character sequence of unlimited length. The character sequence must begin with a Java identifier start character, and all other characters must be Java identifier part characters. An identifier start character is any character for which the method `Character.isJavaIdentifierStart` returns true. This includes the underscore (`_`) character and the dollar sign (`$`) character. An identifier part character is any character for which the method `Character.isJavaIdentifierPart` returns true. The question mark (`?`) character is reserved for use by EJB QL.

The following are the reserved identifiers in EJB QL: *SELECT*, *FROM*, *WHERE*, *DISTINCT*, *OBJECT*, *NULL*, *TRUE*, *FALSE*, *NOT*, *AND*, *OR*, *BETWEEN*, *LIKE*, *IN*, *AS*, *UNKNOWN*^[22], *EMPTY*, *MEMBER*, *OF*, *IS*, *AVG*, *MAX*, *MIN*, *SUM*, *COUNT*, *ORDER*, *BY*, *ASC*, *DESC*, *MOD*.

Reserved identifiers are case insensitive. Reserved identifiers must not be used as identification variables.

It is recommended that the Bean Provider not use other SQL reserved words as identification variables in EJB QL queries because they may be used as EJB QL reserved identifiers in future versions of the EJB specification.

11.2.5.2 Identification Variables

An identification variable is a valid identifier declared in the FROM clause of an EJB QL query. An identification variable may be declared using the special operators IN and, optionally, AS.

All identification variables must be declared in the FROM clause. Identification variables cannot be declared in other clauses.

An identification variable must not be a reserved identifier or have the same name as any of the following:

- abstract-schema-name
- ejb-name^[23]

Identification variables are case insensitive.

An identification variable evaluates to a value of the type of the expression used in declaring the variable. For example, consider the previous finder query for OrderEJB:

```
SELECT DISTINCT OBJECT(o)
FROM Order o, IN(o.lineItems) l
WHERE l.product.product_type = 'office_supplies'
```

In the FROM clause declaration `IN(o.lineItems) l`, the identification variable `l` evaluates to any `LineItem` value directly reachable from `Order`. The cmr-field `lineItems` is a collection of instances of the abstract schema type `LineItem` and the identification variable `l` refers to an element of this collection. The type of `l` is the abstract schema type of `LineItem`.

An identification variable ranges over the abstract schema type of an entity bean. An identification variable designates an instance of an entity bean abstract schema type or an element of a collection of entity bean abstract schema types instances. Identification variables are existentially quantified in an EJB QL query.

An identification variable always designates a reference to a single value. It is declared in one of two ways; as a range variable or as a collection member identification variable:

- A range variable is declared using the abstract schema name of an entity bean.

[22] Not currently used in EJB QL; reserved for future use.

[23] Use of ejb-names in EJB QL is reserved for future use.

- A collection member identification variable is declared using a collection-valued path expression.

The identification variable declarations are evaluated from left to right in the FROM clause. A collection member identification variable declaration can use the result of a preceding identification variable declaration of the query string.

11.2.5.3 Range Variable Declarations

The EJB QL syntax for declaring an identification variable as a range variable is similar to that of SQL; optionally, it uses the AS keyword.

range_variable_declaration ::= abstract_schema_name [AS] identifier

Objects or values that are related to an entity bean are typically obtained by navigation using path expressions. However, navigation does not reach all objects. Range variable declarations allow the Bean Provider to designate a “root” for objects which may not be reachable by navigation.

If the Bean Provider wants to select values by comparing more than one instance of an entity bean abstract schema type, more than one identification variable ranging over the abstract schema type is needed in the FROM clause.

The following finder method query returns orders whose quantity is greater than the order quantity for John Smith. This example illustrates the use of two different identification variables in the FROM clause, both of the abstract schema type Order. The SELECT clause of this query determines that it is the orders with quantities larger than John Smith’s that are returned.

```
SELECT DISTINCT OBJECT(o1)
FROM Order o1, Order o2
WHERE o1.quantity > o2.quantity AND
      o2.customer.lastname = 'Smith' AND
      o2.customer.firstname= 'John'
```

11.2.5.4 Collection Member Declarations

An identification variable declared by a *collection_member_declaration* ranges over values of a collection obtained by navigation using a path expression. Such a path expression represents a navigation involving the cmr-fields of an entity bean abstract schema type. Because a path expression can be based on another path expression, the navigation can use the cmr-fields of related entity beans. Path expressions are discussed in Section 11.2.5.6.

An identification variable of a collection member declaration is declared using a special operator, the reserved identifier IN. The argument to the IN operator is a collection-valued path expression. The path expression evaluates to a collection type specified as a result of navigation to a collection-valued cmr-field of an entity bean abstract schema type.

The syntax for declaring a collection member identification variable is as follows:

collection_member_declaration ::= IN (collection_valued_path_expression) [AS] identifier

For example, the FROM clause for a query defined for OrderEJB might contain the following collection member declaration:

```
IN(o.lineItems) l
```

In this example, `lineItems` is the name of a cmr-field whose value is a collection of instances of the abstract schema type `LineItem` of the `LineItemEJB` entity bean. The identification variable `l` designates a member of this collection, a *single* `LineItem` abstract schema type instance. In this example, `o` is an identification variable of the abstract schema type `Order`.

11.2.5.5 Example

The following FROM clause contains two identification variable declaration clauses. The identification variable declared in the first clause is used in the second clause. The clauses declare the variables `o` and `l` respectively. The range variable declaration `Order AS o` designates the identification variable `o` as a range variable whose type is the abstract schema type, `Order`. The identification variable `l` has the abstract schema type `LineItem`. Because the clauses are evaluated from left to right, the identification variable `l` can utilize the results of the navigation on `o`.

```
FROM Order AS o, IN(o.lineItems) l
```

11.2.5.6 Path Expressions

An identification variable followed by the navigation operator (`.`) and a cmp-field or cmr-field is a path expression. The type of the path expression is the type computed as the result of navigation; that is, the type of the cmp-field or cmr-field to which the expression navigates.

Depending on navigability, a path expression that leads to a cmr-field may be further composed. Path expressions can be composed from other path expressions if the original path expression evaluates to a single-valued type (not a collection) corresponding to a cmr-field. A path expression that ends in a cmp-field is terminal and cannot be further composed.

Path expression navigability is composed using “inner join” semantics. That is, if the value of a non-terminal cmr-field in the path expression is null, the path is considered to have no value, and does not participate in the determination of the result.

The syntax for single-valued path expressions and collection valued path expressions is as follows:

```
cmp_path_expression ::=
    {identification_variable | single_valued_cmr_path_expression}.cmp_field
single_valued_cmr_path_expression ::=
    identification_variable.[single_valued_cmr_field.*]single_valued_cmr_field
single_valued_path_expression ::=
    cmp_path_expression | single_valued_cmr_path_expression
collection_valued_path_expression ::=
    identification_variable.[single_valued_cmr_field.*]collection_valued_cmr_field
```

A *single_valued_cmr_field* is designated by the name of a cmr-field in a one-to-one or many-to-one relationship. The type of a *single_valued_cmr_path_expression* is the abstract schema type of the related entity bean.

A *collection_valued_cmr_field* is designated by the name of a cmr-field in a one-to-many or a many-to-many relationship. The type of a *collection_valued_cmr_field* is a collection of values of the abstract schema type of the related entity bean.

Navigation to a related entity bean results in a value of the related entity bean's abstract schema type.

The evaluation of a path expression terminating in a cmp-field results in the abstract schema type corresponding to the Java type designated by the cmp-field.

It is syntactically illegal to compose a path expression from a path expression that evaluates to a collection. For example, if *o* designates *Order*, the path expression *o.lineItems.product* is illegal since navigation to *lineItems* results in a collection. This case should produce an error when the EJB QL query string is verified. To handle such a navigation, an identification variable must be declared in the *FROM* clause to range over the elements of the *lineItems* collection. Another path expression must be used to navigate over each such element in the *WHERE* clause of the query, as in the following:

```
SELECT OBJECT(o)
FROM Order AS o, IN(o.lineItems) l
WHERE l.product.name = 'widget'
```

11.2.6 WHERE Clause and Conditional Expressions

The *WHERE* clause of a query consists of a conditional expression used to select objects or values that satisfy the expression. The *WHERE* clause thus restricts the result of a query.

A *WHERE* clause is defined as follows:

where_clause ::= **WHERE** *conditional_expression*

The following sections describe the language constructs that can be used in a conditional expression of the *WHERE* clause.

11.2.6.1 Literals

A string literal is enclosed in single quotes—for example: 'literal'. A string literal that includes a single quote is represented by two single quotes—for example: 'literal"s'. EJB QL string literals, like Java *String* literals, use unicode character encoding.

An exact numeric literal is a numeric value without a decimal point, such as 57, -957, +62. Exact numeric literals support numbers in the range of Java *long*. Exact numeric literals use the Java integer literal syntax.

An approximate numeric literal is a numeric value in scientific notation, such as 7E3, -57.9E2, or a numeric value with a decimal, such as 7., -95.7, +6.2. Approximate numeric literals support numbers in the range of Java `double`. Approximate literals use the Java floating point literal syntax.

The Bean Provider may utilize appropriate suffixes to indicate the specific type of the literal in accordance with the Java Language Specification.

The boolean literals are `TRUE` and `FALSE`.

Although predefined reserved literals appear in upper case, they are case insensitive.

11.2.6.2 Identification Variables

All identification variables used in the `WHERE` clause of an EJB QL query must be declared in the `FROM` clause, as described in Section 11.2.5.2.

Identification variables are existentially quantified in the `WHERE` clause. This means that an identification variable represents a member of a collection or an instance of an entity bean's abstract schema type. An identification variable never designates a collection in its entirety.

11.2.6.3 Path Expressions

It is illegal to use a *collection_valued_path_expression* within a `WHERE` clause as part of a conditional expression except in an *empty_collection_comparison_expression* or *collection_member_expression*.

11.2.6.4 Input Parameters

The following rules apply to input parameters. Input parameters can only be used in the `WHERE` clause of a query.

- Input parameters are designated by the question mark (?) prefix followed by an integer. For example: ?1.
- Input parameters are numbered starting from 1.
- The number of distinct input parameters in an EJB QL query must not exceed the number of input parameters for the finder or select method. It is not required that the EJB QL query use all of the input parameters for the finder or select method.
- An input parameter evaluates to the abstract schema type of the corresponding parameter defined in the signature of the finder or select method with which the query is associated. It is the responsibility of the container to map the input parameter to the appropriate abstract schema type value.

Note that if an input parameter value is null, comparison operations or arithmetic operations involving the input parameter will return an unknown value. See Section 11.2.10.

11.2.6.5 Conditional Expression Composition

Conditional expressions are composed of other conditional expressions, comparison operations, logical operations, path expressions that evaluate to boolean values, and boolean literals.

Arithmetic expressions can be used in comparison expressions. Arithmetic expressions are composed of other arithmetic expressions, arithmetic operations, path expressions that evaluate to numeric values, and numeric literals.

Arithmetic operations use numeric promotion.

Standard bracketing () for ordering expression evaluation is supported.

Conditional expressions are defined as follows:

```
conditional_expression ::= conditional_term | conditional_expression OR conditional_term
conditional_term ::= conditional_factor | conditional_term AND conditional_factor
conditional_factor ::= [ NOT ] conditional_primary
conditional_primary ::= simple_cond_expression | (conditional_expression)
simple_cond_expression ::= comparison_expression | between_expression | like_expression |
                        in_expression | null_comparison_expression |
                        empty_collection_comparison_expression |
                        collection_member_expression
```

11.2.6.6 Operators and Operator Precedence

The operators are listed below in order of decreasing precedence.

- Navigation operator (.)
- Arithmetic operators:
 - +, - unary
 - *, / multiplication and division
 - +, - addition and subtraction
- Comparison operators : =, >, >=, <, <=, <> (not equal)
- Logical operators: NOT, AND, OR

The following sections describe other operators used in specific expressions.

11.2.6.7 Between Expressions

The syntax for the use of the comparison operator [NOT] BETWEEN in an conditional expression is as follows:

```
arithmetic_expression [NOT] BETWEEN arithmetic-expression AND arithmetic-expression
```

The BETWEEN expression

`x BETWEEN y AND z`

is semantically equivalent to:

`y <= x AND x <= z`

The rules for unknown and NULL values in comparison operations apply. See Section 11.2.10.

Examples are:

`p.age BETWEEN 15 and 19` is equivalent to `p.age >= 15 AND p.age <= 19`

`p.age NOT BETWEEN 15 and 19` is equivalent to `p.age < 15 OR p.age > 19`

11.2.6.8 In Expressions

The syntax for the use of the comparison operator [NOT] IN in a conditional expression is as follows:

cmp_path_expression [**NOT**] **IN** ({*literal* | *input_parameter*} [, {*literal* | *input_parameter*}]*)

The *cmp_path_expression* must have a string or numeric value. The literal and/or *input_parameter* values must be *like* the same abstract schema type of the *cmp_path_expression* in type. (See Section 11.2.11).

Examples are:

`o.country IN ('UK', 'US', 'France')` is true for UK and false for Peru, and is equivalent to the expression `(o.country = 'UK') OR (o.country = 'US') OR (o.country = 'France')`.

`o.country NOT IN ('UK', 'US', 'France')` is false for UK and true for Peru, and is equivalent to the expression `NOT ((o.country = 'UK') OR (o.country = 'US') OR (o.country = 'France'))`.

There must be at least one element in the comma separated list that defines the set of values for the IN expression.

If the value of a *cmp_path_expression* in an IN or NOT IN expression is NULL or unknown, the value of the expression is unknown.

11.2.6.9 Like Expressions

The syntax for the use of the comparison operator [NOT] LIKE in a conditional expression is as follows:

cmp_path_expression [**NOT**] **LIKE** *pattern_value* [**ESCAPE** *escape_character*]

The *cmp_path_expression* must have a string value. The *pattern_value* is a string literal or a string-valued input parameter in which an underscore (_) stands for any single character, a percent (%) character stands for any sequence of characters (including the empty sequence), and all other characters stand for themselves. The optional *escape_character* is a single-character string literal or a character-valued input parameter (i.e., `char` or `Character`) and is used to escape the special meaning of the underscore and percent characters in *pattern_value*.^[24]

Examples are:

- *address.phone LIKE '12%3'* is true for '123' '12993' and false for '1234'
- *asentence.word LIKE 'l_se'* is true for 'lose' and false for 'loose'
- *aword.underscored LIKE '_%' ESCAPE '\'* is true for '_foo' and false for 'bar'
- *address.phone NOT LIKE '12%3'* is false for '123' and '12993' and true for '1234'

If the value of the *cmp_path_expression* or *pattern_value* is `NULL` or unknown, the value of the `LIKE` expression is unknown. If the *escape_character* is specified and is `NULL`, the value of the `LIKE` expression is unknown.

11.2.6.10 Null Comparison Expressions

The syntax for the use of the comparison operator `IS NULL` in a conditional expression is as follows:

{single_valued_path_expression | input_parameter} IS [NOT] NULL

A null comparison expression tests whether or not the single-valued path expression or input parameter is a `NULL` value.

11.2.6.11 Empty Collection Comparison Expressions

The syntax for the use of the comparison operator `IS EMPTY` in an *empty_collection_comparison_expression* is as follows:

collection_valued_path_expression IS [NOT] EMPTY

This expression tests whether or not the collection designated by the collection-valued path expression is empty (i.e., has no elements).

Note that a collection-valued path expression can only be used in the WHERE clause in an empty collection comparison expression or in a collection member expression.

[24] Refer to [21] for a more precise characterization of these rules.

The collection designated by the collection-valued path expression used in an empty collection comparison expression must not be used in the FROM clause for the declaration of an identification variable. An identification variable declared as a member of a collection implicitly designates the existence of a non-empty relationship; testing whether the same collection is empty is contradictory. Therefore, the following query is invalid.

```
SELECT OBJECT(o)
FROM Order o, IN(o.lineItems) l
WHERE o.lineItems IS EMPTY
```

If the value of the collection-valued path expression in an empty collection comparison expression is unknown, the value of the empty comparison expression is unknown.

11.2.6.12 Collection Member Expressions

The syntax for the use of the comparison operator `MEMBER OF`^[25] in an *collection_member_expression* is as follows:

```
{single_valued_cmr_path_expression | identification_variable | input_parameter }
  [NOT] MEMBER [OF] collection_valued_path_expression
```

This expression tests whether the designated value is a member of the collection specified by the collection-valued path expression.

Note that a collection-valued path expression can only be used in the WHERE clause in an empty collection comparison expression or in a collection member expression.

If the collection valued path expression designates an empty collection, the value of the `MEMBER OF` expression is `FALSE` and the value of the `NOT MEMBER OF` expression is `TRUE`. Otherwise, if the value of the collection-valued path expression or single-valued cmr path expression in the collection member expression is unknown, the value of the collection member expression is unknown.

11.2.6.13 Functional Expressions

EJB QL includes the following built-in functions^[26], which may be used in the WHERE clause of a query.

String Functions:

- `CONCAT(String, String)` returns a `String`
- `SUBSTRING(String, start, length)` returns a `String`
- `LOCATE(String, String [, start])`^[27] returns an `int`

[25] The use of the reserved word `OF` is optional in this expression.

[26] These functions are a subset of the functions defined for JDBC 2.0 and later drivers, as described in Appendix C of the JDBC 3.0 specification [27].

- `LENGTH(String)` returns an `int`

Note that `start` and `length` are integer values. The first position in a string is designated as 1 by these functions.

Arithmetic Functions:

- `ABS(number)` returns a number (`int`, `float`, or `double`) of the same type as the argument to the function
- `SQRT(double)` returns a `double`
- `MOD(int, int)` returns an `int`

Numeric arguments to these functions may correspond to the numeric Java object types as well as the primitive numeric types.

If the value of any argument to a functional expression is null or unknown, the value of the functional expression is unknown.

11.2.7 **SELECT Clause**

The **SELECT** clause denotes the query result. The **SELECT** clause contains either a single range variable that ranges over an entity bean abstract schema type, a single-valued path expression, or an aggregate select expression. In the case of a finder method, the **SELECT** clause is restricted to contain either a single range variable or a single-valued path expression that evaluates to the abstract schema type of the entity bean for which the finder method is defined.

The **SELECT** clause has the following syntax:

select_clause ::= **SELECT** [**DISTINCT**] {*select_expression* | **OBJECT**(*identification_variable*)}

select_expression ::= *single_valued_path_expression* | *aggregate_select_expression*

aggregate_select_expression ::=

{ **AVG** | **MAX** | **MIN** | **SUM** | **COUNT** } ([**DISTINCT**] *cmp_path_expression*) |
COUNT ([**DISTINCT**] *identification_variable* | *single_valued_cmr_path_expression*)

All standalone identification variables in the **SELECT** clause must be qualified by the **OBJECT** operator. The **SELECT** clause must not use the **OBJECT** operator to qualify path expressions.

Note that the **SELECT** clause must be specified to return a single-valued expression. The query below is therefore not valid:

```
SELECT o.lineItems FROM Order AS o
```

[27] Containers and databases may not support the use of the optional, third argument of the **LOCATE** function. Portable applications should therefore avoid use of this argument.

The **DISTINCT** keyword is used to specify that duplicate values must be eliminated from the query result. If **DISTINCT** is not specified, duplicate values are not eliminated unless the query is specified for a method whose result type is `java.util.Set`. If a query is specified for a method whose result type is `java.util.Set`, but does not specify **DISTINCT**, the container must interpret the query as if **SELECT DISTINCT** had been specified. In general, however, the Bean Provider should specify the **DISTINCT** keyword when writing queries for methods that return `java.util.Set`.

The **SELECT** clause determines the type of the values returned by a query. For example, the following query returns a collection of products:

```
SELECT l.product FROM Order AS o, IN(o.lineItems) l
```

It is the responsibility of the container to map the abstract schema types returned by the query to the Java types that are returned by the finder or select method with which the query is associated and to materialize those return types, as described in Section 11.2.9.

11.2.7.1 Null Values in the Query Result

If the result of an EJB QL query corresponds to a `cmr-field` or `cmp-field` whose value is null, the container must include that null value in the result that is returned by the finder or select method. The Bean Provider can use the **IS NOT NULL** construct to eliminate such null values from the result set of the query.

If the finder or select method is a single-object finder or select method, and the result set of the query consists of a single null value, the container must return the null value as the result of the method. If the result set of a query for a single-object finder or select method contains more than one value (whether non-null, null, or a combination), the container must throw the `FinderException`.

Note, however, that `cmp-field` types defined in terms of Java numeric primitive types cannot produce **NULL** values in the query result. An EJB QL query that returns such a `cmp-field` type as a result type must not return a null value. (If the Bean Provider wishes to allow null values for `cmp-fields`, he or she should specify those `cmp-fields` to have the equivalent Java object types instead of primitive types, e.g., `Integer` rather than `int`.)

11.2.7.2 Aggregate Functions in the SELECT Clause

The result of an EJB QL query may be the result of an aggregate function applied to a path expression.

The following aggregate functions can be used in the **SELECT** clause of an EJB QL query: **AVG**, **COUNT**, **MAX**, **MIN**, **SUM**.

For all aggregate functions except **COUNT**, the path expression that is the argument to the aggregate function must terminate in a `cmp-field`. The path expression argument to **COUNT** may terminate in either a `cmp-field` or a `cmr-field`, or the argument to **COUNT** may be an identification variable.

Arguments to the functions **SUM** and **AVG** must be numeric. Arguments to the functions **MAX** and **MIN** must correspond to orderable `cmp-field` types (i.e., numeric types, string types, character types, or date types).

The argument to an aggregate function may be preceded by the keyword **DISTINCT** to specify that duplicate values are to be eliminated before the aggregate function is applied.^[28]

Null values are eliminated before the aggregate function is applied, regardless of whether the keyword **DISTINCT** is specified.

11.2.7.3 Examples

The following example returns all line items related to some order:

```
SELECT OBJECT(l)
FROM Order o, IN(o.lineItems) l
```

The following query returns all line items regardless of whether a line item is related to any order or product:

```
SELECT OBJECT(l)
FROM LineItems AS l
```

The following query returns the average order quantity:

```
SELECT AVG(o.quantity)
FROM Order o
```

The following query returns the total cost of the items that John Smith has ordered.

```
SELECT SUM(l.price)
FROM Order o, IN(o.lineItems) l
WHERE o.customer.lastname = 'Smith' AND o.customer.firstname = 'John'
```

The following query returns the number of items in John Smith's entire order.

```
SELECT COUNT(l)
FROM Order o, IN(o.lineItems) l
WHERE o.customer.lastname = 'Smith' AND o.customer.firstname = 'John'
```

The following query returns the total number of orders.

```
SELECT COUNT(o)
FROM Order o
```

The following query counts the number of items in John Smith's order for which prices have been specified.

```
SELECT COUNT(l.price)
FROM Order o, IN(o.lineItems) l
WHERE o.customer.lastname = 'Smith' AND o.customer.firstname = 'John'
```

[28] It is legal to specify **DISTINCT** with **MAX** or **MIN**, but it does not affect the result.

Note that this is equivalent to:

```
SELECT COUNT(1)
FROM Order o, IN(o.lineItems) l
WHERE o.customer.lastname = 'Smith' AND o.customer.firstname = 'John'
AND l.price IS NOT NULL
```

11.2.8 ORDER BY Clause

The ORDER BY clause allows the objects or values that are returned by the query to be ordered.

The syntax of the ORDER BY clause is

orderby_clause ::= **ORDER BY** *orderby_item* [, *orderby_item*]*
orderby_item ::= *cmp_path_expression* [ASC | DESC]

When the ORDER BY clause is used in an EJB QL query, the SELECT clause of the query must be one of the following:

1. an identification variable *x*, denoted as OBJECT(*x*)
2. a *single_valued_cmr_path_expression*
3. a *cmp_path_expression*

In the first two cases, each *orderby_item* must be an orderable cmp-field of the entity bean abstract schema type value returned by the SELECT clause. In the third case, the *orderby_item* must evaluate to the same cmp-field of the same entity bean abstract schema type as the *cmp_path_expression* in the SELECT clause.

For example, the first two queries below are legal, but the third and fourth are not.

```
SELECT OBJECT(o)
FROM Customer c, IN(c.orders) o
WHERE c.address.state = 'CA'
ORDER BY o.quantity, o.totalcost
```

```
SELECT o.quantity
FROM Customer c, IN(c.orders) o
WHERE c.address.state = 'CA'
ORDER BY o.quantity
```

```
SELECT l.product.product_name
FROM Order o, IN(o.lineItems) l
WHERE o.customer.lastname = 'Smith' AND o.customer.firstname = 'John'
ORDER BY l.product.price
```

```
SELECT l.product.product_name
FROM Order o, IN(o.lineItems) l
WHERE o.customer.lastname = 'Smith' AND o.customer.firstname = 'John'
ORDER BY o.quantity
```

If more than one *orderby_item* is specified, the left-to-right sequence of the *orderby_item* elements determines the precedence, whereby the leftmost *orderby_item* has highest precedence.

The keyword ASC specifies that ascending ordering be used; the keyword DESC specifies that descending ordering be used. Ascending ordering is the default.

SQL rules for the ordering of null values apply: that is, all null values must appear before all non-null values in the ordering or all null values must appear after all non-null values in the ordering, but it is not specified which.

It is the container's responsibility to ensure that the ordering of the query result is preserved in the result of the finder or select method if the ORDER BY clause is used.

11.2.9 Return Value Types

The value of a query result, specified by the SELECT clause, is an entity bean abstract schema type, a cmp-field type, or the result of an aggregate function. The finder or select method with which the query is associated in the deployment descriptor determines how this result is mapped to the Java type that is visible as the result of the query method.

How the result type of a query is mapped depends on whether the query is defined for a finder method on the remote home interface, for a finder method on the local home interface, or for a select method.

- The result type of a query for a finder method must be the entity bean abstract schema type that corresponds to the entity bean type of the entity bean on whose home interface the finder method is defined. If the query is used for a finder method defined on the remote home interface of the bean, the result of the finder method is the entity bean's remote interface (or a collection of objects implementing the entity bean's remote interface). If the finder method is defined on the local home interface, the result is the entity bean's local interface (or a collection of objects implementing the entity bean's local interface).
- If the result type of a query for a select method is an entity bean abstract schema type, the return values for the query method are instances of the entity bean's local interface or instances of the entity bean's remote interface, depending on whether the value of the `result-type-mapping` deployment descriptor element contained in the `query` element for the select method is `Local` or `Remote`. The default value for `result-type-mapping` is `Local`.
- If the result type of a query used for a select method is an abstract schema type corresponding to a cmp-field type (excluding queries whose SELECT clause uses one of the aggregate functions AVG, COUNT, MAX, MIN, SUM), the result type of the select method is as follows:
 - If the Java type of the cmp-field is an object type and the select method is a single-object select method, the result of the select method is an instance of that object type. If the select method is a multi-object select method, the result is a collection of instances of that type.
 - If the Java type of the cmp-field is a primitive Java type (e.g., int), and the select method is a single-object select method, the result of the select method is that primitive type.

- If the Java type of the cmp-field is a primitive Java type (e.g., int), and the select method is a multi-object select method, the result of the select method is a collection of values of the corresponding wrapped type (e.g., Integer).
- If the select method query is an aggregate query, the select method must be a single-object select method.
 - The result type of the select method must be a primitive type, a wrapped type, or an object type that is compatible with the standard JDBC conversion mappings for the type of the cmp-field [27].
 - If the aggregate query uses the SUM, AVG, MAX, or MIN operator, and the result type of the select method is an object type and there are no values to which the aggregate function can be applied, the select method returns null.
 - If the aggregate query uses the SUM, AVG, MAX, or MIN operator, and the result type of the select method is a primitive type and there are no values to which the aggregate function can be applied, the container must throw the `ObjectNotFoundException`.
 - If the aggregate query uses the COUNT operator, the result of the select method should be an exact numeric type. If there are no values to which the COUNT method can be applied, the result of the select method is 0.

The result of a finder or select method may contain a null value if a cmp-field or cmr-field in the query result is null.

11.2.10 Null Values

When the target of a reference does not exist in the persistent store, its value is regarded as NULL. SQL 92 NULL semantics [21] defines the evaluation of conditional expressions containing NULL values.

The following is a brief description of these semantics:

- Comparison or arithmetic operations with a NULL value always yield an unknown value.
- Two NULL values are not considered to be equal, the comparison yields an unknown value.
- Comparison or arithmetic operations with an unknown value always yield an unknown value.
- The IS NULL and IS NOT NULL operators convert a NULL cmp-field or single-valued cmr-field value into the respective TRUE or FALSE value.
- Boolean operators use three valued logic, defined by Table 7, Table 8, and Table 9.

Table 7 Definition of the AND Operator

AND	T	F	U
T	T	F	U
F	F	F	F
U	U	F	U

Table 8 Definition of the OR Operator

OR	T	F	U
T	T	T	T
F	T	F	U
U	T	U	U

Table 9 Definition of the NOT Operator

NOT	
T	F
F	T
U	U

Note: EJB QL defines the empty string, "", as a string with 0 length, which is not equal to a NULL value. However, NULL values and empty strings may not always be distinguished when queries are mapped to some persistent stores. The Bean Provider should therefore not rely on the semantics of EJB QL comparisons involving the empty string and NULL value.

11.2.11 Equality and Comparison Semantics

EJB QL only permits the values of *like* types to be compared. A type is *like* another type if they correspond to the same Java language type, or if one is a primitive Java language type and the other is the wrapped Java class type equivalent (e.g., `int` and `Integer` are like types in this sense). There is one exception to this rule: it is valid to compare numeric values for which the rules of numeric promotion apply. Conditional expressions attempting to compare non-like type values are disallowed except for this numeric case.

Note that EJB QL permits the arithmetic operators and comparison operators to be applied to cmp-fields and input parameters of the wrapped Java class equivalents to the primitive numeric Java types.

Two entity objects of the same abstract schema type are equal if and only if they have the same primary key value.

11.2.12 Restrictions

Date and time values should use the standard Java `long` millisecond value. The standard way to produce millisecond values is to use `java.util.Calendar`.

Although SQL requires support for fixed decimal comparison in arithmetic expressions, EJB QL does not. For this reason EJB QL restricts exact numeric literals to those without a decimal point (and numerics with a decimal point as an alternate representation for approximate numeric values).

Support for the `BigDecimal` and `BigInteger` types is optional for containers in EJB 2.1. Applications that depend on such types in EJB QL queries may not be portable.

Boolean comparison is restricted to `=` and `<>`.

EJB QL does not support the use of comments.

The data model for container-managed persistence does not currently support inheritance. Therefore, entity objects of different types cannot be compared. EJB QL queries that contain such comparisons are invalid.

11.3 Examples

The following examples illustrate the syntax and semantics of EJB QL. These examples are based on the example presented in Section 11.2.4.

11.3.1 Simple Queries

Find all orders:

```
SELECT OBJECT(o)
FROM Order o
```

Find all orders that need to be shipped to California:

```
SELECT OBJECT(o)
FROM Order o
WHERE o.shipping_address.state = 'CA'
```

Find all states for which there are orders:

```
SELECT DISTINCT o.shipping_address.state
FROM Order o
```

11.3.2 Queries with Relationships

Find all orders that have line items:

```
SELECT DISTINCT OBJECT(o)
FROM Order o, IN(o.lineItems) l
```

Note that the result of this query does not include orders with no associated line items. This query can also be written as:

```
SELECT OBJECT(o)
FROM Order o
WHERE o.lineItems IS NOT EMPTY
```

Find all orders that have no line items:

```
SELECT OBJECT(o)
FROM Order o
WHERE o.lineItems IS EMPTY
```

Find all pending orders:

```
SELECT DISTINCT OBJECT(o)
FROM Order o, IN(o.lineItems) l
WHERE l.shipped = FALSE
```

Find all orders in which the shipping address differs from the billing address. This example assumes that the Bean Provider uses two distinct entity beans to designate shipping and billing addresses, as in Figure 38.

```
SELECT OBJECT(o)
FROM Order o
WHERE
NOT (o.shipping_address.state = o.billing_address.state AND
     o.shipping_address.city = o.billing_address.city AND
     o.shipping_address.street = o.billing_address.street)
```

If the Bean Provider uses a single entity bean in two different relationships for both the shipping address and the billing address, the above expression can be simplified based on the equality rules defined in Section 11.2.11. The query can then be written as:

```
SELECT OBJECT(o)
FROM Order o
WHERE o.shipping_address <> o.billing_address
```

The query checks whether the same entity bean abstract schema type instance (identified by its primary key) is related to an order through two distinct relationships.

Find all orders for a book titled 'Applying Enterprise JavaBeans: Component-Based Development for the J2EE Platform':

```
SELECT DISTINCT OBJECT(o)
FROM Order o, IN(o.lineItems) l
WHERE l.product.type = 'book' AND
     l.product.name = 'Applying Enterprise JavaBeans:
     Component-Based Development for the J2EE Platform'
```

11.3.3 Queries Using Input Parameters

The following query finds the orders for a product whose name is designated by an input parameter:

```
SELECT DISTINCT OBJECT(o)
FROM Order o, IN(o.lineItems) l
WHERE l.product.name = ?1
```

For this query, the input parameter must be of the type of the cmp-field name, i.e., a string.

11.3.4 Queries for Select Methods

The following select queries illustrate the selection of values other than entity beans.

The following EJB QL query selects the names of all products that have been ordered.

```
SELECT DISTINCT l.product.name
FROM Order o, IN(o.lineItems) l
```

The following query finds the names of all products in the order specified by a particular order number. The order number is specified by a parameter that corresponds to the primary key of Order. Note that because this query does not specify `DISTINCT` in its `SELECT` clause, if it is specified for a query method whose return type is `java.util.Collection`, the collection that is returned may contain duplicates. In this example, such duplicates correspond to products that have been ordered multiple times in the given order.

```
SELECT l.product.name
FROM Order o, IN(o.lineItems) l
WHERE o.ordernumber = ?1
```

It is the responsibility of the container to interpret the query such that no duplicates are produced if the result type of the query method is `java.util.Set`.

Consider the following query for a select method:

```
SELECT o.shipping_address.city
FROM Order o
```

This query returns the names of all the cities of the shipping addresses of all orders. The result type of the select method, which is either `java.util.Collection` or `java.util.Set`, determines whether the query may return duplicate city names.

11.3.5 EJB QL and SQL

EJB QL, like SQL, treats the `FROM` clause as a cartesian product. The `FROM` clause is similar to that of SQL in that the declared identification variables affect the results of the query even if they are not used in the `WHERE` clause. The Bean Provider should use caution in defining identification variables because the domain of the query can depend on whether there are any values of the declared type.

For example, the `FROM` clause below defines a query over all orders that have line items and existing products. If there are no Product instances in the persistent store, the domain of the query is empty and no order is selected.

```
SELECT OBJECT(o)
FROM Order AS o, IN(o.lineItems) l, Product p
```

The container can represent the abstract schemas of a set of entity beans in an application using a relational database. There are multiple ways to define a mapping to a set of tables in a relational database. Although this area is beyond the scope of this specification, a sample mapping and translation of EJB QL to SQL is described to clarify the semantics of EJB QL.

A mapping strategy from a set of entity beans to a relational database might be to map each entity bean to a separate table. One-to-many relationships may be represented by foreign keys in the related table from the many side and many-to-many relationships may be represented by using an auxiliary table that contains the primary keys of the related objects.

Because the FROM clause represents a cartesian product, the SQL result may contain duplicates. If the query is for a method whose return type is `java.util.Set`, the container would therefore typically utilize a SELECT DISTINCT clause in translating the query to SQL. The query method result may contain duplicates if the return type of the query method is `java.util.Collection` and DISTINCT is not specified in the SELECT clause of the EJB QL query.

The following translation example illustrates the mapping of entity beans to relational database tables. The entity bean OrderEJB is represented by the table ORDER and the entity bean LineItemEJB is represented by the table LINEITEM. The column OKEY represents the primary key for OrderEJB entity bean, FKEY represents the foreign key column of LINEITEM that holds the values of the ORDER primary keys. FKEY is defined in the LINEITEM table to model the one-to-many relationship.

Using this mapping, the following EJB QL finder query

```
SELECT OBJECT(o)
FROM Order o, IN(o.lineItems) l
WHERE l.quantity > 5
```

might be represented in SQL as

```
SELECT DISTINCT o.OKEY
FROM ORDERBEAN o, LINEITEM l
WHERE o.OKEY = l.FKEY AND l.QUANTITY > 5
```

11.4 EJB QL BNF

EJB QL BNF notation summary:

- { ... } grouping
- [...] optional constructs
- **boldface** keywords

The following is the complete BNF notation for EJB QL:

```
EJB QL ::= select_clause from_clause [where_clause] [orderby_clause]
from_clause ::= FROM identification_variable_declaration
               [, identification_variable_declaration]*
identification_variable_declaration ::= collection_member_declaration |
               range_variable_declaration
collection_member_declaration ::= IN (collection_valued_path_expression) [AS] identifier
range_variable_declaration ::= abstract_schema_name [AS] identifier
cmp_path_expression ::=
    {identification_variable | single_valued_cmr_path_expression}.cmp_field
single_valued_cmr_path_expression ::=
    identification_variable.[single_valued_cmr_field.]* single_valued_cmr_field
single_valued_path_expression ::=
```

```

    cmp_path_expression | single_valued_cmr_path_expression
collection_valued_path_expression ::=
    identification_variable.[single_valued_cmr_field.]*collection_valued_cmr_field
select_clause ::= SELECT [DISTINCT] {select_expression | OBJECT(identification_variable)}
select_expression ::= single_valued_path_expression | aggregate_select_expression
aggregate_select_expression ::=
    { AVG | MAX | MIN | SUM | COUNT } ([DISTINCT] cmp_path_expression) |
    COUNT ([DISTINCT] identification_variable | single_valued_cmr_path_expression)
where_clause ::= WHERE conditional_expression
conditional_expression ::= conditional_term | conditional_expression OR conditional_term
conditional_term ::= conditional_factor | conditional_term AND conditional_factor
conditional_factor ::= [NOT] conditional_primary
conditional_primary ::= simple_cond_expression | (conditional_expression)
simple_cond_expression ::= comparison_expression | between_expression | like_expression |
    in_expression | null_comparison_expression |
    empty_collection_comparison_expression |
    collection_member_expression
between_expression ::=
    arithmetic_expression [NOT] BETWEEN
        arithmetic_expression AND arithmetic_expression
in_expression ::=
    cmp_path_expression [NOT] IN
        ( {literal | input_parameter} [, {literal | input_parameter} ]*)
like_expression ::=
    cmp_path_expression [NOT] LIKE pattern_value [ESCAPE escape_character]
null_comparison_expression ::=
    {single_valued_path_expression | input_parameter} IS [NOT] NULL
empty_collection_comparison_expression ::=
    collection_valued_path_expression IS [NOT] EMPTY
collection_member_expression ::=
    {single_valued_cmr_path_expression | identification_variable | input_parameter}
        [NOT] MEMBER [OF] collection_valued_path_expression
comparison_expression ::=
    string_value comparison_operator string_expression |
    boolean_value { = | <> } boolean_expression |
    datetime_value comparison_operator datetime_expression |
    entity_bean_value { = | <> } entity_bean_expression |
    arithmetic_value comparison_operator arithmetic_expression
arithmetic_value ::= cmp_path_expression | functions_returning_numerics
comparison_operator ::=
    = | > | >= | < | <= | <>
arithmetic_expression ::= arithmetic_term | arithmetic_expression { + | - } arithmetic_term
arithmetic_term ::= arithmetic_factor | arithmetic_term { * | / } arithmetic_factor
arithmetic_factor ::= { + | - } arithmetic_primary
arithmetic_primary ::= cmp_path_expression | literal | (arithmetic_expression) |
    input_parameter | functions_returning_numerics
string_value ::= cmp_path_expression | functions_returning_strings
string_expression ::= string_primary | input_parameter
string_primary ::= cmp_path_expression | literal | (string_expression) |
    functions_returning_strings

```

```
datetime_value ::= cmp_path_expression
datetime_expression ::= datetime_value | input_parameter
boolean_value ::= cmp_path_expression
boolean_expression ::= cmp_path_expression | literal | input_parameter
entity_bean_value ::= single_valued_cmr_path_expression | identification_variable
entity_bean_expression ::= entity_bean_value | input_parameter
functions_returning_strings ::= CONCAT(string_expression, string_expression) |
    SUBSTRING(string_expression, arithmetic_expression, arithmetic_expression)
functions_returning_numerics ::=
    LENGTH(string_expression) |
    LOCATE(string_expression, string_expression [, arithmetic_expression]) |
    ABS(arithmetic_expression) |
    SQRT(arithmetic_expression) |
    MOD(arithmetic_expression, arithmetic_expression)
orderby_clause ::= ORDER BY orderby_item [, orderby_item]*
orderby_item ::= cmp_path_expression [ ASC | DESC ]
```


Chapter 12

Entity Bean Component Contract for Bean-Managed Persistence

The entity bean component contract for bean-managed persistence is the contract between an entity bean and its container. It defines the life cycle of the entity bean instances and the model for method delegation of the client-invoked business methods. The main goal of this contract is to ensure that a component using bean-managed persistence is portable across all compliant EJB containers.

This chapter defines the Enterprise Bean Provider's view of this contract and the Container Provider's responsibility for managing the life cycle of the enterprise bean instances. It also describes the Bean Provider's responsibilities when persistence is provided by the Bean Provider.

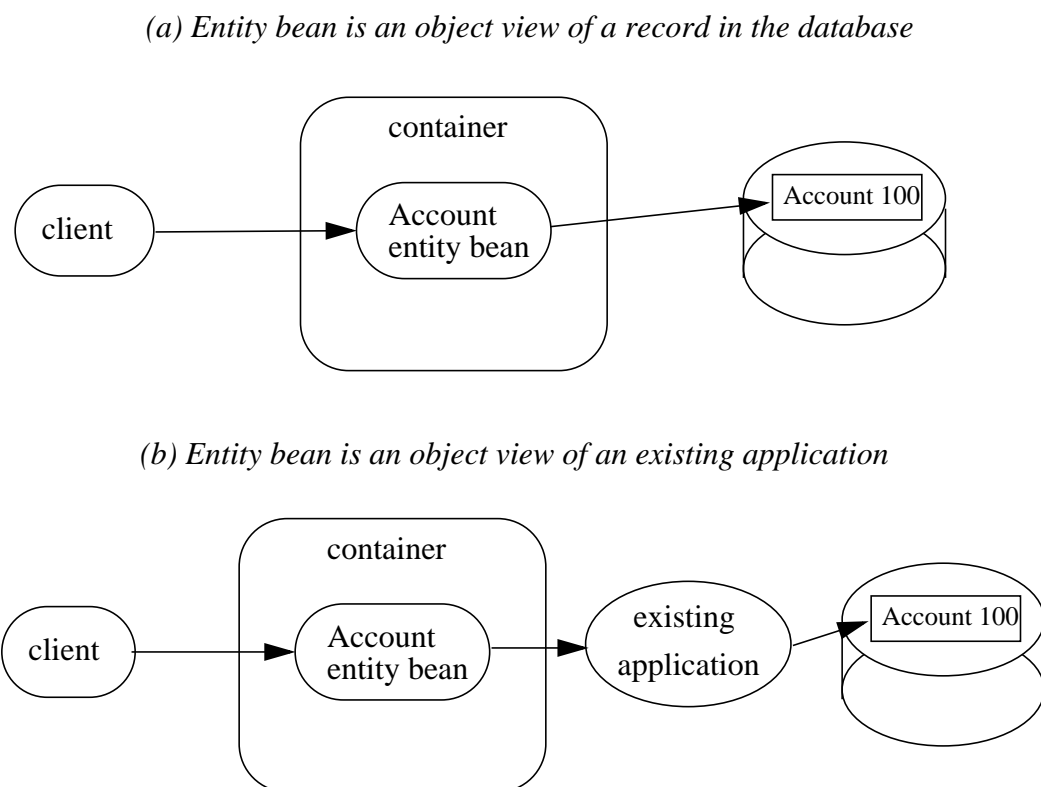
12.1 Overview of Bean-Managed Entity Persistence

An entity bean implements an object view of an entity stored in an underlying database, or an entity implemented by an existing enterprise application (for example, by a mainframe program or by an ERP application). The data access protocol for transferring the state of the entity between the entity bean instances and the underlying database is referred to as object persistence.

The entity bean component protocol for bean-managed persistence allows the entity Bean Provider to implement the entity bean's persistence directly in the entity bean class or in one or more helper classes provided with the entity bean class. This chapter describes the contracts for bean-managed persistence. container-managed persistence, which allows the entity Bean Provider to delegate the management of the entity bean's persistence to the Container Provider, is discussed in Chapter 10.

In many cases, the underlying data source may be an existing application rather than a database.

Figure 39 Client View of Underlying Data Sources Accessed Through Entity Bean



12.1.1 Entity Bean Provider's View of Persistence

Using bean-managed persistence, the entity Bean Provider writes database access calls (e.g. using JDBC™ or SQLJ) directly in the entity bean component. The data access calls are performed in the `ejbCreate<METHOD>`, `ejbRemove`, `ejbFind<METHOD>`, `ejbLoad`, and `ejbStore` methods, and/or in the business methods.

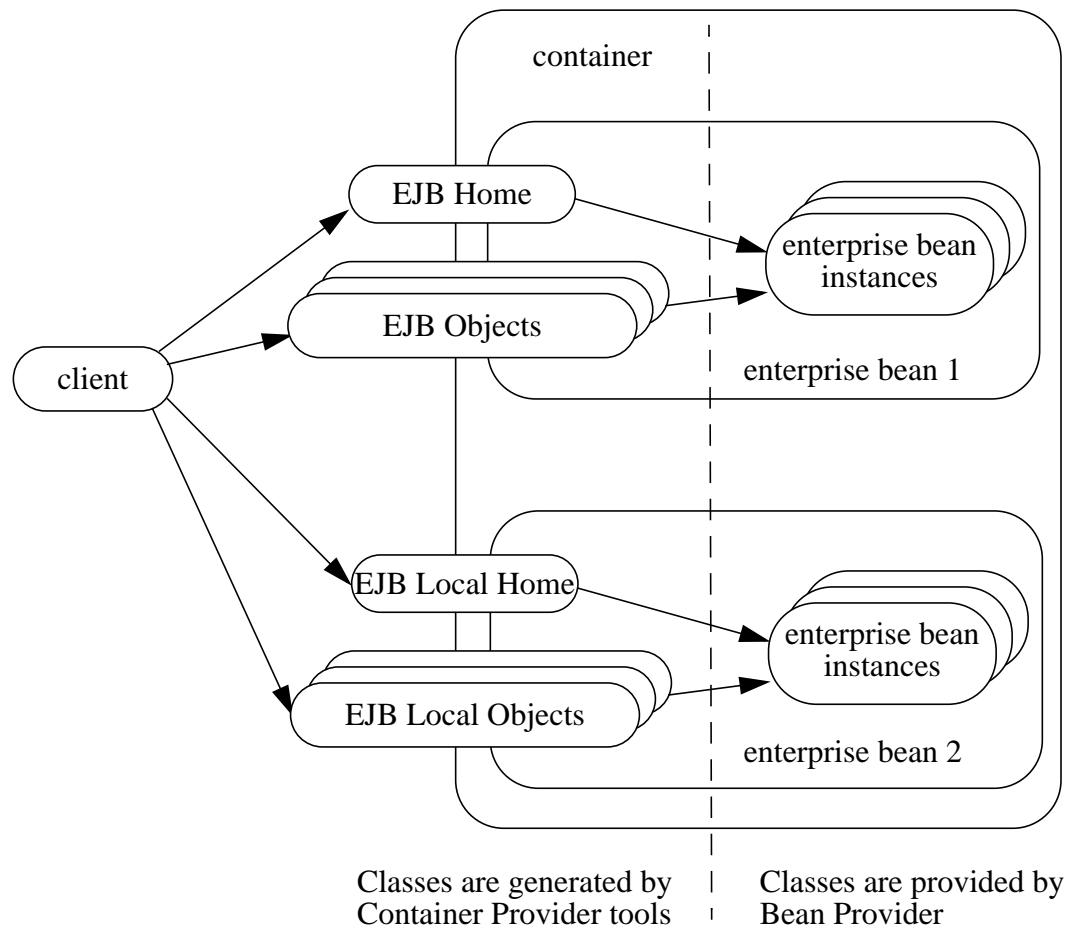
The data access calls can be coded directly into the entity bean class, or they can be encapsulated in a data access component that is part of the entity bean. Directly coding data access calls in the entity bean class may make it more difficult to adapt the entity bean to work with a database that has a different schema, or with a different type of database.

We expect that most enterprise beans with bean-managed persistence will be created by application development tools which will encapsulate data access in components. These data access components will probably not be the same for all tools. Further, if the data access calls are encapsulated in data access components, the data access components may require deployment interfaces to allow adapting data access to different schemas or even to a different database type. This EJB specification does not define the architecture for data access objects, strategies for tailoring and deploying data access components or ensuring portability of these components for bean-managed persistence.

In contrast to container-managed persistence, the entity Bean Provider does not provide a description of the bean's relationships in the deployment descriptor. With bean-managed persistence, it is the responsibility of the Bean Provider to locate related beans and to maintain relationships within the bean code.

12.1.2 Runtime Execution Model

This section describes the runtime model and the classes used in the description of the contract between an entity bean with bean-managed persistence and its container.

Figure 40 Overview of the Entity Bean Runtime Execution Model

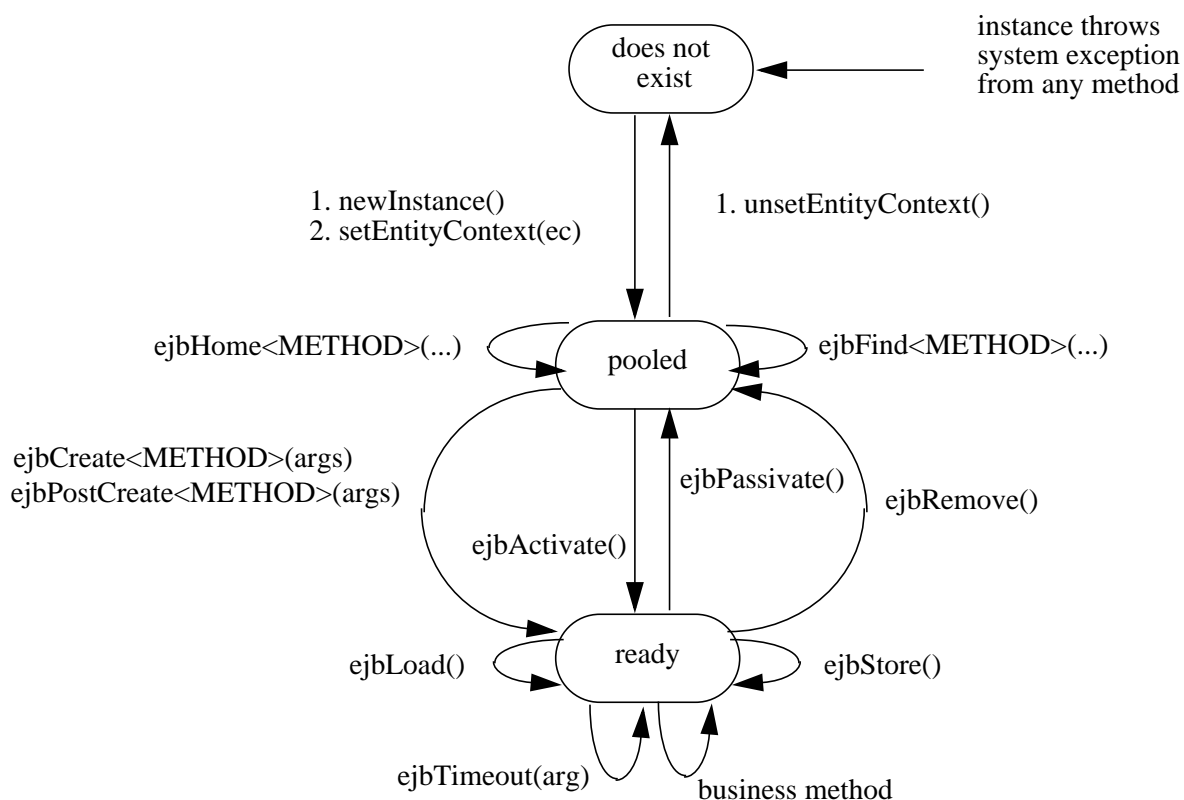
An enterprise bean instance is an object whose class is provided by the Bean Provider.

An entity **EJBObject** or **EJBLocalObject** is an object whose class is generated at deployment time by the Container Provider's tools. The entity **EJBObject** class implements the entity bean's remote interface. The entity **EJBLocalObject** class implements the entity bean's local interface. A client never references an entity bean instance directly—a client always references an entity **EJBObject** or entity **EJBLocalObject** whose class is generated by the Container Provider's tools.

An entity **EJBHome** or **EJBLocalHome** object provides the life cycle operations (create, remove, find) for its entity objects as well as home business methods, which are not specific to an entity bean instance. The class for the entity **EJBHome** or **EJBLocalHome** object is generated by the Container Provider's tools at deployment time. The entity **EJBHome** or **EJBLocalHome** object implements the entity bean's home interface that was defined by the Bean Provider.

12.1.3 Instance Life Cycle

Figure 41 Life Cycle of an Entity Bean Instance.



An entity bean instance is in one of the following three states:

- It does not exist.
- Pooled state. An instance in the pooled state is not associated with any particular entity object identity.
- Ready state. An instance in the ready state is assigned an entity object identity.

The following steps describe the life cycle of an entity bean instance:

- An entity bean instance's life starts when the container creates the instance using `newInstance`. The container then invokes the `setEntityContext` method to pass the instance a reference to the `EntityContext` interface. The `EntityContext` interface allows the

instance to invoke services provided by the container and to obtain the information about the caller of a client-invoked method.

- The instance enters the pool of available instances. Each entity bean has its own pool. While the instance is in the available pool, the instance is not associated with any particular entity object identity. All instances in the pool are considered equivalent, and therefore any instance can be assigned by the container to any entity object identity at the transition to the ready state. While the instance is in the pooled state, the container may use the instance to execute any of the entity bean's finder methods (shown as `ejbFind<METHOD>` in the diagram) or home methods (shown as `ejbHome<METHOD>` in the diagram). The instance does not move to the ready state during the execution of a finder or a home method.
- An instance transitions from the pooled state to the ready state when the container selects that instance to service a client call to an entity object or an `ejbTimeout` method. There are two possible transitions from the pooled to the ready state: through the `ejbCreate<METHOD>` and `ejbPostCreate<METHOD>` methods, or through the `ejbActivate` method. The container invokes the `ejbCreate<METHOD>` and `ejbPostCreate<METHOD>` methods when the instance is assigned to an entity object during entity object creation (i.e., when the client invokes a `create<METHOD>` method on the entity bean's home object). The container invokes the `ejbActivate` method on an instance when an instance needs to be activated to service an invocation on an existing entity object—this occurs because there is no suitable instance in the ready state to service the client's call or the `ejbTimeout` method.
- When an entity bean instance is in the ready state, the instance is associated with a specific entity object identity. While the instance is in the ready state, the container can invoke the `ejbLoad` and `ejbStore` methods zero or more times. A business method can be invoked on the instance zero or more times. The `ejbTimeout` method can be invoked on the instance zero or more times. Invocations of the `ejbLoad` and `ejbStore` methods can be arbitrarily mixed with invocations of business methods or the `ejbTimeout` method. The purpose of the `ejbLoad` and `ejbStore` methods is to synchronize the state of the instance with the state of the entity in the underlying data source—the container can invoke these methods whenever it determines a need to synchronize the instance's state.
- The container can choose to passivate an entity bean instance within a transaction. To passivate an instance, the container first invokes the `ejbStore` method to allow the instance to synchronize the database state with the instance's state, and then the container invokes the `ejbPassivate` method to return the instance to the pooled state.
- Eventually, the container will transition the instance to the pooled state. There are three possible transitions from the ready to the pooled state: through the `ejbPassivate` method, through the `ejbRemove` method, and because of a transaction rollback for `ejbCreate`, `ejbPostCreate`, or `ejbRemove` (not shown in Figure 41). The container invokes the `ejbPassivate` method when the container wants to disassociate the instance from the entity object identity without removing the entity object. The container invokes the `ejbRemove` method when the container is removing the entity object (i.e., when the client invoked the `remove` method on the entity object's component interface, or a `remove` method on the entity bean's home interface). If `ejbCreate`, `ejbPostCreate`, or `ejbRemove` is called and the transaction rolls back, the container will transition the bean instance to the pooled state.

- When the instance is put back into the pool, it is no longer associated with an entity object identity. The container can assign the instance to any entity object within the same entity bean home.
- An instance in the pool can be removed by calling the `unsetEntityContext` method on the instance.

Notes:

1. The `EntityContext` interface passed by the container to the instance in the `setEntityContext` method is an interface, not a class that contains static information. For example, the result of the `EntityContext.getPrimaryKey` method might be different each time an instance moves from the pooled state to the ready state, and the result of the `getCallerPrincipal` and `isCallerInRole` methods may be different in each business method.
2. A `RuntimeException` thrown from any method of the entity bean class (including the business methods and the callbacks invoked by the container) results in the transition to the “does not exist” state. The container must not invoke any method on the instance after a `RuntimeException` has been caught. From the client perspective, the corresponding entity object continues to exist. The client can continue accessing the entity object through its component interface because the container can use a different entity bean instance to delegate the client’s requests. Exception handling is described further in Chapter 18.
3. The container is not required to maintain a pool of instances in the pooled state. The pooling approach is an example of a possible implementation, but it is not the required implementation. Whether the container uses a pool or not has no bearing on the entity bean coding style.

12.1.4 The Entity Bean Component Contract

This section specifies the contract between an entity bean with bean-managed persistence and its container.

12.1.4.1 Entity Bean Instance’s View

The following describes the entity bean instance’s view of the contract:

The Bean Provider is responsible for implementing the following methods in the entity bean class:

- A public constructor that takes no arguments. The container uses this constructor to create instances of the entity bean class.
- `public void setEntityContext(EntityContext ic);`

A container uses this method to pass a reference to the `EntityContext` interface to the entity bean instance. If the entity bean instance needs to use the `EntityContext` interface during its lifetime, it must remember the `EntityContext` interface in an instance variable. This method executes with an unspecified transaction context (Refer to Subsection 17.6.5 for how the container executes methods with an unspecified transaction context). An identity of an entity object is not available during this method.

The instance can take advantage of the `setEntityContext` method to allocate any resources that are to be held by the instance for its lifetime. Such resources cannot be specific to an entity object identity because the instance might be reused during its lifetime to serve multiple entity object identities.

- `public void unsetEntityContext();`

A container invokes this method before terminating the life of the instance.

This method executes with an unspecified transaction context. An identity of an entity object is not available during this method.

The instance can take advantage of the `unsetEntityContext` method to free any resources that are held by the instance. (These resources typically had been allocated by the `setEntityContext` method.)

- `public PrimaryKeyClass ejbCreate<METHOD>(...);`

There are zero^[29] or more `ejbCreate<METHOD>` methods, whose signatures match the signatures of the `create<METHOD>` methods of the entity bean home interface. The container invokes an `ejbCreate<METHOD>` method on an entity bean instance when a client invokes a matching `create<METHOD>` method to create an entity object.

The implementation of the `ejbCreate<METHOD>` method typically validates the client-supplied arguments, and inserts a record representing the entity object into the database. The method also initializes the instance's variables. The `ejbCreate<METHOD>` method must return the primary key for the created entity object.

An `ejbCreate<METHOD>` method executes in the transaction context determined by the transaction attribute of the matching `create<METHOD>` method, as described in subsection 17.6.2.

- `public void ejbPostCreate<METHOD>(...);`

For each `ejbCreate<METHOD>` method, there is a matching `ejbPostCreate<METHOD>` method that has the same input parameters but whose return value is `void`. The container invokes the matching `ejbPostCreate<METHOD>` method on an instance after it invokes the `ejbCreate<METHOD>` method with the same arguments. The entity object identity is available during the `ejbPostCreate<METHOD>` method. The instance may, for example, obtain the component interface of the associated entity object and pass it to another enterprise bean as a method argument.

An `ejbPostCreate<METHOD>` method executes in the same transaction context as the previous `ejbCreate<METHOD>` method.

- `public void ejbActivate();`

The container invokes this method on the instance when the container picks the instance from the pool and assigns it to a specific entity object identity. The `ejbActivate` method gives the entity bean instance the chance to acquire additional resources that it needs while it is in the ready state.

[29] An entity Bean has no `ejbCreate<METHOD>` and `ejbPostCreate<METHOD>` methods if it does not define any `create` methods in its home interface. Such an entity bean does not allow the clients to create new entity objects. The entity bean restricts the clients to accessing entities that were created through direct database inserts.

This method executes with an unspecified transaction context. The instance can obtain the identity of the entity object via the `getPrimaryKey`, `getEJBLocalObject`, or `getEJBObject` method on the entity context. The instance can rely on the fact that the primary key and entity object identity will remain associated with the instance until the completion of `ejbPassivate` or `ejbRemove`.

Note that the instance should not use the `ejbActivate` method to read the state of the entity from the database; the instance should load its state only in the `ejbLoad` method.

- `public void ejbPassivate();`

The container invokes this method on an instance when the container decides to disassociate the instance from an entity object identity, and to put the instance back into the pool of available instances. The `ejbPassivate` method gives the instance the chance to release any resources that should not be held while the instance is in the pool. (These resources typically had been allocated during the `ejbActivate` method.)

This method executes with an unspecified transaction context. The instance can still obtain the identity of the entity object via the `getPrimaryKey`, `getEJBLocalObject`, or `getEJBObject` method of the `EntityContext` interface.

Note that an instance should not use the `ejbPassivate` method to write its state to the database; an instance should store its state only in the `ejbStore` method.

- `public void ejbRemove();`

The container invokes this method on an instance as a result of a client's invoking a `remove` method. The instance is in the ready state when `ejbRemove` is invoked and it will be entered into the pool when the method completes.

This method executes in the transaction context determined by the transaction attribute of the `remove` method that triggered the `ejbRemove` method. The instance can still obtain the identity of the entity object via the `getPrimaryKey`, `getEJBLocalObject`, or `getEJBObject` method of the `EntityContext` interface.

The container synchronizes the instance's state before it invokes the `ejbRemove` method. This means that the state of the instance variables at the beginning of the `ejbRemove` method is the same as it would be at the beginning of a business method.

An entity bean instance should use this method to remove the entity object's representation from the database.

Since the instance will be entered into the pool, the state of the instance at the end of this method must be equivalent to the state of a passivated instance. This means that the instance must release any resource that it would normally release in the `ejbPassivate` method.

- `public void ejbLoad();`

The container invokes this method on an instance in the ready state to inform the instance that it must synchronize the entity state cached in its instance variables from the entity state in the database. The instance must be prepared for the container to invoke this method at any time that the instance is in the ready state.

If the instance is caching the entity state (or parts of the entity state), the instance must not use the previously cached state in the subsequent business method. The instance may take advantage of the `ejbLoad` method, for example, to refresh the cached state by reading it from the database.

This method executes in the transaction context determined by the transaction attribute of the business method or `ejbTimeout` method that triggered the `ejbLoad` method.

- `public void ejbStore();`

The container invokes this method on an instance to inform the instance that the instance must synchronize the entity state in the database with the entity state cached in its instance variables. The instance must be prepared for the container to invoke this method at any time that the instance is in the ready state.

An instance must write any updates cached in the instance variables to the database in the `ejbStore` method.

This method executes in the same transaction context as the previous `ejbLoad` or `ejbCreate<METHOD>` method invoked on the instance. All business methods or the `ejbTimeout` method invoked between the previous `ejbLoad` or `ejbCreate<METHOD>` method and this `ejbStore` method are also invoked in the same transaction context.

- `public <primary key type or collection> ejbFind<METHOD>(...);`

The container invokes this method on the instance when the container selects the instance to execute a matching client-invoked `find<METHOD>` method. The instance is in the pooled state (i.e., it is not assigned to any particular entity object identity) when the container selects the instance to execute the `ejbFind<METHOD>` method on it, and it is returned to the pooled state when the execution of the `ejbFind<METHOD>` method completes.

The `ejbFind<METHOD>` method executes in the transaction context determined by the transaction attribute of the matching `find` method, as described in subsection 17.6.2.

The implementation of an `ejbFind<METHOD>` method typically uses the method's arguments to locate the requested entity object or a collection of entity objects in the database. The method must return a primary key or a collection of primary keys to the container (see Subsection 12.1.8).

- `public <type> ejbHome<METHOD>(...);`

The container invokes this method on any instance when the container selects the instance to execute a matching client-invoked `<METHOD>` home method. The instance is in the pooled state (i.e., it is not assigned to any particular entity object identity) when the container selects the instance to execute the `ejbHome<METHOD>` method on it, and it is returned to the pooled state when the execution of the `ejbHome<METHOD>` method completes.

The `ejbHome<METHOD>` method executes in the transaction context determined by the transaction attribute of the matching `<METHOD>` home method, as described in subsection 17.6.2.

- `public void ejbTimeout(...);`

The container invokes the `ejbTimeout` method on an instance when a timer for the instance has expired. The `ejbTimeout` method notifies the instance of the time-based event and allows the instance to execute the business logic to handle it.

The `ejbTimeout` method executes in the transaction context determined by its transaction attribute.

12.1.4.2 Container's View

This subsection describes the container's view of the state management contract. The container must call the following methods:

- `public void setEntityContext(ec);`
The container invokes this method to pass a reference to the `EntityContext` interface to the entity bean instance. The container must invoke this method after it creates the instance, and before it puts the instance into the pool of available instances.
The container invokes this method with an unspecified transaction context. At this point, the `EntityContext` is not associated with any entity object identity.
- `public void unsetEntityContext();`
The container invokes this method when the container wants to reduce the number of instances in the pool. After this method completes, the container must not reuse this instance.
The container invokes this method with an unspecified transaction context.
- `public PrimaryKeyClass ejbCreate<METHOD>(...);`
`public void ejbPostCreate<METHOD>(...);`
The container invokes these two methods during the creation of an entity object as a result of a client invoking a `create<METHOD>` method on the entity bean's home interface.
The container first invokes the `ejbCreate<METHOD>` method whose signature matches the `create<METHOD>` method invoked by the client. The `ejbCreate<METHOD>` method returns a primary key for the created entity object. The container creates an entity `EJBObject` reference and/or `EJBLocalObject` reference for the primary key. The container then invokes a matching `ejbPostCreate<METHOD>` method to allow the instance to fully initialize itself. Finally, the container returns the entity object's remote interface (i.e., a reference to the entity `EJBObject`) to the client if the client is a remote client, or the entity object's local interface (i.e., a reference to the entity `EJBLocalObject`) to the client if the client is a local client.
The container must invoke the `ejbCreate<METHOD>` and `ejbPostCreate<METHOD>` methods in the transaction context determined by the transaction attribute of the matching `create<METHOD>` method, as described in subsection 17.6.2.
- `public void ejbActivate();`
The container invokes this method on an entity bean instance at activation time (i.e., when the instance is taken from the pool and assigned to an entity object identity). The container must ensure that the primary key of the associated entity object is available to the instance if the instance invokes the `getPrimaryKey`, `getEJBLocalObject`, or `getEJBObject` method on its `EntityContext` interface.
The container invokes this method with an unspecified transaction context.
Note that instance is not yet ready for the delivery of a business method. The container must still invoke the `ejbLoad` method prior to a business method or `ejbTimeout` method invocation.
- `public void ejbPassivate();`
The container invokes this method on an entity bean instance at passivation time (i.e., when the instance is being disassociated from an entity object identity and moved into the pool). The

container must ensure that the identity of the associated entity object is still available to the instance if the instance invokes the `getPrimaryKey`, `getEJBLocalObject`, or `getEJBObject` method on its entity context.

The container invokes this method with an unspecified transaction context.

Note that if the instance state has been updated by a transaction, the container must first invoke the `ejbStore` method on the instance before it invokes `ejbPassivate` on it.

- `public void ejbRemove();`

The container invokes this method before it ends the life of an entity object as a result of a client invoking a `remove` operation.

The container invokes this method in the transaction context determined by the transaction attribute of the invoked `remove` method. The container must ensure that the identity of the associated entity object is still available to the instance in the `ejbRemove` method (i.e., the instance can invoke the `getPrimaryKey`, `getEJBLocalObject`, or `getEJBObject` method on its `EntityContext` in the `ejbRemove` method).

The container must ensure that the instance's state is synchronized from the state in the database before invoking the `ejbRemove` method (i.e., if the instance is not already synchronized from the state in the database, the container must invoke `ejbLoad` before it invokes `ejbRemove`).

- `public void ejbLoad();`

The container must invoke this method on the instance whenever it becomes necessary for the instance to synchronize its instance state from its state in the database. The exact times that the container invokes `ejbLoad` depend on the configuration of the component and the container, and are not defined by the EJB architecture. Typically, the container will call `ejbLoad` before the first business method within a transaction or before invoking the `ejbTimeout` method to ensure that the instance can refresh its cached state of the entity object from the database. After the first `ejbLoad` within a transaction, the container is not required to recognize that the state of the entity object in the database has been changed by another transaction, and it is not required to notify the instance of this change via another `ejbLoad` call.

The container must invoke this method in the transaction context determined by the transaction attribute of the business method or `ejbTimeout` method that triggered the `ejbLoad` method.

- `public void ejbStore();`

The container must invoke this method on the instance whenever it becomes necessary for the instance to synchronize its state in the database with the state of the instance's fields. This synchronization always happens at the end of a transaction. However, the container may also invoke this method when it passivates the instance in the middle of a transaction, or when it needs to transfer the most recent state of the entity object to another instance for the same entity object in the same transaction (see Subsection 17.7).

The container must invoke this method in the same transaction context as the previously invoked `ejbLoad`, `ejbCreate<METHOD>`, or `ejbTimeout` method.

- `public <primary key type or collection> ejbFind<METHOD>(...);`

The container invokes the `ejbFind<METHOD>` method on an instance when a client invokes a matching `find<METHOD>` method on the entity bean's home interface. The container must

pick an instance that is in the pooled state (i.e., the instance is not associated with any entity object identity) for the execution of the `ejbFind<METHOD>` method. If there is no instance in the pooled state, the container creates one and calls the `setEntityContext` method on the instance before dispatching the finder method.

Before invoking the `ejbFind<METHOD>` method, the container must first synchronize the state of any entity bean instances that are participating in the same transaction context as is used to execute the `ejbFind<METHOD>` by invoking the `ejbStore` method on those entity bean instances. ^[30]

After the `ejbFind<METHOD>` method completes, the instance remains in the pooled state. The container may, but is not required to, immediately activate the objects that were located by the finder using the transition through the `ejbActivate` method.

The container must invoke the `ejbFind<METHOD>` method in the transaction context determined by the transaction attribute of the matching `find` method, as described in subsection 17.6.2.

If the `ejbFind<METHOD>` method is declared to return a single primary key, the container creates an entity `EJBObject` reference for the primary key and returns it to the client if the client is a remote client. If the client is a local client, the container creates and returns an entity `EJBLocalObject` reference for the primary key. If the `ejbFind<METHOD>` method is declared to return a collection of primary keys, the container creates a collection of entity `EJBObject` or `EJBLocalObject` references for the primary keys returned from `ejbFind<METHOD>`, and returns the collection to the client. (See Subsection 12.1.8 for information on collections.)

- `public <type> ejbHome<METHOD>(...);`

The container invokes the `ejbHome<METHOD>` method on an instance when a client invokes a matching `<METHOD>` home method on the entity bean's home interface. The container must pick an instance that is in the pooled state (i.e., the instance is not associated with any entity object identity) for the execution of the `ejbHome<METHOD>` method. If there is no instance in the pooled state, the container creates one and calls the `setEntityContext` method on the instance before dispatching the home method.

After the `ejbHome<METHOD>` method completes, the instance remains in the pooled state.

The container must invoke the `ejbHome<METHOD>` method in the transaction context determined by the transaction attribute of the matching `<METHOD>` home method, as described in subsection 17.6.2.

- `public void ejbTimeout(...);`

The container invokes the `ejbTimeout` method on the instance when a timer with which the entity has been registered expires. If there is no suitable instance in the ready state, the container must activate an instance, invoking the `ejbActivate` method and transitioning it to the ready state.

The container invokes the `ejbTimeout` method in the context of a transaction determined by its transaction attribute.

[30] The EJB specification does not require the distributed flushing of state. The container in which the `ejbFind<METHOD>` method executes is not required to propagate the flush to a different container.

12.1.5 The EntityContext Interface

A container provides the entity bean instances with an `EntityContext`, which gives the entity bean instance access to the instance's context maintained by the container. The `EntityContext` interface has the following methods:

- The `getEJBObject` method returns the entity bean's remote interface.
- The `getEJBHome` method returns the entity bean's remote home interface.
- The `getEJBLocalObject` method returns the entity bean's local interface.
- The `getEJBLocalHome` method returns the entity bean's local home interface.
- The `getCallerPrincipal` method returns the `java.security.Principal` that identifies the invoker.
- The `isCallerInRole` method tests if the entity bean instance's caller has a particular role.
- The `setRollbackOnly` method allows the instance to mark the current transaction such that the only outcome of the transaction is a rollback.
- The `getRollbackOnly` method allows the instance to test if the current transaction has been marked for rollback.
- The `getPrimaryKey` method returns the entity bean's primary key.
- The `getTimerService` method returns the `javax.ejb.TimerService` interface.
- The `getUserTransaction` method returns the `javax.transaction.UserTransaction` interface. Entity bean instances must not call this method.

12.1.6 Operations Allowed in the Methods of the Entity Bean Class

Table 10 defines the methods of an entity bean class in which the enterprise bean instances can access the methods of the `javax.ejb.EntityContext` interface, the `java:comp/env` environment naming context, resource managers, the `TimerService` and `Timer` methods, and other enterprise beans.

If an entity bean instance attempts to invoke a method of the `EntityContext` interface, and the access is not allowed in Table 10, the container must throw the `java.lang.IllegalStateException`.

If an entity bean instance attempts to invoke a method of the `TimerService` or `Timer` interface and the access is not allowed in Table 10, the container must throw the `java.lang.IllegalStateException`.

If an entity bean instance attempts to access a resource manager or an enterprise bean, and the access is not allowed in Table 10, the behavior is undefined by the EJB architecture.

Table 10 Operations Allowed in the Methods of an Entity Bean

Bean method	Bean method can perform the following operations
constructor	-
setEntityContext unsetEntityContext	EntityContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> JNDI access to java:comp/env
ejbCreate	EntityContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getCallerPrincipal</i> , <i>getRollbackOnly</i> , <i>isCallerInRole</i> , <i>setRollbackOnly</i> , <i>getTimerService</i> JNDI access to java:comp/env Resource manager access Enterprise bean access
ejbPostCreate	EntityContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getCallerPrincipal</i> , <i>getRollbackOnly</i> , <i>isCallerInRole</i> , <i>setRollbackOnly</i> , <i>getEJBObject</i> , <i>getEJBLocalObject</i> , <i>getPrimaryKey</i> , <i>getTimerService</i> JNDI access to java:comp/env Resource manager access Enterprise bean access Timer service or Timer methods
ejbRemove	EntityContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getCallerPrincipal</i> , <i>getRollbackOnly</i> , <i>isCallerInRole</i> , <i>setRollbackOnly</i> , <i>getEJBObject</i> , <i>getEJBLocalObject</i> , <i>getPrimaryKey</i> , <i>getTimerService</i> JNDI access to java:comp/env Resource manager access Enterprise bean access Timer service or Timer methods
ejbFind	EntityContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getCallerPrincipal</i> , <i>getRollbackOnly</i> , <i>isCallerInRole</i> , <i>setRollbackOnly</i> JNDI access to java:comp/env Resource manager access Enterprise bean access
ejbHome	EntityContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getCallerPrincipal</i> , <i>getRollbackOnly</i> , <i>isCallerInRole</i> , <i>setRollbackOnly</i> , <i>getTimerService</i> JNDI access to java:comp/env Resource manager access Enterprise bean access

Table 10 Operations Allowed in the Methods of an Entity Bean

Bean method	Bean method can perform the following operations
ejbActivate ejbPassivate	EntityContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getEJBObject</i> , <i>getEJBLocalObject</i> , <i>getPrimaryKey</i> , <i>getTimerService</i> JNDI access to java:comp/env
ejbLoad ejbStore	EntityContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getCallerPrincipal</i> , <i>getRollbackOnly</i> , <i>isCallerInRole</i> , <i>setRollbackOnly</i> , <i>getEJBObject</i> , <i>getEJBLocalObject</i> , <i>getPrimaryKey</i> , <i>getTimerService</i> JNDI access to java:comp/env Resource manager access Enterprise bean access Timer service or Timer methods
business method from component inter- face	EntityContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getCallerPrincipal</i> , <i>getRollbackOnly</i> , <i>isCallerInRole</i> , <i>setRollbackOnly</i> , <i>getEJBObject</i> , <i>getEJBLocalObject</i> , <i>getPrimaryKey</i> , <i>getTimerService</i> JNDI access to java:comp/env Resource manager access Enterprise bean access Timer service or Timer methods
ejbTimeout	EntityContext methods: <i>getEJBHome</i> , <i>getEJBLocalHome</i> , <i>getCallerPrincipal</i> , <i>isCallerInRole</i> , <i>getRollbackOnly</i> , <i>setRollbackOnly</i> , <i>getEJBObject</i> , <i>getEJBLocalObject</i> , <i>getPrimaryKey</i> , <i>getTimerService</i> JNDI access to java:comp/env Resource manager access Enterprise bean access Timer service or Timer methods

Additional restrictions:

- The *getRollbackOnly* and *setRollbackOnly* methods of the *EntityContext* interface should be used only in the enterprise bean methods that execute in the context of a transaction. The container must throw the *java.lang.IllegalStateException* if the methods are invoked while the instance is not associated with a transaction.

Reasons for disallowing operations:

- Invoking the *getEJBObject*, *getEJBLocalObject*, and *getPrimaryKey* methods is disallowed in the entity bean methods in which there is no entity object identity associated with the instance.
- Invoking the *getEJBObject* and *getEJBHome* methods is disallowed if the entity bean does not define a remote client view.

- Invoking the `getEJBLocalObject` and `getEJBLocalHome` methods is disallowed if the entity bean does not define a local client view.
- Invoking the `getRollbackOnly` and `setRollbackOnly` methods is disallowed in the entity bean methods for which the container does not have a meaningful transaction context. These are the methods that have the `NotSupported`, `Never`, or `Supports` transaction attribute.
- Accessing resource managers and enterprise beans is disallowed in the entity bean methods for which the container does not have a meaningful transaction context or client security context.

12.1.7 Caching of Entity State and the `ejbLoad` and `ejbStore` Methods

An instance of an entity bean with bean-managed persistence can cache the entity object's state between business method invocations. An instance may choose to cache the entire entity object's state, part of the state, or no state at all.

The container-invoked `ejbLoad` and `ejbStore` methods assist the instance with the management of the cached entity object's state. The instance must handle the `ejbLoad` and `ejbStore` methods as follows:

- When the container invokes the `ejbStore` method on the instance, the instance must push all cached updates of the entity object's state to the underlying database. The container invokes the `ejbStore` method at the end of a transaction, and may also invoke it at other times when the instance is in the ready state. (For example the container may invoke `ejbStore` when passivating an instance in the middle of a transaction, or when transferring the instance's state to another instance to support distributed transactions in a multi-process server.)
- When the container invokes the `ejbLoad` method on the instance, the instance must discard any cached entity object's state. The instance may, but is not required to, refresh the cached state by reloading it from the underlying database.

The following examples, which are illustrative but not prescriptive, show how an instance may cache the entity object's state:

- An instance loads the entire entity object's state in the `ejbLoad` method and caches it until the container invokes the `ejbStore` method. The business methods read and write the cached entity state. The `ejbStore` method writes the updated parts of the entity object's state to the database.
- An instance loads the most frequently used part of the entity object's state in the `ejbLoad` method and caches it until the container invokes the `ejbStore` method. Additional parts of the entity object's state are loaded as needed by the business methods. The `ejbStore` method writes the updated parts of the entity object's state to the database.
- An instance does not cache any entity object's state between business methods. The business methods access and modify the entity object's state directly in the database. The `ejbLoad` and `ejbStore` methods have an empty implementation.

We expect that most entity developers will not manually code the cache management and data access calls in the entity bean class. We expect that they will rely on application development tools to provide various data access components that encapsulate data access and provide state caching.

12.1.7.1 `ejbLoad` and `ejbStore` with the `NotSupported` Transaction Attribute

The use of the `ejbLoad` and `ejbStore` methods for caching an entity object's state in the instance works well only if the container can use transaction boundaries to drive the `ejbLoad` and `ejbStore` methods. When the `NotSupported`^[31] transaction attribute is assigned to a component interface method, the corresponding enterprise bean class method executes with an unspecified transaction context (See Subsection 17.6.5). This means that the container does not have any well-defined transaction boundaries to drive the `ejbLoad` and `ejbStore` methods on the instance.

Therefore, the `ejbLoad` and `ejbStore` methods are “unreliable” for the instances that the container uses to dispatch the methods with an unspecified transaction context. The following are the only guarantees that the container provides for the instances that execute the methods with an unspecified transaction context:

- The container invokes at least one `ejbLoad` between `ejbActivate` and the first business method in the instance.
- The container invokes at least one `ejbStore` between the last business method on the instance and the `ejbPassivate` method.

Because the entity object's state accessed between the `ejbLoad` and `ejbStore` method pair is not protected by a transaction boundary for the methods that execute with an unspecified transaction context, the Bean Provider should not attempt to use the `ejbLoad` and `ejbStore` methods to control caching of the entity object's state in the instance. Typically, the implementation of the `ejbLoad` and `ejbStore` methods should be a no-op (i.e., an empty method), and each business method should access the entity object's state directly in the database.

12.1.8 Finder Method Return Type

12.1.8.1 Single-Object Finder

Some finder methods (such as `ejbFindByPrimaryKey`) are designed to return at most one entity object. For single-object finders, the result type of a `find<METHOD>` method defined in the entity bean's remote home interface is the entity bean's remote interface, and the result type of the `find<METHOD>` method defined in the entity bean's local home interface is the entity bean's local interface. The result type of the corresponding `ejbFind<METHOD>` method defined in the entity's implementation class is the entity bean's primary key type.

[31] This applies also to the `Never` and `Supports` attribute.

The following code illustrates the definition of a single-object finder on the remote home interface.

```
// Entity's home interface
public AccountHome extends javax.ejb.EJBHome {
    ...
    Account findByPrimaryKey(AccountPrimaryKey primkey)
        throws FinderException, RemoteException;
    ...
}
```

Note that a finder method defined on the local home interface, however, must not throw the RemoteException.

```
// Entity's implementation class
public AccountBean implements javax.ejb.EntityBean {
    ...
    public AccountPrimaryKey ejbFindByPrimaryKey(
        AccountPrimaryKey primkey)
        throws FinderException
    {
        ...
    }
    ...
}
```

12.1.8.2 Multi-Object Finders

Some finder methods are designed to return multiple entity objects. For multi-object finders defined in the entity bean's remote home interface, the result type of the `find<METHOD>` method is a collection of objects implementing the entity bean's remote interface. For multi-object finders defined in the entity bean's local home interface, the result type is a collection of objects implementing the entity bean's local interface. In either case, the result type of the corresponding `ejbFind<METHOD>` implementation method defined in the entity bean's implementation class is a collection of objects of the entity bean's primary key type.

The Bean Provider can choose two types to define a collection type for a finder:

- the Java™ 2 `java.util.Collection` interface
- the JDK™ 1.1 `java.util.Enumeration` interface

A Bean Provider targeting containers and clients based on Java 2 should use the `java.util.Collection` interface for the finder's result type.

A Bean Provider who wants to ensure that the entity bean is compatible with containers and clients based on JDK 1.1 must use the `java.util.Enumeration` interface for the finder's result type^[32].

[32] The finder will be also compatible with Java 2-based containers and clients.

The Bean Provider must ensure that the objects in the `java.util.Enumeration` or `java.util.Collection` returned from the `ejbFind<METHOD>` method are instances of the entity bean's primary key class.

A client program must use the `PortableRemoteObject.narrow` method to convert the objects contained in the collections returned by a finder method on the entity bean's remote home interface to the entity bean's remote interface type.

The following is an example of a multi-object finder method definition that is compatible with containers and clients based on Java 2:

```
// Entity's remote home interface
public AccountHome extends javax.ejb.EJBHome {
    ...
    java.util.Collection findLargeAccounts(double limit)
        throws FinderException, RemoteException;
    ...
}

// Entity's implementation class
public AccountBean implements javax.ejb.EntityBean {
    ...
    public java.util.Collection ejbFindLargeAccounts(
        double limit) throws FinderException
    {
        ...
    }
    ...
}
```

The following is an example of a multi-object finder method definition compatible with containers and clients that are based on both JDK 1.1 and Java 2:

```
// Entity's remote home interface
public AccountHome extends javax.ejb.EJBHome {
    ...
    java.util.Enumeration findLargeAccounts(double limit)
        throws FinderException, RemoteException;
    ...
}

// Entity's implementation class
public AccountBean implements javax.ejb.EntityBean {
    ...
    public java.util.Enumeration ejbFindLargeAccounts(
        double limit) throws FinderException
    {
        ...
    }
    ...
}
```

12.1.9 Timer Notifications

An entity bean can be registered with the EJB Timer Service for time-based event notifications if it implements the `javax.ejb.TimedObject` interface. The container invokes the bean instance's `ejbTimeout` method when a timer for the bean has expired. See Chapter 22, "Timer Service".

12.1.10 Standard Application Exceptions for Entities

The EJB specification defines the following standard application exceptions:

- `javax.ejb.CreateException`
- `javax.ejb.DuplicateKeyException`
- `javax.ejb.FinderException`
- `javax.ejb.ObjectNotFoundException`
- `javax.ejb.RemoveException`

12.1.10.1 CreateException

From the client's perspective, a `CreateException` (or a subclass of `CreateException`) indicates that an application level error occurred during the `create<METHOD>` operation. If a client receives this exception, the client does not know, in general, whether the entity object was created but not fully initialized, or not created at all. Also, the client does not know whether or not the transaction has been marked for rollback. (However, the client may determine the transaction status using the `UserTransaction` interface or the `setRollbackOnly` method of the `EJBContext` interface.)

The Bean Provider throws the `CreateException` (or subclass of `CreateException`) from the `ejbCreate<METHOD>` and `ejbPostCreate<METHOD>` methods to indicate an application-level error from the create or initialization operation. Optionally, the Bean Provider may mark the transaction for rollback before throwing this exception.

The Bean Provider is encouraged to mark the transaction for rollback only if data integrity would be lost if the transaction were committed by the client. Typically, when a `CreateException` is thrown, it leaves the database in a consistent state, allowing the client to recover. For example, `ejbCreate` may throw the `CreateException` to indicate that the some of the arguments to the `create<METHOD>` method are invalid.

The container treats the `CreateException` as any other application exception. See Section 18.3.

12.1.10.2 DuplicateKeyException

The `DuplicateKeyException` is a subclass of `CreateException`. It is thrown by the `ejbCreate<METHOD>` method to indicate to the client that the entity object cannot be created because an entity object with the same key already exists. The unique key causing the violation may be the primary key, or another key defined in the underlying database.

Normally, the Bean Provider should not mark the transaction for rollback before throwing the exception.

When the client receives the `DuplicateKeyException`, the client knows that the entity was not created, and that the client's transaction has not typically been marked for rollback.

12.1.10.3 `FinderException`

From the client's perspective, a `FinderException` (or a subclass of `FinderException`) indicates that an application level error occurred during the `find` operation. Typically, the client's transaction has not been marked for rollback because of the `FinderException`.

The Bean Provider throws the `FinderException` (or subclass of `FinderException`) from the `ejbFind<METHOD>` method to indicate an application-level error in the finder method. The Bean Provider should not, typically, mark the transaction for rollback before throwing the `FinderException`.

The container treats the `FinderException` as any other application exception. See Section 18.3.

12.1.10.4 `ObjectNotFoundException`

The `ObjectNotFoundException` is a subclass of `FinderException`. It is thrown by the `ejbFind<METHOD>` method to indicate that the requested entity object does not exist.

Only single-object finders (see Subsection 12.1.8) should throw this exception. Multi-object finders must not throw this exception. Multi-object finders should return an empty collection as an indication that no matching objects were found.

12.1.10.5 `RemoveException`

From the client's perspective, a `RemoveException` (or a subclass of `RemoveException`) indicates that an application level error occurred during a `remove` operation. If a client receives this exception, the client does not know, in general, whether the entity object was removed or not. The client also does not know if the transaction has been marked for rollback. (However, the client may determine the transaction status using the `UserTransaction` interface.)

The Bean Provider throws the `RemoveException` (or subclass of `RemoveException`) from the `ejbRemove` method to indicate an application-level error from the entity object removal operation. Optionally, the Bean Provider may mark the transaction for rollback before throwing this exception.

The Bean Provider is encouraged to mark the transaction for rollback only if data integrity would be lost if the transaction were committed by the client. Typically, when a `RemoveException` is thrown, it leaves the database in a consistent state, allowing the client to recover.

The container treats the `RemoveException` as any other application exception. See Section 18.3.

12.1.11 Commit Options

The Entity Bean protocol is designed to give the container the flexibility to select the disposition of the instance state at transaction commit time. This flexibility allows the container to optimally manage the caching of entity object's state and the association of an entity object identity with the enterprise bean instances.

The container can select from the following commit-time options:

- **Option A:** The container caches a “ready” instance between transactions. The container ensures that the instance has exclusive access to the state of the object in the persistent storage. Therefore, the container does not have to synchronize the instance's state from the persistent storage at the beginning of the next transaction.
- **Option B:** The container caches a “ready” instance between transactions. In contrast to Option A, in this option the container does not ensure that the instance has exclusive access to the state of the object in the persistent storage. Therefore, the container must synchronize the instance's state from the persistent storage at the beginning of the next transaction.
- **Option C:** The container does not cache a “ready” instance between transactions. The container returns the instance to the pool of available instances after a transaction has completed.

The following table provides a summary of the commit-time options.

Table 11 Summary of Commit-Time Options

	Write instance state to database	Instance stays ready	Instance state remains valid
Option A	Yes	Yes	Yes
Option B	Yes	Yes	No
Option C	Yes	No	No

Note that the container synchronizes the instance's state with the persistent storage at transaction commit for all three options.

The selection of the commit option is transparent to the entity bean implementation—the entity bean will work correctly regardless of the commit-time option chosen by the container. The Bean Provider writes the entity bean in the same way.

The object interaction diagrams in subsection 12.4.4 illustrate the three alternative commit options in detail.

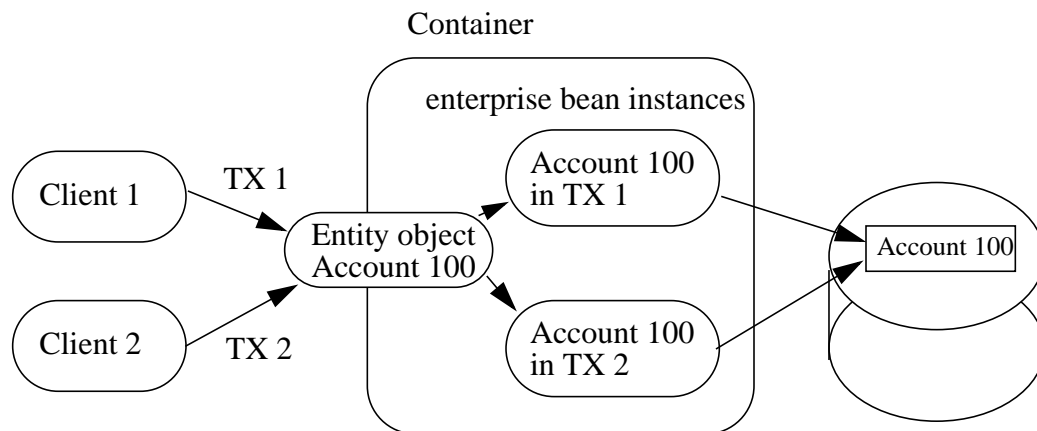
12.1.12 Concurrent Access from Multiple Transactions

When writing the entity bean business methods, the Bean Provider does not have to worry about concurrent access from multiple transactions. The Bean Provider may assume that the container will ensure appropriate synchronization for entity objects that are accessed concurrently from multiple transactions.

The container typically uses one of the following implementation strategies to achieve proper synchronization. (These strategies are illustrative, not prescriptive.)


- The container activates multiple instances of the entity bean, one for each transaction in which the entity object is being accessed. The transaction synchronization is performed automatically by the underlying database during the database access calls performed by the business methods and the `ejbTimeout` method; and by the `ejbLoad`, `ejbCreate<METHOD>`, `ejbStore`, and `ejbRemove` methods. The database system provides all the necessary transaction synchronization; the container does not have to perform any synchronization logic. The commit-time options B and C in Subsection 12.4.4 apply to this type of container.

Figure 42 Multiple Clients Can Access the Same Entity Object Using Multiple Instances



With this strategy, the type of lock acquired by `ejbLoad` leads to a trade-off. If `ejbLoad` acquires an exclusive lock on the instance's state in the database, then throughput of read-only transactions could be impacted. If `ejbLoad` acquires a shared lock and the instance is updated, then `ejbStore` will need to promote the lock to an exclusive lock. This may cause a deadlock if it happens concurrently under multiple transactions.

- The container acquires exclusive access to the entity object's state in the database. The container activates a single instance and serializes the access from multiple transactions to this instance. The commit-time option A in Subsection 12.4.4 applies to this type of container.

Figure 43 Multiple Clients Can Access the Same Entity Object Using Single Instance

12.1.13 Non-reentrant and Re-entrant Instances

An entity Bean Provider can specify that an entity bean is non-reentrant. If an instance of a non-reentrant entity bean executes a client request in a given transaction context, and another request with the same transaction context arrives for the same entity object, the container will throw an exception to the second request. This rule allows the Bean Provider to program the entity bean as single-threaded, non-reentrant code.

The functionality of some entity beans may require loopbacks in the same transaction context. An example of a loopback is when the client calls entity object A, A calls entity object B, and B calls back A in the same transaction context. The entity bean's method invoked by the loopback shares the current execution context (which includes the transaction and security contexts) with the bean's method invoked by the client.

If the entity bean is specified as non-reentrant in the deployment descriptor, the container must reject an attempt to re-enter the instance via the entity bean's component interface while the instance is executing a business method. (This can happen, for example, if the instance has invoked another enterprise bean, and the other enterprise bean tries to make a loopback call.) If the attempt is made to reenter the instance through the remote interface, the container must throw the `java.rmi.RemoteException` to the caller. If the attempt is made to reenter the instance through the local interface, the container must throw the `javax.ejb.EJBException` to the caller. The container must allow the call if the bean's deployment descriptor specifies that the entity bean is re-entrant.

Re-entrant entity beans must be programmed and used with caution. First, the Bean Provider must code the entity bean with the anticipation of a loopback call. Second, since the container cannot, in general, tell a loopback from a concurrent call from a different client, the client programmer must be careful to avoid code that could lead to a concurrent call in the same transaction context.

Concurrent calls in the same transaction context targeted at the same entity object are illegal and may lead to unpredictable results. Since the container cannot, in general, distinguish between an illegal concurrent call and a legal loopback, application programmers are encouraged to avoid using loopbacks. Entity beans that do not need callbacks should be marked as non-reentrant in the deployment descriptor, allowing the container to detect and prevent illegal concurrent calls from clients.

12.2 Responsibilities of the Enterprise Bean Provider

This section describes the responsibilities of a bean-managed persistence entity Bean Provider to ensure that the entity bean can be deployed in any EJB container.

12.2.1 Classes and Interfaces

The Bean Provider is responsible for providing the following class files:

- Entity bean class and any dependent classes
- Primary key class
- Entity bean's remote interface and remote home interface, if the entity bean provides a remote client view
- Entity bean's local interface and local home interface, if the entity bean provides a local client view

The Bean Provider must provide a remote interface and a remote home interface or a local interface and a local home interface for the bean. The Bean Provider may provide a remote interface, remote home interface, local interface, and local home interface for the bean. Other combinations are not allowed.

12.2.2 Enterprise Bean Class

The following are the requirements for an entity bean class:

The class must implement, directly or indirectly, the `javax.ejb.EntityBean` interface.

The class may implement, directly or indirectly, the `javax.ejb.TimedObject` interface.

The class must be defined as `public` and must not be `abstract`. The class must be a top level class.

The class must not be defined as `final`.

The class must define a public constructor that takes no arguments.

The class must not define the `finalize` method.

The class may, but is not required to, implement the entity bean's component interface^[33]. If the class implements the entity bean's component interface, the class must provide no-op implementations of the methods defined in the `javax.ejb.EJBObject` or `javax.ejb.EJBLocalObject` interface. The container will never invoke these methods on the bean instances at runtime.

A no-op implementation of these methods is required to avoid defining the entity bean class as abstract.

The entity bean class must implement the business methods, and the `ejbCreate<METHOD>`, `ejbPostCreate<METHOD>`, `ejbFind<METHOD>`, and `ejbHome<METHOD>` methods as described later in this section.

The entity bean class may have superclasses and/or superinterfaces. If the entity bean has superclasses, the business methods, the `ejbCreate` and `ejbPostCreate` methods, the finder methods, and the methods of the `EntityBean` interface or the `TimedObject` interface may be implemented in the enterprise bean class or in any of its superclasses.

The entity bean class is allowed to implement other methods (for example helper methods invoked internally by the business methods) in addition to the methods required by the EJB specification.

12.2.3 `ejbCreate<METHOD>` Methods

The entity bean class must implement the `ejbCreate<METHOD>` methods that correspond to the `create<METHOD>` methods specified in the entity bean's home interface.

The entity bean class may define zero or more `ejbCreate<METHOD>` methods whose signatures must follow these rules:

The method name must have `ejbCreate` as its prefix.

The method must be declared as `public`.

The method must not be declared as `final` or `static`.

The return type must be the entity bean's primary key type.

The method argument and return value types must be legal types for RMI-IIOP if the `ejbCreate<METHOD>` corresponds to a `create<METHOD>` on the entity bean's remote home interface.

The `throws` clause may define arbitrary application specific exceptions, including the `javax.ejb.CreateException`.

[33] If the entity bean class does implement the component interface, care must be taken to avoid passing of `this` as a method argument or result. This potential error can be avoided by choosing not to implement the component interface in the entity bean class.

Compatibility Note: EJB 1.0 allowed the `ejbCreate` method to throw the `java.rmi.RemoteException` to indicate a non-application exception. This practice was deprecated in EJB 1.1—an EJB 1.1 or EJB 2.0 or later compliant enterprise bean should throw the `javax.ejb.EJBException` or another `java.lang.RuntimeException` to indicate non-application exceptions to the container (see Section 18.2.2). An EJB 2.0 or later enterprise bean should not throw the `java.rmi.RemoteException` from the `ejbCreate` method.

The entity object created by the `ejbCreate<METHOD>` method must have a unique primary key. This means that the primary key must be different from the primary keys of all the existing entity objects within the same home. The `ejbCreate<METHOD>` method should throw the `DuplicateKeyException` on an attempt to create an entity object with a duplicate primary key. However, it is legal to reuse the primary key of a previously removed entity object.

12.2.4 `ejbPostCreate<METHOD>` Methods

For each `ejbCreate<METHOD>` method, the entity bean class must define a matching `ejbPostCreate<METHOD>` method, using the following rules:

The method name must have `ejbPostCreate` as its prefix.

The method must be declared as `public`.

The method must not be declared as `final` or `static`.

The return type must be `void`.

The method arguments must be the same as the arguments of the matching `ejbCreate<METHOD>` method.

The `throws` clause may define arbitrary application specific exceptions, including the `javax.ejb.CreateException`.

Compatibility Note: EJB 1.0 allowed the `ejbPostCreate` method to throw the `java.rmi.RemoteException` to indicate a non-application exception. This practice was deprecated in EJB 1.1—an EJB 1.1 or EJB 2.0 or later compliant enterprise bean should throw the `javax.ejb.EJBException` or another `java.lang.RuntimeException` to indicate non-application exceptions to the container (see Section 18.2.2). An EJB 2.0 or later enterprise bean should not throw the `java.rmi.RemoteException` from the `ejbPostCreate` method.

12.2.5 `ejbFind` Methods

The entity bean class may also define additional `ejbFind<METHOD>` finder methods.

The signatures of the finder methods must follow the following rules:

A finder method name must start with the prefix “`ejbFind`” (e.g. `ejbFindByPrimaryKey`, `ejbFindLargeAccounts`, `ejbFindLateShipments`).

A finder method must be declared as `public`.

The method must not be declared as `final` or `static`.

The method argument types must be legal types for RMI-IIOP if the `ejbFind<METHOD>` method corresponds to a `find<METHOD>` method on the entity bean's remote home interface.

The return type of a finder method must be the entity bean's primary key type, or a collection of primary keys (see Subsection 12.1.8).

The `throws` clause may define arbitrary application specific exceptions, including the `javax.ejb.FinderException`.

Compatibility Note: EJB 1.0 allowed the finder methods to throw the `java.rmi.RemoteException` to indicate a non-application exception. This practice was deprecated in EJB 1.1—an EJB 1.1 or EJB 2.0 or later compliant enterprise bean should throw the `javax.ejb.EJBException` or another `java.lang.RuntimeException` to indicate non-application exceptions to the container (see Section 18.2.2). An EJB 2.0 or later enterprise bean should not throw the `java.rmi.RemoteException` from the `ejbFind` method.

Every entity bean must define the `ejbFindByPrimaryKey` method. The result type for this method must be the primary key type (i.e., the `ejbFindByPrimaryKey` method must be a single-object finder).

12.2.6 `ejbHome<METHOD>` Methods

The entity bean class may define zero or more home methods whose signatures must follow the following rules:

An `ejbHome<METHOD>` method must exist for every home `<METHOD>` method on the entity bean's remote home or local home interface. The method name must have `ejbHome` as its prefix followed by the name of the `<METHOD>` method in which the first character has been uppercased.

The method must be declared as `public`.

The method must not be declared as `static`.

The method argument and return value types must be legal types for RMI-IIOP if the `ejbHome` method corresponds to a method on the entity bean's remote home interface.

The `throws` clause may define arbitrary application specific exceptions. The `throws` clause must not throw the `java.rmi.RemoteException`.

12.2.7 Business Methods

The entity bean class may define zero or more business methods whose signatures must follow these rules:

The method names can be arbitrary, but they must not start with ‘`ejb`’ to avoid conflicts with the call-back methods used by the EJB architecture.

The business method must be declared as `public`.

The method must not be declared as `final` or `static`.

The method argument and return value types must be legal types for RMI-IIOP if the method corresponds to a business method on the entity bean’s remote interface.

The `throws` clause may define arbitrary application specific exceptions.

Compatibility Note: EJB 1.0 allowed the business methods to throw the `java.rmi.RemoteException` to indicate a non-application exception. This practice was deprecated in EJB 1.1—an EJB 1.1 or EJB 2.0 or later compliant enterprise bean should throw the `javax.ejb.EJBException` or another `java.lang.RuntimeException` to indicate non-application exceptions to the container (see Section 18.2.2). An EJB 2.0 or later enterprise bean should not throw the `java.rmi.RemoteException` from a business method.

12.2.8 Entity Bean’s Remote Interface

The following are the requirements for the entity bean’s remote interface:

The interface must extend the `javax.ejb.EJBObject` interface.

The methods defined in the remote interface must follow the rules for RMI-IIOP. This means that their argument and return value types must be valid types for RMI-IIOP, and their `throws` clauses must include the `java.rmi.RemoteException`.

The remote interface is allowed to have superinterfaces. Use of interface inheritance is subject to the RMI-IIOP rules for the definition of remote interfaces.

For each method defined in the remote interface, there must be a matching method in the entity bean’s class. The matching method must have:

- The same name.
- The same number and types of its arguments, and the same return type.
- All the exceptions defined in the `throws` clause of the matching method of the enterprise bean class must be defined in the `throws` clause of the method of the remote interface.

The remote interface methods must not expose local interface types, local home interface types, timers or timer handles, or the managed collection classes that are used for entity beans with container-managed persistence as arguments or results.

12.2.9 Entity Bean's Remote Home Interface

The following are the requirements for the entity bean's remote home interface:

The interface must extend the `javax.ejb.EJBHome` interface.

The methods defined in this interface must follow the rules for RMI-IIOP. This means that their argument and return types must be of valid types for RMI-IIOP, and that their `throws` clauses must include the `java.rmi.RemoteException`.

The remote home interface is allowed to have superinterfaces. Use of interface inheritance is subject to the RMI-IIOP rules for the definition of remote interfaces.

Each method defined in the remote home interface must be one of the following:

- A create method.
- A finder method.
- A home method.

Each `create` method must be named "`create<METHOD>`", and it must match one of the `ejbCreate<METHOD>` methods defined in the enterprise bean class. The matching `ejbCreate<METHOD>` method must have the same number and types of its arguments. (Note that the return type is different.)

The return type for a `create<METHOD>` method must be the entity bean's remote interface type.

All the exceptions defined in the `throws` clause of the matching `ejbCreate<METHOD>` and `ejbPostCreate<METHOD>` methods of the enterprise bean class must be included in the `throws` clause of the matching `create<METHOD>` method of the remote home interface (i.e., the set of exceptions defined for the `create<METHOD>` method must be a superset of the union of exceptions defined for the `ejbCreate<METHOD>` and `ejbPostCreate<METHOD>` methods).

The `throws` clause of a `create<METHOD>` method must include the `javax.ejb.CreateException`.

Each `finder` method must be named "`find<METHOD>`" (e.g. `findLargeAccounts`), and it must match one of the `ejbFind<METHOD>` methods defined in the entity bean class (e.g. `ejbFindLargeAccounts`). The matching `ejbFind<METHOD>` method must have the same number and types of arguments. (Note that the return type may be different.)

The return type for a `find<METHOD>` method must be the entity bean's remote interface type (for a single-object finder), or a collection thereof (for a multi-object finder).

The remote home interface must always include the `findByPrimaryKey` method, which is always a single-object finder. The method must declare the primary key class as the method argument.

All the exceptions defined in the `throws` clause of an `ejbFind` method of the entity bean class must be included in the `throws` clause of the matching `find` method of the remote home interface.

The `throws` clause of a `finder` method must include the `javax.ejb.FinderException`.

Home methods can have arbitrary names, provided that they do not clash with `create`, `find`, and `remove` method names. The matching `ejbHome` method specified in the entity bean class must have the same number and types of arguments and must return the same type as the home method as specified in the remote home interface of the bean.

The remote home interface methods must not expose local interface types, local home interface types, timer handles, or the managed collection classes that are used for entity beans with container-managed persistence as arguments or results.

12.2.10 Entity Bean's Local Interface

The following are the requirements for the entity bean's local interface:

The interface must extend the `javax.ejb.EJBLocalObject` interface.

The `throws` clause of a method defined on the local interface must not include the `java.rmi.RemoteException`.

The local interface is allowed to have superinterfaces.

For each method defined in the local interface, there must be a matching method in the entity bean's class. The matching method must have:

- The same name.
- The same number and types of its arguments, and the same return type.
- All the exceptions defined in the `throws` clause of the matching method of the enterprise Bean class must be defined in the `throws` clause of the method of the local interface.

12.2.11 Entity Bean's Local Home Interface

The following are the requirements for the entity bean's local home interface:

The interface must extend the `javax.ejb.EJBLocalHome` interface.

The `throws` clause of a method on the local home interface must not include the `java.rmi.RemoteException`.

The local home interface is allowed to have superinterfaces.

Each method defined in the local home interface must be one of the following:

- A create method.
- A finder method.
- A home method.

Each `create` method must be named “`create<METHOD>`”, and it must match one of the `ejbCreate<METHOD>` methods defined in the enterprise bean class. The matching `ejbCreate<METHOD>` method must have the same number and types of its arguments. (Note that the return type is different.)

The return type for a `create<METHOD>` method must be the entity bean’s local interface type.

All the exceptions defined in the `throws` clause of the matching `ejbCreate<METHOD>` and `ejbPostCreate<METHOD>` methods of the enterprise bean class must be included in the `throws` clause of the matching `create<METHOD>` method of the local home interface (i.e., the set of exceptions defined for the `create<METHOD>` method must be a superset of the union of exceptions defined for the `ejbCreate<METHOD>` and `ejbPostCreate<METHOD>` methods).

The `throws` clause of a `create<METHOD>` method must include the `javax.ejb.CreateException`.

Each `finder` method must be named “`find<METHOD>`” (e.g. `findLargeAccounts`), and it must match one of the `ejbFind<METHOD>` methods defined in the entity bean class (e.g. `ejbFindLargeAccounts`). The matching `ejbFind<METHOD>` method must have the same number and types of arguments. (Note that the return type may be different.)

The return type for a `find<METHOD>` method must be the entity bean’s local interface type (for a single-object finder), or a collection thereof (for a multi-object finder).

The local home interface must always include the `findByPrimaryKey` method, which is always a single-object finder. The method must declare the primary key class as the method argument.

All the exceptions defined in the `throws` clause of an `ejbFind` method of the entity bean class must be included in the `throws` clause of the matching `find` method of the local home interface.

The `throws` clause of a `finder` method must include the `javax.ejb.FinderException`.

Home methods can have arbitrary names, provided that they do not clash with `create`, `find`, and `remove` method names. The matching `ejbHome` method specified in the entity bean class must have the same number and types of arguments and must return the same type as the home method as specified in the local home interface of the bean.

The `throws` clause of any method on the entity bean’s local home interface must not include the `java.rmi.RemoteException`.

12.2.12 Entity Bean’s Primary Key Class

The Bean Provider must specify a primary key class in the deployment descriptor.

The primary key type must be a legal Value Type in RMI-IIOP.

The class must provide suitable implementation of the `hashCode()` and `equals(Object other)` methods to simplify the management of the primary keys by client code.

12.3 The Responsibilities of the Container Provider

This section describes the responsibilities of the Container Provider to support bean-managed persistence entity beans. The Container Provider is responsible for providing the deployment tools, and for managing entity bean instances at runtime.

Because the EJB specification does not define the API between deployment tools and the container, we assume that the deployment tools are provided by the Container Provider. Alternatively, the deployment tools may be provided by a different vendor who uses the container vendor's specific API.

12.3.1 Generation of Implementation Classes

The deployment tools provided by the Container Provider are responsible for the generation of additional classes when the entity bean is deployed. The tools obtain the information that they need for generation of the additional classes by introspecting the classes and interfaces provided by the entity Bean Provider and by examining the entity bean's deployment descriptor.

The deployment tools must generate the following classes:

- A class that implements the entity bean's remote home interface (i.e., the entity `EJBHome` class).
- A class that implements the entity bean's remote interface (i.e., the entity `EJBObject` class).
- A class that implements the entity bean's local home interface (i.e., the entity `EJBLocalHome` class).
- A class that implements the entity bean's local interface (i.e., the entity `EJBLocalObject` class).

The deployment tools may also generate a class that mixes some container-specific code with the entity bean class. The code may, for example, help the container to manage the entity bean instances at runtime. Tools can use subclassing, delegation, and code generation.

The deployment tools may also allow generation of additional code that wraps the business methods and that is used to customize the business logic for an existing operational environment. For example, a wrapper for a `debit` function on the `Account` bean may check that the debited amount does not exceed a certain limit, or perform security checking that is specific to the operational environment.

12.3.2 Entity EJBHome Class

The entity EJBHome class, which is generated by deployment tools, implements the entity bean's remote home interface. This class implements the methods of the `javax.ejb.EJBHome` interface, and the type-specific create, finder, and home methods specific to the entity bean.

The implementation of each `create<METHOD>` method invokes a matching `ejbCreate<METHOD>` method, followed by the matching `ejbPostCreate<METHOD>` method, passing the `create<METHOD>` parameters to these matching methods.

The implementation of the `remove` methods defined in the `javax.ejb.EJBHome` interface must activate an instance (if an instance is not already in the ready state) and invoke the `ejbRemove` method on the instance.

The implementation of each `find<METHOD>` method invokes a matching `ejbFind<METHOD>` method. The implementation of the `find<METHOD>` method must create an entity object reference for the primary key returned from the `ejbFind<METHOD>` and return the entity object reference (i.e., `EJBObject`) to the client. If the `ejbFind<METHOD>` method returns a collection of primary keys, the implementation of the `find<METHOD>` method must create a collection of entity object references for the primary keys and return the collection to the client.

The implementation of each `<METHOD>` home method invokes a matching `ejbHome<METHOD>` method (in which the first character of `<METHOD>` is uppercased in the name of the `ejbHome<METHOD>` method), passing the `<METHOD>` parameters to the matching method.

12.3.3 Entity EJBObject Class

The entity EJBObject class, which is generated by deployment tools, implements the entity bean's remote interface. It implements the methods of the `javax.ejb.EJBObject` interface and the business methods specific to the entity bean.

The implementation of the `remove` method (defined in the `javax.ejb.EJBObject` interface) must activate an instance (if an instance is not already in the ready state) and invoke the `ejbRemove` method on the instance.

The implementation of each business method must activate an instance (if an instance is not already in the ready state) and invoke the matching business method on the instance.

12.3.4 Entity EJBLocalHome Class

The entity EJBLocalHome class, which is generated by deployment tools, implements the entity bean's local home interface. This class implements the methods of the `javax.ejb.EJBLocalHome` interface, and the type-specific create, finder, and home methods specific to the entity bean.

The implementation of each `create<METHOD>` method invokes a matching `ejbCreate<METHOD>` method, followed by the matching `ejbPostCreate<METHOD>` method, passing the `create<METHOD>` parameters to these matching methods.

The implementation of the `remove` method defined in the `javax.ejb.EJBLocalHome` interface must activate an instance (if an instance is not already in the ready state) and invoke the `ejbRemove` method on the instance.

The implementation of each `find<METHOD>` method invokes a matching `ejbFind<METHOD>` method. The implementation of the `find<METHOD>` method must create an entity object reference for the primary key returned from the `ejbFind<METHOD>` and return the entity object reference (i.e., `EJBLocalObject`) to the client. If the `ejbFind<METHOD>` method returns a collection of primary keys, the implementation of the `find<METHOD>` method must create a collection of entity object references for the primary keys and return the collection to the client.

The implementation of each `<METHOD>` home method invokes a matching `ejbHome<METHOD>` method (in which the first character of `<METHOD>` is uppercased in the name of the `ejbHome<METHOD>` method), passing the `<METHOD>` parameters to the matching method.

12.3.5 Entity EJBLocalObject Class

The entity `EJBLocalObject` class, which is generated by deployment tools, implements the entity bean's local interface. It implements the methods of the `javax.ejb.EJBLocalObject` interface and the business methods specific to the entity bean.

The implementation of the `remove` method (defined in the `javax.ejb.EJBLocalObject` interface) must activate an instance (if an instance is not already in the ready state) and invoke the `ejbRemove` method on the instance.

The implementation of each business method must activate an instance (if an instance is not already in the ready state) and invoke the matching business method on the instance.

12.3.6 Handle Class

The deployment tools are responsible for implementing the handle class for the entity bean. The handle class must be serializable by the Java Serialization protocol.

As the handle class is not entity-bean specific, the container may, but is not required to, use a single class for all deployed entity beans.

12.3.7 Home Handle Class

The deployment tools responsible for implementing the home handle class for the entity bean. The handle class must be serializable by the Java Serialization protocol.

Because the home handle class is not entity-bean specific, the container may, but is not required to, use a single class for the home handles of all deployed entity beans.

12.3.8 Metadata Class

The deployment tools are responsible for implementing the class that provides metadata information to the remote client view contract. The class must be a valid RMI-IIOP Value Type, and must implement the `javax.ejb.EJBMetaData` interface.

Because the metadata class is not entity-bean specific, the container may, but is not required to, use a single class for all deployed enterprise beans.

12.3.9 Instance's Re-entrance

The container runtime must enforce the rules defined in Section 12.1.13.

12.3.10 Transaction Scoping, Security, Exceptions

The container runtime must follow the rules on transaction scoping, security checking, and exception handling described in Chapters 17, 21, and 18.

12.3.11 Implementation of Object References

The container should implement the distribution protocol between the client and the container such that the object references of the remote home and remote interfaces used by entity bean clients are usable for a long period of time. Ideally, a client should be able to use an object reference across a server crash and restart. An object reference should become invalid only when the entity object has been removed, or after a reconfiguration of the server environment (for example, when the entity bean is moved to a different EJB server or container).

The motivation for this is to simplify the programming model for the entity bean client. While the client code needs to have a recovery handler for the system exceptions thrown from the individual method invocations on the home and remote interface, the client should not be forced to re-obtain the object references.

12.3.12 EntityContext

The container must implement the `EntityContext.getEJBObject` method such that the bean instance can use the Java language cast to convert the returned value to the entity bean's remote interface type. Specifically, the bean instance does not have to use the `PortableRemoteObject.narrow` method for the type conversion.

12.4 Object Interaction Diagrams

This section uses object interaction diagrams to illustrate the interactions between a bean-managed persistence entity bean instance and its container.

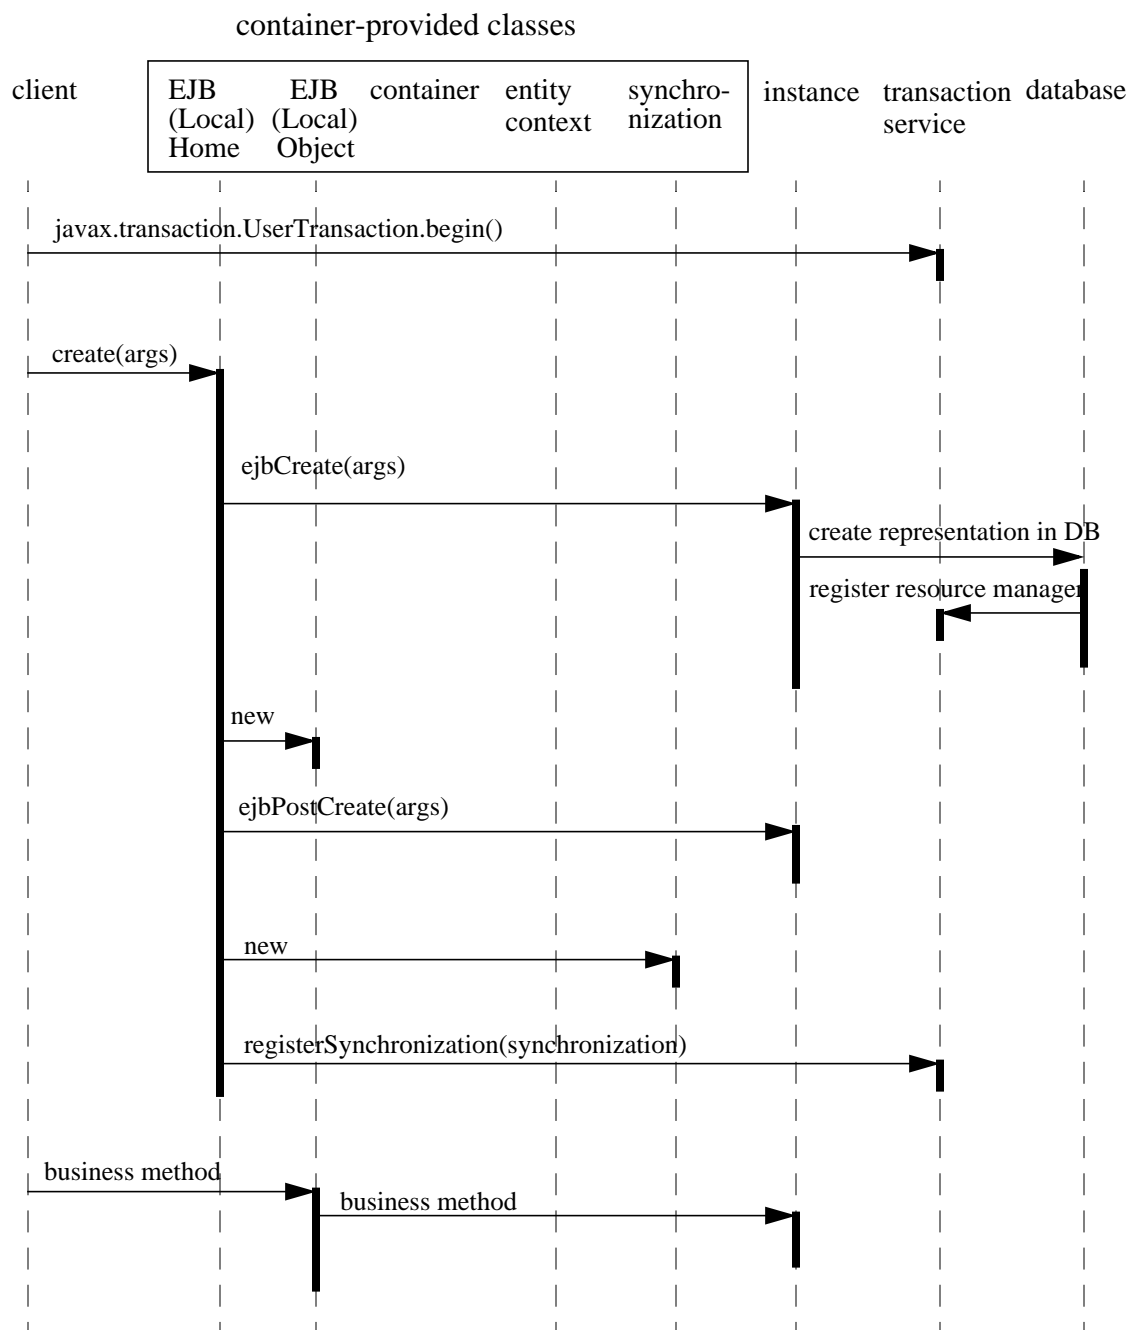
12.4.1 Notes

The object interaction diagrams illustrate a box labeled “container-provided classes.” These classes are either part of the container or are generated by the container tools. These classes communicate with each other through protocols that are container implementation-specific. Therefore, the communication between these classes is not shown in the diagrams.

The classes shown in the diagrams should be considered as an illustrative implementation rather than as a prescriptive one.

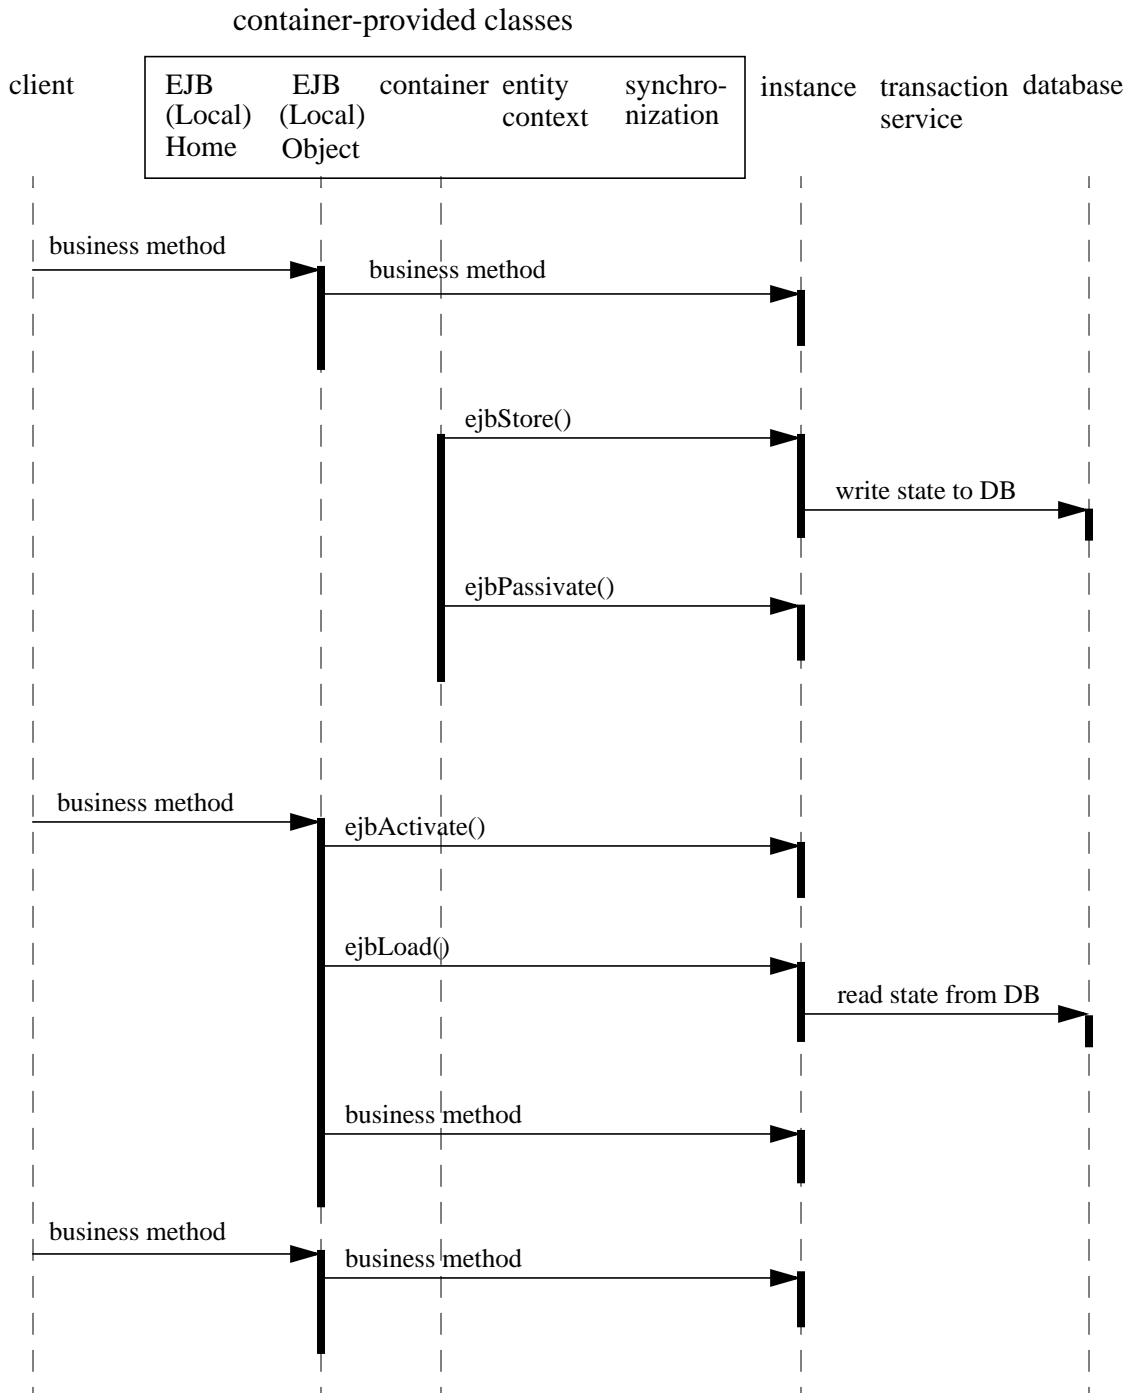
12.4.2 Creating an Entity Object

Figure 44 OID of Creation of an Entity Object with Bean-Managed Persistence



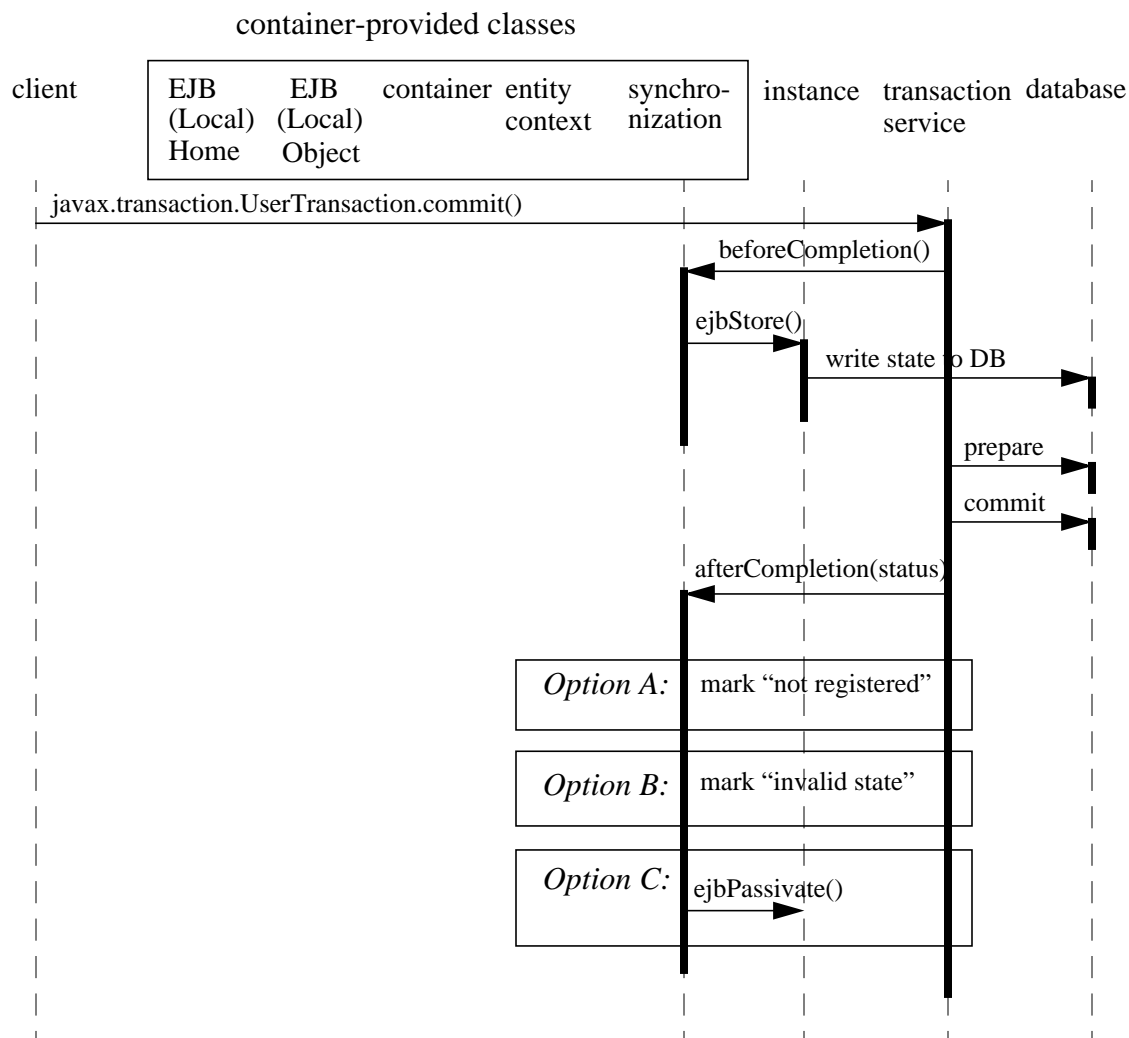
12.4.3 Passivating and Activating an Instance in a Transaction

Figure 45 OID of Passivation and Reactivation of an Entity Bean Instance with Bean-Managed Persistence



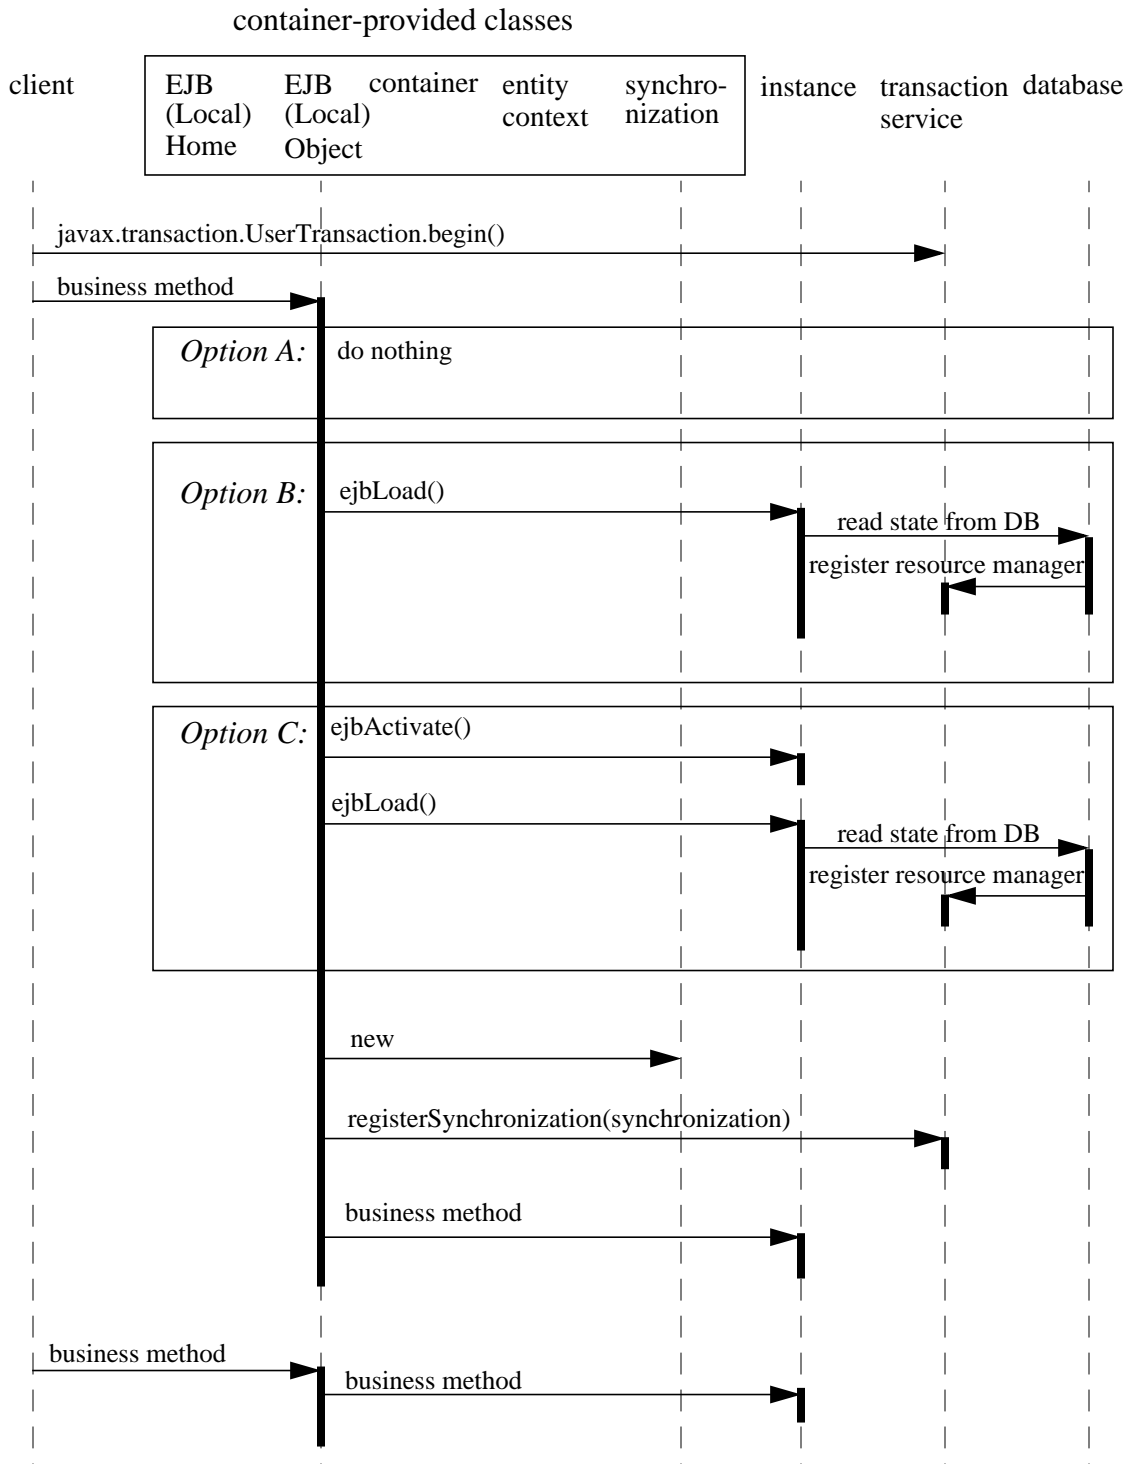
12.4.4 Committing a Transaction

Figure 46 OID of Transaction Commit Protocol for an Entity Bean Instance with Bean-Managed Persistence



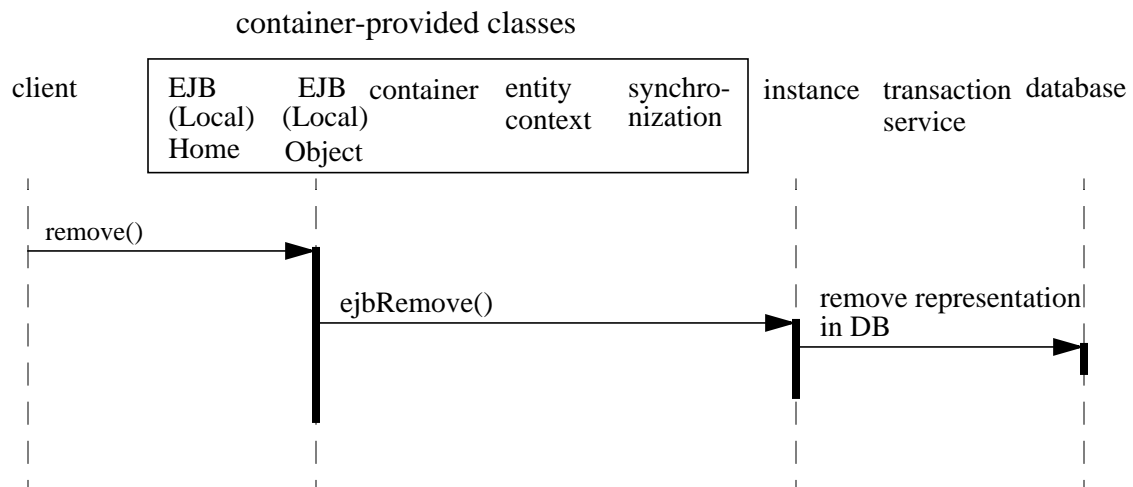
12.4.5 Starting the Next Transaction

The following diagram illustrates the protocol performed for an entity bean instance with bean-managed persistence at the beginning of a new transaction. The three options illustrated in the diagram correspond to the three commit options in the previous subsection.

Figure 47 OID of Start of Transaction for an Entity Bean Instance with Bean-Managed Persistence

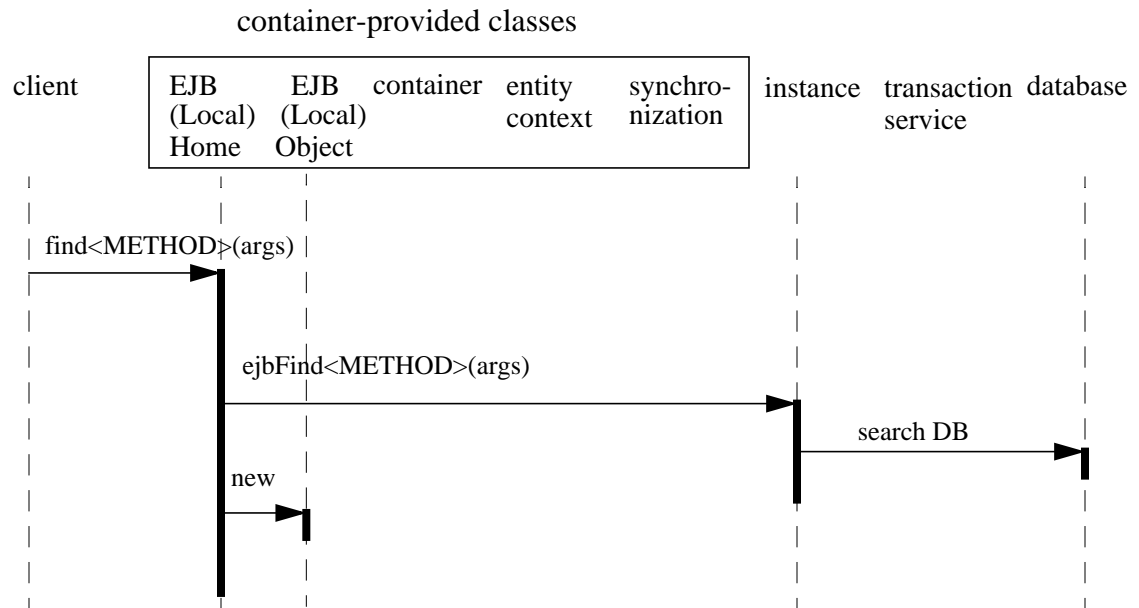
12.4.6 Removing an Entity Object

Figure 48 OID of Removal of an Entity Bean Object with Bean-Managed Persistence



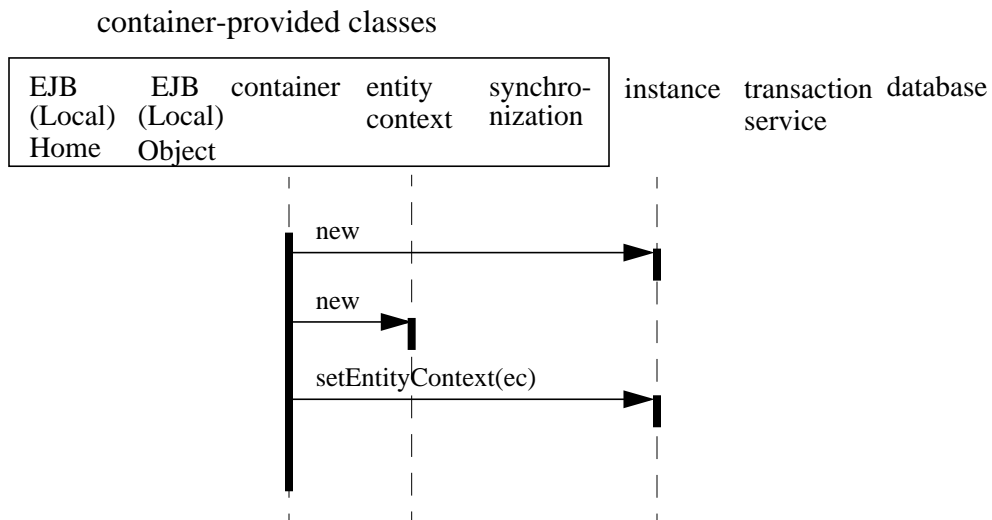
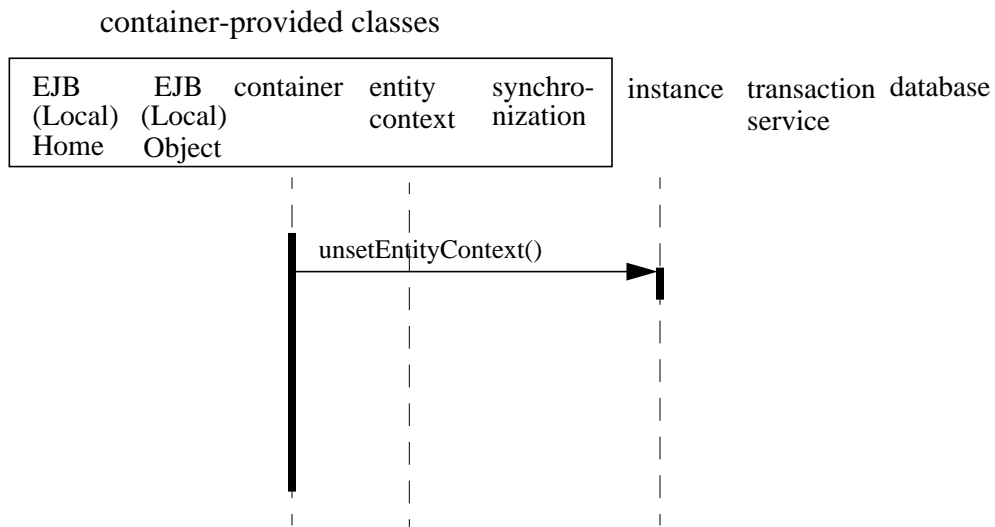
12.4.7 Finding an Entity Object

Figure 49 OID of Execution of a Finder Method on an Entity Bean Instance with Bean-Managed Persistence



12.4.8 Adding and Removing an Instance from the Pool

The diagrams in Subsections 12.4.2 through 12.4.7 did not show the sequences between the “does not exist” and “pooled” state (see the diagram in Section 12.1.3).

Figure 50 OID of a Container Adding an Instance to the Pool**Figure 51** OID of a Container Removing an Instance from the Pool

Example Bean-Managed Persistence Entity Scenario

This chapter describes an example development and deployment scenario for an entity bean using bean-managed persistence that provides a remote client view. We use the scenario to explain the responsibilities of the entity Bean Provider and those of the Container Provider.

The classes generated by the Container Provider's tools in this scenario should be considered illustrative rather than prescriptive. Container providers are free to implement the contract between an entity bean and its container in a different way that achieves an equivalent effect (from the perspectives of the entity Bean Provider and the client-side programmer).

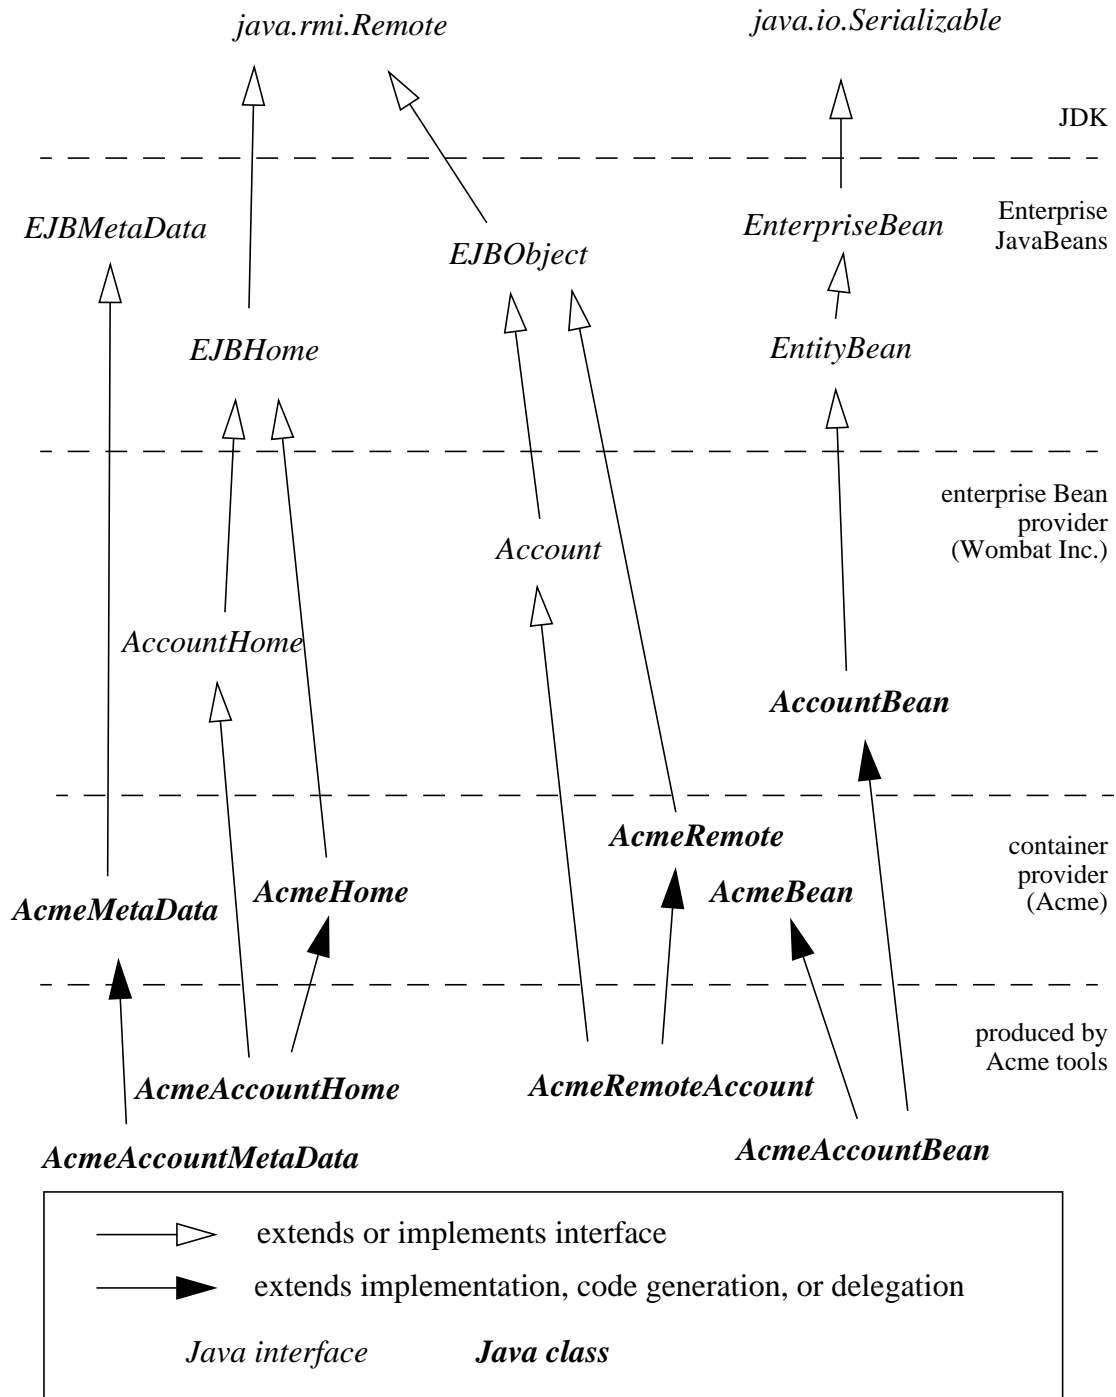
13.1 Overview

Wombat Inc. has developed the AccountBean entity bean. The AccountBean entity bean is deployed in a container provided by the Acme Corporation.

13.2 Inheritance Relationship

Figure 52

Example of the Inheritance Relationship Between the Remote Interfaces and Classes:



13.2.1 What the Entity Bean Provider is Responsible For

Wombat Inc. is responsible for providing the following:

- *Define the entity bean's remote interface (`Account`). The remote interface defines the business methods callable by a client. The remote interface must extend the `javax.ejb.EJBObject` interface, and follow the standard rules for a RMI-IIOP remote interface. The remote interface must be defined as public.*
- *Write the business logic in the entity bean class (`AccountBean`). The entity bean class may, but is not required to, implement the entity bean's remote interface (`Account`). The entity bean must implement the methods of the `javax.ejb.EntityBean` interface, the `ejbCreate<METHOD>` and `ejbPostCreate<METHOD>` methods invoked at entity object creation, and the finder methods.*
- *Define a remote home interface (`AccountHome`) for the entity bean. The remote home interface defines the entity bean's specific create and finder methods. The remote home interface must be defined as public, extend the `javax.ejb.EJBHome` interface, and follow the standard rules for RMI-IIOP remote interfaces.*
- *Define a deployment descriptor that specifies any declarative information that the entity Bean Provider wishes to pass with the entity bean to the next stage of the development/deployment workflow.*

13.2.2 Classes Supplied by Container Provider

The following classes are supplied by the Container Provider, Acme Corp:

- *The `AcmeHome` class provides the Acme implementation of the `javax.ejb.EJBHome` methods.*
- *The `AcmeRemote` class provides the Acme implementation of the `javax.ejb.EJBObject` methods.*
- *The `AcmeBean` class provides additional state and methods to allow Acme's container to manage its entity bean instances. For example, if Acme's container uses an LRU algorithm, then `AcmeBean` may include the clock count and methods to use it.*
- *The `AcmeMetaData` class provides the Acme implementation of the `javax.ejb.EJBMetaData` methods.*

13.2.3 What the Container Provider is Responsible For

The tools provided by Acme Corporation are responsible for the following:

- *Generate the entity `EJBObject` class (`AcmeRemoteAccount`) that implements the entity bean's remote interface. The tools also generate the classes that implement the communication protocol specific artifacts for the remote interface.*

- *Generate the implementation of the entity bean class suitable for the Acme container (AcmeAccountBean). AcmeAccountBean includes the business logic from the AccountBean class mixed with the services defined in the AcmeBean class. Acme tools can use inheritance, delegation, and code generation to achieve mix-in of the two classes.*
- *Generate the entity EJBHome class (AcmeAccountHome) for the entity bean, that implements the entity bean's remote home interface (AccountHome). The tools also generate the classes that implement the communication protocol specific artifacts for the remote home interface.*
- *Generate a class (AcmeAccountMetaData) that implements the javax.ejb.EJBMetaData interface for the Account Bean.*

The above classes and tools are container-specific (i.e., they reflect the way Acme Corp implemented them). Other container providers may use different mechanisms to produce their runtime classes, and the generated classes most likely will be different from those generated by Acme's tools.

EJB 1.1 Entity Bean Component Contract for Container-Managed Persistence

This chapter specifies the EJB 1.1 entity bean component contract for container-managed persistence.

While we require container providers to support backward compatibility for EJB 1.1 entity beans with container-managed persistence by the implementation of this contract, we highly recommend that Bean Providers use the Entity Bean Component Contract for Container-Managed Persistence specified in Chapter 10 for the development of new entity beans because of the more complete functionality that it provides.

14.1 EJB 1.1 Entity Beans with Container-Managed Persistence

Chapter 12, “Entity Bean Component Contract for Bean-Managed Persistence” describes the component contract for entity beans with bean-managed persistence. The contract for an EJB 1.1 entity bean with container-managed persistence is the same as the contract for an entity bean with bean-managed persistence as described in Chapter 12, except for the differences described in this chapter.

An EJB 1.1 entity bean with container-managed persistence cannot have a local interface or local home interface. Use of the local interfaces of other enterprise beans is not supported for an EJB 1.1 entity bean with container-managed persistence.

Use of the EJB Timer Service is not supported for an EJB 1.1 entity bean with container-managed persistence. An EJB 1.1 entity bean with container-managed persistence should not implement the `javax.ejb.TimedObject` interface.

14.1.1 Container-Managed Fields

An EJB 1.1 entity bean with container-managed persistence relies on the Container Provider's tools to generate methods that perform data access on behalf of the entity bean instances. The generated methods transfer data between the entity bean instance's variables and the underlying resource manager at the times defined by the EJB specification. The generated methods also implement the creation, removal, and lookup of the entity object in the underlying database.

An entity bean with container-manager persistence must not code explicit data access—all data access must be deferred to the container.

The EJB 1.1 entity Bean Provider is responsible for using the `cmp-field` elements of the deployment descriptor to declare the instance's fields that the container must load and store at the defined times. The fields must be defined in the entity bean class as `public`, and must not be defined as `transient`.

The container is responsible for transferring data between the entity bean's instance variables and the underlying data source before or after the execution of the `ejbCreate`, `ejbRemove`, `ejbLoad`, and `ejbStore` methods, as described in the following subsections. The container is also responsible for the implementation of the finder methods.

The EJB 2.0 or later deployment descriptor for an EJB 1.1 entity bean with container-managed persistence indicates that the entity bean uses container-managed persistence and that the value of its `cmp-version` element is `1.x`.

The EJB 1.1 component contract does not architect support for relationships for entity beans with container-managed persistence. The EJB 2.0 and later specifications do not support the use of the `cmr-field`, `ejb-relation`, or `query` deployment descriptor elements or their subelements for EJB 1.1 entity beans.

The following requirements ensure that an EJB 1.1 entity bean with container-managed persistence can be deployed in any compliant container.

- The Bean Provider must ensure that the Java types assigned to the container-managed fields are restricted to the following: Java primitive types, Java serializable types, and references of enterprise beans' remote or remote home interfaces.
- The Container Provider may, but is not required to, use Java Serialization to store the container-managed fields in the database. If the container chooses a different approach, the effect should be equivalent to that of Java Serialization. The container must also be capable of persisting references to enterprise beans' remote and remote home interfaces (for example, by storing their handle or primary key).

Although the above requirements allow the Bean Provider to specify almost any arbitrary type for the container-managed fields, we expect that in practice the Bean Provider of EJB 1.1 entity beans with container-managed persistence will use relatively simple Java types, and that most containers will be able to map these simple Java types to columns in a database schema to externalize the entity state in the database, rather than use Java serialization.

If the Bean Provider expects that the container-managed fields will be mapped to database fields, he or she should provide mapping instructions to the Deployer. The mapping between the instance's container-managed fields and the schema of the underlying database manager will be then realized by the data access classes generated by the Container Provider's tools. Because entity beans are typically coarse-grained objects, the content of the container-managed fields may be stored in multiple rows, possibly spread across multiple database tables. These mapping techniques are beyond the scope of the EJB specification, and do not have to be supported by an EJB compliant container. (The container may simply use the Java serialization protocol in all cases).

Because a compliant EJB container is not required to provide any support for mapping the container-managed fields to a database schema, a Bean Provider of entity beans that need a particular mapping to an underlying database schema should use bean-managed persistence or the container-managed persistence contract specified in Chapter 10 of this specification instead.

The provider of EJB 1.1 entity beans with container-managed persistence should take into account the following limitations of the EJB 1.1 container-managed persistence protocol:

- Data aliasing problems. If container-managed fields of multiple entity beans map to the same data item in the underlying database, the entity beans may see an inconsistent view of the data item if the multiple entity beans are invoked in the same transaction. (That is, an update of the data item done through a container-managed field of one entity bean may not be visible to another entity bean in the same transaction if the other entity bean maps to the same data item.)
- Eager loading of state. The container loads the entire entity object state into the container-managed fields before invoking the `ejbLoad` method. This approach may not be optimal for entity objects with large state if most business methods require access to only parts of the state.

An entity bean designer who runs into the limitations of EJB 1.1 container-managed persistence should use the container-managed persistence contracts specified in Chapter 10 of this specification instead.

14.1.2 ejbCreate, ejbPostCreate

With bean-managed persistence, the entity Bean Provider is responsible for writing the code that inserts a record into the database in the `ejbCreate` methods. However, with container-managed persistence, the container performs the database insert after the `ejbCreate` method completes.

The container must ensure that the values of the container-managed fields are set to the Java language defaults (e.g. 0 for integer, null for pointers) prior to invoking an `ejbCreate` method on an instance.

The EJB 1.1 entity Bean Provider's responsibility is to initialize the container-managed fields in the `ejbCreate` methods from the input arguments such that when an `ejbCreate` method returns, the container can extract the container-managed fields from the instance and insert them into the database.

The `ejbCreate` methods must be defined to return the primary key class type. The implementation of the `ejbCreate` methods should be coded to return a `null`. The returned value is ignored by the container.

Note: The above requirement is to allow the creation of an entity bean with bean-managed persistence by subclassing an EJB 1.1 entity bean with container-managed persistence. The Java language rules for overriding methods in subclasses requires the signatures of the `ejbCreate` methods in the subclass and the superclass be the same.

The container is responsible for creating the entity object's representation in the underlying database, extracting the primary key fields of the newly created entity object representation in the database, and for creating an entity `EJBObject` reference for the newly created entity object. The container must establish the primary key before it invokes the `ejbPostCreate` method. The container may create the representation of the entity in the database immediately after `ejbCreate` returns, or it can defer it to a later time (for example to the time after the matching `ejbPostCreate` has been called, or to the end of the transaction).

The container then invokes the matching `ejbPostCreate` method on the instance. The instance can discover the primary key by calling the `getPrimaryKey` method on its entity context object.

The container must invoke `ejbCreate`, perform the database insert operation, and invoke `ejbPostCreate` in the transaction context determined by the transaction attribute of the matching `create` method, as described in subsection 17.6.2.

The container throws the `DuplicateKeyException` if the newly created entity object would have the same primary key as one of the existing entity objects within the same home.

14.1.3 ejbRemove

The container invokes the `ejbRemove` method on an entity bean instance with container-managed persistence in response to a client-invoked `remove` operation on the entity bean's remote home or remote interface.

The entity Bean Provider can use the `ejbRemove` method to implement any actions that must be done before the entity object's representation is removed from the database.

The container synchronizes the instance's state before it invokes the `ejbRemove` method. This means that the state of the instance variables at the beginning of the `ejbRemove` method is the same as it would be at the beginning of a business method.

After `ejbRemove` returns, the container removes the entity object's representation from the database.

The container must perform `ejbRemove` and the database delete operation in the transaction context determined by the transaction attribute of the invoked `remove` method, as described in subsection 17.6.2.

14.1.4 ejbLoad

When the container needs to synchronize the state of an enterprise bean instance with the entity object's state in the database, the container reads the entity object's state from the database into the container-managed fields and then it invokes the `ejbLoad` method on the instance.

The entity Bean Provider can rely on the container's having loaded the container-managed fields from the database just before the container invokes the `ejbLoad` method. The entity bean can use the `ejbLoad` method, for instance, to perform some computation on the values of the fields that were read by the container (for example, uncompressing text fields).

14.1.5 ejbStore

When the container needs to synchronize the state of the entity object in the database with the state of the enterprise bean instance, the container first calls the `ejbStore` method on the instance, and then it extracts the container-managed fields and writes them to the database.

The entity Bean Provider should use the `ejbStore` method to set up the values of the container-managed fields just before the container writes them to the database. For example, the `ejbStore` method may perform compression of text before the text is stored in the database.

14.1.6 Finder Methods

The entity Bean Provider does not write the finder (`ejbFind<METHOD>`) methods.

The finder methods are generated at the entity bean deployment time using the Container Provider's tools. The tools can, for example, create a subclass of the entity bean class that implements the `ejbFind<METHOD>` methods, or the tools can generate the implementation of the finder methods directly in the class that implements the entity bean's remote home interface.

Note that the `ejbFind<METHOD>` names and parameter signatures of EJB 1.1 entity beans do not provide the container tools with sufficient information for automatically generating the implementation of the finder methods for methods other than `ejbFindByPrimaryKey`. Therefore, the Bean Provider is responsible for providing a description of each finder method. The entity bean Deployer uses container tools to generate the implementation of the finder methods based in the description supplied by the Bean Provider.

The EJB1.1 component contract for container-managed persistence does not specify the format of the finder method description. A Bean Provider of entity beans that needs this functionality should use the container-managed persistence contract specified in Chapter 10 of this specification instead.

14.1.7 Home Methods

The EJB1.1 entity bean contract does not support `ejbHome` methods. A Bean Provider of entity beans that need the home method functionality should use the container-managed persistence contracts specified in Chapter 10 of this specification instead.

14.1.8 Create Methods

The EJB1.1 entity bean contract does not support `create<METHOD>` methods. A Bean Provider of entity beans that needs the flexibility in method naming that `create<METHOD>` methods provide should use the container-managed persistence contracts specified in Chapter 10 of this specification instead.

14.1.9 Primary Key Type

The container must be able to manipulate the primary key type. Therefore, the primary key type for an entity bean with container-managed persistence must follow the rules in this subsection, in addition to those specified in Subsection 12.2.12.

There are two ways to specify a primary key class for an entity bean with container-managed persistence:

- Primary key that maps to a single field in the entity bean class.
- Primary key that maps to multiple fields in the entity bean class.

The second method is necessary for implementing compound keys, and the first method is convenient for single-field keys. Without the first method, simple types such as `String` would have to be wrapped in a user-defined class.

14.1.9.1 Primary Key that Maps to a Single Field in the Entity Bean Class

The Bean Provider uses the `primkey-field` element of the deployment descriptor to specify the container-managed field of the entity bean class that contains the primary key. The field's type must be the primary key type.

14.1.9.2 Primary Key that Maps to Multiple Fields in the Entity Bean Class

The primary key class must be `public`, and must have a `public` constructor with no parameters.

All fields in the primary key class must be declared as `public`.

The names of the fields in the primary key class must be a subset of the names of the container-managed fields. (This allows the container to extract the primary key fields from an instance's container-managed fields, and vice versa.)

14.1.9.3 Special Case: Unknown Primary Key Class

In special situations, the entity Bean Provider may choose not to specify the primary key class for an entity bean with container-managed persistence. This case usually happens when the entity bean does not have a natural primary key, and the Bean Provider wants to allow the Deployer to select the primary key fields at deployment time. The entity bean's primary key type will usually be derived from the primary key type used by the underlying database system that stores the entity objects. The primary key used by the database system may not be known to the Bean Provider.

When defining the primary key for the enterprise bean, the Deployer may sometimes need to subclass the entity bean class to add additional container-managed fields (this typically happens for entity beans that do not have a natural primary key, and the primary keys are system-generated by the underlying database system that stores the entity objects).

In this special case, the type of the argument of the `findByPrimaryKey` method must be declared as `java.lang.Object`, and the return value of `ejbCreate` must be declared as `java.lang.Object`. The Bean Provider must specify the primary key class in the deployment descriptor as of the type `java.lang.Object`.

The primary key class is specified at deployment time in the situations when the Bean Provider develops an entity bean that is intended to be used with multiple back-ends that provide persistence, and when these multiple back-ends require different primary key structures.

Use of entity beans with a deferred primary key type specification limits the client application programming model, because the clients written prior to deployment of the entity bean may not use, in general, the methods that rely on the knowledge of the primary key type.

The implementation of the enterprise bean class methods must be done carefully. For example, the methods should not depend on the type of the object returned from `EntityContext.getPrimaryKey`, because the return type is determined by the Deployer after the EJB class has been written.

14.2 Object Interaction Diagrams

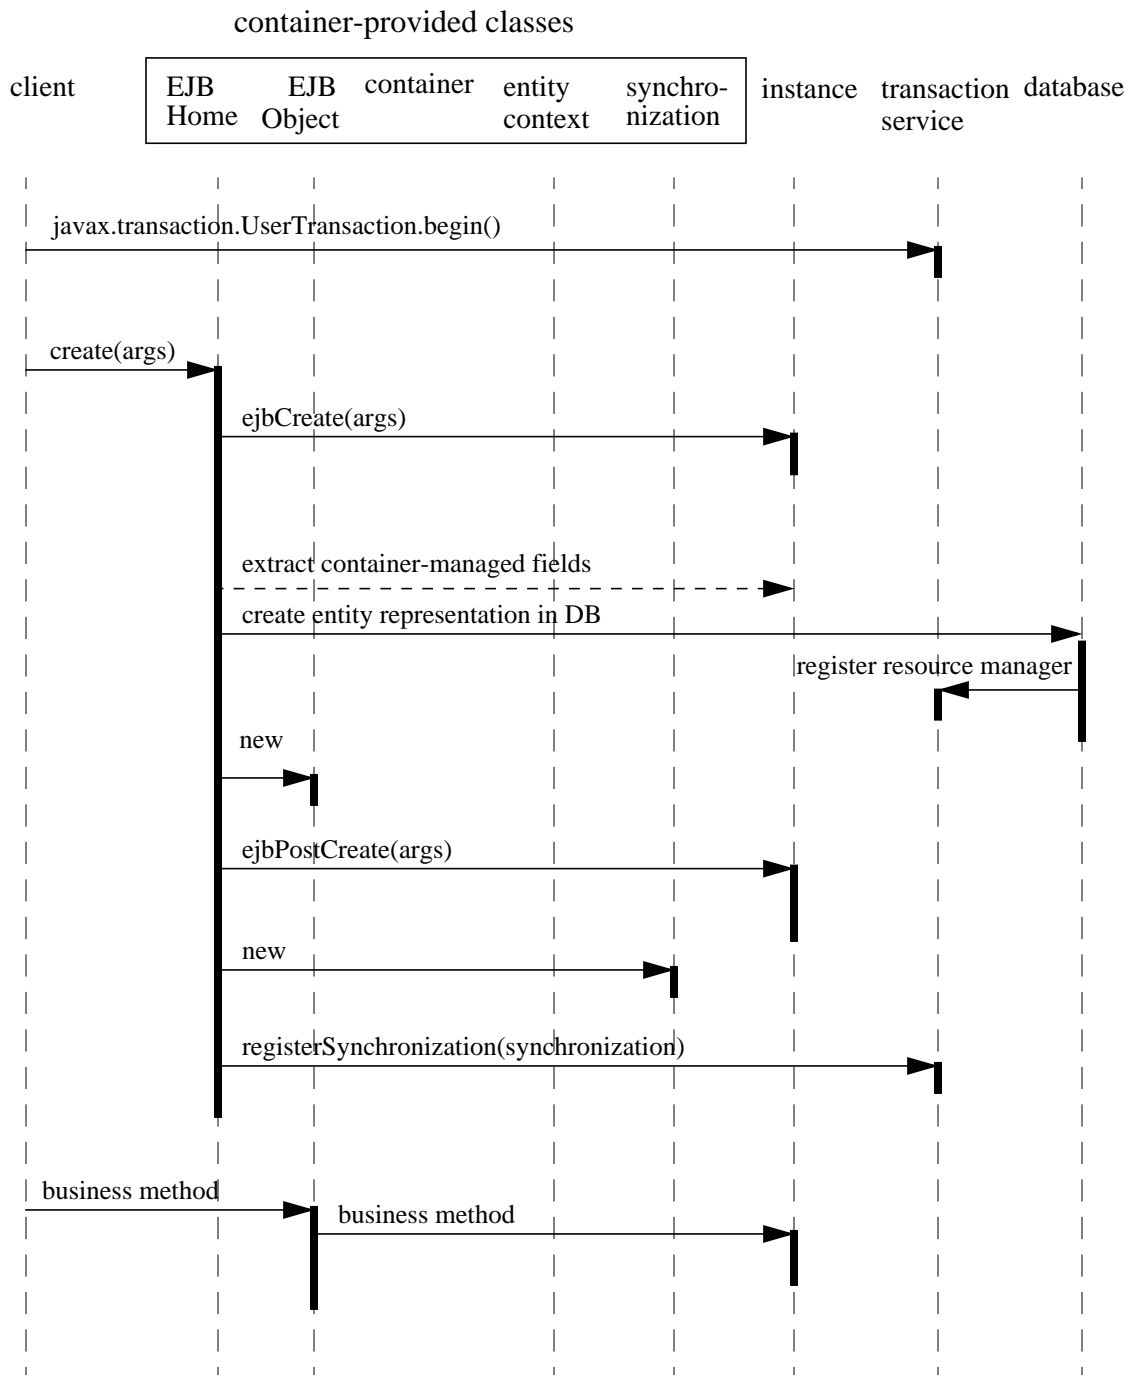
This section uses object interaction diagrams to illustrate the interactions between an EJB 1.1 entity bean instance and its container.

14.2.1 Notes

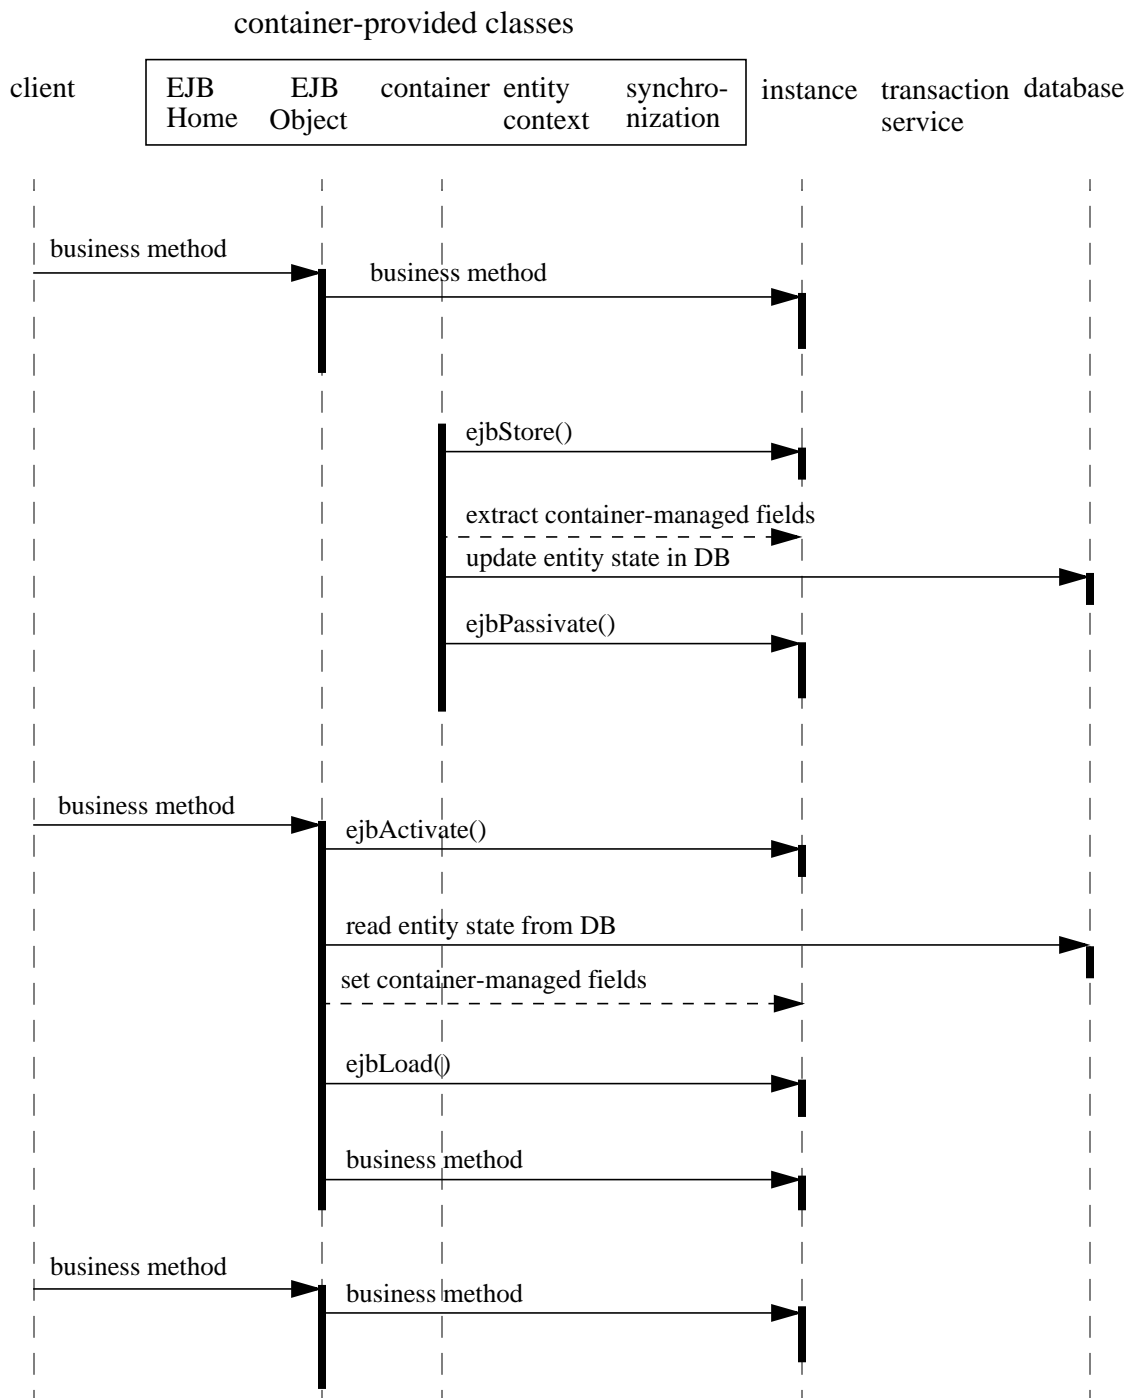
The object interaction diagrams illustrate a box labeled "container-provided classes." These classes are either part of the container or are generated by the container tools. These classes communicate with each other through protocols that are container implementation-specific. Therefore, the communication between these classes is not shown in the diagrams.

The classes shown in the diagrams should be considered as an illustrative implementation rather than as a prescriptive one

14.2.2 Creating an Entity Object


Figure 53 OID of Creation of an Entity Object with EJB 1.1 Container-Managed Persistence

14.2.3 Passivating and Activating an Instance in a Transaction

Figure 54 OID of Passivation and Reactivation of an Entity Bean Instance with EJB 1.1 CMP

14.2.4 Committing a Transaction

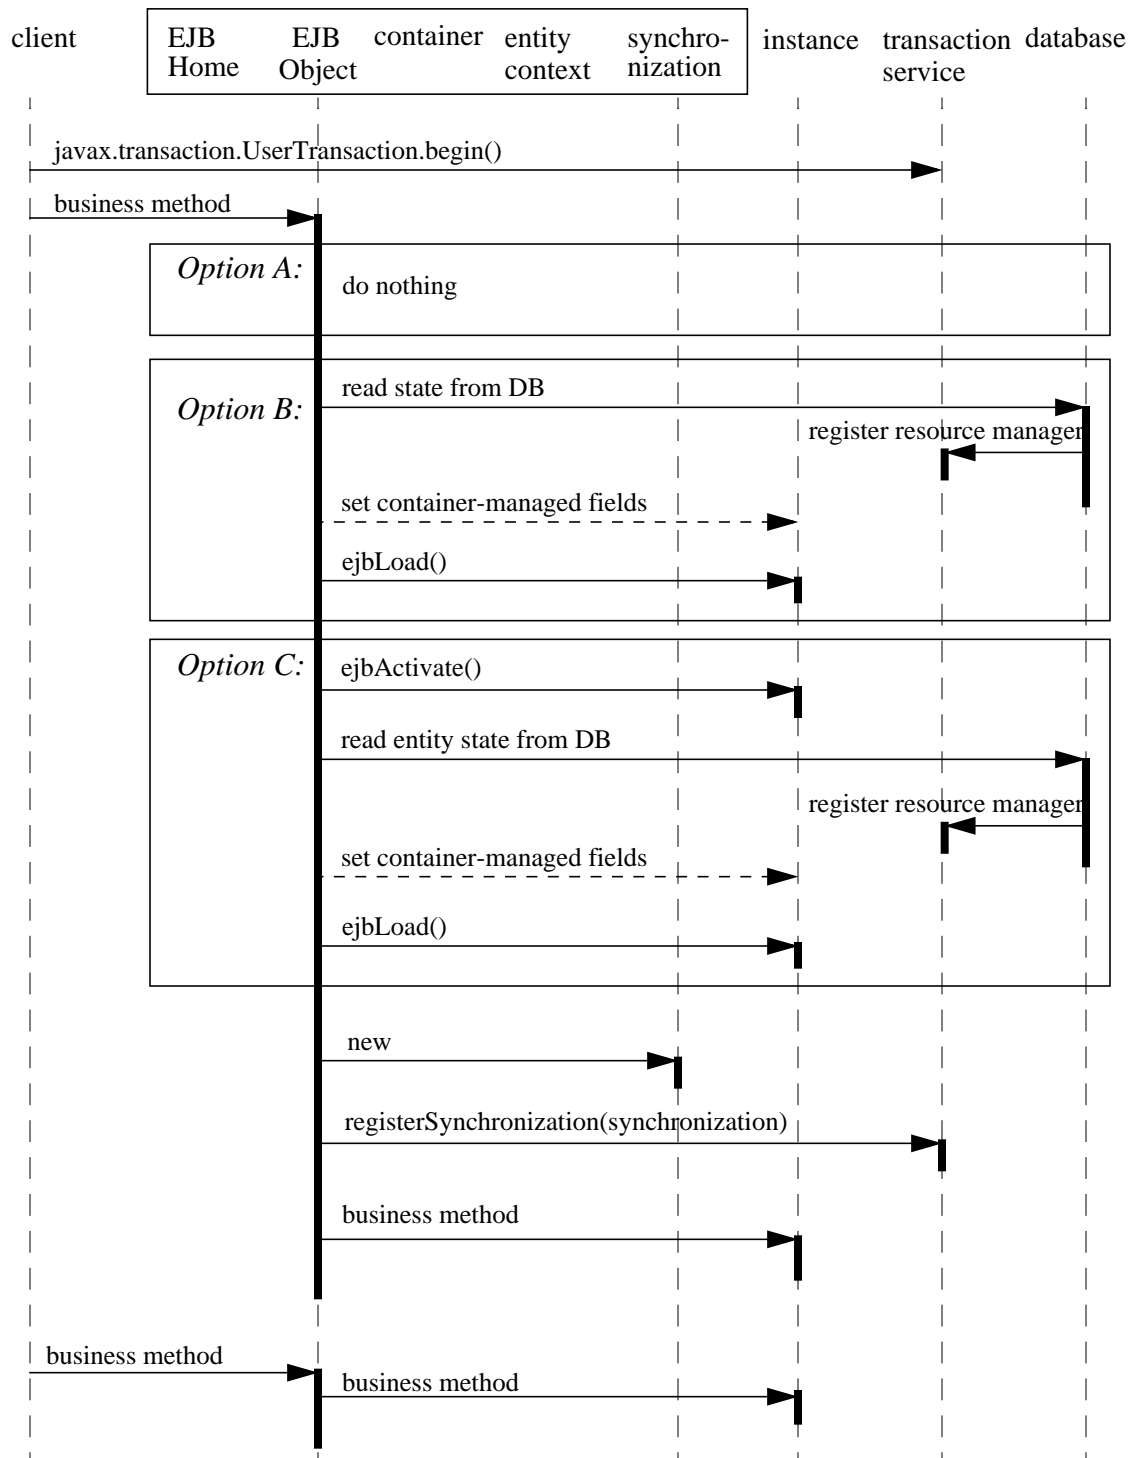
Figure 55 OID of Transaction Commit Protocol for an Entity Bean Instance with EJB 1.1 Container-Managed Persistence



14.2.5 Starting the Next Transaction

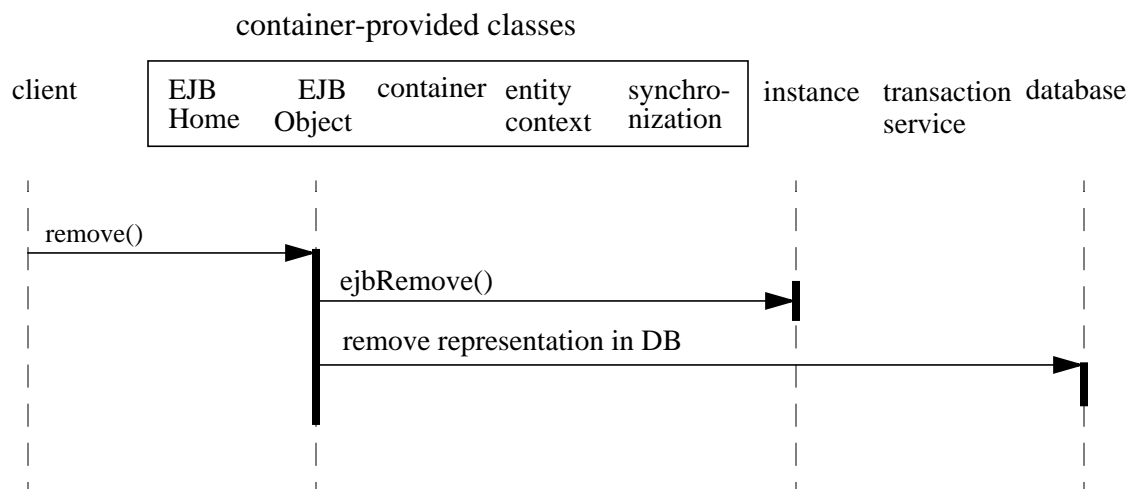
The following diagram illustrates the protocol performed for an entity bean instance with EJB 1.1 container-managed persistence at the beginning of a new transaction. The three options illustrated in the diagram correspond to the three commit options in the previous subsection.

Figure 56 OID of Start of Transaction for an Entity Bean Instance with EJB 1.1 Container-Managed Persistence container-provided classes



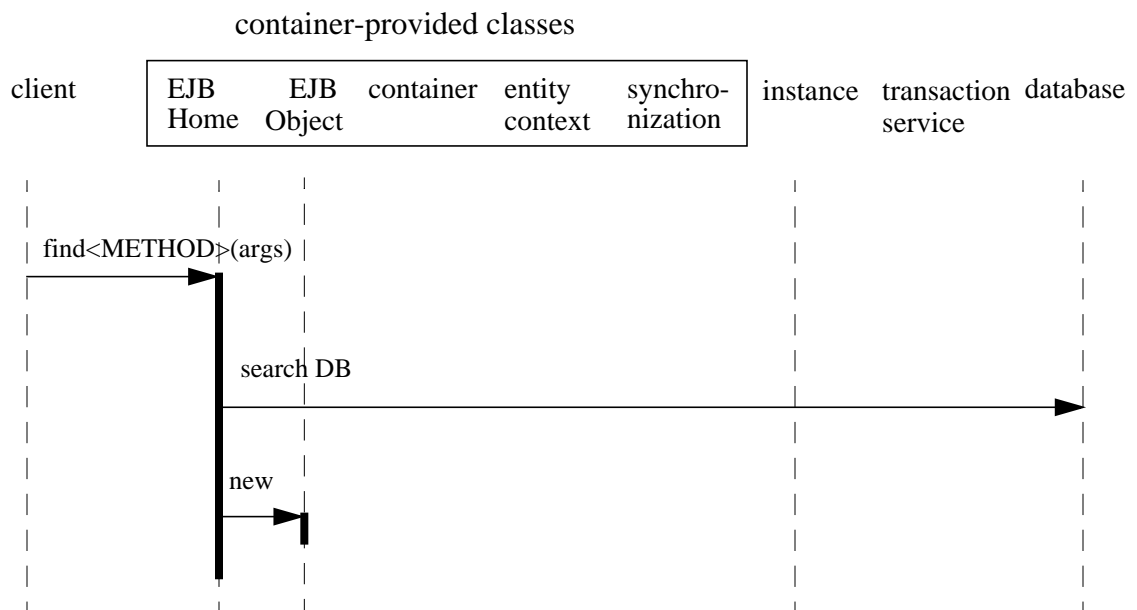
14.2.6 Removing an Entity Object

Figure 57 OID of Removal of an Entity Bean Object with EJB 1.1 Container-Managed Persistence



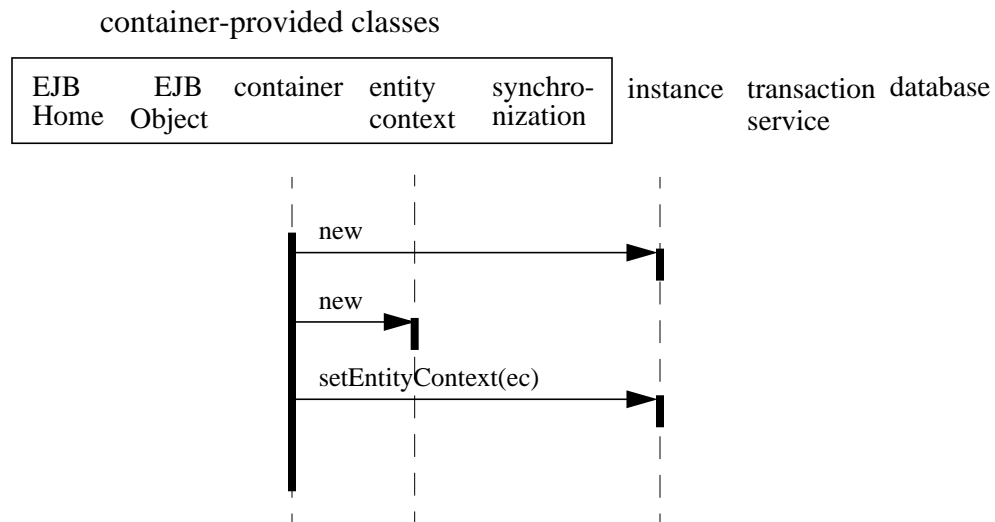
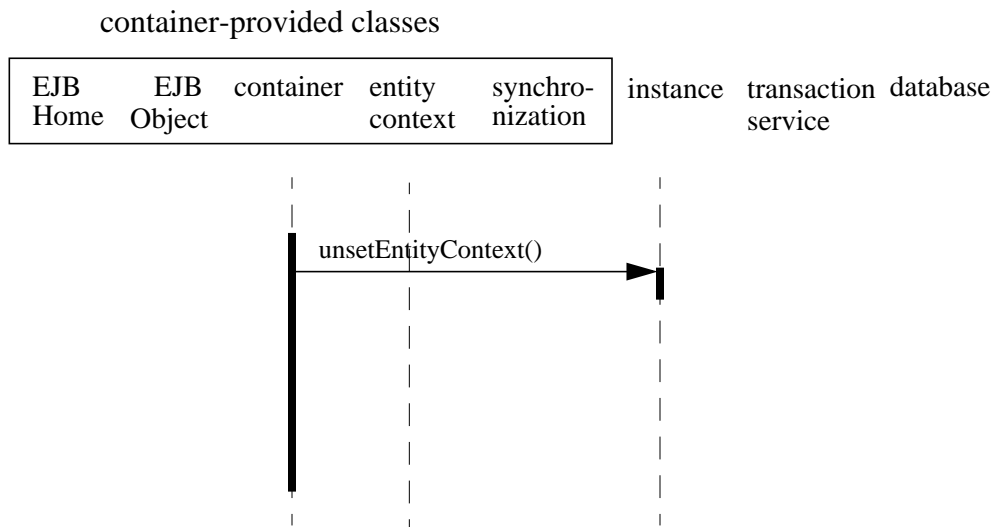
14.2.7 Finding an Entity Object

Figure 58 OID of Execution of a Finder Method on an Entity Bean Instance with EJB 1.1 Container-Managed Persistence



14.2.8 Adding and Removing an Instance from the Pool

The diagrams in Subsections 14.2.7 through 14.2.7 did not show the sequences between the “does not exist” and “pooled” state (see the diagram in Section 12.1.3).

Figure 59 OID of a Container Adding an Instance to the Pool**Figure 60** OID of a Container Removing an Instance from the Pool

Message-Driven Bean Component Contract

This chapter specifies the contract between a message-driven bean and its container. It defines the life cycle of the message-driven bean instances.

This chapter defines the developer's view of message-driven bean state management and the container's responsibility for managing it.

15.1 Overview

A message-driven bean is an asynchronous message consumer. A message-driven bean is invoked by the container as a result of the arrival of a message at the destination or endpoint that is serviced by the message-driven bean. A message-driven bean has neither a home nor a component interface. A message-driven bean instance is an instance of a message-driven bean class.

The EJB 2.0 specification supported only JMS message-driven beans. The EJB 2.1 specification extends the message-driven bean contracts to support other messaging types in addition to JMS^[34]. A message-driven bean is defined for a single messaging type, in accordance with the message listener interface it employs.

To a client, a message-driven bean is a message consumer that implements some business logic running on the server. A client accesses a message-driven bean by sending messages to the destination or endpoint for which the message-driven bean class is the message listener.

Message-driven beans are anonymous. They have no client-visible identity.

Message-driven bean instances have no conversational state. This means that all bean instances are equivalent when they are not involved in servicing a client message.

A message-driven bean instance is created by the container to handle the processing of the messages for which the message-driven bean is the consumer. Its lifetime is controlled by the container.

A message-driven bean instance has no state for a specific client. However, the instance variables of the message-driven bean instance can contain state across the handling of client messages. Examples of such state include an open database connection and an object reference to an EJB object.

15.2 Goals

The goal of the message-driven bean model is to make developing an enterprise bean that is asynchronously invoked to handle the processing of incoming messages as simple as developing the same functionality in any other message listener.


A further goal of the message-driven bean model is to allow for the concurrent processing of a stream of messages by means of container-provided pooling of message-driven bean instances.

15.3 Client View of a Message-Driven Bean

To a client, a message-driven bean is simply a message consumer. The client sends messages to the destination or endpoint for which the message-driven bean is the message listener just as it would to any other destination or endpoint. The message-driven bean, as a message consumer, handles the processing of the messages.

From the perspective of the client, the existence of a message-driven bean is completely hidden behind the destination or endpoint for which the message-driven bean is the message listener. The following diagram illustrates the view that is provided to a message-driven bean's clients.

[34] The pluggable integration of the messaging provider with the container is through a resource adapter, as described in [12].

Figure 61 Client view of Message-Driven Beans Deployed in a Container

A client's JNDI name space may be configured to include the destinations or endpoints of message-driven beans installed in multiple EJB containers located on multiple machines on a network. The actual locations of an enterprise bean and EJB container are, in general, transparent to the client using the enterprise bean.

For example, the queue for the `StockInfo` JMS message-driven bean can be located using the following code segment:

```
Context initialContext = new InitialContext();
Queue stockInfoQueue = (javax.jms.Queue)initialContext.lookup
    ("java:comp/env/jms/stockInfoQueue");
```

The remainder of this section describes the message-driven bean life cycle in detail and the protocol between the message-driven bean and its container.

15.4 Protocol Between a Message-Driven Bean Instance and its Container

From its creation until destruction, a message-driven bean instance lives in a container. The container provides security, concurrency, transactions, and other services for the message-driven bean. The container manages the life cycle of the message-driven bean instances, notifying the instances when bean action may be necessary, and providing a full range of services to ensure that the message-driven bean implementation is scalable and can support the concurrent processing of a large number of messages.

From the Bean Provider's point of view, a message-driven bean exists as long as its container does. It is the container's responsibility to ensure that the message-driven bean comes into existence when the container is started up and that instances of the bean are ready to receive an asynchronous message delivery before the delivery of messages is started.

Containers themselves make no actual service demands on the message-driven bean instances. The calls a container makes on a bean instance provide it with access to container services and deliver notifications issued by the container.

Since all instances of a message-driven bean are equivalent, a client message can be delivered to any available instance.

15.4.1 The Required MessageDrivenBean Interface

All message-driven beans must implement the `MessageDrivenBean` interface.

The `setMessageDrivenContext` method is called by the bean's container to associate a message-driven bean instance with its context maintained by the container. Typically a message-driven bean instance retains its message-driven context as part of its state.

The `ejbRemove` notification signals that the instance is in the process of being removed by the container. In the `ejbRemove` method, the instance releases the resources that it is holding.

15.4.2 The Required Message Listener Interface

The message-driven bean class must implement the appropriate message listener interface for the messaging type that the message-driven bean supports. The specific message listener interface that is implemented by a message-driven bean class distinguishes the messaging type that the message-driven bean supports.

The message-driven bean class's implementation of the `javax.jms.MessageListener` interface distinguishes the message-driven bean as a JMS message-driven bean.

The bean's message listener method (e.g., `onMessage` in the case of `javax.jms.MessageListener`) is called by the container when a message has arrived for the bean to service. The message listener method contains the business logic that handles the processing of the message.

A bean's message listener interface may define more than one message listener method. If the message listener interface contains more than one method, it is the resource adapter that determines which method is invoked. See [12].

Only message-driven beans can asynchronously receive messages. Session and entity beans are not permitted to be message listeners.

15.4.3 The Optional TimedObject Interface

A message driven bean can be registered with the EJB timer service for time-based event notifications if it implements the `javax.ejb.TimedObject` interface in addition to the message listener interface. The container invokes the bean instance's `ejbTimeout` method when a timer for the bean has expired. See Chapter 22, "Timer Service".

15.4.4 The MessageDrivenContext Interface

The container provides the message-driven bean instance with a `MessageDrivenContext`. This gives the message-driven bean instance access to the instance's context maintained by the container. The `MessageDrivenContext` interface has the following methods:

- The `setRollbackOnly` method allows the instance to mark the current transaction such that the only outcome of the transaction is a rollback. Only instances of a message-driven bean with container-managed transaction demarcation can use this method.
- The `getRollbackOnly` method allows the instance to test if the current transaction has been marked for rollback. Only instances of a message-driven bean with container-managed transaction demarcation can use this method.
- The `getUserTransaction` method returns the `javax.transaction.UserTransaction` interface that the instance can use to demarcate transactions, and to obtain transaction status. Only instances of a message-driven bean with bean-managed transaction demarcation can use this method.
- The `getTimerService` method returns the `javax.ejb.TimerService` interface.
- The `getCallerPrincipal` method returns the `java.security.Principal` that is associated with the invocation.
- The `isCallerInRole` method is inherited from the `EJBContext` interface. Message-driven bean instances must not call this method.
- The `getEJBHome` and `getEJBLocalHome` methods are inherited from the `EJBContext` interface. Message-driven bean instances must not call these methods.

15.4.5 Message-Driven Bean's ejbCreate Method

The container creates an instance of a message-driven bean in three steps. First, the container calls the bean class' `newInstance` method to create a new message-driven bean instance. Second, the container calls the `setMessageDrivenContext` method to pass the context object to the instance. Third, the container calls the instance's `ejbCreate` method.

Each message-driven bean class must have one `ejbCreate` method, with no arguments.

15.4.6 Serializing Message-Driven Bean Methods

The container serializes calls to each message-driven bean instance. Most containers will support many instances of a message-driven bean executing concurrently; however, each instance sees only a serialized sequence of method calls. Therefore, a message-driven bean does not have to be coded as reentrant.

The container must serialize all the container-invoked callbacks (e.g., `ejbRemove`, `ejbTimeout` methods), and it must serialize these callbacks with the message listener method calls.

15.4.7 Concurrency of Message Processing

A container allows many instances of a message-driven bean class to be executing concurrently, thus allowing for the concurrent processing of a stream of messages. No guarantees are made as to the exact order in which messages are delivered to the instances of the message-driven bean class, although the container should attempt to deliver messages in order when it does not impair the concurrency of message processing. Message-driven beans should therefore be prepared to handle messages that are out of sequence: for example, the message to cancel a reservation may be delivered before the message to make the reservation.

15.4.8 Transaction Context of Message-Driven Bean Methods

A bean's message listener method is invoked in the scope of a transaction determined by the transaction attribute specified in the deployment descriptor. If the bean is specified as using container-managed transaction demarcation, either the `Required` or the `NotSupported` transaction attribute must be used.

When a message-driven bean using bean-managed transaction demarcation uses the `javax.transaction.UserTransaction` interface to demarcate transactions, the message receipt that causes the bean to be invoked is not part of the transaction. If the message receipt is to be part of the transaction, container-managed transaction demarcation with the `Required` transaction attribute must be used.

A message-driven bean's `newInstance`, `setMessageDrivenContext`, `ejbCreate`, and `ejbRemove` methods are called with an unspecified transaction context. Refer to Subsection 17.6.5 for how the container executes methods with an unspecified transaction context.

15.4.9 Activation Configuration Properties

The Bean Provider may provide information to the Deployer about the configuration of the message-driven bean in its operational environment. This may include information about message acknowledgement modes, message selectors, expected destination or endpoint types, etc.

Activation configuration properties for JMS message-driven beans are described in Sections 15.4.10 through 15.4.12.

15.4.10 Message Acknowledgment for JMS Message-Driven Beans

JMS message-driven beans should not attempt to use the JMS API for message acknowledgment. Message acknowledgment is automatically handled by the container. If the message-driven bean uses container-managed transaction demarcation, message acknowledgment is handled automatically as a part of the transaction commit. If bean-managed transaction demarcation is used, the message receipt cannot be part of the bean-managed transaction, and, in this case, the receipt is acknowledged by the container. If bean-managed transaction demarcation is used, the Bean Provider can indicate in the `activation-config-property` deployment descriptor element whose `activation-config-property-name` is `acknowledgeMode` whether JMS `AUTO_ACKNOWLEDGE` semantics or `DUPS_OK_ACKNOWLEDGE` semantics should apply. If the `acknowledgeMode` property is not specified, JMS `AUTO_ACKNOWLEDGE` semantics are assumed. The `activation-config-property-value` for this property must be either `Auto-acknowledge` or `Dups-ok-acknowledge` for a JMS message-driven bean.

15.4.11 Message Selectors for JMS Message-Driven Beans

The Bean Provider may declare the JMS message selector to be used in determining which messages a JMS message-driven bean is to receive. If the Bean Provider wishes to restrict the messages that a JMS message-driven bean receives, the Bean Provider can specify the value of the message selector the using the `activation-config-property` deployment descriptor element whose `activation-config-property-name` is `messageSelector`.

For example:

```
<activation-config>
<activation-config-property>
<activation-config-property-name>messageSelector</activation-con-
fig-property-name>
<activation-config-property-value>JMSType = 'car' AND color = 'blue'
AND weight > 2500</activation-config-property-value>
</activation-config-property>
</activation-config>
```

The Application Assembler may further restrict, but not replace, the value of the `messageSelector` property of a JMS message-driven bean.

15.4.12 Association of a Message-Driven Bean with a Destination or Endpoint

A message-driven bean is associated with a destination or endpoint when the bean is deployed in the container. It is the responsibility of the Deployer to associate the message-driven bean with a destination or endpoint.

15.4.12.1 JMS Message-Driven Beans

A JMS message-driven bean is associated with a JMS Destination (Queue or Topic) when the bean is deployed in the container. It is the responsibility of the Deployer to associate the message-driven bean with a Queue or Topic.

The Bean Provider may provide advice to the Deployer as to whether a message-driven bean is intended to be associated with a queue or a topic by using the `activation-config-property` deployment element whose `activation-config-property-name` is `destinationType`. The `activation-config-property-value` for this property must be either `javax.jms.Queue` or `javax.jms.Topic` for a JMS message-driven bean.

If the message-driven bean is intended to be used with a topic, the Bean Provider may further indicate whether a durable or non-durable subscription should be used by using the `activation-config-property` deployment descriptor element whose `activation-config-property-name` is `subscriptionDurability`. The `activation-config-property-value` for this property must be either `Durable` or `NonDurable` for a JMS message-driven bean. If a topic subscription is specified and `subscriptionDurability` is not specified, a non-durable subscription is assumed.

- Durable topic subscriptions, as well as queues, ensure that messages are not missed even if the EJB server is not running. Reliable applications will typically make use of queues or durable topic subscriptions rather than non-durable topic subscriptions.
- If a non-durable topic subscription is used, it is the container's responsibility to make sure that the message driven bean subscription is active (i.e., that there is a message driven bean available to service the message) in order to ensure that messages are not missed as long as the EJB server is running. Messages may be missed, however, when a bean is not available to service them. This will occur, for example, if the EJB server goes down for any period of time.

The Deployer should avoid associating more than one message-driven bean with the same JMS Queue. If there are multiple JMS consumers for a queue, JMS does not define how messages are distributed between the queue receivers.

15.4.13 Dealing with Exceptions

A message-driven bean's message listener method must not throw the `java.rmi.RemoteException`.

Message-driven beans should not, in general, throw `RuntimeExceptions`.

A `RuntimeException` thrown from any method of the message-driven bean class (including a message listener method and the callbacks invoked by the container) results in the transition to the “does not exist” state. If a message-driven bean uses bean-managed transaction demarcation and throws a `RuntimeException`, the container should not acknowledge the message. Exception handling is described in detail in Chapter 18.

From the client perspective, the message consumer continues to exist. If the client continues sending messages to the destination or endpoint associated with the bean, the container can delegate the client’s messages to another instance.

The message listener methods of some messaging types may throw application exceptions. An application exception is propagated by the container to the resource adapter.

15.4.14 Missed ejbRemove Calls

The Bean Provider cannot assume that the container will always invoke the `ejbRemove` method on a message-driven bean instance. The following scenarios result in `ejbRemove` not being called on an instance:

- A crash of the EJB container.
- A system exception thrown from the instance’s method to the container.

If the message-driven bean instance allocates resources in the `ejbCreate` method and/or in the message listener method, and releases normally the resources in the `ejbRemove` method, these resources will not be automatically released in the above scenarios. The application using the message-driven bean should provide some clean up mechanism to periodically clean up the unreleased resources.

15.4.15 Replying to a JMS Message

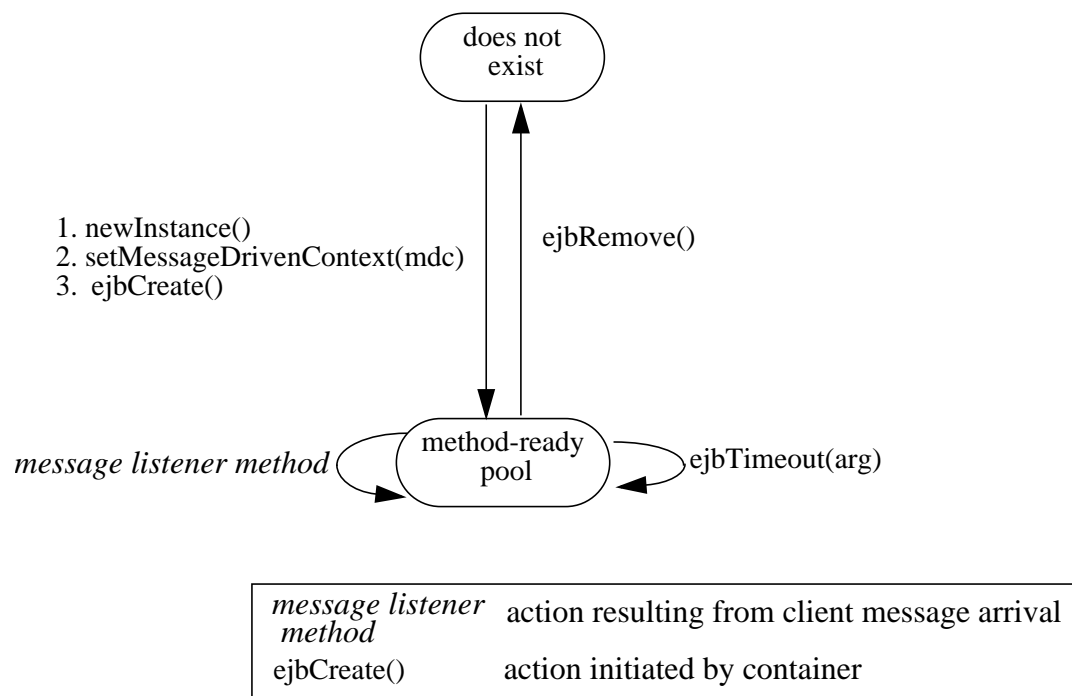
In standard JMS usage scenarios, the messaging mode of a message’s `JMSReplyTo` destination (Queue or Topic) is the same as the mode of the destination to which the message has been sent. Although a message-driven bean is not directly dependent on the mode of the JMS destination from which it is consuming messages, it may contain code that depends on the mode of its message’s `JMSReplyTo` destination. In particular, if a message-driven bean replies to a message, the mode of the reply’s message producer and the mode of the `JMSReplyTo` destination must be the same. In order to implement a message-driven bean that is independent of `JMSReplyTo` mode, the Bean Provider should use `instanceOf` to test whether a `JMSReplyTo` destination is a Queue or Topic, and then use a matching message producer for the reply.

15.5 Message-Driven Bean State Diagram

When a client sends a message to a Destination for which a message-driven bean is the consumer, the container selects one of its method-ready instances and invokes the instance’s message listener method.

The following figure illustrates the life cycle of a message-driven bean instance.

Figure 62 Life Cycle of a Message-Driven Bean.



The following steps describe the life cycle of a message-driven bean instance:

- A message-driven bean instance's life starts when the container invokes `newInstance` on the message-driven bean class to create a new instance. Next, the container calls `setMessageDrivenContext` followed by `ejbCreate` on the instance.
- The message-driven bean instance is now ready to be delivered a message sent to its associated destination or endpoint by any client or a call from the container to the `ejbTimeout` method.
- When the container no longer needs the instance (which usually happens when the container wants to reduce the number of instances in the method-ready pool), the container invokes `ejbRemove` on it. This ends the life of the message-driven bean instance.

15.5.1 Operations Allowed in the Methods of a Message-Driven Bean Class

Table 12 defines the methods of a message-driven bean class in which the message-driven bean instances can access the methods of the `javax.ejb.MessageDrivenContext` interface, the `java:comp/env` environment naming context, resource managers, `TimerService` and `Timer` methods, and other enterprise beans.

If a message-driven bean instance attempts to invoke a method of the `MessageDrivenContext` interface, and the access is not allowed in Table 12, the container must throw and log the `java.lang.IllegalStateException`.

If a message-driven bean instance attempts to invoke a method of the `TimerService` or `Timer` interface and the access is not allowed in Table 12, the container must throw the `java.lang.IllegalStateException`.

If a bean instance attempts to access a resource manager or an enterprise bean and the access is not allowed in Table 12, the behavior is undefined by the EJB architecture.

Table 12 Operations Allowed in the Methods of a Message-Driven Bean

Bean method	Bean method can perform the following operations	
	Container-managed transaction demarcation	Bean-managed transaction demarcation
constructor	-	-
setMessageDrivenContext	JNDI access to java:comp/env	JNDI access to java:comp/env
ejbCreate ejbRemove	MessageDrivenContext methods: <i>getTimerService</i> JNDI access to java:comp/env	MessageDrivenContext methods: <i>getUserTransaction, getTimerService</i> JNDI access to java:comp/env
message listener method	MessageDrivenContext methods: <i>getRollbackOnly, setRollbackOnly, getCallerPrincipal, getTimerService</i> JNDI access to java:comp/env Resource manager access Enterprise bean access Timer service or Timer methods	MessageDrivenContext methods: <i>getUserTransaction, getCallerPrincipal, getTimerService</i> UserTransaction methods JNDI access to java:comp/env Resource manager access Enterprise bean access Timer service or Timer methods
ejbTimeout	MessageDrivenContext methods: <i>getRollbackOnly, setRollbackOnly, getCallerPrincipal, getTimerService</i> JNDI access to java:comp/env Resource manager access Enterprise bean access Timer service or Timer methods	MessageDrivenContext methods: <i>getUserTransaction, getCallerPrincipal, getTimerService</i> UserTransaction methods JNDI access to java:comp/env Resource manager access Enterprise bean access Timer service or Timer methods

Additional restrictions:

- The `getRollbackOnly` and `setRollbackOnly` methods of the `MessageDrivenContext` interface should be used only in the message-driven bean methods that execute in the context of a transaction. The container must throw the `java.lang.IllegalStateException` if the methods are invoked while the instance is not associated with a transaction.

The reasons for disallowing operations in Table 12:

- Invoking the `getRollbackOnly` and `setRollbackOnly` methods is disallowed in the message-driven bean methods for which the container does not have a meaningful transaction context, and for all message-driven beans with bean-managed transaction demarcation.

- The `UserTransaction` interface is unavailable to message-driven beans with container-managed transaction demarcation.
- Invoking `getEJBHome` or `getEJBLocalHome` is disallowed in message-driven bean methods because there are no `EJBHome` or `EJBLocalHome` objects for message-driven beans. The container must throw and log the `java.lang.IllegalStateException` if these methods are invoked.

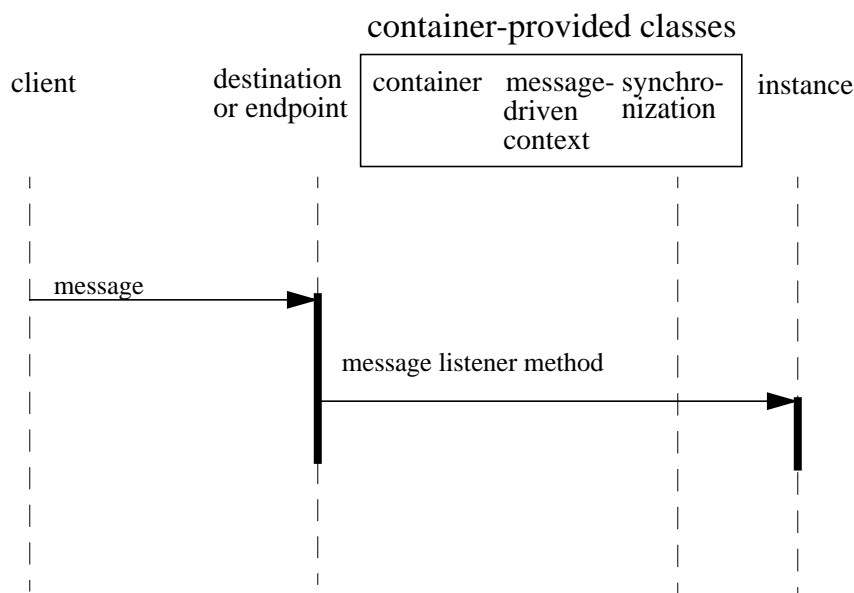
15.6 Object Interaction Diagrams for a Message-Driven Bean

This section contains object interaction diagrams that illustrate the interaction of the classes.

15.6.1 Message Receipt: Message Listener Method Invocation

The following diagram illustrates message listener method invocation.

Figure 63 OID for Invocation of Message Listener Method on Message-Driven Bean Instance




15.6.2 Adding an Instance to the Pool

The following diagram illustrates the sequence for a container adding an instance to the method-ready pool.

Figure 64 OID for Container Adding Instance of a Message-Driven Bean to a Method-Ready Pool

container-provided classes

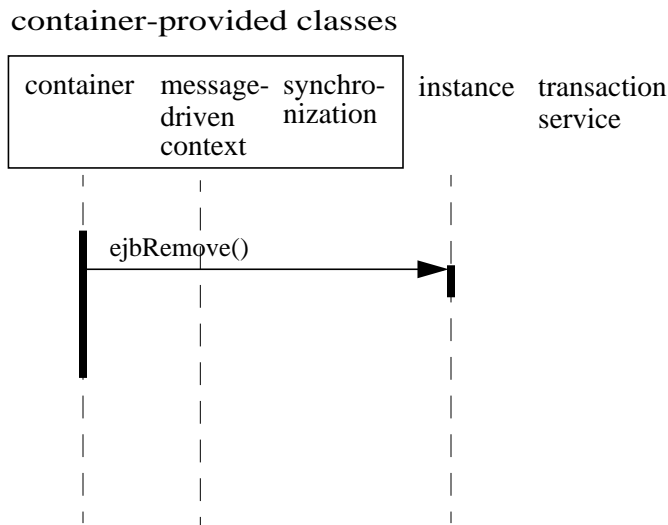


15.6.3 Removing an Instance from the Pool

The following diagram illustrates the sequence for a container removing an instance from the method-ready pool.

Figure 65

OID for a Container Removing an Instance of Message-Driven Bean from Ready Pool



15.7 The Responsibilities of the Bean Provider

This section describes the responsibilities of the message-driven Bean Provider to ensure that a message-driven bean can be deployed in any EJB container.

15.7.1 Classes and Interfaces

The message-driven Bean Provider is responsible for providing the following class files:

- Message-driven bean class.

15.7.2 Message-Driven Bean Class

The following are the requirements for the message-driven bean class:

The class must implement, directly or indirectly, the `javax.ejb.MessageDrivenBean` interface.

The class must implement, directly or indirectly, the message listener interface required by the messaging type that it supports. In the case of JMS, this is the `javax.jms.MessageListener` interface.

The class may implement, directly or indirectly, the `javax.ejb.TimedObject` interface.

The class must be defined as `public`, must not be `final`, and must not be `abstract`. The class must be a top level class.

The class must have a `public` constructor that takes no arguments. The container uses this constructor to create instances of the message-driven bean class.

The class must not define the `finalize` method.

The class must implement the `ejbCreate` method.

The message-driven bean class may have superclasses and/or superinterfaces. If the message-driven bean has superclasses, the `ejbCreate` method, the methods of the `MessageDrivenBean` interface, the methods of the message listener interface, and the `ejbTimeout` method may be defined in the message-driven bean class or in any of its superclasses.

The message-driven bean class is allowed to implement other methods (for example, helper methods invoked internally by the message listener method) in addition to the methods required by the EJB specification.

15.7.3 ejbCreate Method

The message-driven bean class must define one `ejbCreate` method whose signature must follow these rules:

The method name must be `ejbCreate`.

The method must be declared as `public`.

The method must not be declared as `final` or `static`.

The return type must be `void`.

The method must have no arguments.

The `throws` clause must not define any application exceptions.

15.7.4 Message Listener Method

The message-driven bean class must define the message listener methods. The signature of a message listener method must follow these rules:

The method must be declared as `public`.

The method must not be declared as `final` or `static`.

15.7.5 ejbRemove Method

The message-driven bean class must define one `ejbRemove` method whose signature must follow these rules:

The method name must be `ejbRemove`.

The method must be declared as `public`.

The method must not be declared as `final` or `static`.

The return type must be `void`.

The method must have no arguments.

The `throws` clause must not define any application exceptions.

15.8 The Responsibilities of the Container Provider

This section describes the responsibilities of the Container Provider to support a message-driven bean. The Container Provider is responsible for providing the deployment tools, and for managing the message-driven bean instances at runtime.

Because the EJB specification does not define the API between deployment tools and the container, we assume that the deployment tools are provided by the Container Provider. Alternatively, the deployment tools may be provided by a different vendor who uses the container vendor's specific API.

15.8.1 Generation of Implementation Classes

The deployment tools provided by the container are responsible for the generation of additional classes when the message-driven bean is deployed. The tools obtain the information that they need for generation of the additional classes by introspecting the classes and interfaces provided by the enterprise Bean Provider and by examining the message-driven bean's deployment descriptor.

The deployment tools may generate a class that mixes some container-specific code with the message-driven bean class. This code may, for example, help the container to manage the bean instances at runtime. Subclassing, delegation, and code generation can be used by the tools.

15.8.2 Deployment of JMS Message-Driven Beans

The Container Provider must support the deployment of a JMS message-driven bean as the consumer of a JMS queue or a durable subscription.

15.8.3 Request/Response Messaging Types

If the message listener supports a request/response messaging type, it is the container's responsibility to deliver the message response.

15.8.4 Non-reentrant Instances

The container must ensure that only one thread can be executing an instance at any time.

15.8.5 Transaction Scoping, Security, Exceptions

The container must follow the rules with respect to transaction scoping, security checking, and exception handling, as described in Chapters 17, 21, and 18.

Example Message-Driven Bean Scenario

This chapter describes an example development and deployment scenario of a JMS message-driven bean. We use the scenario to explain the responsibilities of the Bean Provider and those of the Container Provider.

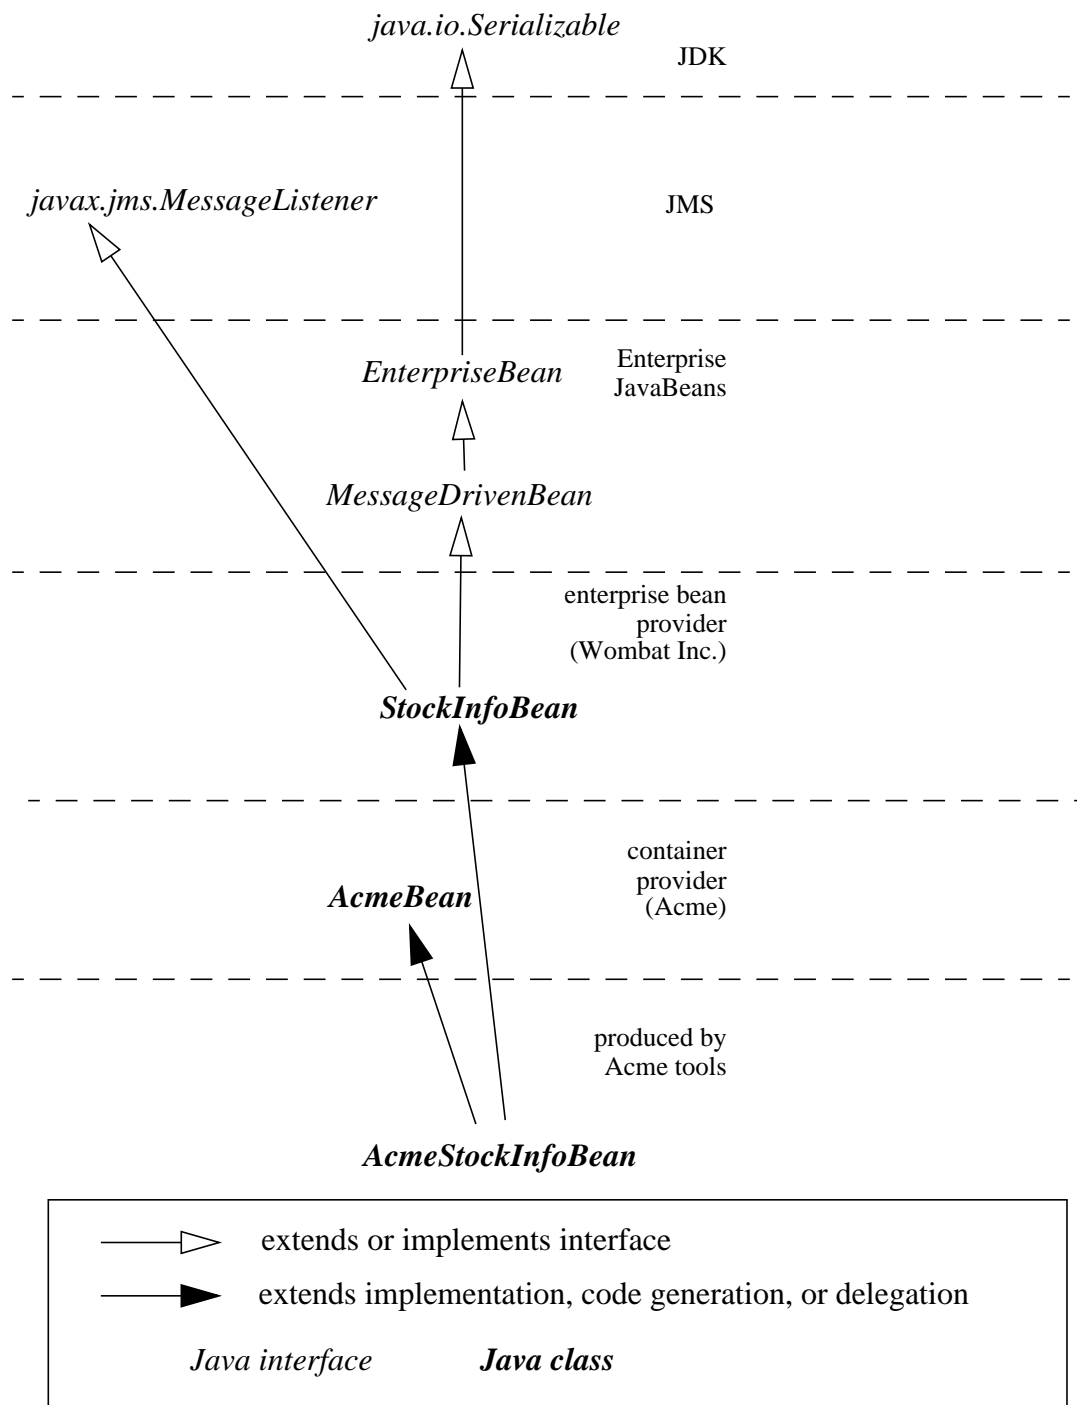
The classes generated by the Container Provider's tools in this scenario should be considered illustrative rather than prescriptive. Container providers are free to implement the contract between a message-driven bean and its container in a different way, provided that it achieves an equivalent effect (from the perspectives of the Bean Provider and the client-side programmer).

16.1 Overview

Wombat Inc. has developed the `StockInfoBean` message-driven Bean. The `StockInfoBean` is deployed in a container provided by the Acme Corporation.

16.2 Inheritance Relationship

An example of the inheritance relationship between the interfaces and classes is illustrated in the following diagram:

Figure 66 Example of Inheritance Relationships Between EJB Classes

16.2.1 What the Message-Driven Bean Provider is Responsible For

Wombat Inc. is responsible for providing the following:

- *Write the business logic in the message-driven bean class (`StockInfoBean`), defining the `onMessage` message listener method that is invoked when the bean is to service a JMS message. The message-driven bean must implement the `javax.ejb.MessageDrivenBean` and `javax.jms.MessageListener` interfaces, and define the `ejbCreate` method invoked at an EJB object creation.*
- *Define a deployment descriptor that specifies any declarative metadata that the message-driven Bean Provider wishes to pass with the bean to the next stage of the development/deployment workflow.*

16.2.2 Classes Supplied by Container Provider

The following classes are supplied by the Container Provider Acme Corp:

- *The `AcmeBean` class provides additional state and methods to allow Acme's container to manage its message-driven bean instances.*

16.2.3 What the Container Provider is Responsible For

The tools provided by Acme Corporation are responsible for the following:

- *Generate the implementation of the message-driven bean class suitable for the Acme container (`AcmeStockInfoBean`). `AcmeStockInfoBean` includes the business logic from the `StockInfoBean` class mixed with the services defined in the `AcmeBean` class. Acme tools can use inheritance, delegation, and code generation to achieve a mix-in of the two classes.*

Many of the above classes and tools are container-specific (i.e., they reflect the way Acme Corp implemented them). Other container providers may use different mechanisms to produce their runtime classes, which will likely be different from those generated by Acme's tools.

Support for Transactions

One of the key features of the Enterprise JavaBeans™ architecture is support for distributed transactions. The Enterprise JavaBeans architecture allows an application developer to write an application that atomically updates data in multiple databases which may be distributed across multiple sites. The sites may use EJB servers from different vendors.

17.1 Overview

This section provides a brief overview of transactions and illustrates a number of transaction scenarios in EJB.

17.1.1 Transactions

Transactions are a proven technique for simplifying application programming. Transactions free the application programmer from dealing with the complex issues of failure recovery and multi-user programming. If the application programmer uses transactions, the programmer divides the application's work into units called transactions. The transactional system ensures that a unit of work either fully completes, or the work is fully rolled back. Furthermore, transactions make it possible for the programmer to design the application as if it ran in an environment that executes units of work serially.

Support for transactions is an essential component of the Enterprise JavaBeans architecture. The Enterprise Bean Provider and the client application programmer are not exposed to the complexity of distributed transactions. The Bean Provider can choose between using programmatic transaction demarcation in the enterprise bean code (this style is called *bean-managed transaction demarcation*) or declarative transaction demarcation performed automatically by the EJB container (this style is called *container-managed transaction demarcation*).

With bean-managed transaction demarcation, the enterprise bean code demarcates transactions using the `javax.transaction.UserTransaction` interface. All resource manager^[35] accesses between the `UserTransaction.begin` and `UserTransaction.commit` calls are part of a transaction.

With container-managed transaction demarcation, the container demarcates transactions per instructions provided by the Application Assembler in the deployment descriptor. These instructions, called *transaction attributes*, tell the container whether it should include the work performed by an enterprise bean method in a client's transaction, run the enterprise bean method in a new transaction started by the container, or run the method with "no transaction" (Refer to Subsection 17.6.5 for the description of the "no transaction" case).

Regardless of whether an enterprise bean uses bean-managed or container-managed transaction demarcation, the burden of implementing transaction management is on the EJB Container and Server Provider. The EJB container and server implement the necessary low-level transaction protocols, such as the two-phase commit protocol between a transaction manager and a database system or messaging provider, transaction context propagation, and distributed two-phase commit.

Many applications will consist of one or several enterprise beans that all use a single resource manager (typically a relational database management system). The EJB container can make use of resource manager local transactions as an optimization technique for enterprise beans for which distributed transactions are not needed. A resource manager local transaction does not involve control or coordination by an external transaction manager. The container's use of local transactions as an optimization technique for enterprise beans with either container-managed transaction demarcation or bean-managed transaction demarcation is not visible to the enterprise beans. For a discussion of the use of resource manager local transactions as a container optimization strategy, refer to [9] and [12].

17.1.2 Transaction Model

The Enterprise JavaBeans architecture supports flat transactions. A flat transaction cannot have any child (nested) transactions.

Note: The decision not to support nested transactions allows vendors of existing transaction processing and database management systems to incorporate support for Enterprise JavaBeans. If these vendors provide support for nested transactions in the future, Enterprise JavaBeans may be enhanced to take advantage of nested transactions.

[35] The terms *resource* and *resource manager* used in this chapter refer to the resources declared in the enterprise bean's deployment descriptor using the `resource-ref` element. This includes not only database resources, but also other resources, such as JMS Connections. These resources are considered to be "managed" by the container.

17.1.3 Relationship to JTA and JTS

The Java™ Transaction API (JTA) [5] is a specification of the interfaces between a transaction manager and the other parties involved in a distributed transaction processing system: the application programs, the resource managers, and the application server.

The Java Transaction Service (JTS) [6] API is a Java binding of the CORBA Object Transaction Service (OTS) 1.1 specification. JTS provides transaction interoperability using the standard IIOP protocol for transaction propagation between servers. The JTS API is intended for vendors who implement transaction processing infrastructure for enterprise middleware. For example, an EJB server vendor may use a JTS implementation as the underlying transaction manager.

The EJB architecture does not require the EJB container to support the JTS interfaces. The EJB architecture requires that the EJB container support the JTA API defined in [5] and the Connector APIs defined in [12].


17.2 Sample Scenarios

This section describes several scenarios that illustrate the distributed transaction capabilities of the Enterprise JavaBeans architecture.

17.2.1 Update of Multiple Databases

The Enterprise JavaBeans architecture makes it possible for an application program to update data in multiple databases in a single transaction.

In the following figure, a client invokes the enterprise bean X. Bean X updates data using two database connections that the Deployer configured to connect with two different databases, A and B. Then X calls another enterprise bean, Y. Bean Y updates data in database C. The EJB server ensures that the updates to databases A, B, and C are either all committed or all rolled back.


Figure 67 Updates to Simultaneous Databases

The application programmer does not have to do anything to ensure transactional semantics. The enterprise beans X and Y perform the database updates using the standard JDBC™ API. Behind the scenes, the EJB server enlists the database connections as part of the transaction. When the transaction commits, the EJB server and the database systems perform a two-phase commit protocol to ensure atomic updates across all three databases.

17.2.2 Messages Sent or Received Over JMS Sessions and Update of Multiple Databases


The Enterprise JavaBeans architecture makes it possible for an application program to send messages to or receive messages from one or more JMS Destinations and/or to update data in one or more databases in a single transaction.

In the following figure, a client invokes the enterprise bean X. Bean X sends a message to a JMS queue A and updates data in a database B using connections that the Deployer configured to connect with a JMS provider and a database. Then X calls another enterprise bean, Y. Bean Y updates data in database C. The EJB server ensures that the operations on A, B, and C are either all committed, or all rolled back.

Figure 68 Message Sent to JMS Queue and Updates to Multiple Databases

The application programmer does not have to do anything to ensure transactional semantics. The enterprise beans X and Y perform the message send and database updates using the standard JMS and JDBC™ APIs. Behind the scenes, the EJB server enlists the session on the connection to the JMS provider and the database connections as part of the transaction. When the transaction commits, the EJB server and the messaging and database systems perform a two-phase commit protocol to ensure atomic updates across all the three resources.


In the following figure, a client sends a message to the JMS queue A serviced by the message-driven bean X. Bean X updates data using two database connections that the Deployer configured to connect with two different databases, B and C. The EJB server ensures that the dequeuing of the JMS message, its receipt by bean X, and the updates to databases B and C are either all committed or all rolled back.

Figure 69 Message Sent to JMS Queue Served by Message-Driven Bean and Updates to Multiple Databases

17.2.3 Update of Databases via Multiple EJB Servers

The Enterprise JavaBeans architecture allows updates of data at multiple sites to be performed in a single transaction.

In the following figure, a client invokes the enterprise bean X. Bean X updates data in database A, and then calls another enterprise bean Y that is installed in a remote EJB server. Bean Y updates data in database B. The Enterprise JavaBeans architecture makes it possible to perform the updates to databases A and B in a single transaction.

Figure 70 Updates to Multiple Databases in Same Transaction

When X invokes Y, the two EJB servers cooperate to propagate the transaction context from X to Y. This transaction context propagation is transparent to the application-level code.

At transaction commit time, the two EJB servers use a distributed two-phase commit protocol (if the capability exists) to ensure the atomicity of the database updates.


17.2.4 Client-Managed Demarcation

A Java client can use the `javax.transaction.UserTransaction` interface to explicitly demarcate transaction boundaries. The client program obtains the `javax.transaction.UserTransaction` interface using JNDI as defined in the Java 2, Enterprise Edition specification [9].

The EJB specification does not imply that the `javax.transaction.UserTransaction` is available to all Java clients. The Java 2, Enterprise Edition specification specifies the client environments in which the `javax.transaction.UserTransaction` interface is available.

A client program using explicit transaction demarcation may perform, via enterprise beans, atomic updates across multiple databases residing at multiple EJB servers, as illustrated in the following figure.

Figure 71 Updates on Multiple Databases on Multiple Servers



The application programmer demarcates the transaction with `begin` and `commit` calls. If the enterprise beans `X` and `Y` are configured to use a client transaction (i.e., their methods have either the `Required`, `Mandatory`, or `Supports` transaction attribute), the EJB server ensures that the updates to databases `A` and `B` are made as part of the client's transaction.


17.2.5 Container-Managed Demarcation

Whenever a client invokes a method on the home or component interface of an enterprise bean, the container interposes on the method invocation. The interposition allows the container to control transaction demarcation declaratively through the transaction attribute set by the Application Assembler. (See Section 17.4.1 for a description of transaction attributes.)

For example, if an enterprise bean method is configured with the `Required` transaction attribute, the container behaves as follows: If the client request is not associated with a transaction context, the container automatically initiates a transaction whenever a client invokes an enterprise bean method that requires a transaction context. If the client request contains a transaction context, the container includes the enterprise bean method in the client transaction.

The following figure illustrates such a scenario. A non-transactional client invokes the enterprise bean `X`, and the invoked method has the `Required` transaction attribute. Because the message from the client does not include a transaction context, the container starts a new transaction before dispatching the method on `X`. Bean `X`'s work is performed in the context of the transaction. When `X` calls other enterprise beans (`Y` in our example), the work performed by the other enterprise beans is also automatically included in the transaction (subject to the transaction attribute of the other enterprise bean).

Figure 72 Update of Multiple Databases from Non-Transactional Client



The container automatically commits the transaction at the time `X` returns a reply to the client.

If a message-driven bean's message listener method is configured with the `Required` transaction attribute, the container automatically starts a new transaction before the delivery of the message and, hence, before the invocation of the method.^[36]

JMS requires that the transaction be started before the dequeuing of the message. See [10].

The container automatically enlists the resource manager associated with the arriving message and all the resource managers accessed by the message listener method with the transaction.

17.3 Bean Provider's Responsibilities

This section describes the Bean Provider's view of transactions and defines the Bean Provider's responsibilities.

17.3.1 Bean-Managed Versus Container-Managed Transaction Demarcation

When designing an enterprise bean, the Bean Provider must decide whether the enterprise bean will demarcate transactions programmatically in the business methods (bean-managed transaction demarcation), or whether the transaction demarcation is to be performed by the container based on the transaction attributes in the deployment descriptor (container-managed transaction demarcation).

A session bean or a message-driven bean can be designed with bean-managed transaction demarcation or with container-managed transaction demarcation. (But it cannot be both at the same time.)

An entity bean must always be designed with container-managed transaction demarcation.

An enterprise bean instance can access resource managers in a transaction only in the enterprise bean's methods in which there is a transaction context available. An entity bean with container-managed persistence can access its persistent state in a transaction only in the enterprise bean's methods in which there is a transaction context available. Refer to Table 4 on page 197.

17.3.1.1 Non-Transactional Execution

Some enterprise beans may need to access resource managers that do not support an external transaction coordinator. The container cannot manage the transactions for such enterprise beans in the same way that it can for the enterprise beans that access resource managers that support an external transaction coordinator.

If an enterprise bean needs to access a resource manager that does not support an external transaction coordinator, the Bean Provider should design the enterprise bean with container-managed transaction demarcation and assign the `NotSupported` transaction attribute to all the bean's methods. The EJB architecture does not specify the transactional semantics of the enterprise bean methods. See Subsection 17.6.5 for how the container implements this case.

[36] We use the term "container" here to encompass both the container and the messaging provider. When the contracts outlined in [12] are used, it may be the messaging provider that starts the transaction.

17.3.2 Isolation Levels

Transactions not only make completion of a unit of work atomic, but they also isolate the units of work from each other, provided that the system allows concurrent execution of multiple units of work.

The *isolation level* describes the degree to which the access to a resource manager by a transaction is isolated from the access to the resource manager by other concurrently executing transactions.

The following are guidelines for managing isolation levels in enterprise beans.

- The API for managing an isolation level is resource-manager-specific. (Therefore, the EJB architecture does not define an API for managing isolation levels.)
- If an enterprise bean uses multiple resource managers, the Bean Provider may specify the same or different isolation level for each resource manager. This means, for example, that if an enterprise bean accesses multiple resource managers in a transaction, access to each resource manager may be associated with a different isolation level.
- The Bean Provider must take care when setting an isolation level. Most resource managers require that all accesses to the resource manager within a transaction are done with the same isolation level. An attempt to change the isolation level in the middle of a transaction may cause undesirable behavior, such as an implicit sync point (a commit of the changes done so far).
- For session beans and message-driven beans with bean-managed transaction demarcation, the Bean Provider can specify the desirable isolation level programmatically in the enterprise bean's methods, using the resource-manager specific API. For example, the Bean Provider can use the `java.sql.Connection.setTransactionIsolation` method to set the appropriate isolation level for database access.
- The Container Provider should insure that suitable isolation levels are provided to guarantee data consistency for entity beans. Typically this means that an equivalent of a repeatable read or serializable isolation level should be available for applications that require a high degree of isolation.
- For entity beans with container-managed persistence, transaction isolation is managed by the data access classes that are generated by the Container Provider's tools. The tools must ensure that the management of the isolation levels performed by the data access classes will not result in conflicting isolation level requests for a resource manager within a transaction.
- Additional care must be taken if multiple enterprise beans access the same resource manager in the same transaction. Conflicts in the requested isolation levels must be avoided.

17.3.3 Enterprise Beans Using Bean-Managed Transaction Demarcation

This subsection describes the requirements for the Bean Provider of an enterprise bean with bean-managed transaction demarcation.

The enterprise bean with bean-managed transaction demarcation must be a session bean or a message-driven bean.

An instance that starts a transaction must complete the transaction before it starts a new transaction.

The Bean Provider uses the `UserTransaction` interface to demarcate transactions. All updates to the resource managers between the `UserTransaction.begin` and `UserTransaction.commit` methods are performed in a transaction. While an instance is in a transaction, the instance must not attempt to use the resource-manager specific transaction demarcation API (e.g. it must not invoke the `commit` or `rollback` method on the `java.sql.Connection` interface or on the `javax.jms.Session` interface).

A stateful session bean instance may, but is not required to, commit a started transaction before a business method returns. If a transaction has not been completed by the end of a business method, the container retains the association between the transaction and the instance across multiple client calls until the instance eventually completes the transaction.

The bean-managed transaction demarcation programming model presented to the programmer of a stateful session bean is natural because it is the same as that used by a stand-alone Java application.

A stateless session bean instance must commit a transaction before a business method or `ejbTimeout` method returns.

A message-driven bean instance must commit a transaction before a message listener method or `ejbTimeout` method returns.

The following example illustrates a business method that performs a transaction involving two database connections.

```
public class MySessionEJB implements SessionBean {
    EJBContext ejbContext;

    public void someMethod(...) {
        javax.transaction.UserTransaction ut;
        javax.sql.DataSource ds1;
        javax.sql.DataSource ds2;
        java.sql.Connection con1;
        java.sql.Connection con2;
        java.sql.Statement stmt1;
        java.sql.Statement stmt2;

        InitialContext initCtx = new InitialContext();

        // obtain con1 object and set it up for transactions
        ds1 = (javax.sql.DataSource)
            initCtx.lookup("java:comp/env/jdbc/Database1");
        con1 = ds1.getConnection();

        stmt1 = con1.createStatement();

        // obtain con2 object and set it up for transactions
        ds2 = (javax.sql.DataSource)
            initCtx.lookup("java:comp/env/jdbc/Database2");
        con2 = ds2.getConnection();

        stmt2 = con2.createStatement();

        //
        // Now do a transaction that involves con1 and con2.
        //
        ut = ejbContext.getUserTransaction();

        // start the transaction
        ut.begin();

        // Do some updates to both con1 and con2. The container
        // automatically enlists con1 and con2 with the transaction.
        stmt1.executeQuery(...);
        stmt1.executeUpdate(...);
        stmt2.executeQuery(...);
        stmt2.executeUpdate(...);
        stmt1.executeUpdate(...);
        stmt2.executeUpdate(...);

        // commit the transaction
        ut.commit();

        // release connections
        stmt1.close();
        stmt2.close();
        con1.close();
        con2.close();
    }
    ...
}
```

The following example illustrates a business method that performs a transaction involving both a database connection and a JMS connection.

```
public class MySessionEJB implements SessionBean {
    EJBContext ejbContext;

    public void someMethod(...) {
        javax.transaction.UserTransaction ut;
        javax.sql.DataSource ds;
        java.sql.Connection dcon;
        java.sql.Statement stmt;
        javax.jms.QueueConnectionFactory qcf;
        javax.jms.QueueConnection qcon;
        javax.jms.Queue q;
        javax.jms.QueueSession qsession;
        javax.jms.QueueSender qsender;
        javax.jms.Message message;

        InitialContext initCtx = new InitialContext();

        // obtain db conn object and set it up for transactions

        ds = (javax.sql.DataSource)
            initCtx.lookup("java:comp/env/jdbc/Database");
        dcon = ds.getConnection();

        stmt = dcon.createStatement();

        // obtain jms conn object and set up session for transactions
        qcf = (javax.jms.QueueConnectionFactory)
            initCtx.lookup("java:comp/env/jms/qConnFactory");
        qcon = qcf.createQueueConnection();
        qsession = qcon.createQueueSession(true,0);
        q = (javax.jms.Queue)
            initCtx.lookup("java:comp/env/jms/jmsQueue");
        qsender = qsession.createSender(q);
        message = qsession.createTextMessage();
        message.setText("some message");

        //
        // Now do a transaction that involves the two connections.
        //
        ut = ejbContext.getUserTransaction();

        // start the transaction
        ut.begin();

        // Do database updates and send message. The container
        // automatically enlists dcon and qsession with the
        // transaction.
        stmt.executeQuery(...);
        stmt.executeUpdate(...);
        stmt.executeUpdate(...);
        qsender.send(message);

        // commit the transaction
        ut.commit();
    }
}
```

```
        // release connections
        stmt.close();
        qsender.close();
        qsession.close();
        dcon.close();
        qcon.close();
    }
    ...
}
```


The following example illustrates a stateful session bean that retains a transaction across three client calls, invoked in the following order: method1, method2, and method3.^[37]

```
public class MySessionEJB implements SessionBean {
    EJBContext ejbContext;
    javax.sql.DataSource ds1;
    javax.sql.DataSource ds2;
    java.sql.Connection con1;
    java.sql.Connection con2;

    public void method1(...) {
        java.sql.Statement stmt;

        InitialContext initCtx = new InitialContext();

        // obtain user transaction interface
        ut = ejbContext.getUserTransaction();

        // start a transaction
        ut.begin();

        // make some updates on con1
        ds1 = (javax.sql.DataSource)
            initCtx.lookup("java:comp/env/jdbc/Database1");
        con1 = ds1.getConnection();
        stmt = con1.createStatement();
        stmt.executeUpdate(...);
        stmt.executeUpdate(...);

        //
        // The container retains the transaction associated with the
        // instance to the next client call (which is method2(...)).
    }

    public void method2(...) {
        java.sql.Statement stmt;

        InitialContext initCtx = new InitialContext();

        // make some updates on con2
        ds2 = (javax.sql.DataSource)
            initCtx.lookup("java:comp/env/jdbc/Database2");
        con2 = ds2.getConnection();
        stmt = con2.createStatement();
        stmt.executeUpdate(...);
        stmt.executeUpdate(...);

        // The container retains the transaction associated with the
        // instance to the next client call (which is method3(...)).
    }

    public void method3(...) {
        java.sql.Statement stmt;

        // obtain user transaction interface
```

[37] Note that the Bean Provider must use the `ejbPassivate` method here to close the Connections and set the instance variables for the Connection to null.

```
        ut = ejbContext.getUserTransaction();

        // make some more updates on con1 and con2
        stmt = con1.createStatement();
        stmt.executeUpdate(...);
        stmt = con2.createStatement();
        stmt.executeUpdate(...);

        // commit the transaction
        ut.commit();

        // release connections
        stmt.close();
        con1.close();
        con2.close();
    }
    ...
}
```

It is possible for an enterprise bean to open and close a database connection in each business method (rather than hold the connection open until the end of transaction). In the following example, if the client executes the sequence of methods (method1, method2, method2, method2, and method3), all the database updates done by the multiple invocations of method2 are performed in the scope of the same transaction, which is the transaction started in method1 and committed in method3.

```
public class MySessionEJB implements SessionBean {
    EJBContext ejbContext;
    InitialContext initCtx;

    public void method1(...) {
        java.sql.Statement stmt;

        // obtain user transaction interface
        ut = ejbContext.getUserTransaction();

        // start a transaction
        ut.begin();
    }

    public void method2(...) {
        javax.sql.DataSource ds;
        java.sql.Connection con;
        java.sql.Statement stmt;

        // open connection
        ds = (javax.sql.DataSource)
            initCtx.lookup("java:comp/env/jdbc/Database");
        con = ds.getConnection();

        // make some updates on con
        stmt = con.createStatement();
        stmt.executeUpdate(...);
        stmt.executeUpdate(...);

        // close the connection
        stmt.close();
        con.close();
    }

    public void method3(...) {
        // obtain user transaction interface
        ut = ejbContext.getUserTransaction();

        // commit the transaction
        ut.commit();
    }
    ...
}
```

17.3.3.1 getRollbackOnly and setRollbackOnly Methods

An enterprise bean with bean-managed transaction demarcation must not use the `getRollbackOnly` and `setRollbackOnly` methods of the `EJBContext` interface.

An enterprise bean with bean-managed transaction demarcation has no need to use these methods, because of the following reasons:

- An enterprise bean with bean-managed transaction demarcation can obtain the status of a transaction by using the `getStatus` method of the `javax.transaction.UserTransaction` interface.
- An enterprise bean with bean-managed transaction demarcation can rollback a transaction using the `rollback` method of the `javax.transaction.UserTransaction` interface.

17.3.4 Enterprise Beans Using Container-Managed Transaction Demarcation

This subsection describes the requirements for the Bean Provider of an enterprise bean using container-managed transaction demarcation.

The enterprise bean's business methods, message listener methods, or `ejbTimeout` method must not use any resource-manager specific transaction management methods that would interfere with the container's demarcation of transaction boundaries. For example, the enterprise bean methods must not use the following methods of the `java.sql.Connection` interface: `commit`, `setAutoCommit`, and `rollback`; or the following methods of the `javax.jms.Session` interface: `commit` and `rollback`.

The enterprise bean's business methods, message listener methods, or `ejbTimeout` method must not attempt to obtain or use the `javax.transaction.UserTransaction` interface.

The following is an example of a business method in an enterprise bean with container-managed transaction demarcation. The business method updates two databases using JDBC™ connections. The container provides transaction demarcation per the Application Assembler's instructions.

```
public class MySessionEJB implements SessionBean {
    EJBContext ejbContext;

    public void someMethod(...) {
        java.sql.Connection con1;
        java.sql.Connection con2;
        java.sql.Statement stmt1;
        java.sql.Statement stmt2;

        // obtain con1 and con2 connection objects
        con1 = ...;
        con2 = ...;

        stmt1 = con1.createStatement();
        stmt2 = con2.createStatement();

        //
        // Perform some updates on con1 and con2. The container
        // automatically enlists con1 and con2 with the container-
        // managed transaction.
        //
        stmt1.executeQuery(...);
        stmt1.executeUpdate(...);

        stmt2.executeQuery(...);
        stmt2.executeUpdate(...);

        stmt1.executeUpdate(...);
        stmt2.executeUpdate(...);

        // release connections
        con1.close();
        con2.close();
    }
    ...
}
```

17.3.4.1 javax.ejb.SessionSynchronization Interface

A stateful session bean with container-managed transaction demarcation can optionally implement the `javax.ejb.SessionSynchronization` interface. The use of the `SessionSynchronization` interface is described in Subsection 7.5.4.

17.3.4.2 javax.ejb.EJBContext.setRollbackOnly Method

An enterprise bean with container-managed transaction demarcation can use the `setRollbackOnly` method of its `EJBContext` object to mark the transaction such that the transaction can never commit. Typically, an enterprise bean marks a transaction for rollback to protect data integrity before throwing an application exception, because application exceptions do not automatically cause the container to rollback the transaction.

For example, an `AccountTransfer` bean which debits one account and credits another account could mark a transaction for rollback if it successfully performs the debit operation, but encounters a failure during the credit operation.

17.3.4.3 `javax.ejb.EJBContext.getRollbackOnly` method

An enterprise bean with container-managed transaction demarcation can use the `getRollbackOnly` method of its `EJBContext` object to test if the current transaction has been marked for rollback. The transaction might have been marked for rollback by the enterprise bean itself, by other enterprise beans, or by other components (outside of the EJB specification scope) of the transaction processing infrastructure.

17.3.5 Use of JMS APIs in Transactions

The Bean Provider must not make use of the JMS request/reply paradigm (sending of a JMS message, followed by the synchronous receipt of a reply to that message) within a single transaction. Because a JMS message is not delivered to its final destination until the transaction commits, the receipt of the reply within the same transaction will never take place.

Because the container manages the transactional enlistment of JMS sessions on behalf of a bean, the parameters of the `createSession(boolean transacted, int acknowledgeMode)`, `createQueueSession(boolean transacted, int acknowledgeMode)` and `createTopicSession(boolean transacted, int acknowledgeMode)` methods are ignored. It is recommended that the Bean Provider specify that a session is transacted, but provide 0 for the value of the acknowledgment mode.

The Bean Provider should not use the JMS `acknowledge` method either within a transaction or within an unspecified transaction context. Message acknowledgment in an unspecified transaction context is handled by the container. Section 17.6.5 describes some of the techniques that the container can use for the implementation of a method invocation with an unspecified transaction context.

17.3.6 Declaration in Deployment Descriptor

The Bean Provider of a session bean or a message-driven bean must use the `transaction-type` element to declare whether the session bean or message-driven bean is of the bean-managed or container-managed transaction demarcation type. (See Chapter 23 for information about the deployment descriptor.)

The `transaction-type` element is not supported for entity beans because all entity beans must use container-managed transaction demarcation.

The Bean Provider of an enterprise bean with container-managed transaction demarcation may optionally specify the transaction attributes for the enterprise bean's methods. See Subsection 17.4.1.

17.4 Application Assembler's Responsibilities

This section describes the view and responsibilities of the Application Assembler.

There is no mechanism for an Application Assembler to affect enterprise beans with bean-managed transaction demarcation. The Application Assembler must not define transaction attributes for an enterprise bean with bean-managed transaction demarcation.

The Application Assembler can use the transaction attribute mechanism described below to manage transaction demarcation for enterprise beans using container-managed transaction demarcation.

17.4.1 Transaction Attributes

Note: The transaction attributes may be specified either by the Bean Provider or by the Application Assembler.

A transaction attribute is a value associated with a method of a session or entity bean's home or component interface, with a method of a stateless session bean's web service endpoint interface, with a message listener method of a message-driven bean, or with an `ejbTimeout` method. The transaction attribute specifies how the container must manage transactions for a method when a client invokes the method via the enterprise bean's home, component, or web service endpoint interface or when the method is invoked as the result of the arrival of a message or a timer event.

The transaction attribute must be specified for the following methods:

- For a session bean, the transaction attributes must be specified for the methods defined in the bean's component interface and all the direct and indirect superinterfaces of the component interface, excluding the methods of the `javax.ejb.EJBObject` or `javax.ejb.EJBLocalObject` interface; and for the `ejbTimeout` method. Transaction attributes must not be specified for the methods of a session bean's home interface.
- For a stateless session bean with a web service endpoint interface, the transaction attributes must be specified for the methods defined in the bean's web service endpoint interface.
- For an entity bean, the transaction attributes must be specified for the methods defined in the bean's component interface and all the direct and indirect superinterfaces of the component interface, excluding the `getEJBHome`, `getEJBLocalHome`, `getHandle`, `getPrimaryKey`, and `isIdentical` methods; for the methods defined in the bean's home interface and all the direct and indirect superinterfaces of the home interface, excluding the `getEJBMetaData` and `getHomeHandle` methods specific to the remote home interface; and for the `ejbTimeout` method.
- For a message-driven bean, the transaction attributes must be specified for the methods defined in the bean's message listener interface and for the `ejbTimeout` method.

Providing the transaction attributes for an enterprise bean is an optional requirement for the Application Assembler, because, for a given enterprise bean, the Application Assembler must either specify a value of the transaction attribute for *all* the methods for which a transaction attribute must be specified, or the Assembler must specify *none*. If the transaction attributes are not specified for the methods of an enterprise bean, the Deployer will have to specify them.

Enterprise JavaBeans defines the following values for the transaction attribute:

- `NotSupported`
- `Required`
- `Supports`
- `RequiresNew`
- `Mandatory`
- `Never`

Refer to Subsection 17.6.2 for the specification of how the value of the transaction attribute affects the transaction management performed by the container.

For a message-driven bean's message listener interface, only the `Required` and `NotSupported` transaction attributes may be used.

For a stateless session bean's web service endpoint interface, only the `Required`, `RequiresNew`, `Supports`, `Never`, and `NotSupported` attributes may be used.

For entity beans that use EJB 2.0 container-managed persistence, only the `Required`, `RequiresNew`, or `Mandatory` transaction attributes should be used for the methods defined in the bean's component interface and all the direct and indirect superinterfaces of the component interface, excluding the `getEJBHome`, `getEJBLocalHome`, `getHandle`, `getPrimaryKey`, and `isIdentical` methods; and for the methods defined in the bean's home interface and all the direct and indirect superinterfaces of the home interface, excluding the `getEJBMetaData` and `getHomeHandle` methods specific to the remote home interface.

The Bean Provider and Application Assembler must exercise caution when using the `RequiresNew` transaction attribute with the navigation of container-managed relationships. If higher levels of isolation are used, navigating a container-managed relationship in a new transaction context may result in deadlock.

Containers may *optionally* support the use of the `NotSupported`, `Supports`, and `Never` transaction attributes for the methods of entity beans with container-managed persistence. However, entity beans with container-managed persistence that use these transaction attributes will not be portable.

Containers may optionally support the use of the `NotSupported`, `Supports`, and `Never` transaction attributes for the methods of entity beans with container-managed persistence because the use of these transaction modes may be needed to make use of container-managed persistence with non-transactional data stores. In general, however, the Bean Provider and Application Assembler should avoid use of the `NotSupported`, `Supports`, and `Never` transaction attributes for the methods of entity beans with container-managed persistence because it may lead to inconsistent results or to the inconsistent and/or to the partial updating of persistent state and relationships in the event of concurrent use.

For an enterprise bean's `ejbTimeout` method only the `RequiresNew` and `NotSupported` transaction attributes may be used.

If an enterprise bean implements the `javax.ejb.SessionSynchronization` interface, only the following values may be used for the transaction attributes of the bean's methods: `Required`, `RequiresNew`, or `Mandatory`.

The above restriction is necessary to ensure that the enterprise bean is invoked only in a transaction. If the bean were invoked without a transaction, the container would not be able to send the transaction synchronization calls.

The tools used by the Application Assembler can determine if the bean implements the `javax.ejb.SessionSynchronization` interface, for example, by using the Java reflection API on the enterprise bean's class.

The following is the description of the deployment descriptor rules that the Application Assembler uses to specify transaction attributes for the methods of the session and entity beans' home, component, and web service endpoint interfaces; message-driven beans' message listener methods; and `ejbTimeout` methods. (See Section 23.5 for the complete syntax of the deployment descriptor.)

The Application Assembler uses the `container-transaction` elements to define the transaction attributes for the methods of session and entity bean home, component, and web service endpoint interfaces; for the message listener methods of message-driven beans; and for the `ejbTimeout` methods of enterprise beans. Each `container-transaction` element consists of a list of one or more method elements, and the `trans-attribute` element. The `container-transaction` element specifies that all the listed methods are assigned the specified transaction attribute value. It is required that all the methods specified in a single `container-transaction` element be methods of the same enterprise bean.

The method element uses the `ejb-name`, `method-name`, and `method-params` elements to denote one or more methods of an enterprise bean's home and component interfaces. There are three legal styles of composing the method element:

Style 1:

```
<method>
  <ejb-name>EJBNAME</ejb-name>
  <method-name>*</method-name>
</method>
```

This style is used to specify a default value of the transaction attribute for the methods for which there is no Style 2 or Style 3 element specified. There must be at most one `container-transaction` element that uses the Style 1 method element for a given enterprise bean.

Style 2:

```
<method>
  <ejb-name>EJBNAME</ejb-name>
  <method-name>METHOD</method-name>
</method>
```

This style is used for referring to a specified method of a home, component, message listener, or web service endpoint interface of the specified enterprise bean or an `ejbTimeout` method. If there are multiple methods with the same overloaded name, this style refers to all the methods with the same name. There must be at most one `container-transaction` element that uses the Style 2 method element for a given method name. If there is also a `container-transaction` element that uses Style 1 element for the same bean, the value specified by the Style 2 element takes precedence.

Style 3:

```
<method>
  <ejb-name>EJBNAME</ejb-name>
  <method-name>METHOD</method-name>
  <method-params>
    <method-param>PARAMETER_1</method-param>
    ...
    <method-param>PARAMETER_N</method-param>
  </method-params>
</method>
```

This style is used to refer to a single method within a set of methods with an overloaded name. If there is also a `container-transaction` element that uses the Style 2 element for the method name, or the Style 1 element for the bean, the value specified by the Style 3 element takes precedence.

The optional `method-intf` element can be used to differentiate between methods with the same name and signature that are multiply defined across the component, home, and/or web service endpoint interfaces.

The following is an example of the specification of the transaction attributes in the deployment descriptor. The `updatePhoneNumber` method of the `EmployeeRecord` enterprise bean is assigned the transaction attribute `Mandatory`; all other methods of the `EmployeeRecord` bean are assigned the attribute `Required`. All the methods of the enterprise bean `AardvarkPayroll` are assigned the attribute `RequiresNew`.

```
<ejb-jar>
  ...
  <assembly-descriptor>
    ...
    <container-transaction>
      <method>
        <ejb-name>EmployeeRecord</ejb-name>
        <method-name>*</method-name>
      </method>
      <trans-attribute>Required</trans-attribute>
    </container-transaction>

    <container-transaction>
      <method>
        <ejb-name>EmployeeRecord</ejb-name>
        <method-name>updatePhoneNumber</method-name>
      </method>
      <trans-attribute>Mandatory</trans-attribute>
    </container-transaction>

    <container-transaction>
      <method>
        <ejb-name>AardvarkPayroll</ejb-name>
        <method-name>*</method-name>
      </method>
      <trans-attribute>RequiresNew</trans-attribute>
    </container-transaction>
  </assembly-descriptor>
</ejb-jar>
```

17.5 Deployer's Responsibilities

The Deployer is responsible for ensuring that the methods of the deployed enterprise beans with container-managed transaction demarcation have been assigned a transaction attribute. If the transaction attributes have not been assigned previously by the Assembler, they must be assigned by the Deployer.

17.6 Container Provider Responsibilities

This section defines the responsibilities of the Container Provider.

Every client method invocation on a session or entity bean via the bean's home, component, or web service endpoint interface and every invocation of a message listener method on a message-driven bean is interposed by the container, and every connection to a resource manager used by an enterprise bean is obtained via the container. This managed execution environment allows the container to affect the enterprise bean's transaction management.

This does not imply that the container must interpose on every resource manager access performed by the enterprise bean. Typically, the container interposes only on the resource manager connection factory (e.g. a JDBC data source) JNDI look up by registering the container-specific implementation of the resource manager connection factory object. The resource manager connection factory object allows the container to obtain the `javax.transaction.xa.XAResource` interface as described in the JTA specification and pass it to the transaction manager. After the set up is done, the enterprise bean communicates with the resource manager without going through the container.

17.6.1 Bean-Managed Transaction Demarcation

This subsection defines the container's responsibilities for the transaction management of enterprise beans with bean-managed transaction demarcation.

Note that only session and message-driven beans can be used with bean-managed transaction demarcation. A Bean Provider is not allowed to provide an entity bean with bean-managed transaction demarcation.

The container must manage client invocations to an enterprise bean instance with bean-managed transaction demarcation as follows. When a client invokes a business method via one of the enterprise bean's client view interfaces, the container suspends any transaction that may be associated with the client request. If there is a transaction associated with the instance (this would happen if a stateful session bean instance started the transaction in some previous business method), the container associates the method execution with this transaction.

The container must make the `javax.transaction.UserTransaction` interface available to the enterprise bean's business method, message listener method, or `ejbTimeout` method via the `javax.ejb.EJBContext` interface and under the environment entry `java:comp/UserTransaction`. When an instance uses the `javax.transaction.UserTransaction` interface to demarcate a transaction, the container must enlist all the resource managers used by the instance between the `begin` and `commit`—or `rollback`— methods with the transaction. When the instance attempts to commit the transaction, the container is responsible for the global coordination of the transaction commit^[38].

In the case of a *stateful* session bean, it is possible that the business method that started a transaction completes without committing or rolling back the transaction. In such a case, the container must retain the association between the transaction and the instance across multiple client calls until the instance commits or rolls back the transaction. When the client invokes the next business method, the container must invoke the business method in this transaction context.

[38] The container typically relies on a transaction manager that is part of the EJB server to perform the two-phase commit across all the enlisted resource managers. If only a single resource manager is involved in the transaction and the deployment descriptor indicates that connection sharing may be used, the container may use the local transaction optimization. See [9] and [12] for further discussion.

If a *stateless* session bean instance starts a transaction in a business method, it must commit the transaction before the business method returns. The container must detect the case in which a transaction was started, but not completed, in the business method, and handle it as follows:

- Log this as an application error to alert the System Administrator.
- Roll back the started transaction.
- Discard the instance of the session bean.
- Throw the `java.rmi.RemoteException` to the client if the client is a remote client, or throw the `javax.ejb.EJBException` if the client is a local client.

If a message-driven bean instance starts a transaction in a message listener method, it must commit the transaction before the message listener method returns. The container must detect the case in which a transaction was started, but not completed, in a message listener method, and handle it as follows:

- Log this as an application error to alert the System Administrator.
- Roll back the started transaction.
- Discard the instance of the message-driven bean.

If a session bean or message-driven bean instance starts a transaction in an `ejbTimeout` method, it must commit the transaction before the `ejbTimeout` method returns. The container must detect the case in which a transaction was started, but not completed, in an `ejbTimeout` method, and handle it as follows:

- Log this as an application error to alert the System Administrator.
- Roll back the started transaction.
- Discard the instance of the bean.

The actions performed by the container for an instance with bean-managed transaction are summarized by the following table. T1 is a transaction associated with a client request, T2 is a transaction that is currently associated with the instance (i.e. a transaction that was started but not completed by a previous business method).

Table 13 Container's Actions for Methods of Beans with Bean-Managed Transaction

Client's transaction	Transaction currently associated with instance	Transaction associated with the method
none	none	none
T1	none	none
none	T2	T2

Table 13 Container's Actions for Methods of Beans with Bean-Managed Transaction

Client's transaction	Transaction currently associated with instance	Transaction associated with the method
T1	T2	T2

The following items describe each entry in the table:

- If the client request is not associated with a transaction and the instance is not associated with a transaction, or if the bean is a message-driven bean, the container invokes the instance with an unspecified transaction context.
- If the client request is associated with a transaction T1, and the instance is not associated with a transaction, the container suspends the client's transaction association and invokes the method with an unspecified transaction context. The container resumes the client's transaction association (T1) when the method completes. This case can never happen for a message-driven bean or for the invocation of a method on the web service endpoint interface of a stateless session bean.
- If the client request is not associated with a transaction and the instance is already associated with a transaction T2, the container invokes the instance with the transaction that is associated with the instance (T2). This case can never happen for a stateless session bean or a message-driven bean.
- If the client is associated with a transaction T1, and the instance is already associated with a transaction T2, the container suspends the client's transaction association and invokes the method with the transaction context that is associated with the instance (T2). The container resumes the client's transaction association (T1) when the method completes. This case can never happen for a stateless session bean or a message-driven bean.

The container must allow the enterprise bean instance to serially perform several transactions in a method.

When an instance attempts to start a transaction using the `begin` method of the `javax.transaction.UserTransaction` interface while the instance has not committed the previous transaction, the container must throw the `javax.transaction.NotSupportedException` in the `begin` method.

The container must throw the `java.lang.IllegalStateException` if an instance of a bean with bean-managed transaction demarcation attempts to invoke the `setRollbackOnly` or `getRollbackOnly` method of the `javax.ejb.EJBContext` interface.

17.6.2 Container-Managed Transaction Demarcation for Session and Entity Beans

The container is responsible for providing the transaction demarcation for the session and entity beans that the Bean Provider declared with container-managed transaction demarcation. For these enterprise beans, the container must demarcate transactions as specified in the deployment descriptor by the Application Assembler. (See Chapter 23 for more information about the deployment descriptor.)

The following subsections define the responsibilities of the container for managing the invocation of an enterprise bean business method when the method is invoked via the enterprise bean's home, component, or web service endpoint interface. The container's responsibilities depend on the value of the transaction attribute.

17.6.2.1 NotSupported

The container invokes an enterprise bean method whose transaction attribute is set to `NotSupported` with an unspecified transaction context.

If a client calls with a transaction context, the container suspends the association of the transaction context with the current thread before invoking the enterprise bean's business method. The container resumes the suspended association when the business method has completed. The suspended transaction context of the client is not passed to the resource managers or other enterprise bean objects that are invoked from the business method.

If the business method invokes other enterprise beans, the container passes no transaction context with the invocation.

Refer to Subsection 17.6.5 for more details of how the container can implement this case.

17.6.2.2 Required

The container must invoke an enterprise bean method whose transaction attribute is set to `Required` with a valid transaction context.

If a client invokes the enterprise bean's method while the client is associated with a transaction context, the container invokes the enterprise bean's method in the client's transaction context.

If the client invokes the enterprise bean's method while the client is not associated with a transaction context, the container automatically starts a new transaction before delegating a method call to the enterprise bean business method. The container automatically enlists all the resource managers accessed by the business method with the transaction. If the business method invokes other enterprise beans, the container passes the transaction context with the invocation. The container attempts to commit the transaction when the business method has completed. The container performs the commit protocol before the method result is sent to the client.

17.6.2.3 Supports

The container invokes an enterprise bean method whose transaction attribute is set to `Supports` as follows.

- If the client calls with a transaction context, the container performs the same steps as described in the `Required` case.
- If the client calls without a transaction context, the container performs the same steps as described in the `NotSupported` case.

The `Supports` transaction attribute must be used with caution. This is because of the different transactional semantics provided by the two possible modes of execution. Only the enterprise beans that will execute correctly in both modes should use the `Supports` transaction attribute.

17.6.2.4 RequiresNew

The container must invoke an enterprise bean method whose transaction attribute is set to `RequiresNew` with a new transaction context.

If the client invokes the enterprise bean's method while the client is not associated with a transaction context, the container automatically starts a new transaction before delegating a method call to the enterprise bean business method. The container automatically enlists all the resource managers accessed by the business method with the transaction. If the business method invokes other enterprise beans, the container passes the transaction context with the invocation. The container attempts to commit the transaction when the business method has completed. The container performs the commit protocol before the method result is sent to the client.

If a client calls with a transaction context, the container suspends the association of the transaction context with the current thread before starting the new transaction and invoking the business method. The container resumes the suspended transaction association after the business method and the new transaction have been completed.

17.6.2.5 Mandatory

The container must invoke an enterprise bean method whose transaction attribute is set to `Mandatory` in a client's transaction context. The client is required to call with a transaction context.

- If the client calls with a transaction context, the container performs the same steps as described in the `Required` case.
- If the client calls without a transaction context, the container throws the `javax.transaction.TransactionRequiredException` exception if the client is a remote client, or the `javax.ejb.TransactionRequiredLocalException` if the client is a local client.

17.6.2.6 Never

The container invokes an enterprise bean method whose transaction attribute is set to `Never` without a transaction context defined by the EJB specification. The client is required to call without a transaction context.

- If the client calls with a transaction context, the container throws the `java.rmi.RemoteException` exception if the client is a remote client, or the `javax.ejb.EJBException` if the client is a local client.
- If the client calls without a transaction context, the container performs the same steps as described in the `NotSupported` case.

17.6.2.7 Transaction Attribute Summary

The following table provides a summary of the transaction context that the container passes to the business method and resource managers used by the business method, as a function of the transaction attribute and the client's transaction context. T1 is a transaction passed with the client request, while T2 is a transaction initiated by the container.

Table 14

Transaction Attribute Summary

Transaction attribute	Client's transaction	Transaction associated with business method	Transaction associated with resource managers
NotSupported	none	none	none
	T1	none	none
Required	none	T2	T2
	T1	T1	T1
Supports	none	none	none
	T1	T1	T1
RequiresNew	none	T2	T2
	T1	T2	T2
Mandatory	none	error	N/A
	T1	T1	T1
Never	none	none	none
	T1	error	N/A

If the enterprise bean's business method invokes other enterprise beans via their home and component interfaces, the transaction indicated in the column "Transaction associated with business method" will be passed as part of the client context to the target enterprise bean.

See Subsection 17.6.5 for how the container handles the "none" case in Table 14.

17.6.2.8 Handling of `setRollbackOnly` Method

The container must handle the `EJBContext.setRollbackOnly` method invoked from a business method executing with the `Required`, `RequiresNew`, or `Mandatory` transaction attribute as follows:

- The container must ensure that the transaction will never commit. Typically, the container instructs the transaction manager to mark the transaction for rollback.
- If the container initiated the transaction immediately before dispatching the business method to the instance (as opposed to the transaction being inherited from the caller), the container must note that the instance has invoked the `setRollbackOnly` method. When the business method invocation completes, the container must roll back rather than commit the transaction. If the business method has returned normally or with an application exception, the container must pass the method result or the application exception to the client after the container performed the rollback.

The container must throw the `java.lang.IllegalStateException` if the `EJBContext.setRollbackOnly` method is invoked from a business method executing with the `Supports`, `NotSupported`, or `Never` transaction attribute.

17.6.2.9 Handling of `getRollbackOnly` Method

The container must handle the `EJBContext.getRollbackOnly` method invoked from a business method executing with the `Required`, `RequiresNew`, or `Mandatory` transaction attribute.

The container must throw the `java.lang.IllegalStateException` if the `EJBContext.getRollbackOnly` method is invoked from a business method executing with the `Supports`, `NotSupported`, or `Never` transaction attribute.

17.6.2.10 Handling of `getUserTransaction` Method

If an instance of an enterprise bean with container-managed transaction demarcation attempts to invoke the `getUserTransaction` method of the `EJBContext` interface, the container must throw the `java.lang.IllegalStateException`.

17.6.2.11 `javax.ejb.SessionSynchronization` Callbacks

If a session bean class implements the `javax.ejb.SessionSynchronization` interface, the container must invoke the `afterBegin`, `beforeCompletion`, and `afterCompletion` callbacks on the instance as part of the transaction commit protocol.

The container invokes the `afterBegin` method on an instance before it invokes the first business method in a transaction.

The container invokes the `beforeCompletion` method to give the enterprise bean instance the last chance to cause the transaction to rollback. The instance may cause the transaction to roll back by invoking the `EJBContext.setRollbackOnly` method.

The container invokes the `afterCompletion(Boolean committed)` method after the completion of the transaction commit protocol to notify the enterprise bean instance of the transaction outcome.

17.6.3 Container-Managed Transaction Demarcation for Message-Driven Beans

The container is responsible for providing the transaction demarcation for the message-driven beans that the Bean Provider declared as with container-managed transaction demarcation. For these enterprise beans, the container must demarcate transactions as specified in the deployment descriptor by the Application Assembler. (See Chapter 23 for more information about the deployment descriptor.)

The following subsections define the responsibilities of the container for managing the invocation of a message-driven bean's message listener method. The container's responsibilities depend on the value of the transaction attribute.

Only the `NotSupported` and `Required` transaction attributes may be used for message-driven bean message listener methods. The use of the other transaction attributes is not meaningful for message-driven bean message listener methods because there is no pre-existing client transaction context (`RequiresNew`, `Supports`) and no client to handle exceptions (`Mandatory`, `Never`).

17.6.3.1 NotSupported

The container invokes a message-driven bean message listener method whose transaction attribute is set to `NotSupported` with an unspecified transaction context.

If the message listener method invokes other enterprise beans, the container passes no transaction context with the invocation.

17.6.3.2 Required

The container must invoke a message-driven bean message listener method whose transaction attribute is set to `Required` with a valid transaction context. The resource managers accessed by the message listener method within the transaction are enlisted with the transaction. If the message listener method invokes other enterprise beans, the container passes the transaction context with the invocation. The container attempts to commit the transaction when the message listener method has completed.

Messaging systems may differ in quality of service with regard to reliability and transactionality of the dequeuing of messages.

The requirement for JMS are as follows:

A transaction must be started before the dequeuing of the JMS message and, hence, before the invocation of the message-driven bean's `onMessage` method. The resource manager associated with the arriving message is enlisted with the transaction as well as all the resource managers accessed by the `onMessage` method within the transaction. If the `onMessage` method invokes other enterprise beans, the container passes the transaction context with the invocation. The transaction is committed when the `onMessage` method has completed. If the `onMessage` method does not successfully complete or the transaction is rolled back, message redelivery semantics apply.

17.6.3.3 Handling of `setRollbackOnly` Method

The container must handle the `EJBContext.setRollbackOnly` method invoked from a message listener method executing with the `Required` transaction attribute as follows:

- The container must ensure that the transaction will never commit. Typically, the container instructs the transaction manager to mark the transaction for rollback.
- The container must note that the instance has invoked the `setRollbackOnly` method. When the method invocation completes, the container must roll back rather than commit the transaction.

The container must throw and log the `java.lang.IllegalStateException` if the `EJBContext.setRollbackOnly` method is invoked from a message listener method executing with the `NotSupported` transaction attribute.

17.6.3.4 Handling of `getRollbackOnly` Method

The container must handle the `EJBContext.getRollbackOnly()` method invoked from a message listener method executing with the `Required` transaction attribute.

The container must throw and log the `java.lang.IllegalStateException` if the `EJBContext.getRollbackOnly` method is invoked from a message listener method executing with the `NotSupported` transaction attribute.

17.6.3.5 Handling of `getUserTransaction` Method

If an instance of a message-driven bean with container-managed transaction demarcation attempts to invoke the `getUserTransaction` method of the `EJBContext` interface, the container must throw and log the `java.lang.IllegalStateException`.

17.6.4 Local Transaction Optimization

The container may use a local transaction optimization for enterprise beans whose deployment descriptor indicates that connections to a resource manager are shareable (see Section 20.5.1.2, “Declaration of Resource Manager Connection Factory References in Deployment Descriptor”). The container manages the use of the local transaction optimization transparent to the application.

The container may use the optimization for transactions initiated by the container for a bean with container-managed transaction demarcation and for transactions initiated by a bean with bean-managed transaction demarcation with the `UserTransaction` interface. The container cannot apply the optimization for transactions imported from a different container.

The use of local transaction optimization approach is discussed in [9] and [12].

17.6.5 Handling of Methods that Run with “an unspecified transaction context”

The term “an unspecified transaction context” is used in the EJB specification to refer to the cases in which the EJB architecture does not fully define the transaction semantics of an enterprise bean method execution.

This includes the following cases:

- The execution of a method of an enterprise bean with container-managed transaction demarcation for which the value of the transaction attribute is `NotSupported`, `Never`, or `Supports`^[39].
- The execution of the `ejbCreate<METHOD>`, `ejbRemove`, `ejbPassivate`, and `ejbActivate` methods of a session bean with container-managed transaction demarcation.
- The execution of the `ejbCreate` and `ejbRemove` methods of a message-driven bean with container-managed transaction demarcation.

The EJB specification does not prescribe how the container should manage the execution of a method with an unspecified transaction context—the transaction semantics are left to the container implementation. Some techniques for how the container may choose to implement the execution of a method with an unspecified transaction context are as follows (the list is not inclusive of all possible strategies):

- The container may execute the method and access the underlying resource managers without a transaction context.
- The container may treat each call of an instance to a resource manager as a single transaction (e.g. the container may set the auto-commit option on a JDBC connection).
- The container may merge multiple calls of an instance to a resource manager into a single transaction.
- The container may merge multiple calls of an instance to multiple resource managers into a single transaction.
- If an instance invokes methods on other enterprise beans, and the invoked methods are also designated to run with an unspecified transaction context, the container may merge the resource manager calls from the multiple instances into a single transaction.
- Any combination of the above.

Since the enterprise bean does not know which technique the container implements, the enterprise bean must be written conservatively not to rely on any particular container behavior.

[39] For the `Supports` attribute, the handling described in this section applies only to the case when the client calls without a transaction context.

A failure that occurs in the middle of the execution of a method that runs with an unspecified transaction context may leave the resource managers accessed from the method in an unpredictable state. The EJB architecture does not define how the application should recover the resource managers' state after such a failure.

17.7 Access from Multiple Clients in the Same Transaction Context

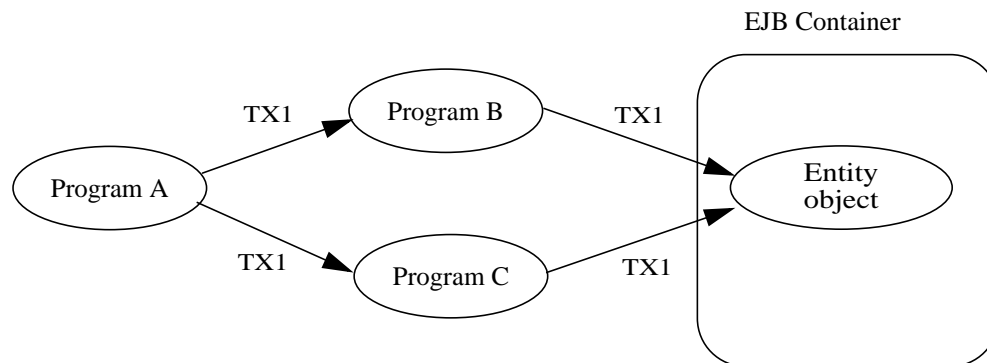
This section describes a more complex distributed transaction scenario, and specifies the container's behavior required for this scenario.

17.7.1 Transaction "Diamond" Scenario with an Entity Object

An entity object may be accessed by multiple clients in the same transaction. For example, program A may start a transaction, call program B and program C in the transaction context, and then commit the transaction. If programs B and C access the same entity object, the topology of the transaction creates a diamond.

Figure 73

Transaction Diamond Scenario with Entity Object



An example (not realistic in practice) is a client program that tries to perform two purchases at two different stores within the same transaction. At each store, the program that is processing the client's purchase request debits the client's bank account.

It is difficult to implement an EJB server that handles the case in which programs B and C access an entity object through different network paths. This case is challenging because many EJB servers implement the EJB container as a collection of multiple processes, running on the same or multiple machines. Each client is typically connected to a single process. If clients B and C connect to different EJB container processes, and both B and C need to access the same entity object in the same transaction, the issue is how the container can make it possible for B and C to see a consistent state of the entity object within the same transaction^[40].

The above example illustrates a simple diamond. We use the term diamond to refer to any distributed transaction scenario in which an entity object is accessed in the same transaction through multiple network paths.

Note that in the diamond scenario the clients B and C access the entity object serially. Concurrent access to an entity object in the same transaction context would be considered an application programming error, and it would be handled in a container-specific way.

Note that the issue of handling diamonds is not unique to the EJB architecture. This issue exists in all distributed transaction processing systems.

The following subsections define the responsibilities of the EJB Roles when handling distributed transaction topologies that may lead to a diamond involving an entity object.

17.7.2 Container Provider's Responsibilities

This Subsection specifies the EJB container's responsibilities with respect to the diamond case involving an entity object.

The EJB specification requires that the container provide support for local diamonds. In a local diamond, components A, B, C, and D are deployed in the same EJB container.

The EJB specification does not require an EJB container to support distributed diamonds. In a distributed diamond, a target entity object is accessed from multiple clients in the same transaction through multiple network paths, and the clients (programs B and C) are not enterprise beans deployed in the same EJB container as the target entity object.

If the Container Provider chooses not to support distributed diamonds, and if the container can detect that a client invocation would lead to a diamond, the container should throw the `java.rmi.RemoteException` to the client.

If the Container Provider chooses to support distributed diamonds, it should provide a consistent view of the entity state within a transaction. The Container Provider can implement the support in several ways. (The options that follow are illustrative, not prescriptive.)

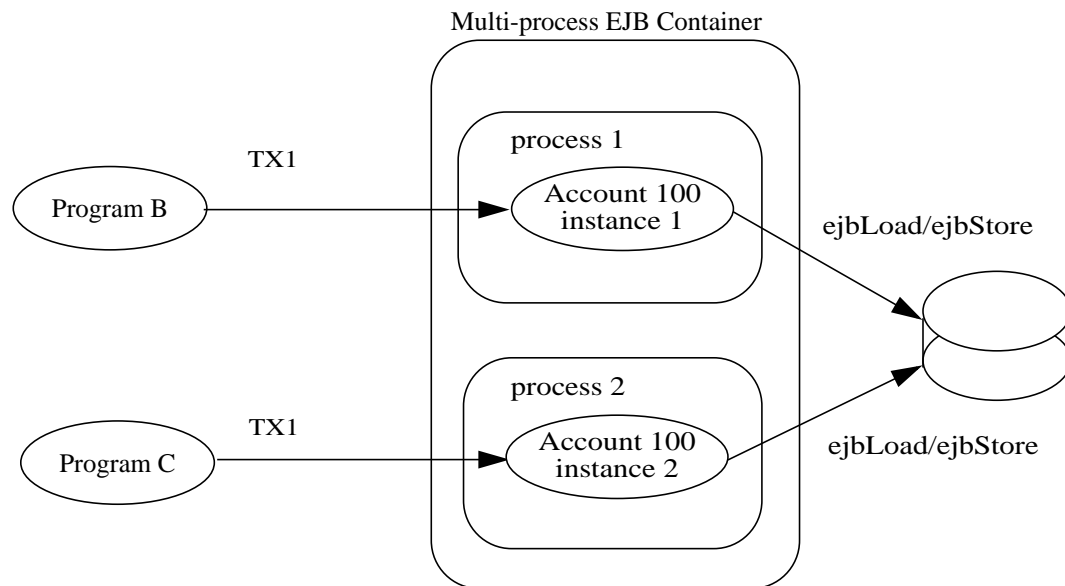
- Always instantiate the entity bean instance for a given entity object in the same process, and route all clients' requests to this process. Within the process, the container routes all the requests within the same transaction to the same enterprise bean instance.

[40] This diamond problem applies only to the case when B and C are in the same transaction.

- Instantiate the entity bean instance for a given entity object in multiple processes, and use the `ejbStore` and `ejbLoad` methods to synchronize the state of the instances within the same transaction. For example, the container can issue `ejbStore` after each business method, and issue `ejbLoad` before the start of the next business method. This technique ensures that the instance used by a one client sees the updates done by other clients within the same transaction.

An illustration of the second approach follows. The illustration is illustrative, not prescriptive for the container implementors.

Figure 74 Handling of Diamonds by a Multi-Process Container



Program B makes a call to an entity object representing Account 100. The request is routed to an instance in process 1. The container invokes `ejbLoad` on the instance. The instance loads the state from the database in the `ejbLoad` method. The instance updates the state in the business method. When the method completes, the container invokes `ejbStore`. The instance writes the updated state to the database in the `ejbStore` method.

Now program C makes a call to the same entity object in the same transaction. The request is routed to a different process (2). The container invokes `ejbLoad` on the instance. The instance loads the state from the database in the `ejbLoad` method. The loaded state was written by the instance in process 1. The instance updates the state in the business method. When the method completes, the container invokes `ejbStore`. The instance writes the updated state to the database in the `ejbStore` method.

In the above scenario, the container presents the business methods operating on the entity object Account 100 with a consistent view of the entity object's state within the transaction.

Another implementation of the EJB container might avoid calling `ejbLoad` and `ejbStore` on each business method by using a distributed lock manager.

17.7.3 Bean Provider's Responsibilities

This Subsection specifies the Bean Provider's responsibilities with respect to the diamond case involving an entity object.

The diamond case is transparent to the Bean Provider—the Bean Provider does not have to code the enterprise bean differently for the bean to participate in a diamond. Any solution to the diamond problem implemented by the container is transparent to the bean and does not change the semantics of the bean.

17.7.4 Application Assembler and Deployer's Responsibilities

This Subsection specifies the Application Assembler and Deployer's responsibilities with respect to the diamond case involving an entity object.


The Application Assembler and Deployer should be aware that distributed diamonds might occur. In general, the Application Assembler should try to avoid creating unnecessary distributed diamonds.

If a distributed diamond is necessary, the Deployer should advise the container (using a container-specific API) that an entity objects of the entity bean may be involved in distributed diamond scenarios.

17.7.5 Transaction Diamonds involving Session Objects

While it is illegal for two clients to access the same session object, it is possible for applications that use session beans to encounter the diamond case. For example, program A starts a transaction and then invokes two different session objects.

Figure 75 Transaction Diamond Scenario with a Session Bean



If the session bean instances cache the same data item (e.g. the current balance of Account 100) across method invocations in the same transaction, most likely the program is going to produce incorrect results.

The problem may exist regardless of whether the two session objects are the same or different session beans. The problem may exist (and may be harder to discover) if there are intermediate objects between the transaction initiator and the session objects that cache the data.

There are no requirements for the Container Provider because it is impossible for the container to detect this problem.

The Bean Provider and Application Assembler must avoid creating applications that would result in inconsistent caching of data in the same transaction by multiple session objects.

Exception Handling

18.1 Overview and Concepts

18.1.1 Application Exceptions

An *application exception* is an exception defined in the `throws` clause of a method of an enterprise bean's home, component, message listener, or web service endpoint interface, other than the `java.rmi.RemoteException`.

Enterprise bean business methods use application exceptions to inform the client of abnormal application-level conditions, such as unacceptable values of the input arguments to a business method. A client can typically recover from an application exception. Application exceptions are not intended for reporting system-level problems.

For example, the `Account` enterprise bean may throw an application exception to report that a debit operation cannot be performed because of an insufficient balance. The `Account` bean should not use an application exception to report, for example, the failure to obtain a database connection.

The `javax.ejb.CreateException`, `javax.ejb.RemoveException`, `javax.ejb.FinderException`, and subclasses thereof are considered to be application exceptions. These exceptions are used as standard application exceptions to report errors to the client from the `create`, `remove`, and `finder` methods (see Subsections 10.5.9 and 12.1.10). These exceptions are covered by the rules on application exceptions that are defined in this chapter.

An application exception class must be a subclass (direct or indirect) of `java.lang.Exception`. An application exception class must not be defined as a subclass of the `java.lang.RuntimeException` or of the `java.rmi.RemoteException`.

18.1.2 Goals for Exception Handling

The EJB specification for exception handling is designed to meet these high-level goals:

- An application exception thrown by an enterprise bean instance should be reported to the client *precisely* (i.e., the client gets the same exception)^[41].
- An application exception thrown by an enterprise bean instance should not automatically roll-back a client's transaction. The client should typically be given a chance to recover a transaction from an application exception.
- An unexpected exception that may have left the instance's state variables and/or underlying persistent data in an inconsistent state can be handled safely.

18.2 Bean Provider's Responsibilities

This section describes the view and responsibilities of the Bean Provider with respect to exception handling.

18.2.1 Application Exceptions

The Bean Provider defines the application exceptions in the `throws` clauses of the methods of the home, component, and web service endpoint interfaces. Because application exceptions are intended to be handled by the client, and not by the System Administrator, they should be used only for reporting business logic exceptions, not for reporting system level problems.

Certain messaging types may define application exceptions in their message listener interfaces. The resource adapter in use for the particular messaging type determines how the exception is processed. See [12].

[41] This may not be the case where web services protocols are used. See [22].

The Bean Provider is responsible for throwing the appropriate application exception from the business method to report a business logic exception to the client. Because the application exception does not automatically result in marking the transaction for rollback, the Bean Provider must do one of the following to ensure data integrity before throwing an application exception from an enterprise bean instance:

- Ensure that the instance is in a state such that a client's attempt to continue and/or commit the transaction does not result in loss of data integrity. For example, the instance throws an application exception indicating that the value of an input parameter was invalid before the instance performed any database updates.
- Mark the transaction for rollback using the `EJBContext.setRollbackOnly` method before throwing an application exception. Marking the transaction for rollback will ensure that the transaction can never commit.

An application exception class must be a subclass (direct or indirect) of `java.lang.Exception`. An application exception class must not be defined as a subclass of the `java.lang.RuntimeException`, or of the `java.rmi.RemoteException`. These are reserved for system exceptions. (See Section 18.2.2).

The Bean Provider is also responsible for using the standard EJB application exceptions (`javax.ejb.CreateException`, `javax.ejb.RemoveException`, `javax.ejb.FinderException`, and subclasses thereof) as described in Subsections 10.5.9 and 12.1.10.

Bean Providers may define subclasses of the standard EJB application exceptions and throw instances of the subclasses in the enterprise bean methods. A subclass will typically provide more information to the client that catches the exception.

18.2.2 System Exceptions

This subsection describes how the Bean Provider should handle various system-level exceptions and errors that an enterprise bean instance may encounter during the execution of a session or entity bean business method, a message-driven bean message listener method, or a container callback method (e.g. `ejbLoad`).

An enterprise bean business method, message listener method, or container callback method may encounter various exceptions or errors that prevent the method from successfully completing. Typically, this happens because the exception or error is unexpected, or the exception is expected but the EJB Provider does not know how to recover from it. Examples of such exceptions and errors are: failure to obtain a database connection, JNDI exceptions, unexpected `RemoteException` from invocation of other enterprise beans^[42], unexpected `RuntimeException`, JVM errors, and so on.

[42] Note that the enterprise bean business method may attempt to recover from a `RemoteException`. The text in this subsection applies only to the case when the business method does not wish to recover from the `RemoteException`.

If the enterprise bean method encounters a system-level exception or error that does not allow the method to successfully complete, the method should throw a suitable non-application exception that is compatible with the method's `throws` clause. While the EJB specification does not prescribe the exact usage of the exception, it encourages the Bean Provider to follow these guidelines:

- If the bean method encounters a `RuntimeException` or error, it should simply propagate the error from the bean method to the container (i.e., the bean method does not have to catch the exception).
- If the bean method performs an operation that results in a checked exception^[43] that the bean method cannot recover, the bean method should throw the `javax.ejb.EJBException` that wraps the original exception.
- Any other unexpected error conditions should be reported using the `javax.ejb.EJBException`.

Note that the `javax.ejb.EJBException` is a subclass of the `java.lang.RuntimeException`, and therefore it does not have to be listed in the `throws` clauses of the business methods.

The container catches a non-application exception; logs it (which can result in alerting the System Administrator); and, unless the bean is a message-driven bean, throws the `java.rmi.RemoteException` (or subclass thereof) to the client if the client is a remote client or web service client, or throws the `javax.ejb.EJBException` (or subclass thereof) to the client if the client is a local client. In the case of a message-driven bean, the container logs the exception and then throws a `javax.ejb.EJBException` that wraps the original exception to the resource adapter. (See [12]).

The Bean Provider can rely on the container to perform the following tasks when catching a non-application exception:

- The transaction in which the bean method participated will be rolled back.
- No other method will be invoked on an instance that threw a non-application exception.

This means that the Bean Provider does not have to perform any cleanup actions before throwing a non-application exception. It is the container that is responsible for the cleanup.

18.2.2.1 `javax.ejb.NoSuchEntityException`

The `NoSuchEntityException` is a subclass of `EJBException`. It should be thrown by the entity bean class methods to indicate that the underlying entity has been removed from the database.

A bean-managed persistence entity bean class typically throws this exception from the `ejbLoad` and `ejbStore` methods, and from the methods that implement the business methods defined in the component interface.

[43] A checked exception is one that is not a subclass of `java.lang.RuntimeException`.

18.3 Container Provider Responsibilities

This section describes the responsibilities of the Container Provider for handling exceptions. The EJB architecture specifies the container's behavior for the following exceptions:

- Exceptions from the business methods of session and entity beans.
- Exceptions from message-driven bean methods
- Exceptions from `ejbTimeout` methods
- Exceptions from other container-invoked callbacks on the enterprise bean.
- Exceptions from management of container-managed transaction demarcation.

18.3.1 Exceptions from a Session or Entity Bean's Business Methods

Business methods are considered to be the methods defined in the enterprise bean's home, component, or web service endpoint interfaces (including all their superinterfaces); and the following session bean or entity bean methods: `ejbCreate<METHOD>`, `ejbPostCreate<METHOD>`, `ejbRemove`, `ejbHome<METHOD>`, and `ejbFind<METHOD>` methods.

Table 15 specifies how the container must handle the exceptions thrown by the business methods for beans with container-managed transaction demarcation. The table specifies the container's action as a function of the condition under which the business method executes and the exception thrown by the business method. The table also illustrates the exception that the client will receive and how the client can recover from the exception. (Section 18.4 describes the client's view of exceptions in detail.)

Table 15 Handling of Exceptions Thrown by a Business Method of a Bean with Container-Managed Transaction Demarcation

Method condition	Method exception	Container's action	Client's view
Bean method runs in the context of the caller's transaction [Note A]. This case may happen with Required, Mandatory, and Supports attributes.	AppException	Re-throw AppException	Receives AppException. Can attempt to continue computation in the transaction, and eventually commit the transaction (the commit would fail if the instance called setRollbackOnly).
	all other exceptions and errors	Log the exception or error [Note B]. Mark the transaction for rollback. Discard instance [Note C]. Throw javax.transaction.TransactionRolledbackException to remote client; throw javax.ejb.TransactionRolledbackLocalException to local client.	Receives javax.transaction.TransactionRolledbackException or javax.ejb.TransactionRolledbackLocalException Continuing transaction is fruitless.
Bean method runs in the context of a transaction that the container started immediately before dispatching the business method. This case may happen with Required and RequiresNew attributes.	AppException	If the instance called setRollbackOnly(), then rollback the transaction, and re-throw AppException. Otherwise, attempt to commit the transaction, and then re-throw AppException.	Receives AppException. If the client executes in a transaction, the client's transaction is not marked for rollback, and client can continue its work.
	all other exceptions	Log the exception or error. Rollback the container-started transaction. Discard instance. Throw RemoteException to remote or web service client [Note D]; throw EJBException to local client.	Receives RemoteException or EJBException. If the client executes in a transaction, the client's transaction may or may not be marked for rollback.

Table 15 Handling of Exceptions Thrown by a Business Method of a Bean with Container-Managed Transaction Demarcation

Method condition	Method exception	Container's action	Client's view
Bean method runs with an unspecified transaction context. This case may happen with the NotSupported, Never, and Supports attributes.	AppException	Re-throw AppException.	Receives AppException. If the client executes in a transaction, the client's transaction is not marked for rollback, and client can continue its work.
	all other exceptions	Log the exception or error. Discard instance. Throw RemoteException to remote or web service client; throw EJBException to local client.	Receives RemoteException OR EJBException. If the client executes in a transaction, the client's transaction may or may not be marked for rollback.

Notes:

- [A] The caller can be another enterprise bean or an arbitrary client program. This case is not applicable for methods on the web service endpoint interface.
- [B] *Log the exception or error* means that the container logs the exception or error so that the System Administrator is alerted of the problem.
- [C] *Discard instance* means that the container must not invoke any business methods or container callbacks on the instance.
- [D] Throw RemoteException to web service client means that the container maps the RemoteException to the appropriate SOAP fault. See [22].

Table 16 specifies how the container must handle the exceptions thrown by the business methods for beans with bean-managed transaction demarcation. The table specifies the container's action as a function of the condition under which the business method executes and the exception thrown by the business method. The table also illustrates the exception that the client will receive and how the client can recover from the exception. (Section 18.4 describes the client's view of exceptions in detail.)

Table 16 Handling of Exceptions Thrown by a Business Method of a Session Bean with Bean-Managed Transaction Demarcation

Bean method condition	Bean method exception	Container action	Client receives
Bean is stateful or stateless session.	AppException	Re-throw AppException	Receives AppException.
	all other exceptions	Log the exception or error. Mark for rollback a transaction that has been started, but not yet completed, by the instance. Discard instance. Throw RemoteException to remote or web service client [Note A]; throw EJBException to local client.	Receives RemoteException or EJBException.

Notes:

[A] Throw RemoteException to web service client means that the container maps the RemoteException to the appropriate SOAP fault. See [22].

18.3.2 Exceptions from ejbCreate and ejbRemove Methods of a Stateless Session Bean with Web Service Client View

Table 17 specifies how the container must handle the exceptions thrown by the ejbCreate and ejbRemove methods for stateless session beans with a web service client view.

Table 17 Handling of Exceptions Thrown by an ejbCreate or ejbRemove Method of a Stateless Session Bean with Web Service Client View.

Bean method condition	Bean method exception	Container action
Bean is stateless session bean with web service client view	system exceptions	Log the exception or error. Discard instance.

18.3.3 Exceptions from Message-Driven Bean Message Listener Methods

This section specifies the container's handling of exceptions thrown from a message-driven bean's message listener method.

Table 18 specifies how the container must handle the exceptions thrown by a message listener method of a message-driven bean with container-managed transaction demarcation. The table specifies the container's action as a function of the condition under which the method executes and the exception thrown by the method.

Table 18

Handling of Exceptions Thrown by a Message Listener Method of a Message-Driven Bean with Container-Managed Transaction Demarcation.

Method condition	Method exception	Container's action
Bean method runs in the context of a transaction that the container started immediately before dispatching the method. This case happens with <code>Required</code> attribute.	<code>AppException</code>	If the instance called <code>setRollbackOnly</code> , rollback the transaction and re-throw <code>AppException</code> to resource adapter. Otherwise, attempt to commit the transaction and re-throw <code>AppException</code> to resource adapter.
	system exceptions	Log the exception or error[Note A]. Rollback the container-started transaction. Discard instance[Note B]. Throw <code>EJBException</code> that wraps the original exception to resource adapter.
Bean method runs with an unspecified transaction context. This case happens with the <code>NotSupported</code> attribute.	<code>AppException</code>	Re-throw <code>AppException</code> to resource adapter.
	system exceptions	Log the exception or error. Discard instance. Throw <code>EJBException</code> that wraps the original exception to resource adapter

Notes:

- [A] *Log the exception or error* means that the container logs the exception or error so that the System Administrator is alerted of the problem.
- [B] *Discard instance* means that the container must not invoke any methods on the instance.

Table 19 specifies how the container must handle the exceptions thrown by a message listener method of a message-driven bean with bean-managed transaction demarcation. The table specifies the container's action as a function of the condition under which the method executes and the exception thrown by the method.

Table 19 Handling of Exceptions Thrown by a Message Listener Method of a Message-Driven Bean with Bean-Managed Transaction Demarcation.

Bean method condition	Bean method exception	Container action
Bean is message-driven bean	AppException	Re-throw AppException to resource adapter.
	system exceptions	Log the exception or error. Mark for rollback a transaction that has been started, but not yet completed, by the instance. Discard instance. Throw EJBException that wraps the original exception to resource adapter.

18.3.4 Exceptions from ejbCreate and ejbRemove Methods of a Message-Driven Bean

Table 20 specifies how the container must handle the exceptions thrown by the ejbCreate and ejbRemove methods of message-driven beans.

Table 20 Handling of Exceptions Thrown by an ejbCreate or ejbRemove Method of a Message-Driven Bean.

Bean method condition	Bean method exception	Container action
Bean is message-driven bean	system exceptions	Log the exception or error. Discard instance.

18.3.5 Exceptions from an Enterprise Bean's ejbTimeout Method

This section specifies the container's handling of exceptions thrown from an enterprise bean's ejbTimeout method.

Table 21 and Table 22 specify how the container must handle the exceptions thrown by the `ejbTimeout` method of an enterprise bean that implements the `TimedObject` interface. The `ejbTimeout` method does not throw application exceptions and cannot throw exceptions to the client.

Table 21 Handling of Exceptions Thrown by the `ejbTimeout` Method of an Enterprise Bean with Container-Managed Transaction Demarcation.

Method condition	Method exception	Container's action
Bean implements <code>TimedObject</code> interface. Bean method runs in the context of a transaction that the container started immediately before dispatching the method.	system exceptions	Log the exception or error[Note A]. Rollback the container-started transaction. Discard instance[Note B].

Notes:

- [A] *Log the exception or error* means that the container logs the exception or error so that the System Administrator is alerted of the problem.
- [B] *Discard instance* means that the container must not invoke any methods on the instance.

Table 22 Handling of Exceptions Thrown by the `ejbTimeout` Method of an Enterprise Bean with Bean-Managed Transaction Demarcation.

Method condition	Method exception	Container's action
Bean implements <code>TimedObject</code> interface. Bean method The bean method may make use of <code>UserTransaction</code> .	system exceptions	Log the exception or error[Note A]. Mark for rollback a transaction that has been started, but not yet completed, by the instance. Discard instance[Note B].

Notes:

- [A] *Log the exception or error* means that the container logs the exception or error so that the System Administrator is alerted of the problem.
- [B] *Discard instance* means that the container must not invoke any methods on the instance.

18.3.6 Exceptions from Container-invoked Callbacks

This subsection specifies the container's handling of exceptions thrown from the container-invoked callbacks on the enterprise bean. This subsection applies to the following callback methods:

- The `ejbActivate`, `ejbLoad`, `ejbPassivate`, `ejbStore`, `setEntityContext`, and `unsetEntityContext` methods of the `EntityBean` interface.
- The `ejbActivate`, `ejbPassivate`, and `setSessionContext` methods of the `SessionBean` interface.
- The `setMessageDrivenContext` method of the `MessageDrivenBean` interface.
- The `afterBegin`, `beforeCompletion` and `afterCompletion` methods of the `SessionSynchronization` interface.

The container must handle all exceptions or errors from these methods as follows:

- Log the exception or error to bring the problem to the attention of the System Administrator.
- If the instance is in a transaction, mark the transaction for rollback.
- Discard the instance (i.e., the container must not invoke any business methods or container callbacks on the instance).
- If the exception or error happened during the processing of a client invoked method, throw the `java.rmi.RemoteException` to the client if the client is a remote client or throw the `javax.ejb.EJBException` to the client if the client is a local client. If the instance executed in the client's transaction, the container should throw the `javax.transaction.TransactionRolledbackException` to a remote client or the `javax.ejb.TransactionRolledbackLocalException` to a local client, because it provides more information to the client. (The client knows that it is fruitless to continue the transaction.)

18.3.7 javax.ejb.NoSuchEntityException

The `NoSuchEntityException` is a subclass of `EJBException`. If it is thrown by a method of an entity bean class, the container must handle the exception using the rules for `EJBException` described in Sections 18.3.1, 18.3.3, and 18.3.6.

To give the client a better indication of the cause of the error, the container should throw the `java.rmi.NoSuchObjectException` (which is a subclass of `java.rmi.RemoteException`) to a remote client, or the `javax.ejb.NoSuchObjectLocalException` to a local client.

18.3.8 Non-existing Stateful Session or Entity Object

If a client makes a call to a stateful session or entity object that has been removed, the container should throw the `java.rmi.NoSuchObjectException` (which is a subclass of `java.rmi.RemoteException`) to a remote client, or the `javax.ejb.NoSuchObjectLocalException` to a local client.

18.3.9 Exceptions from the Management of Container-Managed Transactions

The container is responsible for starting and committing the container-managed transactions, as described in Subsection 17.6.2. This subsection specifies how the container must deal with the exceptions that may be thrown by the transaction start and commit operations.

If the container fails to start or commit a container-managed transaction, the container must throw the `java.rmi.RemoteException` to a remote or web service client and the `javax.ejb.EJBException` to a local client. In the case where the container fails to start or commit a container-managed transaction on behalf of a message-driven bean or an `ejbTimeout` method, the container must throw and log the `javax.ejb.EJBException`.

However, the container should not throw the `java.rmi.RemoteException` or `javax.ejb.EJBException` if the container performs a transaction rollback because the instance has invoked the `setRollbackOnly` method on its `EJBContext` object. In this case, the container must rollback the transaction and pass the business method result or the application exception thrown by the business method to the client.

Note that some implementations of the container may retry a failed transaction transparently to the client and enterprise bean code. Such a container would throw the `java.rmi.RemoteException` or `javax.ejb.EJBException` after a number of unsuccessful tries.

18.3.10 Release of Resources

When the container discards an instance because of a system exception, the container should release all the resources held by the instance that were acquired through the resource factories declared in the enterprise bean environment (See Subsection 20.5).

Note: While the container should release the connections to the resource managers that the instance acquired through the resource factories declared in the enterprise bean environment, the container cannot, in general, release “unmanaged” resources that the instance may have acquired through the JDK APIs. For example, if the instance has opened a TCP/IP connection, most container implementations will not be able to release the connection. The connection will be eventually released by the JVM garbage collector mechanism.

18.3.11 Support for Deprecated Use of `java.rmi.RemoteException`

The EJB 1.0 specification allowed the business methods, `ejbCreate`, `ejbPostCreate`, `ejbFind<METHOD>`, `ejbRemove`, and the container-invoked callbacks (i.e., the methods defined in the `EntityBean`, `SessionBean`, and `SessionSynchronization` interfaces) implemented in the enterprise bean class to use the `java.rmi.RemoteException` to report non-application exceptions to the container.

This use of the `java.rmi.RemoteException` was deprecated in EJB 1.1—enterprise beans written for the EJB 1.1 specification should use the `javax.ejb.EJBException` instead, and enterprise beans written for the EJB 2.0 or later specification must use the `javax.ejb.EJBException` instead.

The EJB 1.1 and EJB 2.0 or later specifications require that a container support the deprecated use of the `java.rmi.RemoteException`. The container should treat the `java.rmi.RemoteException` thrown by an enterprise bean method in the same way as it is specified for the `javax.ejb.EJBException`.

Note: The use of the `java.rmi.RemoteException` is deprecated only in the above-mentioned methods. The methods of the remote and remote home interface still must use the `java.rmi.RemoteException` as required by the EJB specification.

18.4 Client's View of Exceptions

This section describes the client's view of exceptions received from an enterprise bean invocation.

A client accesses an enterprise bean either through the enterprise bean's remote home and remote interfaces, through the enterprise bean's local home and local interfaces, or through the enterprise bean's web service client view depending on whether the client is a remote client, a local client, or a web service client.

The remote home interface, the remote interface, and the web service endpoint interface are Java RMI interfaces, and therefore the `throws` clauses of all their methods (including those inherited from superinterfaces) include the mandatory `java.rmi.RemoteException`. The `throws` clauses may include an arbitrary number of application exceptions.

The local home and local interfaces are both Java local interfaces, and the `throws` clauses of all their methods (including those inherited from superinterfaces) must not include the `java.rmi.RemoteException`. The `throws` clauses may include an arbitrary number of application exceptions.

18.4.1 Application Exception

18.4.1.1 Local and Remote Clients

If a client program receives an application exception from an enterprise bean invocation, the client can continue calling the enterprise bean. An application exception does not result in the removal of the EJB object.

If a client program receives an application exception from an enterprise bean invocation while the client is associated with a transaction, the client can typically continue the transaction because an application exception does not automatically cause the container to mark the transaction for rollback.

For example, if a client receives the `ExceedLimitException` application exception from the `debit` method of an `Account` bean, the client may invoke the `debit` method again, possibly with a lower `debit` amount parameter. If the client executed in a transaction context, throwing the `ExceedLimitException` exception would not automatically result in rolling back, or marking for rollback, the client's transaction.

Although the container does not automatically mark for rollback a transaction because of a thrown application exception, the transaction might have been marked for rollback by the enterprise bean instance before it threw the application exception. There are two ways to learn if a particular application exception results in transaction rollback or not:

- Statically. Programmers can check the documentation of the enterprise bean's home or component interface. The Bean Provider may have specified (although he or she is not required to) the application exceptions for which the enterprise bean marks the transaction for rollback before throwing the exception.
- Dynamically. Clients that are enterprise beans with container-managed transaction demarcation can use the `getRollbackOnly` method of the `javax.ejb.EJBContext` object to learn if the current transaction has been marked for rollback; other clients may use the `getStatus` method of the `javax.transaction.UserTransaction` interface to obtain the transaction status.

18.4.1.2 Web Service Clients

If a stateless session bean throws an application exception from one of the methods on its web service endpoint interface, it is the responsibility of the container to map the exception to the SOAP fault specified in the WSDL that describes the port type that the stateless session bean implements. For Java clients, the exceptions received by the client are described by the mapping rules in [22].

18.4.2 java.rmi.RemoteException and javax.ejb.EJBException

A remote or a Java web service client receives the `java.rmi.RemoteException` and a local client receives the `javax.ejb.EJBException` as an indication of a failure to invoke an enterprise bean method or to properly complete its invocation. The exception can be thrown by the container or by the communication subsystem between the client and the container.

If the client receives the `java.rmi.RemoteException` or the `javax.ejb.EJBException` exception from a method invocation, the client, in general, does not know if the enterprise bean's method has been completed or not.

If the client executes in the context of a transaction, the client's transaction may, or may not, have been marked for rollback by the communication subsystem or target bean's container.

For example, the transaction would be marked for rollback if the underlying transaction service or the target bean's container doubted the integrity of the data because the business method may have been partially completed. Partial completion could happen, for example, when the target bean's method returned with a `RuntimeException` exception, or if the remote server crashed in the middle of executing the business method.

The transaction may not necessarily be marked for rollback. This might occur, for example, when the communication subsystem on the client-side has not been able to send the request to the server.

When a client executing in a transaction context receives a `RemoteException` or an `EJBException` from an enterprise bean invocation, the client may use either of the following strategies to deal with the exception:

- Discontinue the transaction. If the client is the transaction originator, it may simply rollback its transaction. If the client is not the transaction originator, it can mark the transaction for rollback or perform an action that will cause a rollback. For example, if the client is an enterprise bean, the enterprise bean may throw a `RuntimeException` which will cause the container to rollback the transaction.
- Continue the transaction. The client may perform additional operations on the same or other enterprise beans, and eventually attempt to commit the transaction. If the transaction was marked for rollback at the time the `RemoteException` or `EJBException` was thrown to the client, the commit will fail.

If the client chooses to continue the transaction, the client can first inquire about the transaction status to avoid fruitless computation on a transaction that has been marked for rollback. A client that is an enterprise bean with container-managed transaction demarcation can use the `EJBContext.getRollbackOnly` method to test if the transaction has been marked for rollback; a client that is an enterprise bean with bean-managed transaction demarcation, and other client types, can use the `UserTransaction.getStatus` method to obtain the status of the transaction.

Some implementations of EJB servers and containers may provide more detailed exception reporting by throwing an appropriate subclass of the `java.rmi.RemoteException` or `javax.ejb.EJBException` to the client. The following subsections describe the several subclasses of the `java.rmi.RemoteException` and `javax.ejb.EJBException` that may be thrown by the container to give the client more information.

18.4.2.1 `javax.transaction.TransactionRolledbackException` and `javax.ejb.TransactionRolledbackLocalException`

The `javax.transaction.TransactionRolledbackException` is a subclass of the `java.rmi.RemoteException`. It is defined in the JTA standard extension. The `javax.ejb.TransactionRolledbackLocalException` is a subclass of the `javax.ejb.EJBException`.

If a client receives the `javax.transaction.TransactionRolledbackException` or the `javax.ejb.TransactionRolledbackLocalException`, the client knows for certain that the transaction has been marked for rollback. It would be fruitless for the client to continue the transaction because the transaction can never commit.

18.4.2.2 `javax.transaction.TransactionRequiredException` and `javax.ejb.TransactionRequiredLocalException`

The `javax.transaction.TransactionRequiredException` is a subclass of the `java.rmi.RemoteException`. It is defined in the JTA standard extension. The `javax.ejb.TransactionRequiredLocalException` is a subclass of the `javax.ejb.EJBException`.

The `javax.transaction.TransactionRequiredException` or `javax.ejb.TransactionRequiredLocalException` informs the client that the target enterprise bean must be invoked in a client's transaction, and that the client invoked the enterprise bean without a transaction context.

This error usually indicates that the application was not properly formed.

18.4.2.3 `java.rmi.NoSuchObjectException` and `javax.ejb.NoSuchObjectLocalException`

The `java.rmi.NoSuchObjectException` is a subclass of the `java.rmi.RemoteException`. It is thrown to the client if a remote business method cannot complete because the EJB object no longer exists.

The `javax.ejb.NoSuchObjectLocalException` is a subclass of the `javax.ejb.EJBException`. It is thrown to the client if a local business method cannot complete because the EJB object no longer exists.

18.5 System Administrator's Responsibilities

The System Administrator is responsible for monitoring the log of the non-application exceptions and errors logged by the container, and for taking actions to correct the problems that caused these exceptions and errors.

18.6 Differences from EJB 1.0

The EJB 2.0 and EJB 1.1 specification of exception handling preserve the rules defined in the EJB 1.0 specification, with the following exceptions:

- EJB 1.0 specified that the enterprise bean business methods and container-invoked callbacks use the `java.rmi.RemoteException` to report non-application exceptions. This practice was deprecated in EJB 1.1—the enterprise bean methods should use the `javax.ejb.EJBException` or other suitable `RuntimeException` to report non-application exceptions.
- In EJB 2.0 and 1.1, all non-application exceptions thrown by the instance result in the rollback of the transaction in which the instance executed, and in discarding the instance. In EJB 1.0, the container would not rollback a transaction and discard the instance if the instance threw the `java.rmi.RemoteException`.
- In EJB 2.0 and 1.1, an application exception does not cause the container to automatically rollback a transaction. In EJB 1.0, the container was required to rollback a transaction when an application exception was passed through a transaction boundary started by the container. In EJB 1.1, the container performs the rollback only if the instance have invoked the `setRollbackOnly` method on its `EJBContext` object.

Support for Distribution and Interoperability

This chapter describes the support for accessing an enterprise bean through the remote client view from clients distributed over a network, and the interoperability requirements for invocations on enterprise beans from remote clients that are Java 2 Platform, Enterprise Edition (J2EE) components.

19.1 Support for Distribution

The remote home and remote interfaces of an enterprise bean's remote client view are defined as Java™ RMI [3] interfaces. This allows the container to implement the remote home and remote interfaces as *distributed objects*. A client using the remote home and remote interfaces can reside on a different machine than the enterprise bean (location transparency), and the object references of the remote home and remote interfaces can be passed over the network to other applications.


The EJB specification further constrains the Java RMI types that can be used by enterprise beans to be legal RMI-IIOP types [7]. This makes it possible for EJB container implementors to use RMI-IIOP as the object distribution protocol.

19.1.1 Client-Side Objects in a Distributed Environment

When the RMI-IIOP protocol or similar distribution protocols are used, the remote client communicates with the enterprise bean using *stubs* for the server-side objects. The stubs implement the remote home and remote interfaces.

Figure 76

Location of EJB Client Stubs.




The communication stubs used on the client side are artifacts generated at the enterprise bean's deployment time by the Container Provider's tools. The stubs used on the client are specific to the wire protocol used for the remote invocation.

19.2 Interoperability Overview

Session beans and entity beans that are deployed in one vendor's server product often need to be accessed from J2EE client components that are deployed in another vendor's product. EJB defines a standard interoperability protocol based on CORBA/IIOP to address this need.

The interoperability protocols described here must be supported by compatible EJB products. Additional vendor-specific protocols may also be supported.

Figure 77 shows a heterogeneous environment that includes systems from several vendors to illustrate the interoperability enabled by EJB.

Figure 77 Heterogeneous EJB Environment

The following sections in this chapter

- describe the goals for EJB invocation interoperability
- provide illustrative scenarios
- describe the interoperability requirements for remote invocations, transactions, naming, and security.

19.2.1 Interoperability Goals

The goals of the interoperability requirements specified in this chapter are as follows:

- To allow clients in one application deployed in J2EE containers from one server provider to access services from session and entity beans in another application that is deployed in an EJB container from a different server provider. For example, web components (JavaServer Pages and servlets) that are deployed on a J2EE-compliant web server provided by one server provider must be able to invoke the business methods of enterprise beans that are deployed on a J2EE-compliant EJB server from another server provider.
- To achieve interoperability without any new requirements on the J2EE application developer.

- To ensure out-of-the-box interoperability between compliant J2EE products. It must be possible for an enterprise customer to install multiple J2EE server products from different server providers (on potentially different operating systems), deploy applications in the J2EE servers, and have the multiple applications interoperate.
- To leverage the interoperability work done by standards bodies (including the IETF, W3C, and OMG) where possible, so that customers can work with industry standards and use standard protocols to access enterprise beans.

This specification does not address interoperability issues between enterprise beans and non-J2EE components. The J2EE platform specification [9] and the JAX-RPC specification [22] describe requirements for interoperability with Internet clients (using HTTP and XML) and interoperability with enterprise information systems (using the J2EE Connector architecture [12]).

Since the interoperability protocol described here is based on CORBA/IIOP, CORBA clients written in Java, C++, or other languages can also invoke methods on enterprise beans.

This chapter subsumes the previous EJB1.1-to-CORBA mapping document [13].

19.3 Interoperability Scenarios

This section presents a number of interoperability scenarios that motivate the interoperability mechanisms described in later sections of this chapter. These scenarios are illustrative rather than prescriptive. This section does not specify requirements for a J2EE product to support these scenarios in exactly the manner described here.

J2EE applications are multi-tier, web-enabled applications. Each application consists of one or more components, which are deployed in containers. The four types of containers are:

- EJB containers, which host enterprise beans.
- Web containers, which host JavaServer Pages (JSPs) and servlet components as well as static documents, including HTML pages.
- Application client containers, which host standalone applications.
- Applet containers, which host applets which may be downloaded from a web site. At this time, there is no requirement for an applet to be able to directly invoke the remote methods of enterprise beans.

The scenarios below describe interactions between components hosted in these various container types.

19.3.1 Interactions Between Web Containers and EJB Containers for E-Commerce Applications

This scenario occurs for business-to-business and business-to-consumer interactions over the Internet.

Scenario 1: A customer wants to buy a book from an Internet bookstore. The bookstore's web site consists of a J2EE application containing JSPs that form the presentation layer, and another J2EE application containing enterprise beans that have the business logic and database access code. The JSPs and enterprise beans are deployed in containers from different vendors.

At deployment time: The enterprise beans are deployed, and their EJBHome objects are published in the EJB server's name service. The Deployer links the EJB reference in the JSP's deployment descriptor to the URL of the enterprise bean's EJBHome object, which can be looked up from the name service. The transaction attribute specified in the enterprise bean's deployment descriptor is `RequiresNew` for all business methods. Because the "checkout" JSP requires secure access to set up payments for purchases, the bookstore's administrator configures the "checkout" JSP to require access over HTTPS with only server authentication. Customer authentication is done using form-based login. The "book search" JSP is accessed over normal HTTP. Both JSPs talk with enterprise beans that access the book database. The web and EJB containers use the same customer realm and have a trust relationship with each other. The network between the web and EJB servers is not guaranteed to be secure from attacks.

At runtime: The customer accesses the book search JSP using a browser. The JSP looks up the enterprise bean's EJBHome object in a name service, and calls `findBooks(...)` with the search criteria as parameters. The web container establishes a secure session with the EJB container with mutual authentication between the containers, and invokes the enterprise bean. The customer then decides to buy a book, and accesses the "checkout" JSP. The customer enters the necessary information in the login form, which is used by the web server to authenticate the customer. The JSP invokes the enterprise bean to update the book and customer databases. The customer's principal is propagated to the EJB container and used for authorization checks. The enterprise bean completes the updates and commits the transaction. The JSP sends back a confirmation page to the customer.

19.3.2 Interactions Between Application Client Containers and EJB Containers Within an Enterprise's Intranet

Scenario 2.1: An enterprise has an expense accounting application used by employees from their desktops. The server-side consists of a J2EE application containing enterprise beans that are deployed on one vendor's J2EE product, which is hosted in a datacenter. The client side consists of another J2EE application containing an application client deployed using another vendor's J2EE infrastructure. The network between the application client and the EJB container is insecure and needs to be protected against spoofing and other attacks.

At deployment time: The enterprise beans are deployed and their EJBHome objects are published in the enterprise's name service. The application clients are configured with the names of the EJBHome objects. The Deployer maps employees to roles that are allowed access to the enterprise beans. The System Administrator configures the security settings of the application client and EJB container to use client and server authentication and message protection. The System Administrator also does the necessary client-side configuration to allow client authentication.

At runtime: The employee logs on using username and password. The application client container may interact with the enterprise's authentication service infrastructure to set up the employee's credentials. The client application does a remote invocation to the name server to look up the enterprise bean's EJBHome object, and creates the enterprise beans. The application client container uses a secure transport protocol to interact with the name server and EJB server, which does mutual authentication and also guarantees the confidentiality and integrity of messages. The employee then enters the expense information and submits it. This causes remote business methods of the enterprise beans to be invoked. The EJB container performs authorization checks and, if they succeed, executes the business methods.

Scenario 2.2: This is the same as Scenario 2.1, except that there is no client-side authentication infrastructure set up by the System Administrator which can authenticate at the transport protocol layer. At runtime the client container needs to send the user's password to the server during the method invocation to authenticate the employee.

19.3.3 Interactions Between Two EJB Containers in an Enterprise's Intranet

Scenario 3: An enterprise has an expense accounting application which needs to communicate with a payroll application. The applications use enterprise beans and are deployed on J2EE servers from different vendors. The J2EE servers and naming/authentication services may be in the enterprise's data-center with a physically secure private network between them, or they may need to communicate across the intranet, which may be less secure. The applications need to update accounts and payroll databases. The employee (client) accesses the expense accounting application as described in Scenario 2.

At deployment time: The Deployer configures both applications with the appropriate database resources. The accounts application is configured with the name of the EJBHome object of the payroll application. The payroll bean's deployment descriptor specifies the `RequiresNew` transaction attribute for all methods. The applications use the same principal-to-role mappings (e.g. the roles may be `Employee`, `PayrollDept`, `AccountsDept`). The Deployer of these two applications has administratively set up a trust relationship between the two EJB containers, so that the containers do not need to authenticate principals propagated on calls to enterprise beans from the other container. The System Administrator also sets up the message protection parameters of the two containers if the network is not physically secure.

At runtime: An employee makes a request to the accounts application which requires it to access the payroll application. The accounts application does a lookup of the payroll application's EJBHome object in the naming/directory service and creates enterprise beans. It updates the accounts database and invokes a remote method of the payroll bean. The accounts bean's container propagates the employee's principal on the method call. The payroll bean's container maps the propagated employee principal to a role, does authorization checks, and sets up the payroll bean's transaction context. The container starts a new transaction, then the payroll bean updates the payroll database, and the container commits the transaction. The accounts bean receives a status reply from the payroll bean. If an error occurs in the payroll bean, the accounts bean executes code to recover from the error and restore the databases to a consistent state.

19.3.4 Intranet Application Interactions Between Web Containers and EJB Containers

***Scenario 4:** This is the same as scenario 2.1, except that instead of using a “fat-client” desktop application to access the enterprise’s expense accounting application, employees use a web browser and connect to a web server in the intranet that hosts JSPs. The JSPs gather input from the user (e.g., through an HTML form), invoke enterprise beans that contain the actual business logic, and format the results returned by the enterprise beans (using HTML).*

***At deployment time:** The enterprise Deployer configures its expense accounting JSPs to require access over HTTPS with mutual authentication. The web and EJB containers use the same customer realm and have a trust relationship with each other.*

***At run-time:** The employee logs in to the client desktop, starts the browser, and accesses the expense accounting JSP. The browser establishes an HTTPS session with the web server. Client authentication is performed (for example) using the employee’s credentials which have been established by the operating system at login time (the browser interacts with the operating system to obtain the employee’s credentials). The JSP looks up the enterprise bean’s EJBHome object in a name service. The web container establishes a secure session with the EJB container with mutual authentication and integrity/confidentiality protection between the containers, and invokes methods on the enterprise beans.*

19.4 Overview of Interoperability Requirements

The interoperability requirements used to support the above scenarios are:

1. Remote method invocation on an enterprise bean’s EJBObject and EJBHome object references (scenarios 1,2,3,4), described in section 19.5.
2. Name service lookup of the enterprise bean’s EJBHome object (scenarios 1,2,3,4), described in section 19.7.
3. Integrity and confidentiality protection of messages (scenarios 1,2,3,4), described in section 19.8.
4. Authentication between an application client and EJB container (described in section 19.8):
 - 4.1 Mutual authentication at the transport protocol layer when there is client-side authentication infrastructure such as certificates (scenario 2.1).
 - 4.2 Transfer of the user’s authentication data from application client to EJB container to allow the EJB container to authenticate the client when there is no client-side authentication infrastructure (scenario 2.2).
5. Mutual authentication between two EJB containers or between a web and EJB container to establish trust before principals are propagated (scenarios 1,3,4), described in section 19.8.

6. Propagation of the Internet or intranet user's principal name for invocations on enterprise beans from web or EJB containers when the client and server containers have a trust relationship (scenarios 1,3,4), described in section 19.8.

EJB, web, and application client containers must support the above requirements separately as well as in combinations.

19.5 Remote Invocation Interoperability

This section describes the interoperability mechanisms that enable remote invocations on EJBObject and EJBHome object references when client containers and EJB containers are provided by different vendors. This is needed to satisfy interoperability requirement (1) in section 19.4.

All EJB, web, and application client containers must support the IIOP 1.2 protocol for remote invocations on EJBObject and EJBHome references. EJB containers must be capable of servicing IIOP 1.2 based invocations on EJBObject and EJBHome objects. IIOP 1.2 is part of the CORBA 2.3.1 specification [14] from the OMG^[44]. Containers may additionally support vendor-specific protocols.

CORBA Interoperable Object References (IORs) for EJBObject and EJBHome object references must include the GIOP version number 1.2. The IIOP infrastructure in all J2EE containers must be able to accept fragmented GIOP messages, although sending fragmented messages is optional. Bidirectional GIOP messages may optionally be supported by J2EE clients and servers: if a J2EE server receives an IIOP message from a client which contains the `BiDirIIOPServiceContext` structure, it may or may not use the same connection for sending requests back to the client.

Since Java applications use Unicode characters by default, J2EE containers are required to support the Unicode UTF16 code set for transmission of character and string data (in the IDL `wchar` and `wstring` datatypes). J2EE containers may optionally support additional code sets. EJBObject and EJBHome IORs must have the `TAG_CODE_SETS` tagged component which declares the codesets supported by the EJB container. IIOP messages which include `wchar` and `wstring` datatypes must have the code sets service context field. The CORBA 2.3.1 requirements for code set support must be followed by J2EE containers.

EJB containers are required to translate Java types to their on-the-wire representation in IIOP messages using the Java Language to IDL mapping specification [7] with the wire formats for IDL types as described in the GIOP specification in CORBA 2.3. The following subsections describe the mapping details for Java types.

[44] CORBA APIs and earlier versions of the IIOP protocol are already included in the J2SE 1.2, J2SE 1.3 and J2EE 1.2 platforms through JavaIDL and RMI-IIOP.

19.5.1 Mapping Java Remote Interfaces to IDL

For each session bean or entity bean that is deployed in a container, there are two Java RMI remote interfaces in the remote client view—the bean's remote home interface and the bean's remote interface. The Java Language to IDL Mapping specification [7] describes precisely how these remote interfaces are mapped to IDL. This mapping to IDL is typically implicit when Java RMI over IIOP is used to invoke enterprise beans. J2EE clients use only the Java RMI APIs to invoke enterprise beans. The client container may use the CORBA portable Stub APIs for the client-side stubs. EJB containers may create CORBA Tie objects for each EJBObject or EJBHome object.

19.5.2 Mapping Value Objects to IDL

The Java interfaces that are passed by value during remote invocations on enterprise beans are `javax.ejb.Handle`, `javax.ejb.HomeHandle`, and `javax.ejb.EJBMetaData`. The `Enumeration` or `Collection` objects returned by entity bean finder methods are value types. There may also be application-specific value types that are passed as parameters or return values on enterprise bean invocations. In addition, several Java exception classes that are thrown by remote methods also result in concrete IDL value types. All these value types are mapped to IDL abstract value types or abstract interfaces using the rules in the Java Language to IDL Mapping.

19.5.3 Mapping of System Exceptions

Java system exceptions, including the `java.rmi.RemoteException` and its subclasses, may be thrown by the EJB container. If the client's invocation was made over IIOP, the EJB server is required to map these exceptions to CORBA system exceptions and send them in the IIOP reply message to the client, as specified in the following table

System exception thrown by EJB container	CORBA system exception received by client ORB
<code>javax.transaction.TransactionRolledbackException</code>	<code>TRANSACTION_ROLLEDBACK</code>
<code>javax.transaction.TransactionRequiredException</code>	<code>TRANSACTION_REQUIRED</code>
<code>javax.transaction.InvalidTransactionException</code>	<code>INVALID_TRANSACTION</code>
<code>java.rmi.NoSuchObjectException</code>	<code>OBJECT_NOT_EXIST</code>
<code>java.rmi.AccessException</code>	<code>NO_PERMISSION</code>
<code>java.rmi.MarshalException</code>	<code>MARSHAL</code>
<code>java.rmi.RemoteException</code>	<code>UNKNOWN</code>

For EJB clients, the ORB's unmarshaling machinery maps CORBA system exceptions received in the IIOP reply message to the appropriate Java exception as specified in the Java Language to IDL mapping. This results in the original Java exception being received by the client J2EE component.

19.5.4 Obtaining Stub and Client View Classes

When a J2EE component (application client, JSP, servlet or enterprise bean) receives a reference to an EJBObject or EJBHome object through JNDI lookup or as a parameter or return value of an invocation on an enterprise bean, an instance of an RMI-IIOP stub class (proxy) for the enterprise bean's remote home or remote RMI interface needs to be created. When a component receives a value object as a parameter or return value of an enterprise bean invocation, an instance of the value class needs to be created. The stub class, value class, and other client view classes must be available to the referencing container (the container hosting the component that receives the reference or value type).

The client view classes, including application value classes, must be packaged with the referencing component's application, as described in Section 24.3.

Stubs for invoking on EJBHome and EJBObject references must be provided by the referencing container, for example, by generating stub classes at deployment time for the EJBHome and EJBObject interfaces of the referenced beans that are packaged with the referencing component's application. Stub classes may or may not follow the standard RMI-IIOP portable stub architecture.

Containers may optionally support run-time downloading of stub and value classes needed by the referencing container. The CORBA 2.3.1 specification and the Java Language to IDL Mapping specify how stub and value type implementations are to be downloaded: using codebase URLs that are either embedded in the EJBObject or EJBHome's IOR, or sent in the IIOP message service context, or marshalled with the value type. The URLs for downloading may optionally include an HTTPS URL for secure downloading.

19.5.5 System Value Classes

System value classes are serializable value classes implementing the `javax.ejb.Handle`, `javax.ejb.HomeHandle`, `javax.ejb.EJBMetaData`, `java.util.Enumeration`, `java.util.Collection`, and `java.util.Iterator` interfaces. These value classes are provided by the EJB container vendor. They must be provided in the form of a JAR file by the container hosting the referenced bean. For interoperability scenarios, if a referencing component would use such system value classes at runtime, the Deployer must ensure that these system value classes provided by the container hosting the referenced bean are available to the referencing component. This may be done, for example, by including these system value classes in the classpath of the referencing container, or by deploying the system value classes with the referencing component's application by providing them to the deployment tool.

Implementations of these system value classes must be portable (they must use only J2SE and J2EE APIs) so that they can be instantiated in another vendor's container. If the system value class implementation needs to load application-specific classes (such as remote home or remote interfaces) at runtime, it must use the thread context class loader. The referencing container must make application-specific classes available to the system value class instance at runtime through the thread context class loader.

19.5.5.1 HandleDelegate SPI

The `javax.ejb.spi.HandleDelegate` service provider interface defines methods that enable portable implementations of `Handle` and `HomeHandle` that are instantiated in a different vendor's container to serialize and deserialize `EJBObject` and `EJBHome` references. The `HandleDelegate` interface is not used by enterprise beans or J2EE application components directly.

EJB, web and application client containers must provide implementations of the `HandleDelegate` interface. The `HandleDelegate` object must be accessible in the client J2EE component's JNDI namespace at the reserved name "java:comp/HandleDelegate". The `HandleDelegate` object is not exported outside the container that provides it.

Portable implementations of `Handle` and `HomeHandle` must look up the `HandleDelegate` object of the container in which they are instantiated using JNDI at the name "java:comp/HandleDelegate" and use the `HandleDelegate` object to serialize and deserialize `EJBObject` and `EJBHome` references as follows:

- `Handle` and `HomeHandle` implementation classes must define `writeObject` and `readObject` methods to control their serialization and deserialization. These methods must not wrap or substitute the stream objects that are passed to the `HandleDelegate` methods.
- The `writeObject` method of `Handle` implementations must call `HandleDelegate.writeEJBObject` with the `Handle`'s `EJBObject` reference and the serialization output stream object as parameters. The `HandleDelegate` implementation (which is provided by the client container in which the `Handle` was instantiated, potentially from a different vendor) then writes the `EJBObject` to the output stream. If the output stream corresponds to an IIOP message, the `HandleDelegate` must use the standard IIOP abstract interface format for writing the `EJBObject` reference.
- The `readObject` method of `Handle` implementations must call `HandleDelegate.readEJBObject` with the serialization input stream object as parameter, and with the stream positioned at the location where the `EJBObject` can be read. The `HandleDelegate` implementation then reads the `EJBObject` from the input stream and returns it to the `Handle`. If the input stream corresponds to an IIOP message, the `HandleDelegate` must use the standard abstract interface format for reading the `EJBObject` reference. The `HandleDelegate` must ensure that the `EJBObject` reference is capable of performing invocations immediately after deserialization. The `Handle` maintains a reference to the `EJBObject` as a transient instance variable and returns it when the J2EE component calls `Handle.getEJBObject`.
- The `writeObject` and `readObject` methods of `HomeHandle` implementation classes must be implemented similarly, by using `HandleDelegate.writeEJBHome` and `HandleDelegate.readEJBHome` respectively.

19.6 Transaction Interoperability

Transaction interoperability between containers provided by different vendors is an optional feature in this version of the EJB specification. Vendors may choose to not implement transaction interoperability. However, vendors who choose to implement transaction interoperability must follow the requirements in sections 19.6.1 and 19.6.2, and vendors who choose not to implement transaction interoperability must follow the requirements in section 19.6.2.

19.6.1 Transaction Interoperability Requirements

A distributed transaction started by a web or EJB container must be able to propagate in a remote invocation to an enterprise bean in an EJB container provided by a different vendor, and the containers must participate in the distributed two-phase commit protocol.

19.6.1.1 Transaction Context Wire Format

Transaction context propagation from client to EJB container uses the implicit propagation mechanism described in the CORBA Object Transaction Service (OTS) v1.2 specification [8].

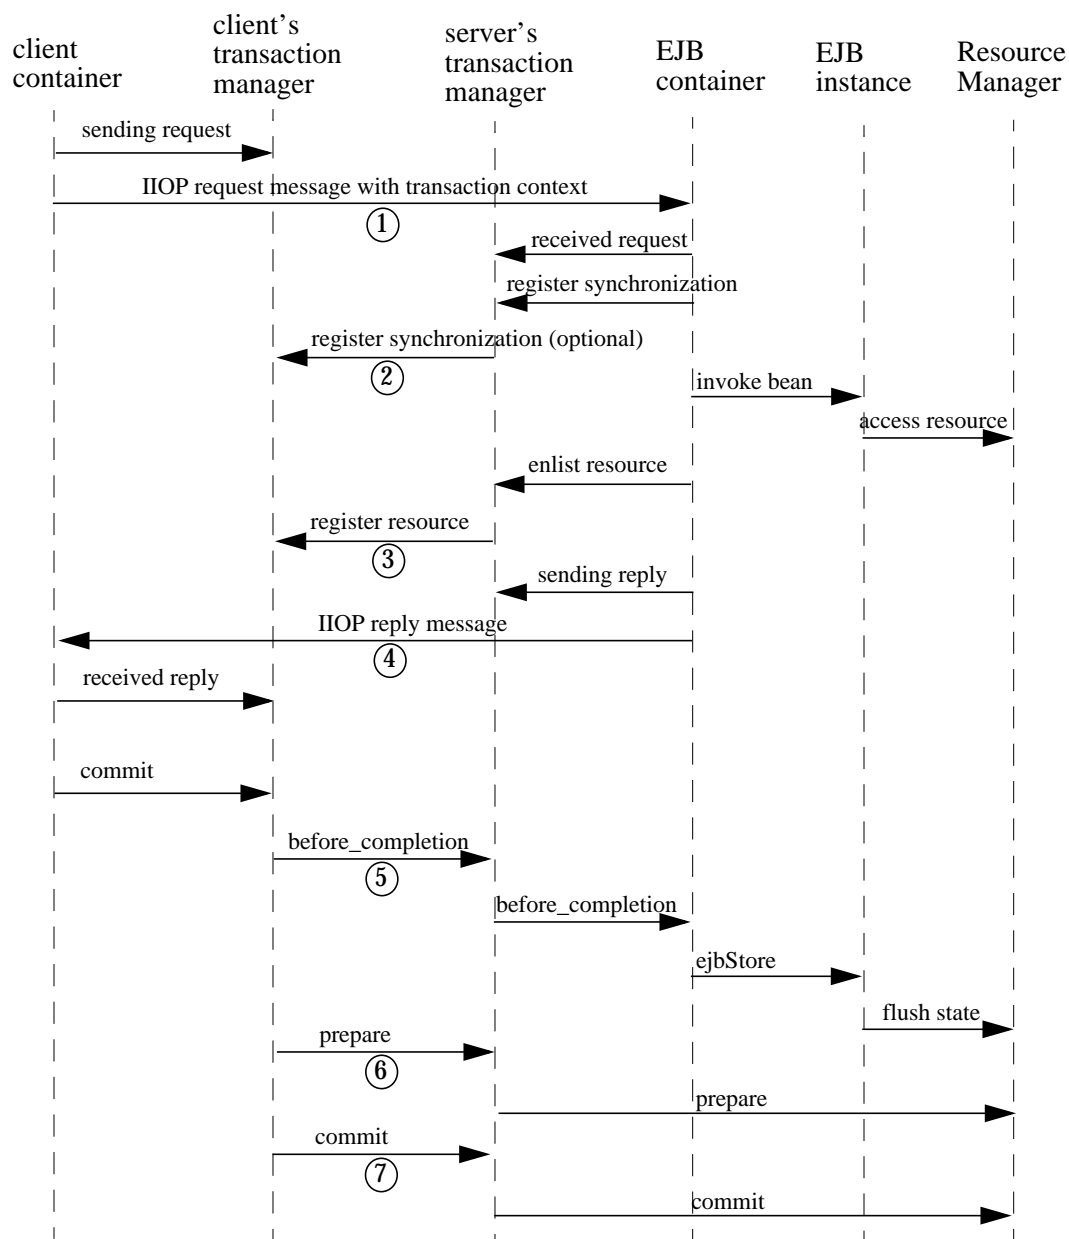
The transaction context format in IIOP messages is specified in the `CorbaTransactions::PropagationContext` structure described in the OTS specification. EJB containers that support transaction interoperability are required to be capable of producing and consuming transaction contexts in IIOP messages in the format described in the OTS specification. Web containers that support transaction interoperability are required to include client-side libraries which can produce the OTS transaction context for sending over IIOP.

Note that it is not necessary for containers to include the Java mappings of the OTS APIs. A container may implement the requirements in the OTS specification in any manner, for example using a non-Java OTS implementation, or an on-the-wire bridge between an existing transaction manager and the OTS protocol, or an OTS wrapper around an existing transaction manager.

The `CorbaTransactions::PropagationContext` structure must be included in IIOP messages sent by web or EJB containers when required by the rules described in the OTS 1.2 specification. The target EJB container must process IIOP invocations based on the transaction policies of `EJBObject` or `EJBHome` references using the rules described in the OTS 1.2 specification [8].

19.6.1.2 Two-Phase Commit Protocol

The object interaction diagram in Figure 78 illustrates the interactions between the client and server transaction managers for transaction context propagation, resource and synchronization object registration, and two-phase commit. This diagram is an example of the interactions between the various entities; it is not intended to be prescriptive.

Figure 78 Transaction Context Propagation

Containers that perform transactional work within the scope of a transaction must register an OTS Resource object with the transaction coordinator whose object reference is included in the propagated transaction context (step 3), and may also register an OTS Synchronization object (step 2). If the server container does not register an OTS Synchronization object, it must still ensure that the `beforeCompletion` method of session beans and `ejbStore` method of entity beans are called with the proper transaction context. Containers must participate in the two-phase commit and recovery procedures performed by the transaction coordinator / terminator (steps 6,7), as described by the OTS specification.

Compliant J2EE containers must not use nested transactions in interoperability scenarios.

19.6.1.3 Transactional Policies of Enterprise Bean References

The OTS1.2 specification describes the `CoTransactions::OTSPolicy` and `CoTransactions::InvocationPolicy` structures that are encoded in IORs as tagged components. EJBObject and EJBHome references must contain these tagged components^[45] with policy values as described below.

The transaction attributes of enterprise beans can be specified per method, while in OTS the entire CORBA object has the same OTS transaction policy. The rules below ensure that the transaction context will be propagated if any method of an enterprise bean needs to execute in the client's transaction context. However, in some cases there may be extra performance overhead of propagating the client's transaction context even if it will not be used by the enterprise bean method.

EJBObject and EJBHome references may have the `InvocationPolicy` value as either `CoTransactions::SHARED` or `CoTransactions::EITHER`^[46].

All EJBObject and EJBHome references must have the `OTSPolicy` value as `CoTransactions::ADAPTS`. This is necessary to allow clients to invoke methods of the `javax.ejb.EJBObject` and `javax.ejb.EJBHome` with or without a transaction.

The `CoTransactions::Synchronization` object registered by the EJB container with the transaction coordinator should have the `OTSPolicy` value `CoTransactions::ADAPTS` and `InvocationPolicy` value `CoTransactions::SHARED` to allow enterprise beans to do transactional work during the `beforeCompletion` notification from the transaction coordinator.

19.6.1.4 Exception Handling Behavior

The exception handling behavior described in the OTS1.2 specification must be followed. In particular, if an application exception (an exception which is not a CORBA system exception and does not extend `java.rmi.RemoteException`) is returned by the server, the request is defined as being successful; hence the client-side OTS library must not roll back the transaction. This allows application exceptions to be propagated back to the client without rolling back the transaction, as required by the exception handling rules in Chapter 18.

19.6.2 Interoperating with Containers that do not Implement Transaction

[45] One way to include the tagged components in IORs is to create the object references using a Portable Object Adapter (POA) which is initialized with the appropriate transaction policies. Note that POA APIs are not required to be supported by server containers.

[46] If the `InvocationPolicy` is not present in the IOR, it is interpreted by the client as if the policy value was `CoTransactions::EITHER`.

Interoperability

The requirements in this subsection are designed to ensure that when a J2EE container does not support transaction interoperability, the failure modes are well defined so that the integrity of an application's data is not compromised: at worst the transaction is rolled back. When a J2EE client component expects the client's transaction to propagate to the enterprise bean but the client or EJB container cannot satisfy this expectation, a `java.rmi.RemoteException` or subclass is thrown, which ensures that the client's transaction will roll back.

In addition, the requirements below allow a container that does not support transaction propagation to interoperate with a container that does support transaction propagation in the cases where the enterprise bean method's transaction attribute indicates that the method would not be executed in the client's transaction.

19.6.2.1 Client Container Requirements

If the client in another container invokes an enterprise bean's method when there is no active global transaction associated with the client's thread, the client container does not include a transaction context in the IIOP request message to the EJB server, i.e., there is no `CosTransactions::PropagationContext` structure in the IIOP request header.

The client application component expects a global transaction to be propagated to the server only if the client's thread has an active global transaction. In this scenario, if the client container does not support transaction interoperability, it has two options:

1. If the client container does not support transaction propagation or uses a non-OTS protocol, it must include the OTS `CosTransactions::PropagationContext` structure in the IIOP request to the server (step 1 in the object interaction diagram above), with the `CosTransactions::Coordinator` and `CosTransactions::Terminator` object references as null. The remaining fields in this "null transaction context," such as the transaction identifier, are not interpreted and may have any value. The "null transaction context" indicates that there is a global client transaction active but the client container is not capable of propagating it to the server. The presence of this "null transaction context" allows the EJB container to determine whether the J2EE client component expects the client's global transaction to propagate to the server.
2. Client containers that use the OTS transaction context format but still do not support transaction interoperability with other vendor's containers must reject the `Coordinator::register_resource` call (step 3 in the object interaction diagram above) by throwing a CORBA system exception if the server's Resource object reference indicates that it belongs to another vendor's container.

19.6.2.2 EJB container requirements

All EJB containers (including those that do not support transaction propagation) must include the `CosTransactions::OTSPolicy` and optionally the `CosTransactions::InvocationPolicy` tagged component in the IOR for `EJBObject` and `EJBHome` references as described in section 19.6.1.3.

19.6.2.2.1 Requirements for EJB Containers Supporting Transaction Interoperability

When an EJB container that supports transaction propagation receives the IIOP request message, it must behave as follows:

- If there is no OTS transaction context in the IIOP message, the container must follow the behavior described in Section 17.6.
- If there is a valid, complete OTS transaction context in the IIOP message, the container must follow the behavior described in Section 17.6.
- If there is a null transaction context (as defined in section 19.6.2.1 above) in the IIOP message, the container's required behavior is described in the table below. The entry "throw RemoteException" indicates that the EJB container must throw the `java.rmi.RemoteException` to the client after the "received request" interaction with the server's transaction manager (after step 1 in the object interaction diagram above).

EJB method's Transaction Attribute	EJB container behavior on receiving null OTS transaction context
Mandatory	throw RemoteException
Required	throw RemoteException
RequiresNew	follow Section 17.6
Supports	throw RemoteException
NotSupported	follow Section 17.6
Never	follow Section 17.6
Bean Managed	follow Section 17.6

19.6.2.2.2 Requirements for EJB Containers not Supporting Transaction Interoperability

When an EJB container that does not support transaction interoperability receives the IIOP request message, it must behave as follows:

- If there is no OTS transaction context in the IIOP message, the container must follow the behavior described in Section 17.6.
- If there is a valid, complete OTS transaction context in the IIOP message, the container's required behavior is described in the table below.

- If there is a null transaction context (as defined in section 19.6.2.1) in the IIOP message, the container's required behavior is described in the table below. Note that the container may not know whether the received transaction context in the IIOP message is valid or null.

EJB method's Transaction Attribute	EJB container behavior on receiving null or valid OTS transaction context
Mandatory	throw RemoteException
Required	throw RemoteException
RequiresNew	follow Section 17.6
Supports	throw RemoteException
NotSupported	follow Section 17.6
Never	follow Section 17.6
Bean Managed	follow Section 17.6

EJB containers that accept the OTS transaction context format but still do not support interoperability with other vendors' client containers must follow the column in the table above for "null or valid OTS transaction context" if the transaction identity or the Coordinator object reference in the propagated client transaction context indicate that the client belongs to a different vendor's container.

19.7 Naming Interoperability

This section describes the requirements for supporting interoperable access to naming services for looking up EJBHome object references (interoperability requirement two in section 19.4).

EJB containers are required to be able to publish EJBHome object references in a CORBA CosNaming service [15]. The CosNaming service must implement the IDL interfaces in the CosNaming module defined in [15] and allow clients to invoke the `resolve` and `list` operations over IIOP.

The CosNaming service must follow the requirements in the CORBA Interoperable Name Service specification [16] for providing the host, port, and object key for its root NamingContext object. The CosNaming service must be able to service IIOP invocations on the root NamingContext at the advertised host, port, and object key.

Client containers (i.e., EJB, web, or application client containers) are required to include a JNDI CosNaming service provider that uses the mechanisms defined in the Interoperable Name Service specification to contact the server's CosNaming service, and to resolve the EJBHome object using standard CosNaming APIs. The JNDI CosNaming service provider may or may not use the JNDI SPI architecture. The JNDI CosNaming service provider must access the root NamingContext of the server's CosNaming service by creating an object reference from the URL `corbaloc:iiop:1.2@<host>:<port>/<objectkey>` (where `<host>`, `<port>`, and `<objectkey>` are the values corresponding to the root NamingContext advertised by the server's CosNaming service), or by using an equivalent mechanism.

At deployment time, the Deployer of the client container should obtain the host, port and object key of the server's CosNaming service and the CosNaming name of the server EJBHome object (e.g. by browsing the server's namespace) for each `ejb-ref` element in the client component's deployment descriptor. The `ejb-ref-name` (which is used by the client code in the JNDI lookup call) should then be linked to the EJBHome object's CosNaming name. At run-time, the client component's JNDI lookup call uses the CosNaming service provider, which contacts the server's CosNaming service, resolves the CosNaming name, and returns the EJBHome object reference to the client component.

Since the EJBHome object's name is scoped within the namespace of the CosNaming service that is accessible at the provided host and port, it is not necessary to federate the namespaces of the client and server containers.

The advantage of using CosNaming is better integration with the IIOP infrastructure that is already required for interoperability, as well as interoperability with non-J2EE CORBA clients and servers. Since CosNaming stores only CORBA objects it is likely that vendors will use other enterprise directory services for storing other resources.

Security of CosNaming service access is achieved using the security interoperability protocol described in Section 19.8. The CosNaming service must support this protocol. Clients which construct the root NamingContext object reference from a URL should send an IIOP `LocateRequest` message to the CosNaming service to obtain the complete IOR (with SSL information) of the root NamingContext, and then initiate an SSL session with the CosNaming service, as determined by the client policy.

19.8 Security Interoperability

This section describes the interoperable mechanisms that support secure invocations on enterprise beans in intranets. These mechanisms are based on the CORBA/IIOP protocol.

EJB containers are required to follow the protocol rules prescribed by the CSIv2 specification Conformance Level 0.


19.8.1 Introduction

The goal of the secure invocation mechanisms is to support the interoperability requirements described earlier in this chapter, as well as be capable of supporting security technologies that are expected to be widely deployed in enterprises, including Kerberos-based secret key mechanisms and X.509 certificate-based public key mechanisms.

The authentication identity (i.e. principal) associated with a J2EE component is usually that of the user on whose behalf the component is executing^[47]. The principal under which an enterprise bean invocation is performed is either that of the bean's caller or the run-as principal which was configured by the Deployer. When there is a chain of invocations across a web component and enterprise beans, an intermediate component may use the principal of the caller (the initiating client) or the intermediate component may use its run-as principal to perform an invocation on the callee, depending on the security identity specified for the intermediate component in its deployment descriptor.

The security principal associated with a container depends on the type of container. Application client containers usually do not have a separate principal associated with them (they operate under the user's principal). Web and EJB containers are typically associated with a security principal of their own (e.g., the operating system user for the container's process) which may be configured by the System Administrator at deployment time. When the client is a web or EJB container, the difference between the client component's principal and the client container's principal is significant for interoperability considerations.

19.8.1.1 Trust Relationships Between Containers, Principal Propagation



When there is a chain of multiple invocations across web components and enterprise beans, intermediate components may not have access to the authentication data of the initiating client to provide proof of the client's identity to the target. In such cases, the target's authentication requirements can be satisfied if the target container trusts the intermediate container to vouch for the authenticity of the propagated principal. The call is made using the intermediate container's principal and authentication data, while also carrying the propagated principal of the initiating client. The invocation on the target enterprise bean is authorized and performed using the propagated principal. This procedure also avoids the overhead associated with authentication of clients on every remote invocation in a chain.

EJB containers are required to provide the Deployer or Administrator with the tools to configure trust relationships for interactions with intermediate web or EJB containers^[48]. If a trust relationship is set up, the containers are usually configured to perform mutual authentication, unless the security of the network can be ensured by some physical means. If the network is physically secure, the target EJB container may be configured to trust all client containers. After a trust relationship is set up, the target EJB container does not need to independently authenticate the initiating client principal sent by the intermediate container on invocations. Thus only the principal name of the initiating client (which may include a realm) needs to be propagated. After a trust relationship has been established, the target EJB container must be able to accept invocations carrying only the principal name of the initiating client.

[47] When there are concurrent invocations on a component from multiple clients, a different principal may be associated with the thread of execution for each invocation.

[48] One way to achieve this is to configure a "trusted container list" for each EJB container which contains the list of intermediate client containers that are trusted. If the list is empty, then the target EJB container does not have a trust relationship with any intermediate container.

For the current interoperability needs of J2EE, it is assumed that trust relationships are transitive, such that if a target container trusts an intermediate container, it implicitly trusts all containers trusted by the intermediate container.

If no trust relationship has been set up between a target EJB container and an intermediate web or EJB container, the target container must not accept principals propagated from that intermediate container, hence the target container needs to have access to and independently verify the initiating client principal's authentication data.

Web and EJB containers are required to support caller propagation mode (where the initiating client's principal is propagated down the chain of calls on enterprise beans) and run-as mode (where the web/EJB component's run-as identity is propagated). This is needed for scenarios 1, 3 and 4 where the internet or intranet user's principal needs to be propagated to the target EJB container.

19.8.1.2 Application Client Authentication

Application client containers that have authentication infrastructure (such as certificates, Kerberos) can:

- authenticate the user by interacting with an authentication service (e.g. the Kerberos KDC) in the enterprise
- inherit an authentication context which was established at system login time from the operating system process, or
- obtain the user's certificate from a client-side store.

These may be achieved by plugging in a Java Authentication and Authorization Service (JAAS) login module for the particular authentication service. After authentication is completed, a credential is associated with the client's thread of execution, which is used for all invocations on enterprise beans made from that thread.

If there is no authentication infrastructure installed in the client's environment, or the authentication infrastructure is not capable of authenticating at the transport protocol layer, the client may send its private credentials (e.g. password) over a secure connection to the EJB server, which authenticates the user by interacting with an authentication service (e.g. a secure user/password database). This is similar to the basic authentication feature of HTTP.

19.8.2 Securing EJB Invocations

This subsection describes the interoperable protocol requirements for providing authentication, protection of integrity and confidentiality, and principal propagation for invocations on enterprise beans. The invocation takes place over an enterprise's intranet as described in the scenarios in section 19.3. Since EJB invocations use the IIOP protocol, we need to secure IIOP messages between client and server containers. The client container may be any of the J2EE containers; the server container is an EJB container.

The secure interoperability requirements for EJB 2.0 (and later) and other J2EE 1.3 (and later) containers are based on Conformance Level 0 of the Common Secure Interoperability version 2 (CSiv2) Final Available specification [20], which was developed by the OMG. EJB, web, and application client containers must support all requirements of Conformance Level 0 of the CSiv2 specification. The following subsections describe how the CSiv2 features are used to realize the scenarios described in section 19.3.

19.8.2.1 Secure Transport Protocol

The Secure Sockets Layer (SSL 3.0) protocol [19] and the related IETF standard Transport Layer Security (TLS 1.0) protocol [17] provide authentication and message protection (that is, integrity and/or confidentiality) at the transport layer. The original SSL and TLS specifications supported only X.509 certificates for authenticating principals. Recently, Kerberos-based authentication mechanisms and cipher suites have been defined for TLS (RFC 2712 [18]). Thus the TLS specification is capable of supporting the two main security technologies that are expected to be widely deployed in enterprises.

EJB, web and application client containers are required to support both SSL 3.0 and TLS 1.0 as security protocols for IIOP. This satisfies interoperability requirement 3 in section 19.4. Compliant containers must be capable of using the following public key SSL/TLS ciphersuites based on policies set by the System Administrator:

- TLS_RSA_WITH_RC4_128_MD5
- SSL_RSA_WITH_RC4_128_MD5
- TLS_DHE_DSS_WITH_3DES_EDE_CBC_SHA^[49]
- SSL_DHE_DSS_WITH_3DES_EDE_CBC_SHA
- TLS_RSA_EXPORT_WITH_RC4_40_MD5
- SSL_RSA_EXPORT_WITH_RC4_40_MD5
- TLS_DHE_DSS_EXPORT_WITH_DES40_CBC_SHA
- SSL_DHE_DSS_EXPORT_WITH_DES40_CBC_SHA

Support for Kerberos ciphersuites is not specified.

When using IIOP over SSL, a secure channel between client and server containers is established at the SSL layer. The SSL handshake layer handles authentication (either mutual or server-only) between containers, negotiation of cipher suite for bulk data encryption, and optionally provides a compression method. The SSL record layer performs confidentiality and integrity protection on application data. Since compliant J2EE products are already required to support SSL (HTTPS for Internet communication), the use of SSL/TLS provides a relatively easy route to interoperable security at the transport layer.

[49] This ciphersuite is mandatory for compliant TLS implementations as specified in [17].

19.8.2.2 Security Information in IORs

Before initiating a secure connection to the EJB container, the client needs to know the hostname and port number at which the server is listening for SSL connections, and the security protocols supported or required by the server object. This information is obtained from the EJBObject or EJBHome reference's IOR.

The CSiv2 specification [20] describes the TAG_CSI_SEC_MECH_LIST tagged component which is included in the IORs of secured objects. This component contains a sequence of CSI_IOP::CompoundSecMech structures (in decreasing order of the server's preference) that contain the target object's security information for transport layer and service context layer mechanisms. This information includes the server's SSL/TLS port, its security principal and supported/required security mechanisms.


EJB containers must be capable of inserting the CSiv2 tagged components into the IORs for EJBObject and EJBHome references, based on the Deployer or System Administrator's security policy settings. Compliant EJB containers must follow the Conformance Level 0 rules described in the CSiv2 specification for constructing these IORs.

EJB containers must also be capable of creating IORs that allow access to enterprise beans over unprotected IIOP, based on the security policies set by the Deployer or System Administrator.

19.8.2.3 Propagating Principals and Authentication Data in IIOP Messages

In scenarios where client authentication does not occur at the transport layer it is necessary to support transfer of authentication data between two containers in the IIOP message service context. When an intermediate client container does not have authentication data for the initiating client, it is necessary to support propagation of client principals in the IIOP message service context.

It is assumed that all information exchanged between client and server at the transport layer is also available to the containers: e.g. the certificates used for authentication at the SSL layer may be used by the server container for authorization.



The following cases are required to be supported:

1. Application client invocations on enterprise beans with mutual authentication between the application client and EJB container (C and S1) at the SSL layer (scenario 2.1 in section 19.3.2, interoperability requirement 4.1 in section 19.4). For example, this is possible when the enterprise has a Kerberos-based authentication infrastructure or when client-side certificates have been installed. In this case the security context of the IIOP message sent from C to S1 should not contain any additional information.

2. Application client invocations on enterprise beans with server-only authentication between the application client and EJB container (C and S1) at the SSL layer (scenario 2.2 in section 19.3.2, interoperability requirement 4.2 in section 19.4). This usually happens when the client cannot authenticate in the transport. In this case, the client container must be capable of inserting into the IIOP message a CSIV2 security context with a client authentication token that contains the client C's authentication data. Once the EJB container S1 has authenticated the client, it may or may not maintain state about the client, so subsequent invocations from the client on the same network connection may need to be authenticated again. The client and server containers must follow the Conformance Level 0 rules in the CSIV2 specification for client authentication. In particular, support for the GSSUP username-password authentication mechanism is required. Support for other GSSAPI mechanisms (such as Kerberos) to perform client authentication at the IIOP layer is optional.
3. Invocations from Web/EJB clients to enterprise beans with a trust relationship between the client container S1 and server container S2 (scenarios 1,3 and 4 in section 19.3.3, interoperability requirements five and six in section 19.4). S2 does not need to independently authenticate the initiating client C. In this case the client container S1 must insert into the IIOP message a security context with an identity token in the format described in the CSIV2 specification. The principal may be propagated as an X.509 certificate chain or as a X.501 distinguished name or as a principal name encoded in the GSS exported name format, as described in the CSIV2 specification. The identity propagated is determined as follows:
 - If the client Web/EJB component is configured to use caller identity, and the caller C authenticated itself to S1, then the identity token contains the initiating client C's identity.
 - If the client component is configured to use caller identity, and the caller C did not authenticate itself to S1, then the identity token contains the anonymous type.
 - If the client component is configured to use a run-as identity then the identity token contains the run-as identity.

J2EE containers are required to support the stateless mode of propagating principal and authentication information defined in CSIV2 (where the server does not store any state for a particular client principal across invocations), and may optionally support the stateful mode.

The caller principal String provided by `EJBContext.getCallerPrincipal().getName()` is defined as follows:

- For case one, the principal should be derived from the distinguished name obtained from the first X.509 certificate in the client's certificate chain that was provided to the server during SSL mutual authentication.
- For case two, the principal should be derived from the username obtained from the client authentication token in the CSIV2 security context of the IIOP message. For the GSSUP username-password mechanism, the principal should be derived from the username in the `GSSUP::InitialContextToken` structure.
- For case three, the principal depends on the identity token type in the CSIV2 security context:
 - If the type is X.509 certificate chain, then the principal should be derived from the distinguished name from the first certificate in the chain.

- If the type is distinguished name, then the principal should be derived from the distinguished name.
- If the type is principal name propagated as a GSS exported name, then the principal should be derived from the mechanism-specific principal name.
- If the anonymous principal type was propagated or the identity token was absent, then `EJBContext.getCallerPrincipal().getName()` returns a product-specific unauthenticated principal name.

19.8.2.4 Security Configuration for Containers

Since the interoperability scenarios involve IIOP/SSL usage in intranets, it is assumed that client and server container administrators cooperatively configure a consistent set of security policies for the enterprise.

At product installation or application deployment time, client and server container administrators may optionally configure the container and SSL infrastructure as described below. These preferences may be specified at any level of granularity (e.g. per host or per container process or per enterprise bean).

- Configure the list of supported SSL cipher suites in preference order.
- For server containers, configure a list of trusted client container principals with whom the server has a trust relationship.
- Configure authentication preferences and requirements (e.g. if the server prefers authenticated clients to anonymous clients). In particular, if a trust relationship has been configured between two servers, then mutual authentication should be required unless there is physical network security.
- If the client and server are using certificates for authentication, configure a trusted common certificate authority for both client and server. If using Kerberos, configure the client and server with the same KDC or cooperating KDCs.
- Configure a restricted list of trusted server principals that a client container is allowed to interact with, to prevent the client's private credentials such as password from being sent to untrusted servers.

19.8.2.5 Runtime Behavior

Client containers determine whether to use SSL for an enterprise bean invocation by using the security policies configured by the client administrator for interactions with the target host or enterprise bean, and the `target_requires` information in the CSIV2 tagged component in the target enterprise bean's IOR. If either the client configuration requires secure interactions with the enterprise bean, or the enterprise bean requires a secure transport, the client should initiate an SSL connection to the server. The client must follow the rules described in the CSIV2 specification Conformance Level 0 for interpreting security information in IORs and including security context information in IIOP messages.

When an EJB container receives an IIOP message, its behavior depends on deployment time configuration, run-time information exchanged with the client at the SSL layer, and principal/authentication data contained in the IIOP message service context. EJB containers are required to follow the protocol rules prescribed by the CSIV2 specification Conformance Level 0.

When the System Administrator changes the security policies associated with an enterprise bean, the IORs for EJB references should be updated. When the bean has existing clients holding IORs, it is recommended that the security policy change should be handled by the client and server containers transparently to the client application if the old security policy is compatible with the new one. This may be done by using interoperable GIOP 1.2 forwarding mechanisms.

Enterprise Bean Environment

This chapter specifies the interfaces for accessing the enterprise bean environment.

20.1 Overview

The Application Assembler and Deployer should be able to customize an enterprise bean's business logic without accessing the enterprise bean's source code.

In addition, ISVs typically develop enterprise beans that are, to a large degree, independent from the operational environment in which the application will be deployed. Most enterprise beans must access resource managers and external information. The key issue is how enterprise beans can locate external information without prior knowledge of how the external information is named and organized in the target operational environment.

The enterprise bean environment mechanism attempts to address both of the above issues.

This chapter is organized as follows.

- Section 20.2 defines the interfaces that specify and access the enterprise bean's environment. The section illustrates the use of the enterprise bean's environment for generic customization of the enterprise bean's business logic.

- Section 20.3 defines the interfaces for obtaining the home interface of another enterprise bean using an *EJB reference*. An EJB reference is a special entry in the enterprise bean's environment.
- Section 20.4 defines the interfaces for obtaining the web service interface using a *web service reference*. A web service reference is a special entry in the enterprise bean's environment.
- Section 20.5 defines the interfaces for obtaining a resource manager connection factory using a *resource manager connection factory reference*. A resource manager connection factory reference is a special entry in the enterprise bean's environment.
- Section 20.6 defines the interfaces for obtaining an administered object that is associated with a resource (e.g., a CCI *InteractionSpec*) using a *resource environment reference*. A resource environment reference is a special entry in the enterprise bean's environment.
- Section 20.7 defines the interfaces for obtaining a message destination associated with a resource using a *message destination reference*. Message destination references allow the flow of messages within an application to be specified. A message destination reference is a special entry in the enterprise bean's environment.

20.2 Enterprise Bean's Environment as a JNDI Naming Context

The enterprise bean's environment is a mechanism that allows customization of the enterprise bean's business logic during deployment or assembly. The enterprise bean's environment allows the enterprise bean to be customized without the need to access or change the enterprise bean's source code.

The container implements the enterprise bean's environment, and provides it to the enterprise bean instance through the JNDI interfaces^[50]. The enterprise bean's environment is used as follows:

1. The enterprise bean's business methods access the environment using the JNDI interfaces. The Bean Provider declares in the deployment descriptor all the environment entries that the enterprise bean expects to be provided in its environment at runtime.
2. The container provides an implementation of the JNDI naming context that stores the enterprise bean environment. The container also provides the tools that allow the Deployer to create and manage the environment of each enterprise bean.
3. The Deployer uses the tools provided by the container to create the environment entries that are declared in the enterprise bean's deployment descriptor. The Deployer can set and modify the values of the environment entries.
4. The container makes the environment naming context available to the enterprise bean instances at runtime. The enterprise bean's instances use the JNDI interfaces to obtain the values of the environment entries.

[50] The container must make the enterprise bean's environment available to any JAX-RPC message handlers for the bean as well. See [25].

Each enterprise bean defines its own set of environment entries. All instances of an enterprise bean within the same home share the same environment entries; the environment entries are not shared with other enterprise beans. Enterprise bean instances are not allowed to modify the bean's environment at runtime.

If an enterprise bean is deployed multiple times in the same container, each deployment results in the creation of a distinct home. The Deployer may set different values for the enterprise bean environment entries for each home.

In general, lookups of objects in the JNDI `java: namespace` are required to return a new instance of the requested object every time. Exceptions are allowed for the following:

- The container knows the object is immutable (for example, objects of type `java.lang.String`), or knows that the application can't change the state of the object.
- The object is defined to be a singleton, such that only one instance of the object may exist in the JVM.
- The name used for the lookup is defined to return an instance of the object that might be shared. The name `java:comp/ORB` is such a name.

In these cases, a shared instance of the object may be returned. In all other cases, a new instance of the requested object must be returned on each lookup. Note that, in the case of resource adapter connection objects, it is the resource adapter's `ManagedConnectionFactory` implementation that is responsible for satisfying this requirement.

Terminology warning: The enterprise bean's "environment" should not be confused with the "environment properties" defined in the JNDI documentation.

The following subsections describe the responsibilities of each EJB Role.

20.2.1 Bean Provider's Responsibilities

This section describes the Bean Provider's view of the enterprise bean's environment, and defines his or her responsibilities.

20.2.1.1 Access to Enterprise Bean's Environment

An enterprise bean instance locates the environment naming context using the JNDI interfaces. An instance creates a `javax.naming.InitialContext` object by using the constructor with no arguments, and looks up the environment naming via the `InitialContext` under the name `java:comp/env`. The enterprise bean's environment entries are stored directly in the environment naming context, or in any of its direct or indirect subcontexts.

The value of an environment entry is of the Java type declared by the Bean Provider in the deployment descriptor.

The following code example illustrates how an enterprise bean accesses its environment entries.

```
public class EmployeeServiceBean implements SessionBean {
    ...
    public void setTaxInfo(int numberOfExemptions, ...)
        throws InvalidNumberOfExemptionsException {
        ...

        // Obtain the enterprise bean's environment naming context.
        Context initCtx = new InitialContext();
        Context myEnv = (Context)initCtx.lookup("java:comp/env");

        // Obtain the maximum number of tax exemptions
        // configured by the Deployer.
        Integer max = (Integer)myEnv.lookup("maxExemptions");

        // Obtain the minimum number of tax exemptions
        // configured by the Deployer.
        Integer min = (Integer)myEnv.lookup("minExemptions");

        // Use the environment entries to customize business logic.
        if (numberOfExemptions > Integer.intValue(max) ||
            numberOfExemptions < Integer.intValue(min))
            throw new InvalidNumberOfExemptionsException();

        // Get some more environment entries. These environment
        // entries are stored in subcontexts.
        String val1 = (String)myEnv.lookup("foo/name1");
        Boolean val2 = (Boolean)myEnv.lookup("foo/bar/name2");

        // The enterprise bean can also lookup using full pathnames.
        Integer val3 = (Integer)
            initCtx.lookup("java:comp/env/name3");
        Integer val4 = (Integer)
            initCtx.lookup("java:comp/env/foo/name4");
        ...
    }
}
```

20.2.1.2 Declaration of Environment Entries

The Bean Provider must declare all the environment entries accessed from the enterprise bean's code. The environment entries are declared using the `env-entry` elements in the deployment descriptor.

Each `env-entry` element describes a single environment entry. The `env-entry` element consists of an optional description of the environment entry, the environment entry name relative to the `java:comp/env` context, the expected Java type of the environment entry value (i.e., the type of the object returned from the JNDI lookup method), and an optional environment entry value.

An environment entry is scoped to the enterprise bean whose declaration contains the `env-entry` element. This means that the environment entry is inaccessible from other enterprise beans at runtime, and that other enterprise beans may define `env-entry` elements with the same `env-entry-name` without causing a name conflict.

The environment entry values may be one of the following Java types: `String`, `Character`, `Integer`, `Boolean`, `Double`, `Byte`, `Short`, `Long`, and `Float`.

If the Bean Provider provides a value for an environment entry using the `env-entry-value` element, the value can be changed later by the Application Assembler or Deployer. The value must be a string that is valid for the constructor of the specified type that takes a single `String` parameter, or for `java.lang.Character`, a single character.

The following example is the declaration of environment entries used by the `EmployeeServiceBean` whose code was illustrated in the previous subsection.

```
<enterprise-beans>
  <session>
    ...
    <ejb-name>EmployeeService</ejb-name>
    <ejb-class>com.wombat.empl.EmployeeServiceBean</ejb-class>
    ...
    <env-entry>
      <description>
        The maximum number of tax exemptions
        allowed to be set.
      </description>
      <env-entry-name>maxExemptions</env-entry-name>
      <env-entry-type>java.lang.Integer</env-entry-type>
      <env-entry-value>15</env-entry-value>
    </env-entry>
    <env-entry>
      <description>
        The minimum number of tax exemptions
        allowed to be set.
      </description>
      <env-entry-name>minExemptions</env-entry-name>
      <env-entry-type>java.lang.Integer</env-entry-type>
      <env-entry-value>1</env-entry-value>
    </env-entry>
    <env-entry>
      <env-entry-name>foo/name1</env-entry-name>
      <env-entry-type>java.lang.String</env-entry-type>
      <env-entry-value>value1</env-entry-value>
    </env-entry>
    <env-entry>
      <env-entry-name>foo/bar/name2</env-entry-name>
      <env-entry-type>java.lang.Boolean</env-entry-type>
      <env-entry-value>true</env-entry-value>
    </env-entry>
    <env-entry>
      <description>Some description.</description>
      <env-entry-name>name3</env-entry-name>
      <env-entry-type>java.lang.Integer</env-entry-type>
    </env-entry>
    <env-entry>
      <env-entry-name>foo/name4</env-entry-name>
      <env-entry-type>java.lang.Integer</env-entry-type>
      <env-entry-value>10</env-entry-value>
    </env-entry>
    ...
  </session>
</enterprise-beans>
...
```

20.2.2 Application Assembler's Responsibility

The Application Assembler is allowed to modify the values of the environment entries set by the Bean Provider, and is allowed to set the values of those environment entries for which the Bean Provider has not specified any initial values.

20.2.3 Deployer's Responsibility

The Deployer must ensure that the values of all the environment entries declared by an enterprise bean are set to meaningful values.

The Deployer can modify the values of the environment entries that have been previously set by the Bean Provider and/or Application Assembler, and must set the values of those environment entries for which no value has been specified.

The `description` elements provided by the Bean Provider or Application Assembler help the Deployer with this task.

20.2.4 Container Provider Responsibility

The Container Provider has the following responsibilities:

- Provide a deployment tool that allows the Deployer to set and modify the values of the enterprise bean's environment entries.
- Implement the `java:comp/env` environment naming context, and provide it to the enterprise bean instances at runtime. The naming context must include all the environment entries declared by the Bean Provider, with their values supplied in the deployment descriptor or set by the Deployer. The environment naming context must allow the Deployer to create subcontexts if they are needed by an enterprise bean.
- The container must ensure that the enterprise bean instances have only read access to their environment variables. The container must throw the `javax.naming.OperationNotSupportedException` from all the methods of the `javax.naming.Context` interface that modify the environment naming context and its subcontexts.

20.3 EJB References

This section describes the programming and deployment descriptor interfaces that allow the Bean Provider to refer to the homes of other enterprise beans using “logical” names called *EJB references*. The EJB references are special entries in the enterprise bean's environment. The Deployer binds the EJB references to the enterprise bean homes in the target operational environment.

The deployment descriptor also allows the Application Assembler to link an EJB reference declared in one enterprise bean to another enterprise bean contained in the same `ejb-jar` file, or in another `ejb-jar` file in the same J2EE application unit. The link is an instruction to the tools used by the Deployer that the EJB reference should be bound to the home of the specified target enterprise bean.

20.3.1 Bean Provider's Responsibilities

This subsection describes the Bean Provider's view and responsibilities with respect to EJB references.

20.3.1.1 EJB Reference Programming Interfaces

The Bean Provider must use EJB references to locate the home interfaces of other enterprise beans as follows.

- Assign an entry in the enterprise bean's environment to the reference. (See subsection 20.3.1.2 for information on how EJB references are declared in the deployment descriptor.)
- *The EJB specification recommends, but does not require, that all references to other enterprise beans be organized in the `ejb` subcontext of the bean's environment (i.e., in the `java:comp/env/ejb` JNDI context).*
- Look up the home interface of the referenced enterprise bean in the enterprise bean's environment using JNDI.

The following example illustrates how an enterprise bean uses an EJB reference to locate the remote home interface of another enterprise bean.

```
public class EmployeeServiceBean implements SessionBean {  
    public void changePhoneNumber(...) {  
        ...  
        // Obtain the default initial JNDI context.  
        Context initCtx = new InitialContext();  
  
        // Look up the home interface of the EmployeeRecord  
        // enterprise bean in the environment.  
        Object result = initCtx.lookup(  
            "java:comp/env/ejb/EmplRecord");  
  
        // Convert the result to the proper type.  
        EmployeeRecordHome emplRecordHome = (EmployeeRecordHome)  
            javax.rmi.PortableRemoteObject.narrow(result,  
                EmployeeRecordHome.class);  
        ...  
    }  
}
```

In the example, the Bean Provider of the `EmployeeServiceBean` enterprise bean assigned the environment entry `ejb/EmplRecord` as the EJB reference name to refer to the remote home of another enterprise bean.

20.3.1.2 Declaration of EJB References in Deployment Descriptor

Although the EJB reference is an entry in the enterprise bean's environment, the Bean Provider must not use a `env-entry` element to declare it. Instead, the Bean Provider must declare all the EJB references using the `ejb-ref` and `ejb-local-ref` elements of the deployment descriptor. This allows the `ejb-jar` consumer (i.e. Application Assembler or Deployer) to discover all the EJB references used by the enterprise bean.

Each `ejb-ref` or `ejb-local-ref` element describes the interface requirements that the referencing enterprise bean has for the referenced enterprise bean. The `ejb-ref` element is used for referencing an enterprise bean that is accessed through its remote home and remote interfaces. The `ejb-local-ref` element is used for referencing an enterprise bean that is accessed through its local home and local interfaces.

The `ejb-ref` element contains an optional `description` element, and the mandatory `ejb-ref-name`, `ejb-ref-type`, `home`, and `remote` elements.

The `ejb-local-ref` element contains an optional `description` element, and the mandatory `ejb-ref-name`, `ejb-ref-type`, `local-home`, and `local` elements.

The `ejb-ref-name` element specifies the EJB reference name: its value is the environment entry name used in the enterprise bean code. The `ejb-ref-type` element specifies the expected type of the enterprise bean: its value must be either `Entity` or `Session`. The `home` and `remote` or `local-home` and `local` elements specify the expected Java types of the referenced enterprise bean's home and component interfaces.

An EJB reference is scoped to the enterprise bean whose declaration contains the `ejb-ref` or `ejb-local-ref` element. This means that the EJB reference is not accessible to other enterprise beans at runtime, and that other enterprise beans may define `ejb-ref` and/or `ejb-local-ref` elements with the same `ejb-ref-name` without causing a name conflict.

The following example illustrates the declaration of EJB references in the deployment descriptor.

```

...
<enterprise-beans>
  <session>
    ...
    <ejb-name>EmployeeService</ejb-name>
    <ejb-class>com.wombat.empl.EmployeeServiceBean</ejb-class>
    ...
    <ejb-ref>
      <description>
        This is a reference to the entity bean that
        encapsulates access to employee records.
      </description>
      <ejb-ref-name>ejb/EmplRecord</ejb-ref-name>
      <ejb-ref-type>Entity</ejb-ref-type>
      <home>com.wombat.empl.EmployeeRecordHome</home>
      <remote>com.wombat.empl.EmployeeRecord</remote>
    </ejb-ref>

    <ejb-ref>
      <ejb-ref-name>ejb/Payroll</ejb-ref-name>
      <ejb-ref-type>Entity</ejb-ref-type>
      <home>com.aardvark.payroll.PayrollHome</home>
      <remote>com.aardvark.payroll.Payroll</remote>
    </ejb-ref>

    <ejb-ref>
      <ejb-ref-name>ejb/PensionPlan</ejb-ref-name>
      <ejb-ref-type>Session</ejb-ref-type>
      <home>com.wombat.empl.PensionPlanHome</home>
      <remote>com.wombat.empl.PensionPlan</remote>
    </ejb-ref>
    ...
  </session>
  ...
</enterprise-beans>
...

```

20.3.2 Application Assembler's Responsibilities

The Application Assembler can use the `ejb-link` element in the deployment descriptor to link an EJB reference to a target enterprise bean.

The Application Assembler specifies the link between two enterprise beans as follows:

- The Application Assembler uses the optional `ejb-link` element of the `ejb-ref` or `ejb-local-ref` element of the referencing enterprise bean. The value of the `ejb-link` element is the name of the target enterprise bean. (It is the name defined in the `ejb-name` element of the target enterprise bean.) The target enterprise bean can be in any `ejb-jar` file in the same J2EE application as the referencing application component.
- Alternatively, to avoid the need to rename enterprise beans to have unique names within an entire J2EE application, the Application Assembler may use the following syntax in the `ejb-link` element of the referencing application component. The Application Assembler

specifies the path name of the ejb-jar file containing the referenced enterprise bean and appends the ejb-name of the target bean separated from the path name by #. The path name is relative to the referencing application component jar file. In this manner, multiple beans with the same ejb-name may be uniquely identified when the Application Assembler cannot change ejb-names.

- The Application Assembler must ensure that the target enterprise bean is type-compatible with the declared EJB reference. This means that the target enterprise bean must be of the type indicated in the ejb-ref-type element, and that the home and component interfaces of the target enterprise bean must be Java type-compatible with the interfaces declared in the EJB reference.

The following illustrates an ejb-link in the deployment descriptor.

```
...
<enterprise-beans>
  <session>
    ...
    <ejb-name>EmployeeService</ejb-name>
    <ejb-class>com.wombat.empl.EmployeeServiceBean</ejb-class>
    ...
    <ejb-ref>
      <ejb-ref-name>ejb/EmplRecord</ejb-ref-name>
      <ejb-ref-type>Entity</ejb-ref-type>
      <home>com.wombat.empl.EmployeeRecordHome</home>
      <remote>com.wombat.empl.EmployeeRecord</remote>
      <ejb-link>EmployeeRecord</ejb-link>
    </ejb-ref>
    ...
  </session>
  ...

  <entity>
    <ejb-name>EmployeeRecord</ejb-name>
    <home>com.wombat.empl.EmployeeRecordHome</home>
    <remote>com.wombat.empl.EmployeeRecord</remote>
    ...
  </entity>
  ...
</enterprise-beans>
...
```

The Application Assembler uses the ejb-link element to indicate that the EJB reference EmplRecord declared in the EmployeeService enterprise bean has been linked to the EmployeeRecord enterprise bean.

The following example illustrates using the `ejb-link` element to indicate an enterprise bean reference to the `ProductEJB` enterprise bean that is in the same J2EE application unit but in a different `ejb-jar` file.

```
<entity>
...
<ejb-name>OrderEJB</ejb-name>
<ejb-class>com.wombat.orders.OrderBean</ejb-class>
...
<ejb-ref>
  <ejb-ref-name>ejb/Product</ejb-ref-name>
  <ejb-ref-type>Entity</ejb-ref-type>
  <home>com.acme.orders.ProductHome</home>
  <remote>com.acme.orders.Product</remote>
  <ejb-link>../products/product.jar#ProductEJB</ejb-link>
</ejb-ref>
...
</entity>
```

20.3.3 Deployer's Responsibility

The Deployer is responsible for the following:

- The Deployer must ensure that all the declared EJB references are bound to the homes of enterprise beans that exist in the operational environment. The Deployer may use, for example, the JNDI `LinkRef` mechanism to create a symbolic link to the actual JNDI name of the target enterprise bean's home.
- The Deployer must ensure that the target enterprise bean is type-compatible with the types declared for the EJB reference. This means that the target enterprise bean must be of the type indicated in the `ejb-ref-type` element, and that the home and component interfaces of the target enterprise bean must be Java type-compatible with the home and component interfaces declared in the EJB reference.
- If an EJB reference declaration includes the `ejb-link` element, the Deployer should bind the enterprise bean reference to the home of the enterprise bean specified as the link's target.

20.3.4 Container Provider's Responsibility

The Container Provider must provide the deployment tools that allow the Deployer to perform the tasks described in the previous subsection. The deployment tools provided by the EJB Container Provider must be able to process the information supplied in the `ejb-ref` and `ejb-local-ref` elements in the deployment descriptor.

At the minimum, the tools must be able to:

- Preserve the application assembly information in the `ejb-link` elements by binding an EJB reference to the home interface of the specified target session or entity bean.
- Inform the Deployer of any unresolved EJB references, and allow him or her to resolve an EJB reference by binding it to a specified compatible target session or entity bean.

20.4 Web Service References

This section describes the programming and deployment descriptor interfaces that allow the Bean Provider to refer to external web services using “logical” names called *web service references*. The web service references are special entries in the enterprise bean’s environment. The Deployer binds the web service references to the web service interfaces in the target operational environment.

20.4.1 Bean Provider’s Responsibilities

This subsection describes the Bean Provider’s view and responsibilities with respect to web service references.

20.4.1.1 Web Service Reference Programming Interfaces

The Bean Provider must use web service references to locate web service interfaces as follows.

- Assign an entry in the enterprise bean’s environment to the reference. (See subsection 20.4.1.2 for information on how web service references are declared in the deployment descriptor.)
- *The EJB specification recommends, but does not require, that all references to web services be organized in the service subcontext of the bean’s environment (i.e., in the `java:comp/env/service` JNDI context).*
- Look up the JAX-RPC web service interface of the referenced service in the enterprise bean’s environment using JNDI.

The following example illustrates how an enterprise bean uses an web service reference to locate a web service interface and invoke a method on the service endpoint.

```
public class InvestmentBean implements SessionBean {  
    public void checkPortfolio(...) {  
        ...  
  
        // Obtain the default initial JNDI context.  
        Context initCtx = new InitialContext();  
  
        // Look up the stock quote service in the environment.  
        com.example.StockQuoteService sqs =  
            (com.example.StockQuoteService) initCtx.lookup(  
                "java:comp/env/service/StockQuoteService");  
  
        // Get the stub for the service endpoint  
        com.example.StockQuoteProvider sqp =  
            sqs.getStockQuoteProviderPort();  
  
        // Get a quote  
        float quotePrice = sqp.getLastTradePrice(...);  
        ...  
    }  
}
```

In the example, the Bean Provider of the `InvestmentBean` enterprise bean assigned the environment entry `service/StockQuoteService` as the web service reference name to refer to the `StockQuoteService` JAX-RPC web service interface.

20.4.1.2 Declaration of Web Service References in Deployment Descriptor

Although the web service reference is an entry in the enterprise bean's environment, the Bean Provider must not use a `env-entry` element to declare it. Instead, the Bean Provider must declare all the web service references using the `service-ref` elements of the deployment descriptor. This allows the `ejb-jar` consumer (i.e. Application Assembler or Deployer) to discover all the web service references used by the enterprise bean.

Each `service-ref` element describes the interface requirements that the referencing enterprise bean has for the referenced web service. The `service-ref` element contains an optional `description` element and the mandatory `service-ref-name` and `service-interface` elements. The use of service references and the `service-ref` deployment descriptor element are described in further detail in [25].

The `service-ref-name` element specifies the web service reference name: its value is the environment entry name used in the enterprise bean code. The `service-interface` element specifies the fully qualified name of the JAX-RPC service interface returned by the JNDI lookup.

A web service reference is scoped to the enterprise bean whose declaration contains the `service-ref` element. This means that the web service reference is not accessible to other enterprise beans at runtime, and that other enterprise beans may define `service-ref` elements with the same `service-ref-name` without causing a name conflict.

The following example illustrates the declaration of a web service reference in the deployment descriptor.

```
<session>
  ...
  <ejb-name>InvestmentBean</ejb-name>
  <ejb-class>com.wombat.empl.InvestmentBean</ejb-class>
  ...
  <service-ref>
    <description>
      This is a reference to the stock quote
      service used to estimate portfolio value.
    </description>
    <service-ref-name>service/StockQuoteService</service-ref-name>
    <service-interface>com.example.StockQuoteService
    </service-interface>
  </service-ref>
  ...
</session>
```

20.4.2 Deployer's Responsibility

The Deployer is responsible for the following:

- The Deployer must ensure that all the declared web service references are bound to classes that exist in the operational environment. The Deployer links a service reference to the actual representation and configuration of the corresponding service.
- The Deployer must provide and configure the required information for the service instance and service endpoint proxies, as described in [22].
- The Deployer ensures that the configuration of the service and service proxy is based on the protocol binding specified in the WSDL description of the service.

20.4.3 Container Provider's Responsibility

The Container Provider must provide the deployment tools that allow the Deployer to perform the tasks described in the previous subsection. The deployment tools provided by the EJB Container Provider must be able to process the information supplied in the `service-ref` elements in the deployment descriptor.

20.5 Resource Manager Connection Factory References

A resource manager connection factory is an object that is used to create connections to a resource manager. For example, an object that implements the `javax.sql.DataSource` interface is a resource manager connection factory for `java.sql.Connection` objects which implement connections to a database management system.

This section describes the enterprise bean programming and deployment descriptor interfaces that allow the enterprise bean code to refer to resource factories using logical names called *resource manager connection factory references*. The resource manager connection factory references are special entries in the enterprise bean's environment. The Deployer binds the resource manager connection factory references to the actual resource manager connection factories that are configured in the container. Because these resource manager connection factories allow the container to affect resource management, the connections acquired through the resource manager connection factory references are called *managed resources* (e.g., these resource manager connection factories allow the container to implement connection pooling and automatic enlistment of the connection with a transaction).

20.5.1 Bean Provider's Responsibilities

This subsection describes the Bean Provider's view of locating resource factories and defines his or her responsibilities.

20.5.1.1 Programming Interfaces for Resource Manager Connection Factory References

The Bean Provider must use resource manager connection factory references to obtain connections to resources as follows.

- Assign an entry in the enterprise bean's environment to the resource manager connection factory reference. (See subsection 20.5.1.2 for information on how resource manager connection factory references are declared in the deployment descriptor.)
- *The EJB specification recommends, but does not require, that all resource manager connection factory references be organized in the subcontexts of the bean's environment, using a different subcontext for each resource manager type. For example, all JDBC™ DataSource references might be declared in the `java:comp/env/jdbc` subcontext, and all JMS connection factories in the `java:comp/env/jms` subcontext. Also, all JavaMail connection factories might be declared in the `java:comp/env/mail` subcontext and all URL connection factories in the `java:comp/env/url` subcontext.*
- Lookup the resource manager connection factory object in the enterprise bean's environment using the JNDI interface.
- Invoke the appropriate method on the resource manager connection factory to obtain a connection to the resource. The factory method is specific to the resource type. It is possible to obtain multiple connections by calling the factory object multiple times.

The Bean Provider can control the shareability of the connections acquired from the resource manager connection factory. By default, connections to a resource manager are shareable across other enterprise beans in the application that use the same resource in the same transaction context. The Bean Provider can specify that connections obtained from a resource manager connection factory reference are not shareable by specifying the value of the `res-sharing-scope` deployment descriptor element to be `Unshareable`. The sharing of connections to a resource manager allows the container to optimize the use of connections and enables the container's use of local transaction optimizations.

The Bean Provider has two choices with respect to dealing with associating a principal with the resource manager access:

- Allow the Deployer to set up principal mapping or resource manager sign-on information. In this case, the enterprise bean code invokes a resource manager connection factory method that has no security-related parameters.
- Sign on to the resource manager from the bean code. In this case, the enterprise bean invokes the appropriate resource manager connection factory method that takes the sign-on information as method parameters.

The Bean Provider uses the `res-auth` deployment descriptor element to indicate which of the two resource manager authentication approaches is used.

We expect that the first form (i.e., letting the Deployer set up the resource manager sign-on information) will be the approach used by most enterprise beans.

The following code sample illustrates obtaining a JDBC connection.

```
public class EmployeeServiceBean implements SessionBean {
    EJBContext ejbContext;

    public void changePhoneNumber(...) {
        ...

        // obtain the initial JNDI context
        Context initCtx = new InitialContext();

        // perform JNDI lookup to obtain resource manager
        // connection factory
        javax.sql.DataSource ds = (javax.sql.DataSource)
            initCtx.lookup("java:comp/env/jdbc/EmployeeAppDB");

        // Invoke factory to obtain a connection. The security
        // principal is not given, and therefore
        // it will be configured by the Deployer.
        java.sql.Connection con = ds.getConnection();
        ...
    }
}
```

20.5.1.2 Declaration of Resource Manager Connection Factory References in Deployment Descriptor

Although a resource manager connection factory reference is an entry in the enterprise bean's environment, the Bean Provider must not use an `env-entry` element to declare it.

Instead, the Bean Provider must declare all the resource manager connection factory references in the deployment descriptor using the `resource-ref` elements. This allows the `ejb-jar` consumer (i.e. Application Assembler or Deployer) to discover all the resource manager connection factory references used by an enterprise bean.

Each `resource-ref` element describes a single resource manager connection factory reference. The `resource-ref` element consists of the `description` element; the mandatory `res-ref-name`, `res-type`, and `res-auth` elements; and the optional `res-sharing-scope` element. The `res-ref-name` element contains the name of the environment entry used in the enterprise bean's code. The name of the environment entry is relative to the `java:comp/env` context (e.g., the name should be `jdbc/EmployeeAppDB` rather than `java:comp/env/jdbc/EmployeeAppDB`). The `res-type` element contains the Java type of the resource manager connection factory that the enterprise bean code expects. The `res-auth` element indicates whether the enterprise bean code performs resource manager sign-on programmatically, or whether the container signs on to the resource manager using the principal mapping information supplied by the Deployer. The Bean Provider indicates the sign-on responsibility by setting the value of the `res-auth` element to `Application` or `Container`. The `res-sharing-scope` element indicates whether connections to the resource manager obtained through the given resource manager connection factory reference are to be shared or whether connections are unshareable. The value of the `res-sharing-scope` element is `Shareable` or `Unshareable`. If the `res-sharing-scope` element is not specified, connections are assumed to be shareable.

A resource manager connection factory reference is scoped to the enterprise bean whose declaration contains the `resource-ref` element. This means that the resource manager connection factory reference is not accessible from other enterprise beans at runtime, and that other enterprise beans may define `resource-ref` elements with the same `res-ref-name` without causing a name conflict.

The type declaration allows the Deployer to identify the type of the resource manager connection factory.

Note that the indicated type is the Java type of the resource factory, not the Java type of the resource.

The following example is the declaration of resource manager connection factory references used by the `EmployeeService` enterprise bean illustrated in the previous subsection.

```
...
<enterprise-beans>
  <session>
    ...
    <ejb-name>EmployeeService</ejb-name>
    <ejb-class>com.wombat.empl.EmployeeServiceBean</ejb-class>
    ...
    <resource-ref>
      <description>
        A data source for the database in which
        the EmployeeService enterprise bean will
        record a log of all transactions.
      </description>
      <res-ref-name>jdbc/EmployeeAppDB</res-ref-name>
      <res-type>javax.sql.DataSource</res-type>
      <res-auth>Container</res-auth>
      <res-sharing-scope>Shareable</res-sharing-scope>
    </resource-ref>
    ...
  </session>
</enterprise-beans>
...
```


The following example illustrates the declaration of the JMS resource manager connection factory references used by the example on page 367.

```

...
<enterprise-beans>
  <session>
    ...
    <resource-ref>
      <description>
        A queue connection factory used by the
        MySession enterprise bean to send
        notifications.
      </description>
      <res-ref-name>jms/qConnFactory</res-ref-name>
      <res-type>javax.jms.QueueConnectionFactory</res-type>
      <res-auth>Container</res-auth>
      <res-sharing-scope>Unshareable</res-sharing-scope>
    </resource-ref>
    ...
  </session>
</enterprise-beans>
...

```

20.5.1.3 Standard Resource Manager Connection Factory Types

The Bean Provider must use the `javax.sql.DataSource` resource manager connection factory type for obtaining JDBC connections, and the `javax.jms.ConnectionFactory`, `javax.jms.QueueConnectionFactory`, or `javax.jms.TopicConnectionFactory` for obtaining JMS connections.

The Bean Provider must use the `javax.mail.Session` resource manager connection factory type for obtaining JavaMail connections, and the `java.net.URL` resource manager connection factory type for obtaining URL connections.

It is recommended that the Bean Provider names JDBC data sources in the `java:comp/env/jdbc` subcontext, and JMS connection factories in the `java:comp/env/jms` subcontext. It is also recommended that the Bean Provider names all JavaMail connection factories in the `java:comp/env/mail` subcontext, and all URL connection factories in the `java:comp/env/url` subcontext.

The Connector architecture [12] allows an enterprise bean to use the API described in this section to obtain resource objects that provide access to additional back-end systems.

20.5.2 Deployer's Responsibility

The Deployer uses deployment tools to bind the resource manager connection factory references to the actual resource factories configured in the target operational environment.

The Deployer must perform the following tasks for each resource manager connection factory reference declared in the deployment descriptor:

- Bind the resource manager connection factory reference to a resource manager connection factory that exists in the operational environment. The Deployer may use, for example, the JNDI `LinkRef` mechanism to create a symbolic link to the actual JNDI name of the resource manager connection factory. The resource manager connection factory type must be compatible with the type declared in the `res-type` element.
- Provide any additional configuration information that the resource manager needs for opening and managing the resource. The configuration mechanism is resource-manager specific, and is beyond the scope of this specification.
- If the value of the `res-auth` element is `Container`, the Deployer is responsible for configuring the sign-on information for the resource manager. This is performed in a manner specific to the EJB container and resource manager; it is beyond the scope of this specification.

For example, if principals must be mapped from the security domain and principal realm used at the enterprise beans application level to the security domain and principal realm of the resource manager, the Deployer or System Administrator must define the mapping. The mapping is performed in a manner specific to the EJB container and resource manager; it is beyond the scope of the current EJB specification.

20.5.3 Container Provider Responsibility

The EJB Container Provider is responsible for the following:

- Provide the deployment tools that allow the Deployer to perform the tasks described in the previous subsection.
- Provide the implementation of the resource manager connection factory classes for the resource managers that are configured with the EJB container.
- If the Bean Provider sets the `res-auth` of a resource manager connection factory reference to `Application`, the container must allow the bean to perform explicit programmatic sign-on using the resource manager's API.
- If the Bean Provider sets the `res-sharing-scope` of a resource manager connection factory reference to `Unshareable`, the container must not attempt to share the connections obtained from the resource manager connection factory *reference*^[51]. If the Bean Provider sets the `res-sharing-scope` of a resource manager connection factory reference to `Shareable` or does not specify `res-sharing-scope`, the container must share the connections obtained from the resource manager connection factory according to the requirements defined in [9].
- The container must provide tools that allow the Deployer to set up resource manager sign-on information for the resource manager references whose `res-auth` element is set to `Container`. The minimum requirement is that the Deployer must be able to specify the user/password information for each resource manager connection factory reference declared by the

[51] Connections obtained from the same resource manager connection factory through a different resource manager connection factory reference may be shareable.

enterprise bean, and the container must be able to use the user/password combination for user authentication when obtaining a connection to the resource by invoking the resource manager connection factory.

Although not required by the EJB specification, we expect that containers will support some form of a single sign-on mechanism that spans the application server and the resource managers. The container will allow the Deployer to set up the resource managers such that the EJB caller principal can be propagated (directly or through principal mapping) to a resource manager, if required by the application.

While not required by the EJB specification, most EJB container providers also provide the following features:

- A tool to allow the System Administrator to add, remove, and configure a resource manager for the EJB server.
- A mechanism to pool connections to the resources for the enterprise beans and otherwise manage the use of resources by the container. The pooling must be transparent to the enterprise beans.

20.5.4 System Administrator's Responsibility

The System Administrator is typically responsible for the following:

- Add, remove, and configure resource managers in the EJB server environment.

In some scenarios, these tasks can be performed by the Deployer.

20.6 Resource Environment References

This section describes the programming and deployment descriptor interfaces that allow the Bean Provider to refer to administered objects that are associated with resources (e.g., a Connector CCI `InteractionSpec` instance) by using “logical” names called *resource environment references*. Resource environment references are special entries in the enterprise bean's environment. The Deployer binds the resource environment references to administered objects in the target operational environment.

20.6.1 Bean Provider's Responsibilities

This subsection describes the Bean Provider's view and responsibilities with respect to resource environment references.

20.6.1.1 Resource Environment Reference Programming Interfaces

The Bean Provider must use resource environment references to locate administered objects that are associated with resources, as follows.

- Assign an entry in the enterprise bean's environment to the reference. (See subsection 20.6.1.2 for information on how resource environment references are declared in the deployment descriptor.)
- *The EJB specification recommends, but does not require, that all resource environment references be organized in the appropriate subcontext of the bean's environment for the resource type.*
- Look up the administered object in the enterprise bean's environment using JNDI.

20.6.1.2 Declaration of Resource Environment References in Deployment Descriptor

Although the resource environment reference is an entry in the enterprise bean's environment, the Bean Provider must not use a `env-entry` element to declare it. Instead, the Bean Provider must declare all references to administered objects associated with resources using the `resource-env-ref` elements of the deployment descriptor. This allows the `ejb-jar` consumer to discover all the resource environment references used by the enterprise bean.

Each `resource-env-ref` element describes the requirements that the referencing enterprise bean has for the referenced administered object. The `resource-env-ref` element contains an optional `description` element, and the mandatory `resource-env-ref-name` and `resource-env-ref-type` elements.

The `resource-env-ref-name` element specifies the resource environment reference name: its value is the environment entry name used in the enterprise bean code. The name of the environment entry is relative to the `java:comp/env` context. The `resource-env-ref-type` element specifies the expected type of the referenced object.

A resource environment reference is scoped to the enterprise bean whose declaration contains the `resource-env-ref` element. This means that the resource environment reference is not accessible to other enterprise beans at runtime, and that other enterprise beans may define `resource-env-ref` elements with the same `resource-env-ref-name` without causing a name conflict.

20.6.2 **Deployer's Responsibility**

The Deployer is responsible for the following:

- The Deployer must ensure that all the declared resource environment references are bound to administered objects that exist in the operational environment. The Deployer may use, for

example, the JNDI `LinkRef` mechanism to create a symbolic link to the actual JNDI name of the target object.

- The Deployer must ensure that the target object is type-compatible with the type declared for the resource environment reference. This means that the target object must be of the type indicated in the `resource-env-ref-type` element.

20.6.3 Container Provider's Responsibility

The Container Provider must provide the deployment tools that allow the Deployer to perform the tasks described in the previous subsection. The deployment tools provided by the EJB Container Provider must be able to process the information supplied in the `resource-env-ref` elements in the deployment descriptor.

At the minimum, the tools must be able to inform the Deployer of any unresolved resource environment references, and allow him or her to resolve a resource environment reference by binding it to a specified compatible target object in the environment.

20.7 Message Destination References

This section describes the programming and deployment descriptor interfaces that allow the Bean Provider to refer to message destination objects by using “logical” names called *message destination references*. Message destination references are special entries in the enterprise bean's environment. The Deployer binds the message destination references to administered message destinations in the target operational environment.

The EJB 2.0 specification supported the use of resource environment references to refer to message destination objects in the target operational environment. This usage is still supported by the EJB 2.1 specification. It is recommended, however, that the Bean Provider use message destination references instead of resource environment references when referring to message destinations because of the additional capabilities that message destination references offer in allowing the message flow within an application to be specified. See Section 20.7.2.

20.7.1 Bean Provider's Responsibilities

This subsection describes the Bean Provider's view and responsibilities with respect to message destination references.

20.7.1.1 Message Destination Reference Programming Interfaces

The Bean Provider uses message destination references to locate message destinations, as follows.

- Assign an entry in the enterprise bean's environment to the reference. (See subsection 20.7.1.2 for information on how message destination references are declared in the deployment descriptor.)

- *The EJB specification recommends, but does not require, that all message destination references be organized in the appropriate subcontext of the bean's environment for the messaging resource type (e.g. in the `java:comp/env/jms` JNDI context for JMS Destinations).*
- Look up the destination in the enterprise bean's environment using JNDI.

The following example illustrates how an enterprise bean uses a message destination reference to locate a JMS Destination.

```
public class StockServiceBean implements SessionBean {  
    public void processStockInfo(...) {  
        ...  
        // Obtain the default initial JNDI context.  
        Context initCtx = new InitialContext();  
  
        // Look up the JMS StockQueue in the environment.  
        Object result = initCtx.lookup(  
            "java:comp/env/jms/StockQueue");  
  
        // Convert the result to the proper type.  
        javax.jms.Queue queue = (javax.jms.Queue)result;  
    }  
}
```

In the example, the Bean Provider of the `StockServiceBean` enterprise bean has assigned the environment entry `jms/StockQueue` as the message destination reference name to refer to a JMS queue.

20.7.1.2 Declaration of Message Destination References in Deployment Descriptor

Although the message destination reference is an entry in the enterprise bean's environment, the Bean Provider must not use a `env-entry` element to declare it. Instead, the Bean Provider should declare all references to message destinations using the `message-destination-ref` elements of the deployment descriptor. This allows the `ejb-jar` consumer to discover all the message destination references used by the enterprise bean.

Each `message-destination-ref` element describes the requirements that the referencing enterprise bean has for the referenced destination. The `message-destination-ref` element contains an optional `description` element, and the mandatory `message-destination-ref-name`, `message-destination-type`, and `message-destination-usage` elements.

The `message-destination-ref-name` element specifies the message destination reference name: its value is the environment entry name used in the enterprise bean code. The name of the message destination reference is relative to the `java:comp/env` context (e.g., the name should be `jms/StockQueue` rather than `java:comp/env/jms/StockQueue`). The `message-destination-type` element specifies the expected type of the referenced destination. For example, in the case of a JMS Destination, its value might be `javax.jms.Queue`. The `message-destination-usage` element specifies whether messages are consumed from the message destination, produced for the destination, or both.

A message destination reference is scoped to the enterprise bean whose declaration contains the `message-destination-ref` element. This means that the message destination reference is not accessible to other enterprise beans at runtime, and that other enterprise beans may define `message-destination-ref` elements with the same `message-destination-ref-name` without causing a name conflict.

The following example illustrates the declaration of message destination references in the deployment descriptor.

```
...
<message-destination-ref>
  <description>
    This is a reference to a JMS queue used in processing Stock info
  </description>
  <message-destination-ref-name>jms/StockInfo</message-destination-ref-name>
  <message-destination-type>javax.jms.Queue</message-destination-type>
  <message-destination-usage>Produces</message-destination-usage>
</message-destination-ref>
...
```

20.7.2 Application Assembler's Responsibilities

By means of linking message consumers and producers to one or more common logical destinations specified in the deployment descriptor, the Application Assembler can specify the flow of messages within an application. The Application Assembler uses the `message-destination` element, the `message-destination-link` element of the `message-destination-ref` element, and the `message-destination-link` element of the `message-driven` element to link message destination references to a common logical destination.

The Application Assembler specifies the link between message consumers and producers as follows:

- The Application Assembler uses the `message-destination` element to specify a logical message destination within the application. The `message-destination` element defines a `message-destination-name`, which is used for the purpose of linking.
- The Application Assembler uses the `message-destination-link` element of the `message-destination-ref` element of an enterprise bean that produces messages to link it to the target destination. The value of the `message-destination-link` element is the name of the target destination, as defined in the `message-destination-name` element of the `message-destination` element. The `message-destination` element can be in any module in the same J2EE application as the referencing component. The Application Assembler uses the `message-destination-usage` element of the `message-destination-ref` element to indicate that the referencing enterprise bean produces messages to the referenced destination.
- If the consumer of messages from the common destination is a message-driven bean, the Application Assembler uses the `message-destination-link` element of the `message-driven` element to reference the logical destination. If the Application Assembler links a message-driven bean to its source destination, he or she should use the mes-

sage-destination-type element of the message-driven element to specify the expected destination type.

- If an enterprise bean is otherwise a message consumer, the Application Assembler uses the message-destination-link element of the message-destination-ref element of the enterprise bean that consumes messages to link to the common destination. In the latter case, the Application Assembler uses the message-destination-usage element of the message-destination-ref element to indicate that the enterprise bean consumes messages from the referenced destination.
- To avoid the need to rename message destinations to have unique names within an entire J2EE application, the Application Assembler may use the following syntax in the message-destination-link element of the referencing application component. The Application Assembler specifies the path name of the ejb-jar file containing the referenced message destination and appends the message-destination-name of the target destination separated from the path name by #. The path name is relative to the referencing application component jar file. In this manner, multiple destinations with the same message-destination-name may be uniquely identified.
- When linking message destinations, the Application Assembler must ensure that the consumers and producers for the destination require a message destination of the same or compatible type, as determined by the messaging system.

The following example illustrates the use of message destination linking in the deployment descriptor.

```

...
<enterprise-beans>
<session>
  ...
  <ejb-name>EmployeeService</ejb-name>
  <ejb-class>com.wombat.empl.EmployeeServiceBean</ejb-class>
  ...
  <message-destination-ref>
    <message-destination-ref-name>jms/EmployeeReimbursements</mes-
sage-destination-ref-name>
    <message-destination-type>javax.jms.Queue</message-destina-
tion-type>
    <message-destination-usage>Produces</message-destina-
tion-usage>
    <message-destination-link>ExpenseProcessingQueue</message-desti-
nation-link>
  </message-destination-ref>
</session>
...

<message-driven>
  <ejb-name>ExpenseProcessing</ejb-name>
  <ejb-class>com.wombat.empl.ExpenseProcessingBean</ejb-class>
  <messaging-type>javax.jms.MessageListener</messaging-type>
  ...
  <message-destination-type>javax.jms.Queue</message-destina-
tion-type>
  <message-destination-link>ExpenseProcessingQueue</message-destina-
tion-link>
  ...
</message-driven>
...
</enterprise-beans>
...
<assembly-descriptor>
  ...
  <message-destination>
    <message-destination-name>ExpenseProcessingQueue</message-destina-
tion-name>
  </message-destination>
  ...
</assembly-descriptor>

```

The Application Assembler uses the `message-destination-link` element to indicate that the message destination reference `EmployeeReimbursement` declared in the `EmployeeService` enterprise bean is linked to the `ExpenseProcessing` message-driven bean by means of the common destination `ExpenseProcessingQueue`.

The following example illustrates using the `message-destination-link` element to indicate an enterprise bean reference to the `ExpenseProcessingQueue` that is in the same J2EE application unit but in a different `ejb-jar` file.

```
<session>
...
<ejb-name>EmployeeService</ejb-name>
<ejb-class>com.wombat.empl.EmployeeServiceBean</ejb-class>
...
<message-destination-ref>
  <message-destination-ref-name>jms/EmployeeReimbursements</mes-
sage-destination-ref-name>
  <message-destination-type>javax.jms.Queue
  </message-destination-type>
  <message-destination-usage>Produces</message-destina-
tion-usage>
  <message-destination-link>finance.jar#ExpenseProcessingQueue
  </message-destination-link>
</message-destination-ref>
</session>
```

20.7.3 Deployer's Responsibility

The Deployer is responsible for the following:

- The Deployer must ensure that all the declared message destination references are bound to destination objects that exist in the operational environment. The Deployer may use, for example, the JNDI `LinkRef` mechanism to create a symbolic link to the actual JNDI name of the target object.
- The Deployer must ensure that the target object is type-compatible with the type declared for the message destination reference.
- The Deployer must observe the message destination links specified by the Application Assembler.

20.7.4 Container Provider's Responsibility

The Container Provider must provide the deployment tools that allow the Deployer to perform the tasks described in the previous subsection. The deployment tools provided by the EJB Container Provider must be able to process the information supplied in the `message-destination-ref` and `message-destination-link` elements in the deployment descriptor.

The tools must be able to inform the Deployer of the message flow between consumers and producers sharing common message destinations. The tools must also be able to inform the Deployer of any unresolved message destination references, and allow him or her to resolve a message destination reference by binding it to a specified compatible target object in the environment.

20.8 Deprecated EJBContext.getEnvironment Method

The *environment naming context* introduced in EJB 1.1 replaces the EJB 1.0 concept of *environment properties*.

An EJB 1.1 or later compliant container is not required to implement support for the EJB 1.0 style environment properties. If the container does not implement the functionality, it should throw a `RuntimeException` (or subclass thereof) from the `EJBContext.getEnvironment` method.

If an EJB 1.1 or later compliant container chooses to provide support for the EJB 1.0 style environment properties (so that it can support enterprise beans written to the EJB 1.0 specification), it should implement the support as described below.

When the tools convert the EJB 1.0 deployment descriptor to the EJB 1.1 XML format, they should place the definitions of the environment properties into the `ejb10-properties` subcontext of the environment naming context. The `env-entry` elements should be defined as follows: the `env-entry-name` element contains the name of the environment property, the `env-entry-type` must be `java.lang.String`, and the optional `env-entry-value` contains the environment property value.

For example, an EJB 1.0 enterprise bean with two environment properties `foo` and `bar`, should declare the following `env-entry` elements in its EJB 1.1 format deployment descriptor.

```
...
<env-entry>
  env-entry-name>ejb10-properties/foo</env-entry-name>
  <env-entry-type>java.lang.String</env-entry-type>
</env-entry>
<env-entry>
  <description>bar's description</description>
  <env-entry-name>ejb10-properties/bar</env-entry-name>
  <env-entry-type>java.lang.String</env-entry-type>
  <env-entry-value>bar value</env-entry-value>
</env-entry>
...
```

The container should provide the entries declared in the `ejb10-properties` subcontext to the instances as a `java.util.Properties` object that the instances obtain by invoking the `EJBContext.getEnvironment` method.

The enterprise bean uses the EJB 1.0 API to access the properties, as shown by the following example.

```
public class SomeBean implements SessionBean {
    SessionContext ctx;
    java.util.Properties env;

    public void setSessionContext(SessionContext sc) {
        ctx = sc;
        env = ctx.getEnvironment();
    }

    public someBusinessMethod(...) ... {
        String fooValue = env.getProperty("foo");
        String barValue = env.getProperty("bar");
    }
    ...
}
```

20.9 UserTransaction Interface

Note: The requirement for the container to publish the UserTransaction interface in the enterprise bean's JNDI context was added to make the requirements on UserTransaction uniform with the other Java 2, Enterprise Edition application component types.

The container must make the UserTransaction interface available to the enterprise beans that are allowed to use this interface (only session and message-driven beans with bean-managed transaction demarcation are allowed to use this interface) in JNDI under the name `java:comp/UserTransaction`.

The container must not make the UserTransaction interface available to the enterprise beans that are not allowed to use this interface. The container should throw `javax.naming.NameNotFoundException` if an instance of an enterprise bean that is not allowed to use the UserTransaction interface attempts to look up the interface in JNDI.

The following code example

```
public MySessionBean implements SessionBean {
    ...
    public someMethod()
    {
        Context initCtx = new InitialContext();
        UserTransaction utx = (UserTransaction)initCtx.lookup(
            "java:comp/UserTransaction");
        utx.begin();
        ...
        utx.commit();
    }
    ...
}
```

is functionally equivalent to

```
public MySessionBean implements SessionBean {
    SessionContext ctx;
    ...
    public someMethod()
    {
        UserTransaction utx = ctx.getUserTransaction();
        utx.begin();
        ...
        utx.commit();
    }
    ...
}
```


Security Management

This chapter defines the EJB architecture's support for security management.

21.1 Overview

We set the following goals for the security management in the EJB architecture:

- *Lessen the burden of the application developer (i.e. the Bean Provider) for securing the application by allowing greater coverage from more qualified EJB roles. The EJB Container Provider provides the implementation of the security infrastructure; the Deployer and System Administrator define the security policies.*
- *Allow the security policies to be set by the Application Assembler or Deployer rather than being hard-coded by the Bean Provider at development time.*
- *Allow the enterprise bean applications to be portable across multiple EJB servers that use different security mechanisms.*

The EJB architecture encourages the Bean Provider to implement the enterprise bean class without hard-coding the security policies and mechanisms into the business methods. In most cases, the enterprise bean's business methods should not contain any security-related logic. This allows the Deployer to configure the security policies for the application in a way that is most appropriate for the operational environment of the enterprise.

To make the Deployer's task easier, the Application Assembler (which could be the same party as the Bean Provider) may define *security roles* for an application composed of one or more enterprise beans. A security role is a semantic grouping of permissions that a given type of users of the application must have in order to successfully use the application. The Application Assembler can define (declaratively in the deployment descriptor) *method permissions* for each security role. A method permission is a permission to invoke a specified group of methods of an enterprise bean's home, component, and/or web service endpoint interfaces. The security roles defined by the Application Assembler present a simplified security view of the enterprise beans application to the Deployer—the Deployer's view of the application's security requirements is the small set of security roles rather than a large number of individual methods.

The security principal under which a method invocation is performed is typically that of the component's caller. By specifying a run-as identity, however, it is possible to specify that a different principal be substituted for the execution of the methods of the bean's home, component, and/or web service endpoint interfaces and any methods of other enterprise beans that the bean may call. The Application Assembler specifies in the deployment descriptor whether the caller's security identity or a run-as security identity should be used for the execution of the bean's methods. See Section 21.3.4.

The Application Assembler should specify the requirements for the caller's principal management of enterprise bean invocations by means of the *security-identity* deployment descriptor element and as part of the description. If *use-caller-identity* is specified as the value of the *security-identity* element, the caller principal is propagated from the caller to the callee. (That is, the called enterprise bean will see the same returned value of the `EJBContext.getCallerPrincipal` as the calling enterprise bean.) If the *run-as* element is specified, a security principal that has been assigned to the specified security role will be used for the execution of the bean's methods and will be visible as the caller principal in the bean's callee.

The Deployer is responsible for assigning principals, or groups of principals, which are defined in the target operational environment, to the security roles defined by the Application Assembler for the enterprise beans in the deployment descriptor. The Deployer is also responsible for assigning principals for the run-as identities specified by the Application Assembler. The Deployer is further responsible for configuring other aspects of the security management of the enterprise beans, such as principal mapping for inter-enterprise bean calls, and principal mapping for resource manager access.

At runtime, a client will be allowed to invoke a business method only if the principal associated with the client call has been assigned by the Deployer to have at least one security role that is allowed to invoke the business method or if the Application Assembler has specified that security authorization is not to be checked for the method. See Section 21.3.2.

The Container Provider is responsible for enforcing the security policies at runtime, providing the tools for managing security at runtime, and providing the tools used by the Deployer to manage security during deployment.

Because not all security policies can be expressed declaratively, the EJB architecture provides a simple programmatic interface that the Bean Provider may use to access the security context from the business methods.

The following sections define the responsibilities of the individual EJB roles with respect to security management.

21.2 Bean Provider's Responsibilities

This section defines the Bean Provider's perspective of the EJB architecture support for security, and defines his or her responsibilities.

21.2.1 Invocation of Other Enterprise Beans

An enterprise bean business method can invoke another enterprise bean via the other bean's home or component interface. The EJB architecture provides no programmatic interfaces for the invoking enterprise bean to control the principal passed to the invoked enterprise bean.

The management of caller principals passed on *inter-enterprise* bean invocations (i.e. principal delegation) is set up by the Deployer and System Administrator in a container-specific way. The Bean Provider and Application Assembler should describe all the requirements for the caller's principal management of inter-enterprise bean invocations as part of the description.

21.2.2 Resource Access

Section 20.5 defines the protocol for accessing resource managers, including the requirements for security management.

21.2.3 Access of Underlying OS Resources

The EJB architecture does not define the operating system principal under which enterprise bean methods execute. Therefore, the Bean Provider cannot rely on a specific principal for accessing the underlying OS resources, such as files. (See Subsection 21.6.8 for the reasons behind this rule.)

We believe that most enterprise business applications store information in resource managers such as relational databases rather than in resources at the operating system levels. Therefore, this rule should not affect the portability of most enterprise beans.

21.2.4 Programming Style Recommendations

The Bean Provider should neither implement security mechanisms nor hard-code security policies in the enterprise beans' business methods. Rather, the Bean Provider should rely on the security mechanisms provided by the EJB container, and should let the Application Assembler and Deployer define the appropriate security policies for the application.

The Bean Provider and Application Assembler may use the deployment descriptor to convey security-related information to the Deployer. The information helps the Deployer to set up the appropriate security policy for the enterprise bean application.

21.2.5 Programmatic Access to Caller's Security Context

Note: In general, security management should be enforced by the container in a manner that is transparent to the enterprise beans' business methods. The security API described in this section should be used only in the less frequent situations in which the enterprise bean business methods need to access the security context information.

The `javax.ejb.EJBContext` interface provides two methods (plus two deprecated methods that were defined in EJB 1.0) that allow the Bean Provider to access security information about the enterprise bean's caller.

```
public interface javax.ejb.EJBContext {
    ...

    //
    // The following two methods allow the EJB class
    // to access security information.
    //
    java.security.Principal getCallerPrincipal();
    boolean isCallerInRole(String roleName);

    //
    // The following two EJB 1.0 methods are deprecated.
    //
    java.security.Identity getCallerIdentity();
    boolean isCallerInRole(java.security.Identity role);

    ...
}
```

The Bean Provider can invoke the `getCallerPrincipal` and `isCallerInRole` methods only in the enterprise bean's business methods as specified in Table 2 on page 87, Table 3 on page 100, Table 4 on page 197, Table 10 on page 279, and Table 12 on page 344. If they are otherwise invoked when no security context exists, they should throw the `java.lang.IllegalStateException` runtime exception.

The `getCallerIdentity()` and `isCallerInRole(Identity role)` methods were deprecated in EJB 1.1. The Bean Provider must use the `getCallerPrincipal()` and `isCallerInRole(String roleName)` methods for new enterprise beans.

An EJB 1.1 or later compliant container may choose to implement the two deprecated methods as follows.

- A container that does not want to provide support for this deprecated method should throw a `RuntimeException` (or subclass of `RuntimeException`) from the `getCallerIdentity` method.

- A container that wants to provide support for the `getCallerIdentity` method should return an instance of a subclass of the `java.security.Identity` abstract class from the method. The `getName` method invoked on the returned object must return the same value that `getCallerPrincipal().getName()` would return.
- A container that does not want to provide support for this deprecated method should throw a `RuntimeException` (or subclass of `RuntimeException`) from the `isCallerInRole(Identity identity)` method.
- A container that wants to implement the `isCallerInRole(Identity identity)` method should implement it as follows:

```
public isCallerInRole(Identity identity) {  
    return isCallerInRole(identity.getName());  
}
```

21.2.5.1 Use of `getCallerPrincipal`

The purpose of the `getCallerPrincipal` method is to allow the enterprise bean methods to obtain the current caller principal's name. The methods might, for example, use the name as a key to information in a database.

An enterprise bean can invoke the `getCallerPrincipal` method to obtain a `java.security.Principal` interface representing the current caller. The enterprise bean can then obtain the distinguished name of the caller principal using the `getName` method of the `java.security.Principal` interface. If the security identity has not been established, `getCallerPrincipal` returns the container's representation of the unauthenticated identity.

Note that `getCallerPrincipal` returns the principal that represents the caller of the enterprise bean, not the principal that corresponds to the run-as security identity for the bean, if any.

The meaning of the *current caller*, the Java class that implements the `java.security.Principal` interface, and the realm of the principals returned by the `getCallerPrincipal` method depend on the operational environment and the configuration of the application.

An enterprise may have a complex security infrastructure that includes multiple security domains. The security infrastructure may perform one or more mapping of principals on the path from an EJB caller to the EJB object. For example, an employee accessing his or her company over the Internet may be identified by a userid and password (basic authentication), and the security infrastructure may authenticate the principal and then map the principal to a Kerberos principal that is used on the enterprise's intranet before delivering the method invocation to the EJB object. If the security infrastructure performs principal mapping, the `getCallerPrincipal` method returns the principal that is the result of the mapping, not the original caller principal. (In the previous example, `getCallerPrincipal` would return the Kerberos principal.) The management of the security infrastructure, such as principal mapping, is performed by the System Administrator role; it is beyond the scope of the EJB specification.

The following code sample illustrates the use of the `getCallerPrincipal()` method.

```
public class EmployeeServiceBean implements SessionBean {
    EJBContext ejbContext;

    public void changePhoneNumber(...) {
        ...

        // Obtain the default initial JNDI context.
        Context initCtx = new InitialContext();

        // Look up the remote home interface of the EmployeeRecord
        // enterprise bean in the environment.
        Object result = initCtx.lookup(
            "java:comp/env/ejb/EmplRecord");

        // Convert the result to the proper type.
        EmployeeRecordHome emplRecordHome = (EmployeeRecordHome)
            javax.rmi.PortableRemoteObject.narrow(result,
                EmployeeRecordHome.class);

        // obtain the caller principal.
        callerPrincipal = ejbContext.getCallerPrincipal();

        // obtain the caller principal's name.
        callerKey = callerPrincipal.getName();

        // use callerKey as primary key to EmployeeRecord finder
        EmployeeRecord myEmployeeRecord =
            emplRecordHome.findByPrimaryKey(callerKey);

        // update phone number
        myEmployeeRecord.changePhoneNumber(...);

        ...
    }
}
```

In the previous example, the enterprise bean obtains the principal name of the current caller and uses it as the primary key to locate an `EmployeeRecord` Entity object. This example assumes that application has been deployed such that the current caller principal contains the primary key used for the identification of employees (e.g., employee number).

21.2.5.2 Use of `isCallerInRole`

The main purpose of the `isCallerInRole(String roleName)` method is to allow the Bean Provider to code the security checks that cannot be easily defined declaratively in the deployment descriptor using method permissions. Such a check might impose a role-based limit on a request, or it might depend on information stored in the database.

The enterprise bean code can use the `isCallerInRole` method to test whether the current caller has been assigned to a given security role. Security roles are defined by the Application Assembler in the deployment descriptor (see Subsection 21.3.1), and are assigned to principals or principal groups that exist in the operational environment by the Deployer.

Note that `isCallerInRole(String roleName)` tests the principal that represents the caller of the enterprise bean, not the principal that corresponds to the run-as security identity for the bean, if any.

The following code sample illustrates the use of the `isCallerInRole(String roleName)` method.

```
public class PayrollBean ... {
    EntityContext ejbContext;

    public void updateEmployeeInfo(EmplInfo info) {
        oldInfo = ... read from database;

        // The salary field can be changed only by callers
        // who have the security role "payroll"
        if (info.salary != oldInfo.salary &&
            !ejbContext.isCallerInRole("payroll")) {
            throw new SecurityException(...);
        }
        ...
    }
    ...
}
```

21.2.5.3 Declaration of Security Roles Referenced from the Bean's Code

The Bean Provider is responsible for declaring in the `security-role-ref` elements of the deployment descriptor all the security role names used in the enterprise bean code. Declaring the security roles references in the code allows the Application Assembler or Deployer to link the names of the security roles used in the code to the security roles defined for an assembled application through the `security-role` elements.

The Bean Provider must declare each security role referenced in the code using the `security-role-ref` element as follows:

- Declare the name of the security role using the `role-name` element. The name must be the security role name that is used as a parameter to the `isCallerInRole(String roleName)` method.
- Optionally provide a description of the security role in the `description` element.

A security role reference, including the name defined by the `role-name` element, is scoped to the session or entity bean element whose declaration contains the `security-role-ref` element.

The following example illustrates how an enterprise bean's references to security roles are declared in the deployment descriptor.

```

...
<enterprise-beans>
  ...
  <entity>
    <ejb-name>AardvarkPayroll</ejb-name>
    <ejb-class>com.aardvark.payroll.PayrollBean</ejb-class>
    ...
    <security-role-ref>
      <description>
        This security role should be assigned to the
        employees of the payroll department who are
        allowed to update employees' salaries.
      </description>
      <role-name>payroll</role-name>
    </security-role-ref>
    ...
  </entity>
  ...
</enterprise-beans>
...

```

The deployment descriptor above indicates that the enterprise bean `AardvarkPayroll` makes the security check using `isCallerInRole("payroll")` in its business method.

21.3 Application Assembler's Responsibilities

The Application Assembler (which could be the same party as the Bean Provider) may define a *security view* of the enterprise beans contained in the `ejb-jar` file. Providing the security view in the deployment descriptor is optional for the Bean Provider and Application Assembler.

The main reason for the Application Assembler's providing the security view of the enterprise beans is to simplify the Deployer's job. In the absence of a security view of an application, the Deployer needs detailed knowledge of the application in order to deploy the application securely. For example, the Deployer would have to know what each business method does to determine which users can call it. The security view defined by the Application Assembler presents a more consolidated view to the Deployer, allowing the Deployer to be less familiar with the application.

The security view consists of a set of *security roles*. A security role is a semantic grouping of permissions that a given type of users of an application must have in order to successfully use the application.

The Application Assembler defines *method permissions* for each security role. A method permission is a permission to invoke a specified group of methods of the enterprise beans' home, component, and/or web service endpoint interfaces.

It is important to keep in mind that the security roles are used to define the logical security view of an application. They should not be confused with the user groups, users, principals, and other concepts that exist in the target enterprise's operational environment.

In special cases, a qualified Deployer may change the definition of the security roles for an application, or completely ignore them and secure the application using a different mechanism that is specific to the operational environment.

If the Bean Provider has declared any security role references using the `security-role-ref` elements, the Application Assembler must link all the security role references listed in the `security-role-ref` elements to the security roles defined in the `security-role` elements. This is described in more detail in subsection 21.3.3.

21.3.1 Security Roles

The Application Assembler can define one or more security roles in the deployment descriptor. The Application Assembler then assigns groups of methods of the enterprise beans' home, component, and/or web service endpoint interfaces to the security roles to define the security view of the application.

Because the Application Assembler does not, in general, know the security environment of the operational environment, the security roles are meant to be *logical* roles (or actors), each representing a type of user that should have the same access rights to the application.

The Deployer then assigns user groups and/or user accounts defined in the operational environment to the security roles defined by the Application Assembler.

Defining the security roles in the deployment descriptor is optional^[52] for the Application Assembler. Their omission in the deployment descriptor means that the Application Assembler chose not to pass any security deployment related instructions to the Deployer in the deployment descriptor.

The Application Assembler is responsible for the following:

- Define each security role using a `security-role` element.
- Use the `role-name` element to define the name of the security role.
- Optionally, use the `description` element to provide a description of a security role.

The security roles defined by the `security-role` elements are scoped to the `ejb-jar` file level, and apply to all the enterprise beans in the `ejb-jar` file.

[52] If the Application Assembler does not define security roles in the deployment descriptor, the Deployer will have to define security roles at deployment time.

The following example illustrates a security role definition in a deployment descriptor.

```

...
<assembly-descriptor>
  <security-role>
    <description>
      This role includes the employees of the
      enterprise who are allowed to access the
      employee self-service application. This role
      is allowed only to access his/her own
      information.
    </description>
    <role-name>employee</role-name>
  </security-role>

  <security-role>
    <description>
      This role includes the employees of the human
      resources department. The role is allowed to
      view and update all employee records.
    </description>
    <role-name>hr-department</role-name>
  </security-role>

  <security-role>
    <description>
      This role includes the employees of the payroll
      department. The role is allowed to view and
      update the payroll entry for any employee.
    </description>
    <role-name>payroll-department</role-name>
  </security-role>

  <security-role>
    <description>
      This role should be assigned to the personnel
      authorized to perform administrative functions
      for the employee self-service application.
      This role does not have direct access to
      sensitive employee and payroll information.
    </description>
    <role-name>admin</role-name>
  </security-role>
  ...
</assembly-descriptor>

```

21.3.2 Method Permissions

If the Application Assembler has defined security roles for the enterprise beans in the ejb-jar file, he or she can also specify the methods of the home, component, and/or web service endpoint interfaces that each security role is allowed to invoke.

Method permissions are defined in the deployment descriptor as a binary relation from the set of security roles to the set of methods of the home, component, and/or web service endpoint interfaces of session and entity beans, including all their superinterfaces (including the methods of the `EJBHome` and `EJBObject` interfaces and/or `EJBLocalHome` and `EJBLocalObject` interfaces). The method permissions relation includes the pair (R, M) if and only if the security role R is allowed to invoke the method M .

The Application Assembler defines the method permissions relation in the deployment descriptor using the `method-permission` elements as follows.

- Each `method-permission` element includes a list of one or more security roles and a list of one or more methods. All the listed security roles are allowed to invoke all the listed methods. Each security role in the list is identified by the `role-name` element, and each method (or a set of methods, as described below) is identified by the `method` element. An optional description can be associated with a `method-permission` element using the `description` element.
- The method permissions relation is defined as the union of all the method permissions defined in the individual `method-permission` elements.
- A security role or a method may appear in multiple `method-permission` elements.

The Application Assembler can indicate that some methods should not be checked for authorization prior to invocation by the container. The Application Assembler uses the `unchecked` element instead of a role name in the `method-permission` element to indicate that a method should not be checked for authorization.

If the method permission relation specifies both the `unchecked` element for a given method and one or more security roles, the method should not be checked for authorization.

The Application Assembler can use the `exclude-list` element to indicate the set of methods that should not be called. The Deployer should configure the enterprise bean's security such that no access is permitted to any method contained in the `exclude-list`.

If a given method is specified both in the `exclude-list` element and in the method permission relation, the Deployer should configure the enterprise bean's security such that no access is permitted to the method.

It is possible that some methods are not assigned to any security roles nor contained in the `exclude-list` element. In this case, it is the responsibility of the Deployer to assign method permissions for all of the unspecified methods, either by assigning them to security roles, or by marking them as `unchecked`.

The `method` element uses the `ejb-name`, `method-name`, and `method-params` elements to denote one or more methods of an enterprise bean's home, component, and/or web service endpoint interfaces. There are three legal styles for composing the `method` element:

Style 1:

```
<method>
  <ejb-name>EJBNAME</ejb-name>
  <method-name>*</method-name>
</method>
```

This style is used for referring to all of the home, component, and web service endpoint interface methods of a specified enterprise bean.

Style 2:

```
<method>
  <ejb-name>EJBNAME</ejb-name>
  <method-name>METHOD</method-name>
</method>
```

This style is used for referring to a specified method of the home, component, or web service endpoint interface of the specified enterprise bean. If there are multiple methods with the same overloaded name, this style refers to all of the overloaded methods.

Style 3:

```
<method>
  <ejb-name>EJBNAME</ejb-name>
  <method-name>METHOD</method-name>
  <method-params>
    <method-param>PARAMETER_1</method-param>
    ...
    <method-param>PARAMETER_N</method-param>
  </method-params>
</method>
```

This style is used to refer to a specified method within a set of methods with an overloaded name. The method must be defined in the specified enterprise bean's home, component, or web service endpoint interface. If there are multiple methods with the same overloaded name, however, this style refers to all of the overloaded methods.

The optional `method-intf` element can be used to differentiate methods with the same name and signature that are multiply defined across the home, component, and/or web service endpoint interfaces of an enterprise bean.

The following example illustrates how security roles are assigned method permissions in the deployment descriptor:

```

...
<method-permission>
  <role-name>employee</role-name>
  <method>
    <ejb-name>EmployeeService</ejb-name>
    <method-name>*</method-name>
  </method>
</method-permission>

<method-permission>
  <role-name>employee</role-name>
  <method>
    <ejb-name>AardvarkPayroll</ejb-name>
    <method-name>findByPrimaryKey</method-name>
  </method>
  <method>
    <ejb-name>AardvarkPayroll</ejb-name>
    <method-name>getEmployeeInfo</method-name>
  </method>
  <method>
    <ejb-name>AardvarkPayroll</ejb-name>
    <method-name>updateEmployeeInfo</method-name>
  </method>
</method-permission>

<method-permission>
  <role-name>payroll-department</role-name>
  <method>
    <ejb-name>AardvarkPayroll</ejb-name>
    <method-name>findByPrimaryKey</method-name>
  </method>
  <method>
    <ejb-name>AardvarkPayroll</ejb-name>
    <method-name>getEmployeeInfo</method-name>
  </method>
  <method>
    <ejb-name>AardvarkPayroll</ejb-name>
    <method-name>updateEmployeeInfo</method-name>
  </method>
  <method>
    <ejb-name>AardvarkPayroll</ejb-name>
    <method-name>updateSalary</method-name>
  </method>
</method-permission>

<method-permission>
  <role-name>admin</role-name>
  <method>
    <ejb-name>EmployeeServiceAdmin</ejb-name>
    <method-name>*</method-name>
  </method>
</method-permission>
...

```

21.3.3 Linking Security Role References to Security Roles

If the Application Assembler defines the `security-role` elements in the deployment descriptor, he or she is also responsible for linking all the security role references declared in the `security-role-ref` elements to the security roles defined in the `security-role` elements.

The Application Assembler links each security role reference to a security role using the `role-link` element. The value of the `role-link` element must be the name of one of the security roles defined in a `security-role` element.

A `role-link` element must be used even if the value of `role-name` is the same as the value of the `role-link` reference.

The following deployment descriptor example shows how to link the security role reference named `payroll` to the security role named `payroll-department`.

```
...
<enterprise-beans>
  ...
  <entity>
    <ejb-name>AardvarkPayroll</ejb-name>
    <ejb-class>com.aardvark.payroll.PayrollBean</ejb-class>
    ...
    <security-role-ref>
      <description>
        This role should be assigned to the
        employees of the payroll department.
        Members of this role have access to
        anyone's payroll record.
        The role has been linked to the
        payroll-department role.
      </description>
      <role-name>payroll</role-name>
      <role-link>payroll-department</role-link>
    </security-role-ref>
    ...
  </entity>
  ...
</enterprise-beans>
...
```

21.3.4 Specification of Security Identities in the Deployment Descriptor

The Application Assembler typically specifies whether the caller's security identity should be used for the execution of the methods of an enterprise bean or whether a specific run-as identity should be used.

The Application Assembler uses the `security-identity` deployment descriptor element for this purpose. The value of the `security-identity` element is either `use-caller-identity` or `run-as`.

Defining the security identities in the deployment descriptor is optional for the Application Assembler. Their omission in the deployment descriptor means that the Application Assembler chose not to pass any instructions related to security identities to the Deployer in the deployment descriptor.

21.3.4.1 Run-as

The Application Assembler can use the `run-as` element to define a run-as identity for an enterprise bean in the deployment descriptor. The run-as identity applies to the enterprise bean as a whole, that is, to all methods of the enterprise bean's home, component, web service endpoint interfaces, to the message listener methods of a message-driven bean, and to the `ejbTimeout` method of an enterprise bean that implements the `TimedObject` interface, and all internal methods of the bean that they might in turn call.

Establishing a run-as identity for an enterprise bean does not affect the identities of its callers, which are the identities tested for permission to access the methods of the enterprise bean. The run-as identity establishes the identity the enterprise bean will use when it makes calls.

Because the Application Assembler does not, in general, know the security environment of the operational environment, the run-as identity is designated by a *logical* role-name, which corresponds to one of the security roles defined by the Application Assembler in the deployment descriptor.

The Deployer then assigns a security principal defined in the operational environment to be used as the principal for the run-as identity. The security principal assigned by the Deployer should be a principal that has been assigned to the security role specified by the `role-name` element.

The Application Assembler is responsible for the following in the specification of run-as identities:

- Use the `role-name` element to define the name of the security role.
- Optionally, use the `description` element to provide a description of the principal that is expected to be bound to the run-as identity in terms of its security role.

The following example illustrates the definition of a run-as identity in the deployment descriptor.

```
...
<enterprise-beans>
  ...
  <session>
    <ejb-name>EmployeeService</ejb-name>
    ...
    <security-identity>
      <run-as>
        <role-name>admin</role-name>
      </run-as>
    </security-identity>
    ...
  </session>
  ...
</enterprise-beans>
...
```

21.4 Deployer's Responsibilities

The Deployer is responsible for ensuring that an assembled application is secure after it has been deployed in the target operational environment. This section defines the Deployer's responsibility with respect to EJB security management.

The Deployer uses deployment tools provided by the EJB Container Provider to read the security view of the application supplied by the Application Assembler in the deployment descriptor. The Deployer's job is to map the security view that was specified by the Application Assembler to the mechanisms and policies used by the security domain in the target operational environment. The output of the Deployer's work includes an application security policy descriptor that is specific to the operational environment. The format of this descriptor and the information stored in the descriptor are specific to the EJB container.

The following subsections describe the security related tasks performed by the Deployer.

21.4.1 Security Domain and Principal Realm Assignment

The Deployer is responsible for assigning the security domain and principal realm to an enterprise bean application.

Multiple principal realms within the same security domain may exist, for example, to separate the realms of employees, trading partners, and customers. Multiple security domains may exist, for example, in application hosting scenarios.

21.4.2 Assignment of Security Roles

The Deployer assigns principals and/or groups of principals (such as individual users or user groups) used for managing security in the operational environment to the security roles defined in the `security-role` elements of the deployment descriptor.

Typically, the Deployer does not need to change the method permissions assigned to each security role in the deployment descriptor.

The Application Assembler linked all the security role references used in the bean's code to the security roles defined in the `security-role` elements. The Deployer does not assign principals and/or principal groups to the security role references—the principals and/or principals groups assigned to a security role apply also to all the linked security role references. For example, the Deployer of the `AardvarkPayroll` enterprise bean in subsection 21.3.3 would assign principals and/or principal groups to the security-role `payroll-department`, and the assigned principals and/or principal groups would be implicitly assigned also to the linked security role `payroll`.

The EJB architecture does not specify how an enterprise should implement its security architecture. Therefore, the process of assigning the logical security roles defined in the application's deployment descriptor to the operational environment's security concepts is specific to that operational environment. Typically, the deployment process consists of assigning to each security role one or more user groups (or individual users) defined in the operational environment. This assignment is done on a per-application basis. (That is, if multiple independent ejb-jar files use the same security role name, each may be assigned differently.)

21.4.3 Principal Delegation

The Deployer is responsible for configuring the principal delegation for inter-component calls. The Deployer must follow any instructions supplied by the Application Assembler (for example, provided in the `run-as` elements of the deployment descriptor, in the `description` elements of the deployment descriptor, or in a deployment manual).

If the `use-caller-identity` element is specified, the caller principal is propagated from one component to another (i.e., the caller principal of the first enterprise bean in a call-chain is passed to the enterprise beans down the chain).

If the Application Assembler specifies that a `run-as` identity be used on behalf of a particular enterprise bean, the Deployer must configure the enterprise beans such that the `run-as` principal is used as the caller principal on any calls that the enterprise bean makes to other beans, and that the `run-as` principal is propagated along the call-chain of those other beans (in the absence of the specification of any further `run-as` elements).

21.4.4 Security Management of Resource Access

The Deployer's responsibilities with respect to securing resource managers access are defined in subsection 20.5.2.

21.4.5 General Notes on Deployment Descriptor Processing

The Deployer can use the security view defined in the deployment descriptor by the Bean Provider and Application Assembler merely as "hints" and may change the information whenever necessary to adapt the security policy to the operational environment.

Since providing the security information in the deployment descriptor is optional for the Application Assembler, the Deployer is responsible for performing any tasks that have not been done by the Application Assembler. (For example, if the definition of security roles and method permissions is missing in the deployment descriptor, the Deployer must define the security roles and method permissions for the application.) It is not required that the Deployer store the output of this activity in the standard ejb-jar file format.

21.5 EJB Client Responsibilities

This section defines the rules that the EJB client program must follow to ensure that the security context passed on the client calls, and possibly imported by the enterprise bean, do not conflict with the EJB server's capabilities for association between a security context and transactions.

These rules are:

- A transactional client cannot change its principal association within a transaction. This rule ensures that all calls from the client within a transaction are performed with the same security context.
- A session bean's client must not change its principal association for the duration of the communication with the session object. This rule ensures that the server can associate a security identity with the session instance at instance creation time, and never have to change the security association during the session instance lifetime.
- If transactional requests within a single transaction arrive from multiple clients (this could happen if there are intermediary objects or programs in the transaction call-chain), all requests within the same transaction must be associated with the same security context.

21.6 EJB Container Provider's Responsibilities

This section describes the responsibilities of the EJB Container and Server Provider.

21.6.1 Deployment Tools

The EJB Container Provider is responsible for providing the deployment tools that the Deployer can use to perform the tasks defined in Section 21.4.

The deployment tools read the information from the deployment descriptor and present the information to the Deployer. The tools guide the Deployer through the deployment process, and present him or her with choices for mapping the security information in the deployment descriptor to the security management mechanisms and policies used in the target operational environment.

The deployment tools' output is stored in an EJB container-specific manner, and is available at runtime to the EJB container.

21.6.2 Security Domain(s)

The EJB container provides a security domain and one or more principal realms to the enterprise beans. The EJB architecture does not specify how an EJB server should implement a security domain, and does not define the scope of a security domain.

A security domain can be implemented, managed, and administered by the EJB server. For example, the EJB server may store X509 certificates or it might use an external security provider such as Kerberos.

The EJB specification does not define the scope of the security domain. For example, the scope may be defined by the boundaries of the application, EJB server, operating system, network, or enterprise.

The EJB server can, but is not required to, provide support for multiple security domains, and/or multiple principal realms.

The case of multiple domains on the same EJB server can happen when a large server is used for application hosting. Each hosted application can have its own security domain to ensure security and management isolation between applications owned by multiple organizations.

21.6.3 Security Mechanisms

The EJB Container Provider must provide the security mechanisms necessary to enforce the security policies set by the Deployer. The EJB specification does not specify the exact mechanisms that must be implemented and supported by the EJB server.

The typical security functions provided by the EJB server include:

- *Authentication of principals.*
- *Access authorization for EJB calls and resource manager access.*
- *Secure communication with remote clients (privacy, integrity, etc.).*

21.6.4 Passing Principals on EJB Calls

The EJB Container Provider is responsible for providing the deployment tools that allow the Deployer to configure the principal delegation for calls from one enterprise bean to another. The EJB container is responsible for performing the principal delegation as specified by the Deployer.

The EJB container must be capable of allowing the Deployer to specify that, for all calls from a single application within a single J2EE product, the caller principal is propagated on calls from one enterprise bean to another (i.e., the multiple beans in the call chain will see the same return value from `getCallerPrincipal()`).

This requirement is necessary for applications that need a consistent return value of `getCallerPrincipal` across a chain of calls between enterprise beans.

The EJB container must be capable of allowing the Deployer to specify that a run-as principal be used for the execution of the home, component, and web service endpoint interface methods of a session or entity bean, or for the message listener methods of a message-driven bean.

21.6.5 Security Methods in `javax.ejb.EJBContext`

The EJB container must provide access to the caller's security context information from the enterprise beans' instances via the `getCallerPrincipal()` and `isCallerInRole(String roleName)` methods. The EJB container must provide the caller's security context information during the execution of a business method invoked via the enterprise bean's home, component, message listener, web service endpoint interface, and/or `TimedObject` interface, as defined in Table 2 on page 87,

Table 3 on page 100, Table 4 on page 197, Table 10 on page 279, and Table 12 on page 344. The container must ensure that all enterprise bean method invocations received through these interfaces are associated with some principal. If the security identity of the caller has not been established, the container returns the container's representation of the unauthenticated identity. The container must never return a null from the `getCallerPrincipal` method.

21.6.6 Secure Access to Resource Managers

The EJB Container Provider is responsible for providing secure access to resource managers as described in Subsection 20.5.3.

21.6.7 Principal Mapping

If the application requires that its clients are deployed in a different security domain, or if multiple applications deployed across multiple security domains need to interoperate, the EJB Container Provider is responsible for the mechanism and tools that allow mapping of principals. The tools are used by the System Administrator to configure the security for the application's environment.

21.6.8 System Principal

The EJB specification does not define the "system" principal under which the JVM running an enterprise bean's method executes.

Leaving the principal undefined makes it easier for the EJB container vendors to provide runtime support for EJB on top of their existing server infrastructures. For example, while one EJB container implementation can execute all instances of all enterprise beans in a single JVM, another implementation can use a separate JVM per ejb-jar per client. Some EJB containers may make the system principal the same as the application-level principal. Others may use different principals, potentially from different principal realms and even security domains.

21.6.9 Runtime Security Enforcement

The EJB container is responsible for enforcing the security policies defined by the Deployer. The implementation of the enforcement mechanism is EJB container implementation-specific. The EJB container may, but does not have to, use the Java programming language security as the enforcement mechanism.

For example, to isolate multiple executing enterprise bean instances, the EJB container can load the multiple instances into the same JVM and isolate them via using multiple class loaders, or it can load each instance into its own JVM and rely on the address space protection provided by the operating system.

The general security enforcement requirements for the EJB container follow:

- The EJB container must provide enforcement of the client access control per the policy defined by the Deployer. A caller is allowed to invoke a method if, and only if, the method is specified as unchecked or the caller is assigned *at least one* of the security roles that includes the method in its method permissions definition. (That is, it is not meant that the caller must be assigned *all* the roles associated with the method.) If the container denies a client access to a business method, the container must throw the `java.rmi.RemoteException` (or its sub-

class, the `java.rmi.AccessException`) to the client if the client is a remote client, or the `javax.ejb.EJBException` (or its subclass, the `javax.ejb.AccessLocalException`) if the client is a local client.

- The EJB container must isolate an enterprise bean instance from other instances and other application components running on the server. The EJB container must ensure that other enterprise bean instances and other application components are allowed to access an enterprise bean only via the enterprise bean's component, home, and/or web service endpoint interfaces.
- The EJB container must isolate an enterprise bean instance at runtime such that the instance does not gain unauthorized access to privileged system information. Such information includes the internal implementation classes of the container, the various runtime state and context maintained by the container, object references of other enterprise bean instances, or resource managers used by other enterprise bean instances. The EJB container must ensure that the interactions between the enterprise beans and the container are only through the EJB architected interfaces.
- The EJB container must ensure the security of the persistent state of the enterprise beans.
- The EJB container must manage the mapping of principals on calls to other enterprise beans or on access to resource managers according to the security policy defined by the Deployer.
- The container must allow the same enterprise bean to be deployed independently multiple times, each time with a different security policy^[53]. The container must allow multiple-deployed enterprise beans to co-exist at runtime.

21.6.10 Audit Trail

The EJB container may provide a security audit trail mechanism. A security audit trail mechanism typically logs all `java.security.Exceptions`. It also logs all denials of access to EJB servers, EJB container, EJB component interfaces, EJB home interfaces, and EJB web service endpoint interfaces.

21.7 System Administrator's Responsibilities

This section defines the security-related responsibilities of the System Administrator. Note that some responsibilities may be carried out by the Deployer instead, or may require cooperation of the Deployer and the System Administrator.

21.7.1 Security Domain Administration

The System Administrator is responsible for the administration of principals. Security domain administration is beyond the scope of the EJB specification.

[53] The enterprise bean is installed each time using a different JNDI name.

Typically, the System Administrator is responsible for creating a new user account, adding a user to a user group, removing a user from a user group, and removing or freezing a user account.

21.7.2 Principal Mapping

If the client is in a different security domain than the target enterprise bean, the System Administrator is responsible for mapping the principals used by the client to the principals defined for the enterprise bean. The result of the mapping is available to the Deployer.

The specification of principal mapping techniques is beyond the scope of the EJB architecture.

21.7.3 Audit Trail Review

If the EJB container provides an audit trail facility, the System Administrator is responsible for its management.

Timer Service

This chapter defines the EJB container-managed timer service. The EJB timer service is a container-provided service that allows the Bean Provider to register enterprise beans for timer callbacks to occur at a specified time, after a specified elapsed time, or at specified intervals.

22.1 Overview

Enterprise applications that model workflow-type business processes are dependent on notifications that certain temporal events have occurred in order to manage the semantic state transitions that are intrinsic to the business processes that they model.

The EJB Timer Service is a container-managed service that provides methods to allow callbacks to be scheduled for time-based events. The container provides a reliable and transactional notification service for timed events. Timer notifications may be scheduled to occur at a specific time, after a specific elapsed duration, or at specific recurring intervals.

The Timer Service is implemented by the EJB container. An enterprise bean accesses this service through the `EJBContext` interface.

The EJB Timer Service is a coarse-grained timer notification service that is designed for use in the modeling of application-level processes. It is not intended for the modeling of real-time events.

While timer durations are expressed in millisecond units, this is because the millisecond is the unit of time granularity used by the APIs of the J2SE platform. It is expected that most timed events will correspond to hours, days, or longer periods of time.

The following sections describe the Timer Service with respect to the various individual EJB roles.

22.2 Bean Provider's View of the Timer Service

The EJB Timer Service is a container-provided service that allows enterprise beans to be registered for timer callback methods to occur at a specified time, after a specified elapsed time, or after specified intervals.

An enterprise bean accesses the timer service through its `EJBContext` interface. The timer service provides methods for the creation and cancellation of timers, as well as for locating the timers that are associated with a bean.

A timer is created to schedule timed callbacks. The bean class of an enterprise bean that uses the timer service must implement the `javax.ejb.TimerObject` interface. This interface has a single method, the timer callback method, `ejbTimeout`. Timers can be created for stateless session beans, message-driven beans, and entity beans. Timers cannot be created for stateful session beans.^[54]

A timer that is created for an entity bean is associated with the entity bean's identity. The `ejbTimeout` method invocation for a timer that is created for a stateless session bean or a message-driven bean may be called on any bean instance in the pooled state.

When the time specified at timer creation elapses, the container invokes the `ejbTimeout` method of the bean. A timer may be cancelled by a bean before its expiration. If a timer is cancelled, the `ejbTimeout` method is not called^[55]. A timer is cancelled by calling its `cancel` method.

Invocations of the methods to create and cancel timers and of the `ejbTimeout` method are typically made within a transaction.

The timer service is intended for the modelling of long-lived business processes. Timers survive container crashes and the activation/passivation and load/store cycles of the enterprise beans that are registered with them.

[54] This functionality may be added in a future release of this specification.

[55] In the event of race conditions, extraneous calls to the `ejbTimeout` method may occur.

22.2.1 The Timer Service Interface

The Timer Service is accessed through the `getTimerService` method of the `EJBContext` interface. The `TimerService` interface has the following methods:

```
public interface javax.ejb.TimerService {  
  
    public Timer createTimer(long duration,  
                             java.io.Serializable info);  
  
    public Timer createTimer(long initialDuration,  
                             long intervalDuration, java.io.Serializable info);  
  
    public Timer createTimer(java.util.Date expiration,  
                             java.io.Serializable info);  
  
    public Timer createTimer(java.util.Date initialExpiration,  
                             long intervalDuration, java.io.Serializable info);  
  
    public Collection getTimers();  
  
}
```

The timer creation methods allow a timer to be created as a single-event timer or as an interval timer. The timer expiration (initial expiration in the case of an interval timer) may be expressed either in terms of a duration or as an absolute time.

The bean may pass some client-specific information at timer creation to help it recognize the significance of the timer's expiration. This information is stored by the timer service and available through the timer. The information object must be serializable. ^[56]

The timer duration is expressed in terms of milliseconds. The timer service begins counting down the timer duration upon timer creation.

The `createTimer` methods return a `Timer` object that allows the bean to cancel the timer or to obtain information about the timer prior to its cancellation and/or expiration (if it is a single-event timer).

The `getTimers` method returns the active timers associated with the bean. For an entity bean, the result of `getTimers` is a collection of those timers that are associated with the bean's identity.

[56] There is currently no way to set the information object other than through the `createTimer` method. An API to do this may be added in a future release of this specification.

22.2.2 The TimedObject Interface

The enterprise bean class of a bean that is to be registered with the timer service for timer callbacks must implement the `javax.ejb.TimedObject` interface. This interface has a single method, `ejbTimeout`.

```
public interface javax.ejb.TimedObject {  
    public void ejbTimeout(Timer timer);  
}
```

When the timer expires (i.e., after the number of milliseconds specified at its creation expires or after the absolute time specified has passed), the container calls the `ejbTimeout` method of the bean that was registered for the timer. The `ejbTimeout` method contains the business logic that the Bean Provider supplies to handle the timeout event. The container calls the `ejbTimeout` method with the timer that has expired. The Bean Provider can use the `getInfo` method to retrieve the information that was supplied when the timer was created. This information may be useful in enabling the timed object to recognize the significance of the timer expiration.

The container interleaves calls to the `ejbTimeout` method with the calls to the business methods and the life cycle methods (`ejbLoad`, `ejbStore`, etc.) of the bean. The time at which the `ejbTimeout` method is called may therefore not correspond exactly to the time specified at timer creation. If multiple timers have been created for a bean and will expire at approximately the same times, the Bean Provider must be prepared to handle timeout callbacks that are out of sequence. The Bean Provider must be prepared to handle extraneous calls to the `ejbTimeout` method in the event that a timer expiration is outstanding when a call to the cancellation method has been made.

In general, the `ejbTimeout` method can perform the same operations as business methods from the component interface or methods from the message listener interface. See Tables 3, 4, 10, and 12 for the specification of the operations that may be performed by the `ejbTimeout` method.

Since the `ejbTimeout` method is an internal method of the bean class, it has no client security context. When `getCallerPrincipal` is called from within the `ejbTimeout` method, it returns the container's representation of the unauthenticated identity.

If the timed object needs to make use of the identity of the timer to recognize the significance of the timer expiration, it may use the `equals` method to compare it with any other timer references it might have outstanding.

If the timer is a single-action timer, the container removes the timer after the `ejbTimeout` method has been successfully invoked (e.g., when the transaction that has been started for the invocation of the `ejbTimeout` method commits). If the bean invokes a method on the timer after the termination of the `ejbTimeout` method, the `NoSuchObjectLocalException` is thrown.

22.2.3 The Timer and TimerHandle Interfaces

The `javax.ejb.Timer` interface allows the Bean Provider to cancel a timer and to obtain information about the timer.

The `javax.ejb.TimerHandle` interface allows the Bean Provider to obtain a serializable timer handle that may be persisted. Since timers are local objects, a `TimerHandle` must not be passed through a bean's remote interface or web service interface.

The methods of these interfaces are as follows:

```
public interface javax.ejb.Timer {
    public void cancel();
    public long getTimeRemaining();
    public java.util.Date getNextTimeout();
    public javax.ejb.TimerHandle getHandle();
    public java.io.Serializable getInfo();
}

public interface javax.ejb.TimerHandle extends java.io.Serializable{
    public javax.ejb.Timer getTimer();
}
```

22.2.4 Timer Identity

The Bean Provider cannot rely on the `==` operator to compare timers for "object equality". The Bean Provider must use the `Timer.equals(Object obj)` method.

22.2.5 Transactions

An enterprise bean typically creates a timer within the scope of a transaction. If the transaction is then rolled back, the timer creation is rolled back.

An enterprise bean typically cancels a timer within a transaction. If the transaction is rolled back, the container rescinds the timer cancellation.

The `ejbTimeout` method is typically specified to have transaction attribute `RequiresNew`. If the transaction is rolled back, the container retries the timeout.

22.3 Bean Provider's Responsibilities

This section defines the Bean Provider's responsibilities.

22.3.1 Enterprise Bean Class

An enterprise bean that is to be registered with the Timer Service must implement the `javax.ejb.TimedObject` interface. The enterprise bean class may have superclasses and/or superinterfaces. If the bean class has superclasses, then the `ejbTimeout` method may be defined in the bean class, or in any of its superclasses.

22.3.2 TimerHandle

Since the `TimerHandle` interface extends `java.io.Serializable`, a client may serialize the handle. The serialized handle may be used later to obtain a reference to the timer identified by the handle. A `TimerHandle` is intended to be storable in persistent storage.

A `TimerHandle` must not be passed as an argument or result of an enterprise bean's remote interface or web service endpoint interface.

22.4 Container's Responsibilities

This section describes the responsibilities of the Container Provider to support the EJB Timer Service.

22.4.1 TimerService, Timer, and TimerHandle Interfaces

The container must provide the implementation of the `TimerService`, `Timer`, and `TimerHandle` interfaces.

Timer instances must not be serializable.

The container must implement a timer handle to be usable over the lifetime of the timer.

The container must provide suitable implementations of the `Timer equals(Object obj)` and `hashCode()` methods.

22.4.2 Timer Expiration and ejbTimeout Method

The container must call the `ejbTimeout` method after the timed duration or the absolute time specification in the timer creation method has passed. The timer service must begin to count down the timer duration upon timer creation. The container must call the `ejbTimeout` method with the expired `Timer`.

If container-managed transaction demarcation is used and the `RequiresNew` transaction attribute is specified, the container must begin a new transaction prior to invoking the `ejbTimeout` method. If the transaction fails or is rolled back, the container must retry the timeout at least once.

If a timer for an entity bean expires, and the bean has been passivated, the container must call the `ejbActivate` and `ejbLoad` methods on the entity bean class before calling the `ejbTimeout` method, as described in Sections 10.5.3 and 12.1.4.2.

If the timer is a single-event timer, the container must cause the timer to no longer exist. If a method is subsequently invoked on the timer after the completion of the `ejbTimeout` method, the container must throw the `javax.ejb.NoSuchObjectLocalException`.

If the Bean Provider invokes the `setRollbackOnly` method from within the `ejbTimeout` method, the container must rollback the transaction in which the `ejbTimeout` method is invoked. This has the effect of rescinding the timer expiration. The container must retry the timeout after the transaction rollback.

Timers are persistent objects. In the event of a container crash, any single-event timers that have expired during the intervening time before container restart must cause the `ejbTimeout` method to be invoked upon restart. Any interval timers that have expired during the intervening time must cause the `ejbTimeout` method to be invoked at least once upon restart.

22.4.3 Timer Cancellation

When a timer's `cancel` method has been called, the container must cause the timer to no longer exist. If a method is subsequently invoked on the timer, the container must throw the `javax.ejb.NoSuchObjectLocalException`.

If the transaction in which the timer cancellation occurs is rolled back, the container must restore the duration of the timer to the duration it would have had if it had not been cancelled. If the timer would have expired by the time that the transaction failed, the failure of the transaction should result in the expired timer providing an expiration notification after the transaction rolls back.

22.4.4 Entity Bean Removal

If an entity bean is removed, the container must remove the timers for that bean.

Deployment Descriptor

This chapter defines the deployment descriptor that is part of the ejb-jar file. Section 23.1 provides an overview of the deployment descriptor. Sections 23.2 through 23.4 describe the information in the deployment descriptor from the perspective of the EJB roles responsible for providing the information. Section 23.5 defines the deployment descriptor's XML Schema elements that are specific to the EJB architecture. The XML Schema elements that are common to the J2EE Platform specifications are provided in [9].

23.1 Overview

The deployment descriptor is part of the contract between the ejb-jar file producer and consumer. This contract covers both the passing of enterprise beans from the Bean Provider to the Application Assembler, and from the Application Assembler to the Deployer.

An ejb-jar file produced by the Bean Provider contains one or more enterprise beans and typically does not contain application assembly instructions. An ejb-jar file produced by an Application Assembler contains one or more enterprise beans, plus application assembly information describing how the enterprise beans are combined into a single application deployment unit.

The J2EE specification defines how enterprise beans and other application components contained in multiple ejb-jar files can be assembled into an application.

The role of the deployment descriptor is to capture the declarative information (i.e. information that is not included directly in the enterprise beans' code) that is intended for the consumer of the ejb-jar file.

There are two basic kinds of information in the deployment descriptor:

- *Enterprise beans' structural information.* Structural information describes the structure of an enterprise bean and declares an enterprise bean's external dependencies. Providing structural information in the deployment descriptor is mandatory for the ejb-jar file producer. The structural information cannot, in general, be changed because doing so could break the enterprise bean's function.
- *Application assembly information.* Application assembly information describes how the enterprise bean (or beans) in the ejb-jar file is composed into a larger application deployment unit. Providing assembly information in the deployment descriptor is optional for the ejb-jar file producer. Assembly level information can be changed without breaking the enterprise bean's function, although doing so may alter the behavior of an assembled application.

23.2 Bean Provider's Responsibilities

The Bean Provider is responsible for providing the structural information for each enterprise bean in the deployment descriptor.

The Bean Provider must use the `enterprise-beans` element to list all the enterprise beans in the ejb-jar file.

The Bean Provider must provide the following information for each enterprise bean:

- **Enterprise bean's name.** The Bean Provider must assign a logical name to each enterprise bean in the ejb-jar file. There is no architected relationship between this name and the JNDI name that the Deployer will assign to the enterprise bean. The Bean Provider specifies the enterprise bean's name in the `ejb-name` element.
- **Enterprise bean's class.** The Bean Provider must specify the fully-qualified name of the Java class that implements the enterprise bean's business methods. The Bean Provider specifies the enterprise bean's class name in the `ejb-class` element.
- **Enterprise bean's remote home interface.** The Bean Provider must specify the fully-qualified name of the enterprise bean's remote home interface, if any, in the `home` element.
- **Enterprise bean's remote interface.** The Bean Provider must specify the fully-qualified name of the enterprise bean's remote interface, if any, in the `remote` element.
- **Enterprise bean's local home interface.** The Bean Provider must specify the fully-qualified name of the enterprise bean's local home interface, if any, in the `local-home` element.
- **Enterprise bean's local interface.** The Bean Provider must specify the fully-qualified name of the enterprise bean's local interface, if any, in the `local` element.

- **Enterprise bean's web service endpoint interface.** The Bean Provider must specify the fully-qualified name of the enterprise bean's web service endpoint interface, in any, in the `service-endpoint` element. This element may only be used for stateless session beans.
- **Enterprise bean's type.** The enterprise bean types are: session, entity, and message-driven. The Bean Provider must use the appropriate `session`, `entity`, or `message-driven` element to declare the enterprise bean's structural information.
- **Re-entrancy indication.** The Bean Provider must specify whether an entity bean is re-entrant or not. Session beans and message-driven beans are never re-entrant.
- **Session bean's state management type.** If the enterprise bean is a session bean, the Bean Provider must use the `session-type` element to declare whether the session bean is stateful or stateless.
- **Session or message-driven bean's transaction demarcation type.** If the enterprise bean is a session bean or message-driven bean, the Bean Provider must use the `transaction-type` element to declare whether transaction demarcation is performed by the enterprise bean or by the container.
- **Entity bean's persistence management.** If the enterprise bean is an entity bean, the Bean Provider must use the `persistence-type` element to declare whether persistence management is performed by the enterprise bean or by the container.
- **Entity bean's primary key class.** If the enterprise bean is an entity bean, the Bean Provider specifies the fully-qualified name of the entity bean's primary key class in the `prim-key-class` element. The Bean Provider *must* specify the primary key class for an entity with bean-managed persistence.
- **Entity bean's abstract schema name.** If the enterprise bean is an entity bean with container-managed persistence and `cmp-version 2.x`, the Bean Provider must specify the abstract schema name of the entity bean using the `abstract-schema-name` element.
- **Container-managed fields.** If the enterprise bean is an entity bean with container-managed persistence, the Bean Provider must specify the container-managed fields using the `cmp-field` elements.
- **Container-managed relationships.** If the enterprise bean is an entity bean with container-managed persistence and `cmp-version 2.x`, the Bean Provider must specify the container-managed relationships of the entity bean using the `relationships` element.
- **Finder and select queries.** If the enterprise bean is an entity bean with container-managed persistence and `cmp-version 2.x`, the Bean Provider must use the `query` element to specify any EJB QL finder or select query for the entity bean other than a query for the `findByPrimaryKey` method.
- **Environment entries.** The Bean Provider must declare all the enterprise bean's environment entries as specified in Subsection 20.2.1.

- **Resource manager connection factory references.** The Bean Provider must declare all the enterprise bean's resource manager connection factory references as specified in Subsection 20.5.1.
- **Resource environment references.** The Bean Provider must declare all the enterprise bean's references to administered objects that are associated with resources as specified in Subsection 20.6.1.
- **EJB references.** The Bean Provider must declare all the enterprise bean's references to the remote homes of other enterprise beans as specified in Subsection 20.3.1.
- **EJB local references.** The Bean Provider must declare all the enterprise bean's references to the local homes of other enterprise beans as specified in Subsection 20.3.1.
- **Web service references.** The Bean Provider must declare all the enterprise bean's references to web service interfaces as specified in Subsection 20.4.1.
- **Message destination references.** The Bean Provider must declare all the enterprise bean's references to message destinations as specified in Subsection 20.7.1.
- **Security role references.** The Bean Provider must declare all the enterprise bean's references to security roles as specified in Subsection 21.2.5.3.
- **Message-driven bean's configuration properties.** The Bean Provider may provide input to the Deployer as to how a message-driven bean should be configured upon activation in its operational environment. Activation configuration properties for a JMS message-driven bean include information about a bean's intended destination type, its message selector, and its acknowledgement mode. Other bean types may make use of different properties. See [12].
- **Message-driven bean's destination.** The Bean Provider may provide advice to the Application Assembler as to the destination type to which a message-driven bean should be assigned when linking message destinations.

The deployment descriptor produced by the Bean Provider must conform to the XML Schema definition in Section 23.5 or the DTD definition from a previous version of this specification. The content of the deployment descriptor must conform to the semantics rules specified in the XML Schema or DTD comments and elsewhere in this specification.

23.3 Application Assembler's Responsibility

The Application Assembler assembles enterprise beans into a single deployment unit. The Application Assembler's input is one or more ejb-jar files provided by one or more Bean Providers, and the output is also one or more ejb-jar files. The Application Assembler can combine multiple input ejb-jar files into a single output ejb-jar file, or split an input ejb-jar file into multiple output ejb-jar files. Each output ejb-jar file is either a deployment unit intended for the Deployer, or a partially assembled application that is intended for another Application Assembler.

The Bean Provider and Application Assembler may be the same person or organization. In such a case, the person or organization performs the responsibilities described both in this and the previous sections.

The Application Assembler may modify the following information that was specified by the Bean Provider:

- **Values of environment entries.** The Application Assembler may change existing and/or define new values of environment properties.
- **Description fields.** The Application Assembler may change existing or create new description elements.
- **Relationship names.** If multiple ejb-jar files use the same names for relationships and are merged into a single ejb-jar file, it is the responsibility of the Application Assembler to modify the relationship names defined in the ejb-relation-name elements.
- **Message-driven bean message selector.** The Application Assembler may further restrict, but not replace, the value of the messageSelector activation-config-property element of a JMS message-driven bean.

In general, the Application Assembler should never modify any of the following.

- **Enterprise bean's abstract schema name.** The Application Assembler should not change the enterprise bean's name defined in the abstract-schema-name element since EJB QL queries may depend on the content of this element.
- **Relationship role source element.** The Application Assembler should not change the value of an ejb-name element in the relationship-role-source element.

If any of these elements must be modified by the Application Assembler in order to resolve name clashes during the merging two ejb-jar files into one, the Application Assembler must also modify all ejb-ql query strings that depend on the value of the modified element(s).

The Application Assembler must not, in general, modify any other information listed in Section 23.2 that was provided in the input ejb-jar file.

The Application Assembler may, but is not required to, specify any of the following application assembly information:

- **Binding of enterprise bean references.** The Application Assembler may link an enterprise bean reference to another enterprise bean in the ejb-jar file or in the same J2EE application unit. The Application Assembler creates the link by adding the ejb-link element to the referencing bean. The Application Assembler uses the ejb-name of the referenced bean for the link. If there are multiple enterprise beans with the same ejb-name, the Application Assembler uses the path name specifying the location of the ejb-jar file that contains the referenced component. The path name is relative to the referencing ejb-jar file. The Application Assembler appends the ejb-name of the referenced bean to the path name separated by #. This allows multiple beans with the same name to be uniquely identified.

- **Linking of message destination references.** The Application Assembler may link message consumers and producers through common message destinations specified in the `ejb-jar` file or in the same J2EE application unit. The Application Assembler creates the link by adding the `message-destination-link` element to the referencing bean.
- **Security roles.** The Application Assembler may define one or more security roles. The security roles define the *recommended* security roles for the clients of the enterprise beans. The Application Assembler defines the security roles using the `security-role` elements.
- **Method permissions.** The Application Assembler may define method permissions. Method permission is a binary relation between the security roles and the methods of the home, component, and/or web service endpoint interfaces of the enterprise beans. The Application Assembler defines method permissions using the `method-permission` elements.
- **Linking of security role references.** If the Application Assembler defines security roles in the deployment descriptor, the Application Assembler must link the security role references declared by the Bean Provider to the security roles. The Application Assembler defines these links using the `role-link` element.
- **Security identity.** The Application Assembler may specify whether the caller's security identity should be used for the execution of the methods of an enterprise bean or whether a specific run-as security identity should be used.
- **Transaction attributes.** The Application Assembler may define the value of the transaction attributes for the methods of the home, component, web service endpoint, and `TimedObject` interfaces of the enterprise beans that require container-managed transaction demarcation. All entity beans and the session and message-driven beans declared by the Bean Provider as transaction-type `Container` require container-managed transaction demarcation. The Application Assembler uses the `container-transaction` elements to declare the transaction attributes.

If an input `ejb-jar` file contains application assembly information, the Application Assembler is allowed to change the application assembly information supplied in the input `ejb-jar` file. (This could happen when the input `ejb-jar` file was produced by another Application Assembler.)

The deployment descriptor produced by the Bean Provider must conform to the XML Schema definition in Section 23.5 or the DTD definition from a previous version of this specification. The content of the deployment descriptor must conform to the semantics rules specified in the XML Schema or DTD comments and elsewhere in this specification.

23.4 Container Provider's Responsibilities

The Container Provider provides tools that read and import the information contained in the XML deployment descriptor.

All EJB 2.1 implementations must support EJB 1.1 and EJB 2.0 as well as EJB 2.1 deployment descriptors.

The EJB 2.0 deployment descriptor is defined in Appendix B.

The EJB 1.1 deployment descriptor is defined in Appendix C.

23.5 Deployment Descriptor XML Schema

This section provides the XML Schema for the EJB 2.1 deployment descriptor. The comments in the XML Schema specify additional requirements for the syntax and semantics that cannot be easily expressed by the XML Schema mechanism.

The content of the XML elements is in general case sensitive (i.e., unless stated otherwise). This means, for example, that

```
<transaction-type>Container</transaction-type>
```

must be used, rather than:

```
<transaction-type>container</transaction-type>.
```

All valid ejb-jar deployment descriptors must conform to the XML Schema definition below or to the DTD definition from a previous version of this specification.

```
<?xml version="1.0" encoding="UTF-8"?>
<xsd:schema xmlns="http://www.w3.org/2001/XMLSchema"
  targetNamespace="http://java.sun.com/xml/ns/j2ee"
  xmlns:j2ee="http://java.sun.com/xml/ns/j2ee"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema"
  elementFormDefault="qualified"
  attributeFormDefault="unqualified"
  version="2.1">

  <xsd:annotation>
    <xsd:documentation>

      @(#)ejb-jar_2_1.xsds      1.23 08/01/03

    </xsd:documentation>
  </xsd:annotation>

  <xsd:annotation>
    <xsd:documentation>

      This is the XML Schema for the EJB 2.1 deployment descriptor.
      The deployment descriptor must be named "META-INF/ejb-jar.xml" in
      the EJB's jar file. All EJB deployment descriptors must indicate
      the ejb-jar schema by using the J2EE namespace:

    </xsd:documentation>
  </xsd:annotation>
```

http://java.sun.com/xml/ns/j2ee

and by indicating the version of the schema by
using the version element as shown below:

```
<ejb-jar xmlns="http://java.sun.com/xml/ns/j2ee"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://java.sun.com/xml/ns/j2ee
    http://java.sun.com/xml/ns/j2ee/ebj-jar_2_1.xsd"
  version="2.1">
  ...
</ejb-jar>
```

The instance documents may indicate the published version of the schema using the `xsi:schemaLocation` attribute for the J2EE namespace with the following location:

http://java.sun.com/xml/ns/j2ee/ebj-jar_2_1.xsd

```
</xsd:documentation>
</xsd:annotation>
```

```
<xsd:annotation>
  <xsd:documentation>
```

The following conventions apply to all J2EE deployment descriptor elements unless indicated otherwise.

- In elements that specify a pathname to a file within the same JAR file, relative filenames (i.e., those not starting with "/") are considered relative to the root of the JAR file's namespace. Absolute filenames (i.e., those starting with "/") also specify names in the root of the JAR file's namespace. In general, relative names are preferred. The exception is .war files where absolute names are preferred for consistency with the Servlet API.

```
</xsd:documentation>
</xsd:annotation>
```

```
<xsd:include schemaLocation="j2ee_1_4.xsd"/>
```

```
<!-- ***** -->
```

```
<xsd:element name="ejb-jar" type="j2ee:ejb-jarType">
```

```
  <xsd:annotation>
    <xsd:documentation>
```

This is the root of the ejb-jar deployment descriptor.

</xsd:documentation>

</xsd:annotation>

<xsd:key name="ejb-name-key">

<xsd:annotation>

<xsd:documentation>

The ejb-name element contains the name of an enterprise bean. The name must be unique within the ejb-jar file.

</xsd:documentation>

</xsd:annotation>

<xsd:selector xpath="j2ee:enterprise-beans/*"/>

<xsd:field xpath="j2ee:ejb-name"/>

</xsd:key>

<xsd:keyref name="ejb-name-references"
refer="j2ee:ejb-name-key">

<xsd:annotation>

<xsd:documentation>

The keyref indicates the references from relationship-role-source must be to a specific ejb-name defined within the scope of enterprise-beans element.

</xsd:documentation>

</xsd:annotation>

<xsd:selector

xpath="//j2ee:ejb-relationship-role/j2ee:relationship-role-source"/>

<xsd:field

xpath="j2ee:ejb-name"/>

</xsd:keyref>

<xsd:key name="role-name-key">

<xsd:annotation>

<xsd:documentation>

A role-name-key is specified to allow the references from the security-role-refs.

</xsd:documentation>

</xsd:annotation>

```

    <xsd:selector xpath="j2ee:assembly-descriptor/j2ee:security-role"/>
    <xsd:field xpath="j2ee:role-name"/>
  </xsd:key>
  <xsd:keyref name="role-name-references"
    refer="j2ee:role-name-key">

    <xsd:annotation>
      <xsd:documentation>

        The keyref indicates the references from
        security-role-ref to a specified role-name.

      </xsd:documentation>
    </xsd:annotation>

    <xsd:selector xpath="j2ee:enterprise-beans/*/j2ee:security-role-ref"/>
    <xsd:field xpath="j2ee:role-link"/>
  </xsd:keyref>
</xsd:element>

<!-- ***** -->
<xsd:complexType name="activation-config-propertyType">

  <xsd:annotation>
    <xsd:documentation>

      The activation-config-propertyType contains a name/value
      configuration property pair for a message-driven bean.

      The properties that are recognized for a particular
      message-driven bean are determined by the messaging type.

    </xsd:documentation>
  </xsd:annotation>

  <xsd:sequence>
    <xsd:element name="activation-config-property-name"
      type="j2ee:xsdStringType">

      <xsd:annotation>
        <xsd:documentation>

          The activation-config-property-name element contains
          the name for an activation configuration property of
          a message-driven bean.

          For JMS message-driven beans, the following property
          names are recognized: acknowledgeMode,

```

```

        messageSelector, destinationType, subscriptionDurability

    </xsd:documentation>
</xsd:annotation>

</xsd:element>
<xsd:element name="activation-config-property-value"
    type="j2ee:xsdStringType">

    <xsd:annotation>
        <xsd:documentation>

            The activation-config-property-value element
            contains the value for an activation configuration
            property of a message-driven bean.

        </xsd:documentation>
    </xsd:annotation>

</xsd:element>
</xsd:sequence>
<xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="activation-configType">

    <xsd:annotation>
        <xsd:documentation>

            The activation-configType defines information about the
            expected configuration properties of the message-driven bean
            in its operational environment. This may include information
            about message acknowledgement, message selector, expected
            destination type, etc.

            The configuration information is expressed in terms of
            name/value configuration properties.

            The properties that are recognized for a particular
            message-driven bean are determined by the messaging type.

        </xsd:documentation>
    </xsd:annotation>

    <xsd:sequence>
        <xsd:element name="description"
            type="j2ee:descriptionType"

```

```

        minOccurs="0"
        maxOccurs="unbounded"/>
<xsd:element name="activation-config-property"
        type="j2ee:activation-config-propertyType"
        maxOccurs="unbounded"/>
</xsd:sequence>
<xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="assembly-descriptorType">

<xsd:annotation>
<xsd:documentation>

    The assembly-descriptorType defines
    application-assembly information.

    The application-assembly information consists of the
    following parts: the definition of security roles, the
    definition of method permissions, the definition of
    transaction attributes for enterprise beans with
    container-managed transaction demarcation and a list of
    methods to be excluded from being invoked.

    All the parts are optional in the sense that they are
    omitted if the lists represented by them are empty.

    Providing an assembly-descriptor in the deployment
    descriptor is optional for the ejb-jar file producer.

</xsd:documentation>
</xsd:annotation>

<xsd:sequence>
<xsd:element name="security-role"
        type="j2ee:security-roleType"
        minOccurs="0"
        maxOccurs="unbounded"/>
<xsd:element name="method-permission"
        type="j2ee:method-permissionType"
        minOccurs="0"
        maxOccurs="unbounded"/>
<xsd:element name="container-transaction"
        type="j2ee:container-transactionType"
        minOccurs="0"
        maxOccurs="unbounded"/>
<xsd:element name="message-destination"

```



```

        type="j2ee:message-destinationType"
        minOccurs="0"
        maxOccurs="unbounded"/>
<xsd:element name="exclude-list"
        type="j2ee:exclude-listType"
        minOccurs="0"/>
</xsd:sequence>
<xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="cmp-fieldType">

    <xsd:annotation>
        <xsd:documentation>

            The cmp-fieldType describes a container-managed field. The
            cmp-fieldType contains an optional description of the field,
            and the name of the field.

        </xsd:documentation>
    </xsd:annotation>

    <xsd:sequence>
        <xsd:element name="description"
            type="j2ee:descriptionType"
            minOccurs="0"
            maxOccurs="unbounded"/>
        <xsd:element name="field-name"
            type="j2ee:java-identifierType">

            <xsd:annotation>
                <xsd:documentation>

                    The field-name element specifies the name of a
                    container managed field.

                    The name of the cmp-field of an entity bean with
                    cmp-version 2.x must begin with a lowercase
                    letter. This field is accessed by methods whose
                    names consists of the name of the field specified by
                    field-name in which the first letter is uppercased,
                    prefixed by "get" or "set".

                    The name of the cmp-field of an entity bean with
                    cmp-version 1.x must denote a public field of the
                    enterprise bean class or one of its superclasses.

```

```

        </xsd:documentation>
    </xsd:annotation>

</xsd:element>
</xsd:sequence>
<xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="cmp-versionType">

    <xsd:annotation>
        <xsd:documentation>

            The cmp-versionType specifies the version of an entity bean
            with container-managed persistence. It is used by
            cmp-version elements.

            The value must be one of the two following:

                1.x
                2.x

        </xsd:documentation>
    </xsd:annotation>

    <xsd:simpleContent>
        <xsd:restriction base="j2ee:string">
            <xsd:enumeration value="1.x"/>
            <xsd:enumeration value="2.x"/>
        </xsd:restriction>
    </xsd:simpleContent>
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="cmr-field-typeType">

    <xsd:annotation>
        <xsd:documentation>

            The cmr-field-type element specifies the class of a
            collection-valued logical relationship field in the entity
            bean class. The value of an element using cmr-field-typeType
            must be either: java.util.Collection or java.util.Set.

        </xsd:documentation>
    </xsd:annotation>

```

```

    <xsd:simpleContent>
      <xsd:restriction base="j2ee:string">
        <xsd:enumeration value="java.util.Collection"/>
        <xsd:enumeration value="java.util.Set"/>
      </xsd:restriction>
    </xsd:simpleContent>
  </xsd:complexType>

<!-- ***** -->
  <xsd:complexType name="cmr-fieldType">

    <xsd:annotation>
      <xsd:documentation>

        The cmr-fieldType describes the bean provider's view of
        a relationship. It consists of an optional description, and
        the name and the class type of a field in the source of a
        role of a relationship. The cmr-field-name element
        corresponds to the name used for the get and set accessor
        methods for the relationship. The cmr-field-type element is
        used only for collection-valued cmr-fields. It specifies the
        type of the collection that is used.

      </xsd:documentation>
    </xsd:annotation>

    <xsd:sequence>
      <xsd:element name="description"
        type="j2ee:descriptionType"
        minOccurs="0"
        maxOccurs="unbounded"/>
      <xsd:element name="cmr-field-name"
        type="j2ee:string">

        <xsd:annotation>
          <xsd:documentation>

            The cmr-field-name element specifies the name of a
            logical relationship field in the entity bean
            class. The name of the cmr-field must begin with a
            lowercase letter. This field is accessed by methods
            whose names consist of the name of the field
            specified by cmr-field-name in which the first
            letter is uppercased, prefixed by "get" or "set".

          </xsd:documentation>
        </xsd:annotation>
      </xsd:element>
    </xsd:sequence>
  </xsd:complexType>

```

```

    </xsd:element>
    <xsd:element name="cmr-field-type"
        type="j2ee:cmr-field-typeType"
        minOccurs="0"/>
</xsd:sequence>
<xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="container-transactionType">

    <xsd:annotation>
        <xsd:documentation>

            The container-transactionType specifies how the container
            must manage transaction scopes for the enterprise bean's
            method invocations. It defines an optional description, a
            list of method elements, and a transaction attribute. The
            transaction attribute is to be applied to all the specified
            methods.

        </xsd:documentation>
    </xsd:annotation>

    <xsd:sequence>
        <xsd:element name="description"
            type="j2ee:descriptionType"
            minOccurs="0"
            maxOccurs="unbounded"/>
        <xsd:element name="method"
            type="j2ee:methodType"
            maxOccurs="unbounded"/>
        <xsd:element name="trans-attribute"
            type="j2ee:trans-attributeType"/>
    </xsd:sequence>
    <xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="ejb-classType">

    <xsd:annotation>
        <xsd:documentation>

            The ejb-classType contains the fully-qualified name of the
            enterprise bean's class. It is used by ejb-class elements.

            Example:

```

```

        <ejb-class>com.wombat.empl.EmployeeServiceBean</ejb-class>

    </xsd:documentation>
</xsd:annotation>

<xsd:simpleContent>
    <xsd:restriction base="j2ee:fully-qualified-classType"/>
</xsd:simpleContent>
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="ejb-jarType">

    <xsd:annotation>
        <xsd:documentation>

            The ejb-jarType defines the root element of the EJB
            deployment descriptor. It contains

- an optional description of the ejb-jar file
- an optional display name
- an optional icon that contains a small and a large
                icon file name
- mandatory structural information about all included
                enterprise beans
- a descriptor for container managed relationships,
                if any
- an optional application-assembly descriptor
- an optional name of an ejb-client-jar file for the
                ejb-jar.



        </xsd:documentation>
    </xsd:annotation>

    <xsd:sequence>
        <xsd:group ref="j2ee:descriptionGroup"/>
        <xsd:element name="enterprise-beans"
            type="j2ee:enterprise-beansType"/>
        <xsd:element name="relationships"
            type="j2ee:relationshipsType"
            minOccurs="0">
            <xsd:unique name="relationship-name-uniqueness">

                <xsd:annotation>
                    <xsd:documentation>

                        The ejb-relation-name contains the name of a

```

```

        relation. The name must be unique within
        relationships.

</xsd:documentation>
</xsd:annotation>

<xsd:selector xpath="j2ee:ejb-relation"/>
<xsd:field    xpath="j2ee:ejb-relation-name"/>
</xsd:unique>
</xsd:element>
<xsd:element name="assembly-descriptor"
            type="j2ee:assembly-descriptorType"
            minOccurs="0">

<xsd:annotation>
  <xsd:documentation>

    Providing an assembly-descriptor in the deployment
    descriptor is optional for the ejb-jar file
    producer.

  </xsd:documentation>
</xsd:annotation>

</xsd:element>
<xsd:element name="ejb-client-jar"
            type="j2ee:pathType"
            minOccurs="0">

<xsd:annotation>
  <xsd:documentation>

    The optional ejb-client-jar element specifies a JAR
    file that contains the class files necessary for a
    client program to access the
    enterprise beans in the ejb-jar file.

    Example:

    <ejb-client-jar>employee_service_client.jar
    </ejb-client-jar>

  </xsd:documentation>
</xsd:annotation>

</xsd:element>
</xsd:sequence>
<xsd:attribute name="version"

```

```

        type="j2ee:dewey-versionType"
        fixed="2.1"
        use="required">

<xsd:annotation>
  <xsd:documentation>

    The version specifies the version of the
    EJB specification that the instance document must
    comply with. This information enables deployment tools
    to validate a particular EJB Deployment
    Descriptor with respect to a specific version of the EJB
    schema.

  </xsd:documentation>
</xsd:annotation>

</xsd:attribute>
<xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="ejb-nameType">

  <xsd:annotation>
    <xsd:documentation>

      The ejb-nameType specifies an enterprise bean's name. It is
      used by ejb-name elements. This name is assigned by the
      ejb-jar file producer to name the enterprise bean in the
      ejb-jar file's deployment descriptor. The name must be
      unique among the names of the enterprise beans in the same
      ejb-jar file.

      There is no architected relationship between the used
      ejb-name in the deployment descriptor and the JNDI name that
      the Deployer will assign to the enterprise bean's home.

      The name for an entity bean must conform to the lexical
      rules for an NMTOKEN.

      Example:

      <ejb-name>EmployeeService</ejb-name>

    </xsd:documentation>
  </xsd:annotation>

```

```

    <xsd:simpleContent>
      <xsd:restriction base="j2ee:xsdNMTOKENType"/>
    </xsd:simpleContent>
  </xsd:complexType>

<!-- ***** -->
  <xsd:complexType name="ejb-relationType">

    <xsd:annotation>
      <xsd:documentation>

        The ejb-relationType describes a relationship between two
        entity beans with container-managed persistence. It is used
        by ejb-relation elements. It contains a description; an
        optional ejb-relation-name element; and exactly two
        relationship role declarations, defined by the
        ejb-relationship-role elements. The name of the
        relationship, if specified, is unique within the ejb-jar
        file.

      </xsd:documentation>
    </xsd:annotation>

    <xsd:sequence>
      <xsd:element name="description"
        type="j2ee:descriptionType"
        minOccurs="0"
        maxOccurs="unbounded"/>
      <xsd:element name="ejb-relation-name"
        type="j2ee:string"
        minOccurs="0">

        <xsd:annotation>
          <xsd:documentation>

            The ejb-relation-name element provides a unique name
            within the ejb-jar file for a relationship.

          </xsd:documentation>
        </xsd:annotation>

      </xsd:element>
      <xsd:element name="ejb-relationship-role"
        type="j2ee:ejb-relationship-roleType"/>
      <xsd:element name="ejb-relationship-role"
        type="j2ee:ejb-relationship-roleType"/>
    </xsd:sequence>
  </xsd:complexType>

```



```

    <xsd:attribute name="id" type="xsd:ID"/>
  </xsd:complexType>

<!-- ***** -->
<xsd:complexType name="ejb-relationship-roleType">

  <xsd:annotation>
    <xsd:documentation>

      The ejb-relationship-roleType describes a role within a
      relationship. There are two roles in each relationship.

      The ejb-relationship-roleType contains an optional
      description; an optional name for the relationship role; a
      specification of the multiplicity of the role; an optional
      specification of cascade-delete functionality for the role;
      the role source; and a declaration of the cmr-field, if any,
      by means of which the other side of the relationship is
      accessed from the perspective of the role source.

      The multiplicity and role-source element are mandatory.

      The relationship-role-source element designates an entity
      bean by means of an ejb-name element. For bidirectional
      relationships, both roles of a relationship must declare a
      relationship-role-source element that specifies a cmr-field
      in terms of which the relationship is accessed. The lack of
      a cmr-field element in an ejb-relationship-role specifies
      that the relationship is unidirectional in navigability and
      the entity bean that participates in the relationship is
      "not aware" of the relationship.

      Example:

      <ejb-relation>
        <ejb-relation-name>Product-LineItem</ejb-relation-name>
        <ejb-relationship-role>
          <ejb-relationship-role-name>product-has-lineitems
          </ejb-relationship-role-name>
          <multiplicity>One</multiplicity>
          <relationship-role-source>
            <ejb-name>ProductEJB</ejb-name>
          </relationship-role-source>
        </ejb-relationship-role>
      </ejb-relation>

    </xsd:documentation>
  </xsd:annotation>

```

```
<xsd:sequence>
  <xsd:element name="description"
    type="j2ee:descriptionType"
    minOccurs="0"
    maxOccurs="unbounded"/>
  <xsd:element name="ejb-relationship-role-name"
    type="j2ee:string"
    minOccurs="0">

    <xsd:annotation>
      <xsd:documentation>

        The ejb-relationship-role-name element defines a
        name for a role that is unique within an
        ejb-relation. Different relationships can use the
        same name for a role.

      </xsd:documentation>
    </xsd:annotation>

  </xsd:element>
  <xsd:element name="multiplicity"
    type="j2ee:multiplicityType"/>
  <xsd:element name="cascade-delete"
    type="j2ee:emptyType"
    minOccurs="0">

    <xsd:annotation>
      <xsd:documentation>

        The cascade-delete element specifies that, within a
        particular relationship, the lifetime of one or more
        entity beans is dependent upon the lifetime of
        another entity bean. The cascade-delete element can
        only be specified for an ejb-relationship-role
        element contained in an ejb-relation element in
        which the other ejb-relationship-role
        element specifies a multiplicity of One.

      </xsd:documentation>
    </xsd:annotation>

  </xsd:element>
  <xsd:element name="relationship-role-source"
    type="j2ee:relationship-role-sourceType"/>
  <xsd:element name="cmr-field"
    type="j2ee:cmr-fieldType"
    minOccurs="0"/>
```

```

    </xsd:sequence>
    <xsd:attribute name="id" type="xsd:ID"/>
  </xsd:complexType>

<!-- ***** -->
<xsd:complexType name="enterprise-beansType">

  <xsd:annotation>
    <xsd:documentation>

      The enterprise-beansType declares one or more enterprise
      beans. Each bean can be a session, entity or message-driven
      bean.

    </xsd:documentation>
  </xsd:annotation>

  <xsd:choice maxOccurs="unbounded">
    <xsd:element name="session"
      type="j2ee:session-beanType">
      <xsd:unique name="session-ejb-local-ref-name-uniqueness">

        <xsd:annotation>
          <xsd:documentation>

            The ejb-ref-name element contains the name of
            an EJB reference. The EJB reference is an entry in
            the component's environment and is relative to the
            java:comp/env context. The name must be unique within
            the component.

            It is recommended that name be prefixed with "ejb/".

          </xsd:documentation>
        </xsd:annotation>

        <xsd:selector xpath="j2ee:ejb-local-ref"/>
        <xsd:field xpath="j2ee:ejb-ref-name"/>
      </xsd:unique>
      <xsd:unique name="session-ejb-ref-name-uniqueness">

        <xsd:annotation>
          <xsd:documentation>

            The ejb-ref-name element contains the name of an EJB
            reference. The EJB reference is an entry in the
            component's environment and is relative to the
            java:comp/env context. The name must be unique

```

within the component.

It is recommended that name is prefixed with "ejb/".

```

    </xsd:documentation>
  </xsd:annotation>

  <xsd:selector xpath="j2ee:ejb-ref"/>
  <xsd:field    xpath="j2ee:ejb-ref-name"/>
</xsd:unique>
<xsd:unique name="session-resource-env-ref-uniqueness">

  <xsd:annotation>
    <xsd:documentation>

      The resource-env-ref-name element specifies the name
      of a resource environment reference; its value is
      the environment entry name used in the component
      code. The name is a JNDI name relative to the
      java:comp/env context and must be unique within an
      component.

    </xsd:documentation>
  </xsd:annotation>

  <xsd:selector xpath="j2ee:resource-env-ref"/>
  <xsd:field    xpath="j2ee:resource-env-ref-name"/>
</xsd:unique>
<xsd:unique name="session-message-destination-ref-uniqueness">

  <xsd:annotation>
    <xsd:documentation>

      The message-destination-ref-name element specifies the name
      of a message destination reference; its value is
      the message destination reference name used in the component
      code. The name is a JNDI name relative to the
      java:comp/env context and must be unique within an
      component.

    </xsd:documentation>
  </xsd:annotation>

  <xsd:selector xpath="j2ee:message-destination-ref"/>
  <xsd:field    xpath="j2ee:message-destination-ref-name"/>
</xsd:unique>
<xsd:unique name="session-res-ref-name-uniqueness">

```

```
<xsd:annotation>
  <xsd:documentation>

    The res-ref-name element specifies the name of a
    resource manager connection factory reference. The name
    is a JNDI name relative to the java:comp/env context.
    The name must be unique within an component.

  </xsd:documentation>
</xsd:annotation>

<xsd:selector xpath="j2ee:resource-ref"/>
<xsd:field    xpath="j2ee:res-ref-name"/>
</xsd:unique>
<xsd:unique name="session-env-entry-name-uniqueness">

<xsd:annotation>
  <xsd:documentation>

    The env-entry-name element contains the name of a
    component's environment entry. The name is a JNDI
    name relative to the java:comp/env context. The
    name must be unique within an component.

  </xsd:documentation>
</xsd:annotation>

<xsd:selector xpath="j2ee:env-entry"/>
<xsd:field    xpath="j2ee:env-entry-name"/>
</xsd:unique>
</xsd:element>
<xsd:element name="entity"
             type="j2ee:entity-beanType">
  <xsd:unique name="entity-ejb-local-ref-name-uniqueness">

<xsd:annotation>
  <xsd:documentation>

    The ejb-ref-name element contains the name of
    an EJB reference. The EJB reference is an entry in
    the component's environment and is relative to the
    java:comp/env context. The name must be unique within
    the component.

    It is recommended that name be prefixed with "ejb/".

  </xsd:documentation>
</xsd:annotation>
```

```

    <xsd:selector xpath="j2ee:ejb-local-ref"/>
    <xsd:field xpath="j2ee:ejb-ref-name"/>
</xsd:unique>
<xsd:unique name="entity-ejb-ref-name-uniqueness">

    <xsd:annotation>
      <xsd:documentation>

        The ejb-ref-name element contains the name of an EJB
        reference. The EJB reference is an entry in the
        component's environment and is relative to the
        java:comp/env context. The name must be unique
        within the component.

        It is recommended that name is prefixed with "ejb/".

      </xsd:documentation>
    </xsd:annotation>

    <xsd:selector xpath="j2ee:ejb-ref"/>
    <xsd:field xpath="j2ee:ejb-ref-name"/>
</xsd:unique>
<xsd:unique name="entity-resource-env-ref-uniqueness">

    <xsd:annotation>
      <xsd:documentation>

        The resource-env-ref-name element specifies the name
        of a resource environment reference; its value is
        the environment entry name used in the component
        code. The name is a JNDI name relative to the
        java:comp/env context and must be unique within an
        component.

      </xsd:documentation>
    </xsd:annotation>

    <xsd:selector xpath="j2ee:resource-env-ref"/>
    <xsd:field xpath="j2ee:resource-env-ref-name"/>
</xsd:unique>
<xsd:unique name="entity-message-destination-ref-uniqueness">

    <xsd:annotation>
      <xsd:documentation>

        The message-destination-ref-name element specifies the name
        of a message destination reference; its value is
        the message destination reference name used in the component

```

```

        code. The name is a JNDI name relative to the
        java:comp/env context and must be unique within an
        component.

    </xsd:documentation>
</xsd:annotation>

<xsd:selector xpath="j2ee:message-destination-ref"/>
<xsd:field    xpath="j2ee:message-destination-ref-name"/>
</xsd:unique>
<xsd:unique name="entity-res-ref-name-uniqueness">

<xsd:annotation>
    <xsd:documentation>

        The res-ref-name element specifies the name of a
        resource manager connection factory reference. The name
        is a JNDI name relative to the java:comp/env context.
        The name must be unique within an component.

    </xsd:documentation>
</xsd:annotation>

<xsd:selector xpath="j2ee:resource-ref"/>
<xsd:field    xpath="j2ee:res-ref-name"/>
</xsd:unique>
<xsd:unique name="entity-env-entry-name-uniqueness">

<xsd:annotation>
    <xsd:documentation>

        The env-entry-name element contains the name of a
        component's environment entry. The name is a JNDI
        name relative to the java:comp/env context. The
        name must be unique within an component.

    </xsd:documentation>
</xsd:annotation>

<xsd:selector xpath="j2ee:env-entry"/>
<xsd:field    xpath="j2ee:env-entry-name"/>
</xsd:unique>
</xsd:element>
<xsd:element name="message-driven"
    type="j2ee:message-driven-beanType">
    <xsd:unique name="messaged-ejb-local-ref-name-uniqueness">

```

```

<xsd:annotation>
  <xsd:documentation>

    The ejb-ref-name element contains the name of
    an EJB reference. The EJB reference is an entry in
    the component's environment and is relative to the
    java:comp/env context. The name must be unique within
    the component.

    It is recommended that name be prefixed with "ejb/".

  </xsd:documentation>
</xsd:annotation>

<xsd:selector xpath="j2ee:ejb-local-ref"/>
<xsd:field    xpath="j2ee:ejb-ref-name"/>
</xsd:unique>
<xsd:unique name="messaging-ejb-ref-name-uniqueness">

<xsd:annotation>
  <xsd:documentation>

    The ejb-ref-name element contains the name of an EJB
    reference. The EJB reference is an entry in the
    component's environment and is relative to the
    java:comp/env context. The name must be unique
    within the component.

    It is recommended that name is prefixed with "ejb/".

  </xsd:documentation>
</xsd:annotation>

<xsd:selector xpath="j2ee:ejb-ref"/>
<xsd:field    xpath="j2ee:ejb-ref-name"/>
</xsd:unique>
<xsd:unique name="messaging-resource-env-ref-uniqueness">

<xsd:annotation>
  <xsd:documentation>

    The resource-env-ref-name element specifies the name
    of a resource environment reference; its value is
    the environment entry name used in the component
    code. The name is a JNDI name relative to the
    java:comp/env context and must be unique within an
    component.

```



```

    </xsd:documentation>
  </xsd:annotation>

  <xsd:selector xpath="j2ee:resource-env-ref"/>
  <xsd:field    xpath="j2ee:resource-env-ref-name"/>
</xsd:unique>
<xsd:unique name="messaged-message-destination-ref-uniqueness">

  <xsd:annotation>
    <xsd:documentation>

      The message-destination-ref-name element specifies the name
      of a message destination reference; its value is
      the message destination reference name used in the component
      code. The name is a JNDI name relative to the
      java:comp/env context and must be unique within an
      component.

    </xsd:documentation>
  </xsd:annotation>

  <xsd:selector xpath="j2ee:message-destination-ref"/>
  <xsd:field    xpath="j2ee:message-destination-ref-name"/>
</xsd:unique>
<xsd:unique name="messaged-res-ref-name-uniqueness">

  <xsd:annotation>
    <xsd:documentation>

      The res-ref-name element specifies the name of a
      resource manager connection factory reference. The name
      is a JNDI name relative to the java:comp/env context.
      The name must be unique within an component.

    </xsd:documentation>
  </xsd:annotation>

  <xsd:selector xpath="j2ee:resource-ref"/>
  <xsd:field    xpath="j2ee:res-ref-name"/>
</xsd:unique>
<xsd:unique name="messaged-env-entry-name-uniqueness">

  <xsd:annotation>
    <xsd:documentation>

      The env-entry-name element contains the name of a
      component's environment entry. The name is a JNDI
      name relative to the java:comp/env context. The

```

```

        name must be unique within an component.

    </xsd:documentation>
</xsd:annotation>

    <xsd:selector xpath="j2ee:env-entry"/>
    <xsd:field xpath="j2ee:env-entry-name"/>
</xsd:unique>
</xsd:element>
</xsd:choice>
<xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="entity-beanType">

    <xsd:annotation>
        <xsd:documentation>

            The entity-beanType declares an entity bean. The declaration
            consists of:

            - an optional description
            - an optional display name
            - an optional icon element that contains a small and a large
              icon file name
            - a unique name assigned to the enterprise bean
              in the deployment descriptor
            - the names of the entity bean's remote home
              and remote interfaces, if any
            - the names of the entity bean's local home and local
              interfaces, if any
            - the entity bean's implementation class
            - the entity bean's persistence management type
            - the entity bean's primary key class name
            - an indication of the entity bean's reentrancy
            - an optional specification of the
              entity bean's cmp-version
            - an optional specification of the entity bean's
              abstract schema name
            - an optional list of container-managed fields
            - an optional specification of the primary key
              field
            - an optional declaration of the bean's environment
              entries
            - an optional declaration of the bean's EJB
              references
            - an optional declaration of the bean's local

```

- EJB references*
- an optional declaration of the bean's web service references
- an optional declaration of the security role references
- an optional declaration of the security identity to be used for the execution of the bean's methods
- an optional declaration of the bean's resource manager connection factory references
- an optional declaration of the bean's resource environment references
- an optional declaration of the bean's message destination references
- an optional set of query declarations for finder and select methods for an entity bean with *cmp-version 2.x*.

The optional *abstract-schema-name* element must be specified for an entity bean with container-managed persistence and *cmp-version 2.x*.

The optional *primkey-field* may be present in the descriptor if the entity's *persistence-type* is *Container*.

The optional *cmp-version* element may be present in the descriptor if the entity's *persistence-type* is *Container*. If the *persistence-type* is *Container* and the *cmp-version* element is not specified, its value defaults to *2.x*.

The optional *home* and *remote* elements must be specified if the entity bean *cmp-version* is *1.x*.

The optional *home* and *remote* elements must be specified if the entity bean has a remote home and remote interface.

The optional *local-home* and *local* elements must be specified if the entity bean has a local home and local interface.

Either both the *local-home* and the *local* elements or both the *home* and the *remote* elements must be specified.

The optional query elements must be present if the *persistence-type* is *Container* and the *cmp-version* is *2.x* and query methods other than *findByPrimaryKey* have been defined for the entity bean.

The other elements that are optional are "optional" in the sense that they are omitted if the lists represented by them

are empty.

At least one cmp-field element must be present in the descriptor if the entity's persistence-type is Container and the cmp-version is 1.x, and none must not be present if the entity's persistence-type is Bean.

```

</xsd:documentation>
</xsd:annotation>

<xsd:sequence>
  <xsd:group ref="j2ee:descriptionGroup"/>
  <xsd:element name="ejb-name"
    type="j2ee:ejb-nameType"/>
  <xsd:element name="home"
    type="j2ee:homeType"
    minOccurs="0"/>
  <xsd:element name="remote"
    type="j2ee:remoteType"
    minOccurs="0"/>
  <xsd:element name="local-home"
    type="j2ee:local-homeType"
    minOccurs="0"/>
  <xsd:element name="local"
    type="j2ee:localType"
    minOccurs="0"/>
  <xsd:element name="ejb-class"
    type="j2ee:ejb-classType"/>
  <xsd:element name="persistence-type"
    type="j2ee:persistence-typeType"/>
  <xsd:element name="prim-key-class"
    type="j2ee:fully-qualified-classType">

    <xsd:annotation>
      <xsd:documentation>

        The prim-key-class element contains the
        fully-qualified name of an
        entity bean's primary key class.

        If the definition of the primary key class is
        deferred to deployment time, the prim-key-class
        element should specify java.lang.Object.

      </xsd:documentation>
    </xsd:annotation>

  </xsd:element>

```

```
<xsd:element name="reentrant"
             type="j2ee:true-falseType">

  <xsd:annotation>
    <xsd:documentation>

      The reentrant element specifies whether an entity
      bean is reentrant or not.

      The reentrant element must be one of the two
      following: true or false

    </xsd:documentation>
  </xsd:annotation>

</xsd:element>
<xsd:element name="cmp-version"
             type="j2ee:cmp-versionType"
             minOccurs="0"/>
<xsd:element name="abstract-schema-name"
             type="j2ee:java-identifierType"
             minOccurs="0">

  <xsd:annotation>
    <xsd:documentation>

      The abstract-schema-name element specifies the name
      of the abstract schema type of an entity bean with
      cmp-version 2.x. It is used in EJB QL queries.

      For example, the abstract-schema-name for an entity
      bean whose local interface is
      com.acme.commerce.Order might be Order.

    </xsd:documentation>
  </xsd:annotation>

</xsd:element>
<xsd:element name="cmp-field"
             type="j2ee:cmp-fieldType"
             minOccurs="0"
             maxOccurs="unbounded"/>
<xsd:element name="primkey-field"
             type="j2ee:string"
             minOccurs="0">

  <xsd:annotation>
    <xsd:documentation>
```

The `primkey-field` element is used to specify the name of the primary key field for an entity with container-managed persistence.

The `primkey-field` must be one of the fields declared in the `cmp-field` element, and the type of the field must be the same as the primary key type.

The `primkey-field` element is not used if the primary key maps to multiple container-managed fields (i.e. the key is a compound key). In this case, the fields of the primary key class must be public, and their names must correspond to the field names of the entity bean class that comprise the key.

```
</xsd:documentation>
</xsd:annotation>
```

```
</xsd:element>
<xsd:group ref="j2ee:jndiEnvironmentRefsGroup"/>
<xsd:element name="security-role-ref"
  type="j2ee:security-role-refType"
  minOccurs="0" maxOccurs="unbounded"/>
<xsd:element name="security-identity"
  type="j2ee:security-identityType"
  minOccurs="0"/>
<xsd:element name="query"
  type="j2ee:queryType"
  minOccurs="0" maxOccurs="unbounded"/>
</xsd:sequence>
<xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>
```

```
<!-- ***** -->
```

```
<xsd:complexType name="exclude-listType">
```

```
<xsd:annotation>
  <xsd:documentation>
```

The `exclude-listType` specifies one or more methods which the Assembler marks to be uncallable.

If the method permission relation contains methods that are in the exclude list, the Deployer should consider those methods to be uncallable.

```
</xsd:documentation>
</xsd:annotation>
```

```

<xsd:sequence>
  <xsd:element name="description"
    type="j2ee:descriptionType"
    minOccurs="0"
    maxOccurs="unbounded"/>
  <xsd:element name="method"
    type="j2ee:methodType"
    maxOccurs="unbounded"/>
</xsd:sequence>
<xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="message-driven-beanType">

  <xsd:annotation>
    <xsd:documentation>

      The message-driven element declares a message-driven
      bean. The declaration consists of:

        - an optional description
        - an optional display name
        - an optional icon element that contains a small and a large
          icon file name.
        - a name assigned to the enterprise bean in
          the deployment descriptor
        - the message-driven bean's implementation class
        - an optional declaration of the bean's messaging
          type
        - the message-driven bean's transaction management type
        - an optional declaration of the bean's
          message-destination-type
        - an optional declaration of the bean's
          message-destination-link
        - an optional declaration of the message-driven bean's
          activation configuration properties
        - an optional declaration of the bean's environment
          entries
        - an optional declaration of the bean's EJB references
        - an optional declaration of the bean's local EJB
          references
        - an optional declaration of the bean's web service
          references
        - an optional declaration of the security
          identity to be used for the execution of the bean's
          methods
        - an optional declaration of the bean's

```

```

        resource manager connection factory
        references
        - an optional declaration of the bean's resource
          environment references.
        - an optional declaration of the bean's message
          destination references

</xsd:documentation>
</xsd:annotation>

<xsd:sequence>
  <xsd:group ref="j2ee:descriptionGroup"/>
  <xsd:element name="ejb-name"
    type="j2ee:ejb-nameType"/>
  <xsd:element name="ejb-class"
    type="j2ee:ejb-classType"/>
  <xsd:element name="messaging-type"
    type="j2ee:fully-qualified-classType"
    minOccurs="0">

    <xsd:annotation>
      <xsd:documentation>

        The messaging-type element specifies the message
        listener interface of the message-driven bean. If
        the messaging-type element is not specified, it is
        assumed to be javax.jms.MessageListener.

      </xsd:documentation>
    </xsd:annotation>

  </xsd:element>
  <xsd:element name="transaction-type"
    type="j2ee:transaction-typeType"/>
  <xsd:element name="message-destination-type"
    type="j2ee:message-destination-typeType"
    minOccurs="0"/>
  <xsd:element name="message-destination-link"
    type="j2ee:message-destination-linkType"
    minOccurs="0"/>
  <xsd:element name="activation-config"
    type="j2ee:activation-configType"
    minOccurs="0"/>
  <xsd:group ref="j2ee:jndiEnvironmentRefsGroup"/>
  <xsd:element name="security-identity"
    type="j2ee:security-identityType"
    minOccurs="0"/>
</xsd:sequence>

```



```

    <xsd:attribute name="id" type="xsd:ID"/>
  </xsd:complexType>

<!-- ***** -->
<xsd:complexType name="method-intfType">

  <xsd:annotation>
    <xsd:documentation>

      The method-intf element allows a method element to
      differentiate between the methods with the same name and
      signature that are multiply defined across the home and
      component interfaces (e.g, in both an enterprise bean's
      remote and local interfaces or in both an enterprise bean's
      home and remote interfaces, etc.); the component and web
      service endpoint interfaces, and so on.

      The method-intf element must be one of the following:

        Home
        Remote
        LocalHome
        Local
        ServiceEndpoint

    </xsd:documentation>
  </xsd:annotation>

  <xsd:simpleContent>
    <xsd:restriction base="j2ee:string">
      <xsd:enumeration value="Home"/>
      <xsd:enumeration value="Remote"/>
      <xsd:enumeration value="LocalHome"/>
      <xsd:enumeration value="Local"/>
      <xsd:enumeration value="ServiceEndpoint"/>
    </xsd:restriction>
  </xsd:simpleContent>
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="method-nameType">

  <xsd:annotation>
    <xsd:documentation>

      The method-nameType contains a name of an enterprise
      bean method or the asterisk (*) character. The asterisk is
      used when the element denotes all the methods of an

```

```

        enterprise bean's client view interfaces.

    </xsd:documentation>
</xsd:annotation>

    <xsd:simpleContent>
        <xsd:restriction base="j2ee:string"/>
    </xsd:simpleContent>
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="method-paramsType">

    <xsd:annotation>
        <xsd:documentation>

            The method-paramsType defines a list of the
            fully-qualified Java type names of the method parameters.

        </xsd:documentation>
    </xsd:annotation>

    <xsd:sequence>
        <xsd:element name="method-param"
            type="j2ee:java-typeType"
            minOccurs="0"
            maxOccurs="unbounded">

            <xsd:annotation>
                <xsd:documentation>

                    The method-param element contains a primitive
                    or a fully-qualified Java type name of a method
                    parameter.

                </xsd:documentation>
            </xsd:annotation>

        </xsd:element>
    </xsd:sequence>
    <xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="method-permissionType">

    <xsd:annotation>
        <xsd:documentation>

```

The method-permissionType specifies that one or more security roles are allowed to invoke one or more enterprise bean methods. The method-permissionType consists of an optional description, a list of security role names or an indicator to state that the method is unchecked for authorization, and a list of method elements.

The security roles used in the method-permissionType must be defined in the security-role elements of the deployment descriptor, and the methods must be methods defined in the enterprise bean's home, component and/or web service endpoint interfaces.

```
</xsd:documentation>
</xsd:annotation>

<xsd:sequence>
  <xsd:element name="description"
    type="j2ee:descriptionType"
    minOccurs="0"
    maxOccurs="unbounded"/>
  <xsd:choice>
    <xsd:element name="role-name"
      type="j2ee:role-nameType"
      maxOccurs="unbounded"/>
    <xsd:element name="unchecked"
      type="j2ee:emptyType">

      <xsd:annotation>
        <xsd:documentation>

          The unchecked element specifies that a method is
            not checked for authorization by the container
            prior to invocation of the method.

        </xsd:documentation>
      </xsd:annotation>

    </xsd:element>
  </xsd:choice>
  <xsd:element name="method"
    type="j2ee:methodType"
    maxOccurs="unbounded"/>
</xsd:sequence>
<xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>
```

```
<!-- ***** -->
```

```
<xsd:complexType name="methodType">
```

```
<xsd:annotation>
```

```
<xsd:documentation>
```

The methodType is used to denote a method of an enterprise bean's home, component, and/or web service endpoint interface, or, in the case of a message-driven bean, the bean's message listener method, or a set of such methods. The ejb-name element must be the name of one of the enterprise beans declared in the deployment descriptor; the optional method-intf element allows to distinguish between a method with the same signature that is multiply defined across the home, component, and/or web service endpoint interfaces; the method-name element specifies the method name; and the optional method-params elements identify a single method among multiple methods with an overloaded method name.

There are three possible styles of using methodType element within a method element:

1.

```
<method>
  <ejb-name>EJBNAME</ejb-name>
  <method-name>*</method-name>
</method>
```

This style is used to refer to all the methods of the specified enterprise bean's home, component, and/or web service endpoint interfaces.

2.

```
<method>
  <ejb-name>EJBNAME</ejb-name>
  <method-name>METHOD</method-name>
</method>
```

This style is used to refer to the specified method of the specified enterprise bean. If there are multiple methods with the same overloaded name, the element of this style refers to all the methods with the overloaded name.

3.

```
<method>
  <ejb-name>EJBNAME</ejb-name>
```

```
<method-name>METHOD</method-name>
<method-params>
  <method-param>PARAM-1</method-param>
  <method-param>PARAM-2</method-param>
  ...
  <method-param>PARAM-n</method-param>
</method-params>
</method>
```

This style is used to refer to a single method within a set of methods with an overloaded name. PARAM-1 through PARAM-n are the fully-qualified Java types of the method's input parameters (if the method has no input arguments, the method-params element contains no method-param elements). Arrays are specified by the array element's type, followed by one or more pair of square brackets (e.g. int[][]). If there are multiple methods with the same overloaded name, this style refers to all of the overloaded methods.

Examples:

Style 1: The following method element refers to all the methods of the EmployeeService bean's home, component, and/or web service endpoint interfaces:

```
<method>
  <ejb-name>EmployeeService</ejb-name>
  <method-name>*</method-name>
</method>
```

Style 2: The following method element refers to all the create methods of the EmployeeService bean's home interface(s).

```
<method>
  <ejb-name>EmployeeService</ejb-name>
  <method-name>create</method-name>
</method>
```

Style 3: The following method element refers to the create(String firstName, String LastName) method of the EmployeeService bean's home interface(s).

```
<method>
  <ejb-name>EmployeeService</ejb-name>
  <method-name>create</method-name>
  <method-params>
```

```

        <method-param>java.lang.String</method-param>
        <method-param>java.lang.String</method-param>
    </method-params>
</method>

```

The following example illustrates a Style 3 element with more complex parameter types. The method `foobar(char s, int i, int[] iar, mypackage.MyClass mycl, mypackage.MyClass[][] myclaar)` would be specified as:

```

<method>
  <ejb-name>EmployeeService</ejb-name>
  <method-name>foobar</method-name>
  <method-params>
    <method-param>char</method-param>
    <method-param>int</method-param>
    <method-param>int[]</method-param>
    <method-param>mypackage.MyClass</method-param>
    <method-param>mypackage.MyClass[][]</method-param>
  </method-params>
</method>

```

The optional `method-intf` element can be used when it becomes necessary to differentiate between a method that is multiply defined across the enterprise bean's home, component, and/or web service endpoint interfaces with the same name and signature.

For example, the method element

```

<method>
  <ejb-name>EmployeeService</ejb-name>
  <method-intf>Remote</method-intf>
  <method-name>create</method-name>
  <method-params>
    <method-param>java.lang.String</method-param>
    <method-param>java.lang.String</method-param>
  </method-params>
</method>

```

can be used to differentiate the `create(String, String)` method defined in the remote interface from the `create(String, String)` method defined in the remote home interface, which would be defined as

```

<method>
  <ejb-name>EmployeeService</ejb-name>
  <method-intf>Home</method-intf>

```

```

    <method-name>create</method-name>
    <method-params>
      <method-param>java.lang.String</method-param>
      <method-param>java.lang.String</method-param>
    </method-params>
  </method>

```

and the create method that is defined in the local home interface which would be defined as

```

<method>
  <ejb-name>EmployeeService</ejb-name>
  <method-intf>LocalHome</method-intf>
  <method-name>create</method-name>
  <method-params>
    <method-param>java.lang.String</method-param>
    <method-param>java.lang.String</method-param>
  </method-params>
</method>

```

The method-intf element can be used with all three Styles of the method element usage. For example, the following method element example could be used to refer to all the methods of the EmployeeService bean's remote home interface.

```

<method>
  <ejb-name>EmployeeService</ejb-name>
  <method-intf>Home</method-intf>
  <method-name>*</method-name>
</method>

```

```

</xsd:documentation>
</xsd:annotation>

```

```

<xsd:sequence>
  <xsd:element name="description"
    type="j2ee:descriptionType"
    minOccurs="0" maxOccurs="unbounded"/>
  <xsd:element name="ejb-name"
    type="j2ee:ejb-nameType"/>
  <xsd:element name="method-intf"
    type="j2ee:method-intfType"
    minOccurs="0">
  </xsd:element>
  <xsd:element name="method-name"
    type="j2ee:method-nameType"/>
  <xsd:element name="method-params"
    type="j2ee:method-paramsType"

```

```

        minOccurs="0"/>
    </xsd:sequence>
    <xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="multiplicityType">

    <xsd:annotation>
        <xsd:documentation>

            The multiplicityType describes the multiplicity of the
            role that participates in a relation.

            The value must be one of the two following:

                One
                Many

        </xsd:documentation>
    </xsd:annotation>

    <xsd:simpleContent>
        <xsd:restriction base="j2ee:string">
            <xsd:enumeration value="One"/>
            <xsd:enumeration value="Many"/>
        </xsd:restriction>
    </xsd:simpleContent>
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="persistence-typeType">

    <xsd:annotation>
        <xsd:documentation>

            The persistence-typeType specifies an entity bean's persistence
            management type.

            The persistence-type element must be one of the two following:

                Bean
                Container

        </xsd:documentation>
    </xsd:annotation>

    <xsd:simpleContent>

```



```

        <xsd:restriction base="j2ee:string">
            <xsd:enumeration value="Bean"/>
            <xsd:enumeration value="Container"/>
        </xsd:restriction>
    </xsd:simpleContent>
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="query-methodType">

    <xsd:annotation>
        <xsd:documentation>

            The query-method specifies the method for a finder or select
            query.

            The method-name element specifies the name of a finder or select
            method in the entity bean's implementation class.

            Each method-param must be defined for a query-method using the
            method-params element.

            It is used by the query-method element.

            Example:

            <query>
                <description>Method finds large orders</description>
                <query-method>
                    <method-name>findLargeOrders</method-name>
                    <method-params></method-params>
                </query-method>
            </ejb-ql>
                SELECT OBJECT(o) FROM Order o
                WHERE o.amount > 1000
            </ejb-ql>
        </query>

        </xsd:documentation>
    </xsd:annotation>

    <xsd:sequence>
        <xsd:element name="method-name"
            type="j2ee:method-nameType"/>
        <xsd:element name="method-params"
            type="j2ee:method-paramsType"/>
    </xsd:sequence>

```

```

    <xsd:attribute name="id" type="xsd:ID"/>
  </xsd:complexType>

<!-- ***** -->
<xsd:complexType name="queryType">

  <xsd:annotation>
    <xsd:documentation>

      The queryType defines a finder or select
      query. It contains
      - an optional description of the query
      - the specification of the finder or select
        method it is used by
        - an optional specification of the result type
          mapping, if the query is for a select method
          and entity objects are returned.
        - the EJB QL query string that defines the query.

      Queries that are expressible in EJB QL must use the ejb-ql
      element to specify the query. If a query is not expressible
      in EJB QL, the description element should be used to
      describe the semantics of the query and the ejb-ql element
      should be empty.

      The result-type-mapping is an optional element. It can only
      be present if the query-method specifies a select method
      that returns entity objects. The default value for the
      result-type-mapping element is "Local".

    </xsd:documentation>
  </xsd:annotation>

  <xsd:sequence>
    <xsd:element name="description"
      type="j2ee:descriptionType" minOccurs="0"/>
    <xsd:element name="query-method"
      type="j2ee:query-methodType"/>
    <xsd:element name="result-type-mapping"
      type="j2ee:result-type-mappingType"
      minOccurs="0"/>
    <xsd:element name="ejb-ql"
      type="j2ee:xsdStringType"/>
  </xsd:sequence>
  <xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>

```

```

<!-- ***** -->
<xsd:complexType name="relationship-role-sourceType">

  <xsd:annotation>
    <xsd:documentation>

      The relationship-role-sourceType designates the source of a
      role that participates in a relationship. A
      relationship-role-sourceType is used by
      relationship-role-source elements to uniquely identify an
      entity bean.

    </xsd:documentation>
  </xsd:annotation>

  <xsd:sequence>
    <xsd:element name="description"
      type="j2ee:descriptionType"
      minOccurs="0"
      maxOccurs="unbounded"/>
    <xsd:element name="ejb-name"
      type="j2ee:ejb-nameType"/>
  </xsd:sequence>
  <xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="relationshipsType">

  <xsd:annotation>
    <xsd:documentation>

      The relationshipsType describes the relationships in
      which entity beans with container-managed persistence
      participate. The relationshipsType contains an optional
      description; and a list of ejb-relation elements, which
      specify the container managed relationships.

    </xsd:documentation>
  </xsd:annotation>

  <xsd:sequence>
    <xsd:element name="description"
      type="j2ee:descriptionType"
      minOccurs="0"
      maxOccurs="unbounded"/>
    <xsd:element name="ejb-relation"
      type="j2ee:ejb-relationType"

```

```

        maxOccurs="unbounded">
<xsd:unique name="role-name-uniqueness">

    <xsd:annotation>
        <xsd:documentation>

            The ejb-relationship-role-name contains the name of a
            relationship role. The name must be unique within
            a relationship, but can be reused in different
            relationships.

        </xsd:documentation>
    </xsd:annotation>

    <xsd:selector
        xpath="//j2ee:ejb-relationship-role-name"/>
    <xsd:field
        xpath="."/>
    </xsd:unique>
</xsd:element>
</xsd:sequence>
<xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="result-type-mappingType">

    <xsd:annotation>
        <xsd:documentation>

            The result-type-mappingType is used in the query element to
            specify whether an abstract schema type returned by a query
            for a select method is to be mapped to an EJBLocalObject or
            EJBObject type.

            The value must be one of the following:

                Local
                Remote

            </xsd:documentation>
        </xsd:annotation>

        <xsd:simpleContent>
            <xsd:restriction base="j2ee:string">
                <xsd:enumeration value="Local"/>
                <xsd:enumeration value="Remote"/>
            </xsd:restriction>

```

```

    </xsd:simpleContent>
  </xsd:complexType>

<!-- ***** -->
<xsd:complexType name="security-identityType">

  <xsd:annotation>
    <xsd:documentation>

      The security-identityType specifies whether the caller's
      security identity is to be used for the execution of the
      methods of the enterprise bean or whether a specific run-as
      identity is to be used. It contains an optional description
      and a specification of the security identity to be used.

    </xsd:documentation>
  </xsd:annotation>

  <xsd:sequence>
    <xsd:element name="description"
      type="j2ee:descriptionType"
      minOccurs="0"
      maxOccurs="unbounded"/>
    <xsd:choice>
      <xsd:element name="use-caller-identity"
        type="j2ee:emptyType">

        <xsd:annotation>
          <xsd:documentation>

            The use-caller-identity element specifies that
            the caller's security identity be used as the
            security identity for the execution of the
            enterprise bean's methods.

          </xsd:documentation>
        </xsd:annotation>

      </xsd:element>
      <xsd:element name="run-as"
        type="j2ee:run-asType"/>
    </xsd:choice>
  </xsd:sequence>
  <xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="session-beanType">

```

<xsd:annotation>
<xsd:documentation>

The session-beanType declares an session bean. The declaration consists of:

- an optional description
- an optional display name
- an optional icon element that contains a small and a large icon file name
- a name assigned to the enterprise bean in the deployment description
- the names of the session bean's remote home and remote interfaces, if any
- the names of the session bean's local home and local interfaces, if any
- the name of the session bean's web service endpoint interface, if any
- the session bean's implementation class
- the session bean's state management type
- the session bean's transaction management type
- an optional declaration of the bean's environment entries
- an optional declaration of the bean's EJB references
- an optional declaration of the bean's local EJB references
- an optional declaration of the bean's web service references
- an optional declaration of the security role references
- an optional declaration of the security identity to be used for the execution of the bean's methods
- an optional declaration of the bean's resource manager connection factory references
- an optional declaration of the bean's resource environment references.
- an optional declaration of the bean's message destination references

The elements that are optional are "optional" in the sense that they are omitted when if lists represented by them are empty.

Either both the local-home and the local elements or both the home and the remote elements must be specified for the session bean.

The service-endpoint element may only be specified if the

```
        bean is a stateless session bean.

    </xsd:documentation>
</xsd:annotation>

<xsd:sequence>
  <xsd:group ref="j2ee:descriptionGroup"/>
  <xsd:element name="ejb-name"
    type="j2ee:ejb-nameType"/>
  <xsd:element name="home"
    type="j2ee:homeType"
    minOccurs="0"/>
  <xsd:element name="remote"
    type="j2ee:remoteType"
    minOccurs="0"/>
  <xsd:element name="local-home"
    type="j2ee:local-homeType"
    minOccurs="0"/>
  <xsd:element name="local"
    type="j2ee:localType"
    minOccurs="0"/>
  <xsd:element name="service-endpoint"
    type="j2ee:fully-qualified-classType"
    minOccurs="0">

    <xsd:annotation>
      <xsd:documentation>

        The service-endpoint element contains the
        fully-qualified name of the enterprise bean's web
        service endpoint interface. The service-endpoint
        element may only be specified for a stateless
        session bean. The specified interface must be a
        valid JAX-RPC service endpoint interface.

      </xsd:documentation>
    </xsd:annotation>

  </xsd:element>
  <xsd:element name="ejb-class"
    type="j2ee:ejb-classType"/>
  <xsd:element name="session-type"
    type="j2ee:session-typeType"/>
  <xsd:element name="transaction-type"
    type="j2ee:transaction-typeType"/>
  <xsd:group ref="j2ee:jndiEnvironmentRefsGroup"/>
  <xsd:element name="security-role-ref"
    type="j2ee:security-role-refType"
```

```

        minOccurs="0"
        maxOccurs="unbounded"/>
<xsd:element name="security-identity"
    type="j2ee:security-identityType"
    minOccurs="0"/>
</xsd:sequence>
<xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="session-typeType">

    <xsd:annotation>
        <xsd:documentation>

            The session-typeType describes whether the session bean is a
            stateful session or stateless session. It is used by
            session-type elements.

            The value must be one of the two following:

                Stateful
                Stateless

        </xsd:documentation>
    </xsd:annotation>

    <xsd:simpleContent>
        <xsd:restriction base="j2ee:string">
            <xsd:enumeration value="Stateful"/>
            <xsd:enumeration value="Stateless"/>
        </xsd:restriction>
    </xsd:simpleContent>
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="trans-attributeType">

    <xsd:annotation>
        <xsd:documentation>

            The trans-attributeType specifies how the container must
            manage the transaction boundaries when delegating a method
            invocation to an enterprise bean's business method.

            The value must be one of the following:

                NotSupported


```



```

        Supports
        Required
        RequiresNew
        Mandatory
        Never

    </xsd:documentation>
</xsd:annotation>

<xsd:simpleContent>
    <xsd:restriction base="j2ee:string">
        <xsd:enumeration value="NotSupported"/>
        <xsd:enumeration value="Supports"/>
        <xsd:enumeration value="Required"/>
        <xsd:enumeration value="RequiresNew"/>
        <xsd:enumeration value="Mandatory"/>
        <xsd:enumeration value="Never"/>
    </xsd:restriction>
</xsd:simpleContent>
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="transaction-typeType">

    <xsd:annotation>
        <xsd:documentation>

            The transaction-typeType specifies an enterprise bean's
            transaction management type.

            The transaction-type must be one of the two following:

                Bean
                Container

        </xsd:documentation>
    </xsd:annotation>

    <xsd:simpleContent>
        <xsd:restriction base="j2ee:string">
            <xsd:enumeration value="Bean"/>
            <xsd:enumeration value="Container"/>
        </xsd:restriction>
    </xsd:simpleContent>
</xsd:complexType>

</xsd:schema>

```


Ejb-jar File

The ejb-jar file is the standard format for the packaging of enterprise beans. The ejb-jar file format is used to package un-assembled enterprise beans (the Bean Provider's output), and to package assembled applications (the Application Assembler's output).

24.1 Overview

The ejb-jar file format is the contract between the Bean Provider and the Application Assembler, and between the Application Assembler and the Deployer.

An ejb-jar file produced by the Bean Provider contains one or more enterprise beans that typically do not contain application assembly instructions. The ejb-jar file produced by an Application Assembler (which can be the same person or organization as the Bean Provider) contains one or more enterprise beans, plus application assembly information describing how the enterprise beans are combined into a single application deployment unit.

24.2 Deployment Descriptor

The ejb-jar file must contain the deployment descriptor in the format defined in Chapter 23. The deployment descriptor must be stored with the name `META-INF/ejb-jar.xml` in the ejb-jar file.

24.3 Ejb-jar File Requirements

The ejb-jar file must contain, either by inclusion or by reference, the class files of each enterprise bean as follows:

- The enterprise bean class.
- The enterprise bean home, component, and/or web service endpoint interfaces.
- The primary key class if the bean is an entity bean.

We say that a jar file contains a second file “by reference” if the second file is named in the Class-Path attribute in the Manifest file of the referencing jar file or is contained (either by inclusion or by reference) in another jar file that is named in the Class-Path attribute in the Manifest file of the referencing jar file.

The ejb-jar file must also contain, either by inclusion or by reference, the class files for all the classes and interfaces that each enterprise bean class and the home, component, and/or web service endpoint interfaces depend on, except J2EE and J2SE classes. This includes their superclasses and superinterfaces, dependent classes, and the classes and interfaces used as method parameters, results, and exceptions.

The Application Assembler must not package the stubs of the EJBHome and EJBObject interfaces in the ejb-jar file. This includes the stubs for the enterprise beans whose implementations are provided in the ejb-jar file as well as the referenced enterprise beans. Generating the stubs is the responsibility of the container. The stubs are typically generated by the Container Provider’s deployment tools for each class that extends the EJBHome or EJBObject interfaces, or they may be generated by the container at runtime.

24.4 The Client View and the ejb-client JAR File

The client view of an enterprise bean is defined in Section 4.2.1, and is comprised of the home and component interfaces of the referenced enterprise bean and other classes that these interfaces depend on, such as their superclasses and superinterfaces, the classes and interfaces used as method parameters, results, and exceptions. The serializable application value classes, including the classes which may be used as members of a collection in a remote method call to an enterprise bean, are part of the client view. An example of an application value class might be an `Address` class used as a parameter in a method call.

The `ejb-jar` file producer can create an `ejb-client` JAR file for the `ejb-jar` file. The `ejb-client` JAR file contains all the class files that a client program needs to use the client view of the enterprise beans that are contained in the `ejb-jar` file. If this option is used, it is the responsibility of the Application Assembler to include all the classes necessary to comprise the client view of an enterprise bean in the `ejb-client` JAR file.

The `ejb-client` JAR file is specified in the deployment descriptor of the `ejb-jar` file using the optional `ejb-client-jar` element. The value of the `ejb-client-jar` element is the path name specifying the location of the `ejb-client` JAR file in the containing J2EE Enterprise Application Archive (`.ear`) file. The path name is relative to the location of the referencing `ejb-jar` file.

The EJB specification does not specify whether an `ejb-jar` file should include by copy or by reference the classes that are in an `ejb-client` JAR file, but they must be included either one way or the other. If the by-copy approach is used, the producer simply includes all the class files in the `ejb-client` JAR file also in the `ejb-jar` file. If the by-reference approach is used, the `ejb-jar` file producer does not duplicate the content of the `ejb-client` JAR file in the `ejb-jar` file, but instead uses a Manifest Class-Path entry in the `ejb-jar` file to specify that the `ejb-jar` file depends on the `ejb-client` JAR at runtime. The use of the Class-Path entries in the JAR files is explained in the Java 2 Platform, Enterprise Edition specification [9].

24.5 Requirements for Clients

The Application Assembler must construct the application to insure that the client view classes are available to the client at runtime. The client of an enterprise bean may be another enterprise bean packaged in the same `ejb-jar` or different `ejb-jar` file, or the client may be another J2EE component, such as a web component.

When clients packaged in jar files refer to enterprise beans, the jar file that contains the client, e.g. an `ejb-jar` file, should contain, either by inclusion or by reference, all the client view classes of the referenced beans. The client view classes may have been packaged in an `ejb-client` JAR file. In other words, the jar file that contains the client should contain one of the following:

- a reference to the `ejb-client` JAR file
- a reference to the `ejb-jar` file that contains the client view classes
- a copy of the client view classes

The client may also require the use of system value classes (e.g., the serializable value classes implementing the `javax.ejb.Handle`, `javax.ejb.HomeHandle`, `javax.ejb.EJBMetaData`, `java.util.Enumeration`, `java.util.Collection`, and `java.util.Iterator` interfaces), although these are not packaged with the application. It is the responsibility of the provider of the container hosting the referenced beans to provide the system value classes and make them available for use when the client is deployed. See Section 19.5.5, “System Value Classes”.

24.6 Example

In this example, the Bean Provider has chosen to package the enterprise bean client view classes in a separate jar file and to reference that jar file from the other jar files that need those classes. Those classes are needed both by `ejb2.jar`, packaged in the same application as `ejb1.jar`, and by `ejb3.jar`, packaged in a different application. Those classes are also needed by `ejb1.jar` itself because they define the remote interface of the enterprise beans in `ejb1.jar`, and the Bean Provider has chosen the *by reference* approach to making these classes available.

The deployment descriptor for `ejb1.jar` names the client view jar file in the `ejb-client-jar` element. Because `ejb2.jar` requires these client view classes, it includes a Class-Path reference to `ejb1_client.jar`.

The Class-Path mechanism must be used by components in `app2.ear` to reference the client view jar file that corresponds to the enterprise beans packaged in `ejb1.jar` of `app1.ear`. Those enterprise beans are referenced by enterprise beans in `ejb3.jar`. Note that the client view jar file must be included directly in the `app2.ear` file.

```
app1.ear:
    META-INF/application.xml
    ejb1.jar      Class-Path: ejb1_client.jar
                  deployment descriptor contains:
                  <ejb-client-jar>ejb1_client.jar</ejb-client-jar>
    ejb1_client.jar
    ejb2.jar      Class-Path: ejb1_client.jar

app2.ear:
    META-INF/application.xml
    ejb1_client.jar
    ejb3.jar      Class-Path: ejb1_client.jar
```

24.7 Removed in EJB 1.1

This section describes the deployment information that was defined in EJB 1.0, and was removed in EJB 1.1.

24.7.1 **ejb-jar Manifest**

The JAR Manifest file is no longer used by the EJB architecture to identify the enterprise beans contained in an `ejb-jar` file.

EJB 1.0 used the Manifest file to identify the individual enterprise beans that were included in the `ejb-jar` file. As of EJB 1.1, the enterprise beans are identified in the deployment descriptor, so the information in the Manifest is no longer needed.

24.7.2 Serialized Deployment Descriptor JavaBeans™ Components

The mechanism of using serialized JavaBeans components as deployment descriptors has been replaced by the XML-based deployment descriptor.

Runtime Environment

This chapter defines the application programming interfaces (APIs) that a compliant EJB 2.1 container must make available to the enterprise bean instances at runtime. These APIs can be used by portable enterprise beans because the APIs are guaranteed to be available in all EJB 2.1 containers.

The chapter also defines the restrictions that the EJB 2.1 Container Provider can impose on the functionality that it provides to the enterprise beans. These restrictions are necessary to enforce security and to allow the container to properly manage the runtime environment.

25.1 Bean Provider's Responsibilities

This section describes the view and responsibilities of the Bean Provider.

25.1.1 APIs Provided by Container

The EJB Provider can rely on the EJB 2.1 Container Provider to provide the following APIs:

- Java 2 Platform, Standard Edition, v1.4 (J2SE) APIs, which include the following APIs:
 - JDBC
 - RMI-IIOP

- JNDI
- JAXP 1.1
- Java IDL
- EJB 2.1
- JTA 1.0.1 (the `UserTransaction` interface only)
- JMS 1.1
- JavaMail 1.3 (for sending mail only)
- JAF 1.0^[57]
- JAXP 1.2
- JAXR 1.0
- SAAJ 1.2
- JAX-RPC 1.1
- Connector 1.5
- Web Services for J2EE 1.1

25.1.2 Programming Restrictions

This section describes the programming restrictions that a Bean Provider must follow to ensure that the enterprise bean is *portable* and can be deployed in any compliant EJB 2.1 container. The restrictions apply to the implementation of the business methods. Section 25.2, which describes the container's view of these restrictions, defines the programming environment that all EJB containers must provide.

- An enterprise bean must not use read/write static fields. Using read-only static fields is allowed. Therefore, it is recommended that all static fields in the enterprise bean class be declared as `final`.

This rule is required to ensure consistent runtime semantics because while some EJB containers may use a single JVM to execute all enterprise bean's instances, others may distribute the instances across multiple JVMs.

- An enterprise bean must not use thread synchronization primitives to synchronize execution of multiple instances.

This is for the same reason as above. Synchronization would not work if the EJB container distributed enterprise bean's instances across multiple JVMs.

[57] See [9] for restrictions.

- An enterprise bean must not use the AWT functionality to attempt to output information to a display, or to input information from a keyboard.

Most servers do not allow direct interaction between an application program and a keyboard/display attached to the server system.

- An enterprise bean must not use the `java.io` package to attempt to access files and directories in the file system.

The file system APIs are not well-suited for business components to access data. Business components should use a resource manager API, such as JDBC, to store data.

- An enterprise bean must not attempt to listen on a socket, accept connections on a socket, or use a socket for multicast.

The EJB architecture allows an enterprise bean instance to be a network socket client, but it does not allow it to be a network server. Allowing the instance to become a network server would conflict with the basic function of the enterprise bean-- to serve the EJB clients.

- The enterprise bean must not attempt to query a class to obtain information about the declared members that are not otherwise accessible to the enterprise bean because of the security rules of the Java language. The enterprise bean must not attempt to use the Reflection API to access information that the security rules of the Java programming language make unavailable.

Allowing the enterprise bean to access information about other classes and to access the classes in a manner that is normally disallowed by the Java programming language could compromise security.

- The enterprise bean must not attempt to create a class loader; obtain the current class loader; set the context class loader; set security manager; create a new security manager; stop the JVM; or change the input, output, and error streams.

These functions are reserved for the EJB container. Allowing the enterprise bean to use these functions could compromise security and decrease the container's ability to properly manage the runtime environment.

- The enterprise bean must not attempt to set the socket factory used by `ServerSocket`, `Socket`, or the stream handler factory used by `URL`.

These networking functions are reserved for the EJB container. Allowing the enterprise bean to use these functions could compromise security and decrease the container's ability to properly manage the runtime environment.

- The enterprise bean must not attempt to manage threads. The enterprise bean must not attempt to start, stop, suspend, or resume a thread, or to change a thread's priority or name. The enterprise bean must not attempt to manage thread groups.

These functions are reserved for the EJB container. Allowing the enterprise bean to manage threads would decrease the container's ability to properly manage the runtime environment.

- The enterprise bean must not attempt to directly read or write a file descriptor.

Allowing the enterprise bean to read and write file descriptors directly could compromise security.

- The enterprise bean must not attempt to obtain the security policy information for a particular code source.

Allowing the enterprise bean to access the security policy information would create a security hole.

- The enterprise bean must not attempt to load a native library.

This function is reserved for the EJB container. Allowing the enterprise bean to load native code would create a security hole.

- The enterprise bean must not attempt to gain access to packages and classes that the usual rules of the Java programming language make unavailable to the enterprise bean.

This function is reserved for the EJB container. Allowing the enterprise bean to perform this function would create a security hole.

- The enterprise bean must not attempt to define a class in a package.

This function is reserved for the EJB container. Allowing the enterprise bean to perform this function would create a security hole.

- The enterprise bean must not attempt to access or modify the security configuration objects (Policy, Security, Provider, Signer, and Identity).

These functions are reserved for the EJB container. Allowing the enterprise bean to use these functions could compromise security.

- The enterprise bean must not attempt to use the subclass and object substitution features of the Java Serialization Protocol.

Allowing the enterprise bean to use these functions could compromise security.

- The enterprise bean must not attempt to pass this as an argument or method result. The enterprise bean must pass the result of `SessionContext.getEJBObject`, `SessionContext.getEJBLocalObject`, `EntityContext.getEJBObject`, or `EntityContext.getEJBLocalObject` instead.

To guarantee portability of the enterprise bean's implementation across all compliant EJB 2.1 containers, the Bean Provider should test the enterprise bean using a container with the security settings defined in Table 23. The tables define the minimal functionality that a compliant EJB container must provide to the enterprise bean instances at runtime.

25.2 Container Provider's Responsibility

This section defines the container's responsibilities for providing the runtime environment to the enterprise bean instances. The requirements described here are considered to be the minimal requirements; a container may choose to provide additional functionality that is not required by the EJB specification.

An EJB 2.1 container must make the following APIs available to the enterprise bean instances at runtime:

- Java 2 Platform, Standard Edition v1.4 (J2SE) APIs, which include the following APIs:
 - JDBC
 - RMI-IIOP
 - JNDI
 - JAXP 1.1
 - Java IDL
- EJB 2.1 APIs
- JTA 1.0.1, the `UserTransaction` interface only
- JMS 1.1
- JavaMail 1.3, sending mail only
- JAF 1.0^[58]
- JAXP 1.2
- JAXR 1.0
- JAX-RPC 1.1
- SAAJ 1.2
- Connector 1.5
- Web Services for J2EE 1.1

The following subsections describes the requirements in more detail.

25.2.1 Java 2 APIs Requirements

The container must provide the full set of Java 2 Platform, Standard Edition, v1.4 (J2SE) APIs. The container is not allowed to subset the Java 2 platform APIs.

[58] See [9] for restrictions.

The EJB container is allowed to make certain Java 2 platform functionality unavailable to the enterprise bean instances by using the Java 2 platform security policy mechanism. The primary reason for the container to make certain functions unavailable to enterprise bean instances is to protect the security and integrity of the EJB container environment, and to prevent the enterprise bean instances from interfering with the container's functions.

The following table defines the Java 2 platform security permissions that the EJB container must be able to grant to the enterprise bean instances at runtime. The term “grant” means that the container must be able to grant the permission, the term “deny” means that the container should deny the permission.

Table 23

Java 2 Platform Security Policy for a Standard EJB Container

Permission name	EJB Container policy
java.security.AllPermission	deny
java.awt.AWTPermission	deny
java.io.FilePermission	deny
java.net.NetPermission	deny
java.util.PropertyPermission	grant “read”, “*” deny all other
java.lang.reflect.ReflectPermission	deny
java.lang.RuntimePermission	grant “queuePrintJob”, deny all other
java.lang.SecurityPermission	deny
java.io.SerializablePermission	deny
java.net.SocketPermission	grant “connect”, “*” [Note A], deny all other

Notes:

[A] This permission is necessary, for example, to allow enterprise beans to use the client functionality of the Java IDL and RMI-IIOP packages that are part of the Java 2 platform.

Some containers may allow the Deployer to grant more, or fewer, permissions to the enterprise bean instances than specified in Table 23. Support for this is not required by the EJB specification. Enterprise beans that rely on more or fewer permissions will not be portable across all EJB containers.

25.2.2 EJB 2.1 Requirements

The container must implement the EJB 2.1 interfaces as defined in this specification.

25.2.3 JNDI Requirements

At the minimum, the EJB container must provide a JNDI API name space to the enterprise bean instances. The EJB container must make the name space available to an instance when the instance invokes the `javax.naming.InitialContext` default (no-arg) constructor.

The EJB container must make available at least the following objects in the name space:

- The home interfaces of other enterprise beans.
- The resource factories used by the enterprise beans.
- The web service interfaces used by the enterprise beans.

The EJB specification does not require that all the enterprise beans deployed in a container be presented with the same JNDI API name space. However, all the instances of the same enterprise bean must be presented with the same JNDI API name space.

25.2.4 JTA 1.0.1 Requirements

The EJB container must include the JTA 1.0.1 extension, and it must provide the `javax.transaction.UserTransaction` interface to enterprise beans with bean-managed transaction demarcation through the `javax.ejb.EJBContext` interface, and also in JNDI under the name `java:comp/UserTransaction`, in the cases required by the EJB specification.

The other JTA interfaces are low-level transaction manager and resource manager integration interfaces, and are not intended for direct use by enterprise beans.

25.2.5 JDBC™ 3.0 Extension Requirements

The EJB container must include the JDBC 3.0 extension and provide its functionality to the enterprise bean instances, with the exception of the low-level XA and connection pooling interfaces. These low-level interfaces are intended for integration of a JDBC driver with an application server, not for direct use by enterprise beans.

25.2.6 JMS 1.1 Requirements

The EJB container must include the JMS 1.1 extension and provide its functionality to the enterprise bean instances, with the exception of the low-level interfaces that are intended for integration of a JMS provider with an application server, not for direct use by enterprise beans. These interfaces include: `javax.jms.ServerSession`, `javax.jms.ServerSessionPool`, `javax.jms.ConnectionConsumer`, and all the `javax.jms` XA interfaces.

In addition, the following methods are for use by the container only. Enterprise beans must not call these methods: `javax.jms.Session.setMessageListener`, `javax.jms.Session.getMessageListener`, `javax.jms.Session.run`, `javax.jms.QueueConnection.createConnectionConsumer`, `javax.jms.TopicConnection.createConnectionConsumer`, `javax.jms.TopicConnection.createDurableConnectionConsumer`, `javax.jms.Connection.createConnectionConsumer`, `javax.jms.Connection.createDurableConnectionConsumer`.

The following methods must not be called by enterprise beans because they may interfere with the connection management by the container: `javax.jms.Connection.setExceptionListener`, `javax.jms.Connection.stop`, `javax.jms.Connection.setClientID`.

Enterprise beans must not call the `javax.jms.MessageConsumer.setMessageListener` or `javax.jms.MessageConsumer.getMessageListener` method.

This specification recommends, but does not require, that the container throw the `javax.jms.JMSEException` if enterprise beans call any of the methods listed in this section.

25.2.7 Argument Passing Semantics

An enterprise bean's remote home and remote interfaces are *remote interfaces* for Java RMI. The container must ensure the semantics for passing arguments conforms to Java RMI-IIOP. Non-remote objects must be passed by value.

Specifically, the EJB container is not allowed to pass non-remote objects by reference on inter-EJB invocations when the calling and called enterprise beans are collocated in the same JVM. Doing so could result in the multiple beans sharing the state of a Java object, which would break the enterprise bean's semantics. Any local optimizations of remote interface calls must ensure the semantics for passing arguments conforms to Java RMI-IIOP.

An enterprise bean's local home and local interfaces are *local Java interfaces*. The caller and callee enterprise beans that make use of these local interfaces are typically collocated in the same JVM. The EJB container must ensure the semantics for passing arguments across these interfaces conforms to the standard argument passing semantics of the Java programming language.

Responsibilities of EJB Roles

This chapter provides the summary of the responsibilities of each EJB Role.

26.1 Bean Provider's Responsibilities

This section highlights the requirements for the Bean Provider. Meeting these requirements is necessary to ensure that the enterprise beans developed by the Bean Provider can be deployed in all compliant EJB containers.

26.1.1 API Requirements

The enterprise beans must meet all the API requirements defined in the individual chapters of this document.

26.1.2 Packaging Requirements

The Bean Provider is responsible for packaging the enterprise beans in an ejb-jar file in the format described in Chapter 24.

The deployment descriptor must include the *structural* information described in Section 23.2.

The deployment descriptor may optionally include any of the *application assembly* information as described in Section 23.3.

26.2 Application Assembler's Responsibilities

The requirements for the Application Assembler are in defined in Section 23.3 and Chapter 24.

26.3 EJB Container Provider's Responsibilities

The EJB Container Provider is responsible for providing the deployment tools used by the Deployer to deploy enterprise beans packaged in the `ejb-jar` file. The requirements for the deployment tools are defined in the individual chapters of this document.

The EJB Container Provider is responsible for implementing its part of the EJB contracts, and for providing all the runtime services described in the individual chapters of this document.

26.4 Deployer's Responsibilities

The Deployer uses the deployment tools provided by the EJB Container Provider to deploy `ejb-jar` files produced by the Bean Providers and Application Assemblers.

The individual chapters of this document describe the responsibilities of the Deployer in more detail.

26.5 System Administrator's Responsibilities

The System Administrator is responsible for configuring the EJB container and server, setting up security management, integrating resource managers with the EJB container, and runtime monitoring of deployed enterprise beans applications.

The individual chapters of this document describe the responsibilities of the System Administrator in more detail.

26.6 Client Programmer's Responsibilities

The EJB client programmer writes applications that access enterprise beans via their home and component interfaces, via their web service client view, or via messages.

Enterprise JavaBeans™ API Reference

The following interfaces and classes comprise the Enterprise JavaBeans API:

package javax.ejb

Interfaces:

```
public interface EJBContext
public interface EJBHome
public interface EJBLocalHome
public interface EJBLocalObject
public interface EJBMetaData
public interface EJBObject
public interface EnterpriseBean
public interface EntityBean
public interface EntityContext
public interface Handle
public interface HomeHandle
public interface MessageDrivenBean
public interface MessageDrivenContext
public interface SessionBean
public interface SessionContext
public interface SessionSynchronization
public interface TimedObject
public interface Timer
public interface TimerHandle
public interface TimerService
```

Classes:

```
public class AccessLocalException
public class CreateException
public class DuplicateKeyException
public class EJBException
public class FinderException
public class NoSuchEntityException
public class NoSuchObjectLocalException
public class ObjectNotFoundException
public class RemoveException
public class TransactionRequiredLocalException
public class TransactionRolledbackLocalException
```

package javax.ejb.spi

Interfaces:

```
public interface HandleDelegate
```

package javax.ejb.deployment

The `javax.ejb.deployment` package that was defined in the EJB 1.0 specification was deprecated in EJB 1.1. The EJB 1.0 deployment descriptor format should not be used by `ejb-jar` file producer, and the support for it is not required by EJB 1.1 and later compliant containers.

The Javadoc specification of the EJB interface is included in a ZIP file distributed with this document.

The assertions contained in the Javadoc specification of the EJB interface are required functionality and must be implemented by compliant containers.

Related Documents

- [1] JavaBeans. <http://java.sun.com/beans>.
- [2] Java Naming and Directory Interface (JNDI). <http://java.sun.com/products/jndi>.
- [3] Java Remote Method Invocation (RMI). <http://java.sun.com/products/rmi>.
- [4] Java Security. <http://java.sun.com/security>.
- [5] Java Transaction API (JTA). <http://java.sun.com/products/jta>.
- [6] Java Transaction Service (JTS). <http://java.sun.com/products/jts>.
- [7] Java Language to IDL Mapping Specification. <http://www.omg.org/cgi-bin/doc?ptc/00-01-06>.
- [8] CORBA Object Transaction Service v1.2. <http://www.omg.org/cgi-bin/doc?ptc/2000-11-07>.
- [9] Java 2 Platform, Enterprise Edition, v1.4 (J2EE). <http://java.sun.com/j2ee>.
- [10] Java Message Service (JMS), v 1.1. <http://java.sun.com/products/jms>.
- [11] Java API for XML Messaging (JAXM).
- [12] Java 2 Enterprise Edition Connector Architecture, v1.5 . <http://java.sun.com/j2ee/connector>.
- [13] Enterprise JavaBeans to CORBA Mapping v1.1. <http://java.sun.com/products/ejb/docs.html>.
- [14] CORBA 2.3.1 Specification. <http://www.omg.org/cgi-bin/doc?formal/99-10-07>.
- [15] CORBA COSNaming Service. <http://www.omg.org/cgi-bin/doc?formal/00-06-19>.

- [16] Interoperable Name Service FTF document. <http://www.omg.org/cgi-bin/doc?ptc/00-08-07>.
- [17] RFC 2246: The TLS Protocol. <ftp://ftp.isi.edu/in-notes/rfc2246.txt>.
- [18] RFC 2712: Addition of Kerberos Cipher Suites to Transport Layer Security. <ftp://ftp.isi.edu/in-notes/rfc2712.txt>.
- [19] The SSL Protocol Version 3.0. <http://home.netscape.com/eng/ssl3/draft302.txt>.
- [20] Common Secure Interoperability Version 2 Final Available Specification. <http://www.omg.org/cgi-bin/doc?ptc/2001-06-17>.
- [21] Database Language SQL. ANSI X3.135-1992 or ISO/IEC 9075:1992.
- [22] Java API for XML-based RPC (JAX-RPC) 1.1. <http://jcp.org/en/jsr/detail?id=101>.
- [23] Web Services Description Language (WSDL) 1.1. <http://www.w3.org/TR/wsdl>.
- [24] W3C: SOAP 1.1. <http://www.w3.org/TR/SOAP/>.
- [25] Web Services for J2EE, Version 1.1. <http://jcp.org/en/jsr/detail?id=109>,
<http://jcp.org/en/jsr/detail?id=921>.
- [26] The Java Virtual Machine Specification.
- [27] JDBC 3.0 Specification. <http://java.sun.com/products/jdbc>.

Features Deferred to Future Releases

We plan to provide the following in future releases of the Enterprise JavaBeans specification:

- Use of the Java Programming Language metadata facility
- Further extensions to EJB QL, e.g.: subqueries, multi-valued results, dynamic queries
- Read-only Entity Beans with container-managed persistence
- Support for component-level inheritance

EJB 2.0 Deployment Descriptor

This appendix defines the EJB 2.0 deployment descriptor that is part of the ejb-jar file. All EJB 2.1 compliant implementations must support EJB 2.0 deployment descriptors. Section B.1 provides an overview of the deployment descriptor. Sections B.2 through B.4 describe the information in the deployment descriptor from the perspective of the EJB roles responsible for providing the information. Section B.5 defines the deployment descriptor's XML DTD.

B.1 Overview

The deployment descriptor is part of the contract between the ejb-jar file producer and consumer. This contract covers both the passing of enterprise beans from the Bean Provider to the Application Assembler, and from the Application Assembler to the Deployer.

An ejb-jar file produced by the Bean Provider contains one or more enterprise beans and typically does not contain application assembly instructions. An ejb-jar file produced by an Application Assembler contains one or more enterprise beans, plus application assembly information describing how the enterprise beans are combined into a single application deployment unit.

The J2EE specification defines how enterprise beans and other application components contained in multiple ejb-jar files can be assembled into an application.

The role of the deployment descriptor is to capture the declarative information (i.e. information that is not included directly in the enterprise beans' code) that is intended for the consumer of the ejb-jar file.

There are two basic kinds of information in the deployment descriptor:

- *Enterprise beans' structural information.* Structural information describes the structure of an enterprise bean and declares an enterprise bean's external dependencies. Providing structural information in the deployment descriptor is mandatory for the ejb-jar file producer. The structural information cannot, in general, be changed because doing so could break the enterprise bean's function.
- *Application assembly information.* Application assembly information describes how the enterprise bean (or beans) in the ejb-jar file is composed into a larger application deployment unit. Providing assembly information in the deployment descriptor is optional for the ejb-jar file producer. Assembly level information can be changed without breaking the enterprise bean's function, although doing so may alter the behavior of an assembled application.

B.2 Bean Provider's Responsibilities

The Bean Provider is responsible for providing the structural information for each enterprise bean in the deployment descriptor.

The Bean Provider must use the `enterprise-beans` element to list all the enterprise beans in the ejb-jar file.

The Bean Provider must provide the following information for each enterprise bean:

- **Enterprise bean's name.** The Bean Provider must assign a logical name to each enterprise bean in the ejb-jar file. There is no architected relationship between this name and the JNDI name that the Deployer will assign to the enterprise bean. The Bean Provider specifies the enterprise bean's name in the `ejb-name` element.
- **Enterprise bean's class.** The Bean Provider must specify the fully-qualified name of the Java class that implements the enterprise bean's business methods. The Bean Provider specifies the enterprise bean's class name in the `ejb-class` element.
- **Enterprise bean's remote home interface.** The Bean Provider must specify the fully-qualified name of the enterprise bean's remote home interface, if any, in the `home` element.
- **Enterprise bean's remote interface.** The Bean Provider must specify the fully-qualified name of the enterprise bean's remote interface, if any, in the `remote` element.
- **Enterprise bean's local home interface.** The Bean Provider must specify the fully-qualified name of the enterprise bean's local home interface, if any, in the `local-home` element.
- **Enterprise bean's local interface.** The Bean Provider must specify the fully-qualified name of the enterprise bean's local interface, if any, in the `local` element.

- **Enterprise bean's type.** The enterprise bean types are: session, entity, and message-driven. The Bean Provider must use the appropriate `session`, `entity`, or `message-driven` element to declare the enterprise bean's structural information.
- **Re-entrancy indication.** The Bean Provider must specify whether an entity bean is re-entrant or not. Session beans and message-driven beans are never re-entrant.
- **Session bean's state management type.** If the enterprise bean is a session bean, the Bean Provider must use the `session-type` element to declare whether the session bean is stateful or stateless.
- **Session or message-driven bean's transaction demarcation type.** If the enterprise bean is a session or a message-driven bean, the Bean Provider must use the `transaction-type` element to declare whether transaction demarcation is performed by the enterprise bean or by the container.
- **Entity bean's persistence management.** If the enterprise bean is an entity bean, the Bean Provider must use the `persistence-type` element to declare whether persistence management is performed by the enterprise bean or by the container.
- **Entity bean's primary key class.** If the enterprise bean is an entity bean, the Bean Provider specifies the fully-qualified name of the entity bean's primary key class in the `prim-key-class` element. The Bean Provider must specify the primary key class for an entity bean with bean-managed persistence.
- **Entity Bean's abstract schema name.** If the enterprise bean is an entity bean with container-managed persistence and `cmp-version 2.x`, the Bean Provider must specify the abstract schema name of the entity bean using the `abstract-schema-name` element.
- **Container-managed fields.** If the enterprise bean is an entity bean with container-managed persistence, the Bean Provider must specify the container-managed fields using the `cmp-field` elements.
- **Container-managed relationships.** If the enterprise bean is an entity bean with container-managed persistence and `cmp-version 2.x`, the Bean Provider must specify the container-managed relationships of the entity bean using the `relationships` element.
- **Finder and select queries.** If the enterprise bean is an entity bean with container-managed persistence and `cmp-version 2.x`, the Bean Provider must use the `query` element to specify any EJB QL finder or select query for the entity bean other than a query for the `findByPrimaryKey` method.
- **Environment entries.** The Bean Provider must declare all the enterprise bean's environment entries as specified in Subsection 20.2.1.
- **Resource manager connection factory references.** The Bean Provider must declare all the enterprise bean's resource manager connection factory references as specified in Subsection 20.5.1.

- **Resource environment references.** The Bean Provider must declare all the enterprise bean's references to administered objects that are associated with resources as specified in Subsection 20.6.1.
- **EJB references.** The Bean Provider must declare all the enterprise bean's references to the remote homes of other enterprise beans as specified in Subsection 20.3.1.
- **EJB local references.** The Bean Provider must declare all the enterprise bean's references to the local homes of other enterprise beans as specified in Subsection 20.3.1.
- **Security role references.** The Bean Provider must declare all the enterprise bean's references to security roles as specified in Subsection 21.2.5.3.
- **Message-driven bean's destination.** The Bean Provider may provide advice to the Deployer as to the destination type to which a message-driven bean should be assigned.
- **Message-driven bean's message selector.** The Bean Provider may declare the JMS message selector to be used in determining which messages the message-driven bean is to receive.
- **Message-driven bean's acknowledgment mode.** The Bean Provider may declare the JMS acknowledgment mode option that should be used for a message-driven bean with bean-managed transaction demarcation.

The deployment descriptor produced by the Bean Provider must be well formed in the XML sense, and valid with respect to the DTD in Section B.5. The content of the deployment descriptor must conform to the semantics rules specified in the DTD comments and elsewhere in this specification. The deployment descriptor must refer to the DTD using the following statement:

```
<!DOCTYPE ejb-jar PUBLIC "-//Sun Microsystems, Inc.//DTD Enterprise  
JavaBeans 2.0//EN" "http://java.sun.com/dtd/ejb-jar_2_0.dtd">
```

B.3 Application Assembler's Responsibility

The Application Assembler assembles enterprise beans into a single deployment unit. The Application Assembler's input is one or more ejb-jar files provided by one or more Bean Providers, and the output is also one or more ejb-jar files. The Application Assembler can combine multiple input ejb-jar files into a single output ejb-jar file, or split an input ejb-jar file into multiple output ejb-jar files. Each output ejb-jar file is either a deployment unit intended for the Deployer, or a partially assembled application that is intended for another Application Assembler.

The Bean Provider and Application Assembler may be the same person or organization. In such a case, the person or organization performs the responsibilities described both in this and the previous sections.

The Application Assembler may modify the following information that was specified by the Bean Provider:

- **Values of environment entries.** The Application Assembler may change existing and/or define new values of environment properties.
- **Description fields.** The Application Assembler may change existing or create new description elements.
- **Relationship names.** If multiple ejb-jar files use the same names for relationships and are merged into a single ejb-jar file, it is the responsibility of the Application Assembler to modify the relationship names defined in the `ejb-relation-name` elements.
- **Message-driven bean message selector.** The Application Assembler may further restrict, but not replace, the value of the `message-selector` element of a message-driven bean.

In general, the Application Assembler should never modify any of the following.

- **Enterprise bean's abstract schema name.** The Application Assembler should not change the enterprise bean's name defined in the `abstract-schema-name` element since EJB QL queries may depend on the content of this element.
- **Relationship role source element.** The Application Assembler should not change the value of an `ejb-name` element in the `relationship-role-source` element.

If any of these elements must be modified by the Application Assembler in order to resolve name clashes during the merging two ejb-jar files into one, the Application Assembler must also modify all `ejb-ql` query strings that depend on the value of the modified element(s).

The Application Assembler must not, in general, modify any other information listed in Section B.2 that was provided in the input ejb-jar file.

The Application Assembler may, but is not required to, specify any of the following *application assembly* information:

- **Binding of enterprise bean references.** The Application Assembler may link an enterprise bean reference to another enterprise bean in the ejb-jar file or in the same J2EE application unit. The Application Assembler creates the link by adding the `ejb-link` element to the referencing bean. The Application Assembler uses the `ejb-name` of the referenced bean for the link. If there are multiple enterprise beans with the same `ejb-name`, the Application Assembler uses the path name specifying the location of the ejb-jar file that contains the referenced component. The path name is relative to the referencing ejb-jar file. The Application Assembler appends the `ejb-name` of the referenced bean to the path name separated by `#`. This allows multiple beans with the same name to be uniquely identified.
- **Security roles.** The Application Assembler may define one or more security roles. The security roles define the *recommended* security roles for the clients of the enterprise beans. The Application Assembler defines the security roles using the `security-role` elements.
- **Method permissions.** The Application Assembler may define method permissions. Method permission is a binary relation between the security roles and the methods of the home and component interfaces of the enterprise beans. The Application Assembler defines method permissions using the `method-permission` elements.

- **Linking of security role references.** If the Application Assembler defines security roles in the deployment descriptor, the Application Assembler must link the security role references declared by the Bean Provider to the security roles. The Application Assembler defines these links using the `role-link` element.
- **Security identity.** The Application Assembler may specify whether the caller's security identity should be used for the execution of the methods of an enterprise bean or whether a specific run-as security identity should be used.
- **Transaction attributes.** The Application Assembler may define the value of the transaction attributes for the methods of the home and component interfaces of the enterprise beans that require container-managed transaction demarcation. All entity beans and the session and message-driven beans declared by the Bean Provider as transaction-type `Container` require container-managed transaction demarcation. The Application Assembler uses the `container-transaction` elements to declare the transaction attributes.

If an input `ejb-jar` file contains application assembly information, the Application Assembler is allowed to change the application assembly information supplied in the input `ejb-jar` file. (This could happen when the input `ejb-jar` file was produced by another Application Assembler.)

The deployment descriptor produced by the Bean Provider must be well formed in the XML sense, and valid with respect to the DTD in Section B.5. The content of the deployment descriptor must conform to the semantic rules specified in the DTD comments and elsewhere in this specification. The deployment descriptor must refer to the DTD using the following statement:

```
<!DOCTYPE ejb-jar PUBLIC "-//Sun Microsystems, Inc.//DTD Enterprise
JavaBeans 2.0//EN" "http://java.sun.com/dtd/ejb-jar_2_0.dtd">
```

B.4 Container Provider's Responsibilities

The Container Provider provides tools that read and import the information contained in the XML deployment descriptor.

All EJB 2.0 implementations must support EJB 1.1 as well as EJB 2.0 deployment descriptors.

The EJB 1.1 deployment descriptor is defined in Appendix C.

B.5 Deployment Descriptor DTD

This section defines the XML DTD for the EJB 2.0 deployment descriptor. The comments in the DTD specify additional requirements for the syntax and semantics that cannot be easily expressed by the DTD mechanism.

The content of the XML elements is in general case sensitive. This means, for example, that

```
<reentrant>True</reentrant>
```

must be used, rather than:

```
<reentrant>true</reentrant>.
```

All valid ejb-jar deployment descriptors must contain the following DOCTYPE declaration:

```
<!DOCTYPE ejb-jar PUBLIC "-//Sun Microsystems, Inc.//DTD Enterprise
JavaBeans 2.0//EN" "http://java.sun.com/dtd/ejb-jar_2_0.dtd">
```

```
<!--
This is the XML DTD for the EJB 2.0 deployment descriptor.
-->
```

```
<!--
The abstract-schema-name element specifies the name of the abstract
schema type of an entity bean with cmp-version 2.x. It is used in EJB
QL queries.
```

For example, the abstract-schema-name for an entity bean whose local interface is com.acme.commerce.Order might be Order.

```
Used in: entity.
-->
```

```
<!ELEMENT abstract-schema-name (#PCDATA)>
```

```
<!--
The acknowledge-mode element specifies whether JMS AUTO_ACKNOWLEDGE
or DUPS_OK_ACKNOWLEDGE message acknowledgment semantics should be
used for the onMessage message of a message-driven bean that uses bean
managed transaction demarcation. If the acknowledge-mode deployment
descriptor element is not specified, JMS AUTO_ACKNOWLEDGE semantics
are assumed.
```

The acknowledge-mode element must be one of the two following:

```
<acknowledge-mode>Auto-acknowledge</acknowledge-mode>
<acknowledge-mode>Dups-ok-acknowledge</acknowledgemode>
```

```
Used in: message-driven
-->
```

```
<!ELEMENT acknowledge-mode (#PCDATA)>
```

```
<!--
The assembly-descriptor element contains application-assembly infor-
mation.
```

The application-assembly information consists of the following parts: the definition of security roles, the definition of method permissions the definition of transaction attributes for enterprise beans with container-managed transaction demarcation, and a list of methods that are to be excluded from being invoked.

All the parts are optional in the sense that they are omitted if the lists represented by them are empty.

Providing an assembly-descriptor in the deployment descriptor is optional for the ejb-jar file producer.

```
Used in: ejb-jar
-->
```

```
<!ELEMENT assembly-descriptor (security-role*, method-permission*,
container-transaction*, exclude-list?)>
```

<!--

The cascade-delete element specifies that, within a particular relationship, the lifetime of one or more entity beans is dependent upon the lifetime of another entity bean. The cascade-delete element can only be specified for an ejb-relationship-role element contained in an ejb-relation element in which the other ejb-relationship-role element specifies a multiplicity of One.

Used in: ejb-relationship-role

-->

<!ELEMENT cascade-delete EMPTY>

<!--

The cmp-field element describes a container-managed field. The field element includes an optional description of the field, and the name of the field.

Used in: entity

-->

<!ELEMENT cmp-field (description?, field-name)>

<!--

The cmp-version element specifies the version of an entity bean with container-managed persistence.

The cmp-version element must be one of the two following:

<cmp-version>1.x</cmp-version>

<cmp-version>2.x</cmp-version>

The default value of the cmp-version element is 2.x.

Used in: entity

-->

<!ELEMENT cmp-version (#PCDATA)>

<!--

The cmr-field element describes the bean provider's view of a relationship. It consists of an optional description, and the name and the class type of a field in the source of a role of a relationship. The cmr-field-name element corresponds to the name used for the get and set accessor methods for the relationship. The cmr-field-type element is used only for collection-valued cmr-fields. It specifies the type of the collection that is used.

Used in: ejb-relationship-role

-->

<!ELEMENT cmr-field (description?, cmr-field-name, cmr-field-type?)>

<!--

The cmr-field-name element specifies the name of a logical relationship field in the entity bean class. The name of the cmr-field must begin with a lowercase letter. This field is accessed by methods whose names consist of the name of the field specified by cmr-field-name in which the first letter is uppercased, prefixed by "get" or "set".

Used in: cmr-field

-->

<!ELEMENT cmr-field-name (#PCDATA)>

```

<!--
The cmr-field-type element specifies the class of a collection-valued
logical relationship field in the entity bean class. The value of the
cmr-field-type element must be either: java.util.Collection or
java.util.Set.

Used in: cmr-field
-->
<!--ELEMENT cmr-field-type (#PCDATA)>

<!--
The container-transaction element specifies how the container must
manage transaction scopes for the enterprise bean's method invoca-
tions. The element consists of an optional description, a list of
method elements, and a transaction attribute. The transaction
attribute is to be applied to all the specified methods.

Used in: assembly-descriptor
-->
<!--ELEMENT container-transaction (description?, method+,
trans-attribute)>

<!--
The description element is used by the ejb-jar file producer to pro-
vide text describing the parent element.

The description element should include any information that the
enterprise bean ejb-jar file producer wants to provide to the consumer
of the enterprise bean ejb-jar file (i.e. to the Deployer). Typically,
the tools used by the ejb-jar file consumer will display the descrip-
tion when processing the parent element that contains the descrip-
tion.

Used in: cmp-field, cmr-field, container-transaction, ejb-jar,
ejb-local-ref, ejb-ref, ejb-relation, ejb-relationship-role, entity,
env-entry, exclude-list, message-driven, method, method-permission,
query, relationships, relationship-role-source, resource-env-ref,
resource-ref, run-as, security-identity, security-role, secu-
rity-role-ref, and session.
-->
<!--ELEMENT description (#PCDATA)>

<!--
The destination-type element specifies the type of the JMS destina-
tion. The type is specified by the Java interface expected to be
implemented by the destination.

The destination-type element must be one of the two following:
    <destination-type>javax.jms.Queue</destination-type>
    <destination-type>javax.jms.Topic</destination-type>

Used in: message-driven-destination
-->
<!--ELEMENT destination-type (#PCDATA)>

<!--
The display-name element contains a short name that is intended to be
displayed by tools. The display name need not be unique.

```

Used in: ejb-jar, session, entity, and message-driven

Example:

```
<display-name>Employee Self Service</display-name>
-->
<!ELEMENT display-name (#PCDATA)>
```

<!--

The ejb-class element contains the fully-qualified name of the enterprise bean's class.

Used in: entity, message-driven, and session

Example:

```
<ejb-class>com.wombat.empl.EmployeeServiceBean</ejb-class>
-->
<!ELEMENT ejb-class (#PCDATA)>
```

<!--

The optional ejb-client-jar element specifies a JAR file that contains the class files necessary for a client program to access the enterprise beans in the ejb-jar file.

Used in: ejb-jar

Example:

```
<ejb-client-jar>employee_service_client.jar</ejb-client-jar>
-->
<!ELEMENT ejb-client-jar (#PCDATA)>
```

<!--

The ejb-jar element is the root element of the EJB deployment descriptor. It contains an optional description of the ejb-jar file; an optional display name; an optional small icon file name; an optional large icon file name; mandatory structural information about all included enterprise beans; a descriptor for container-managed relationships, if any; an optional application-assembly descriptor; and an optional name of an ejb-client-jar file for the ejb-jar.

-->

```
<!ELEMENT ejb-jar (description?, display-name?, small-icon?,
    large-icon?, enterprise-beans, relationships?,
    assembly-descriptor?, ejb-client-jar?)>
```

<!--

The ejb-link element is used in the ejb-ref and ejb-local-ref elements to specify that an EJB reference is linked to an enterprise bean.

The value of the ejb-link element must be the ejb-name of an enterprise bean in the same ejb-jar file or in another ejb-jar file in the same J2EE application unit.

Alternatively, the name in the ejb-link element may be composed of a path name specifying the ejb-jar containing the referenced enterprise bean with the ejb-name of the target bean appended and separated from the path name by "#". The path name is relative to the jar file containing the enterprise bean that is referencing the enterprise bean. This allows multiple enterprise beans with the same ejb-name to be uniquely identified.

Used in: ejb-ref, ejb-local-ref

Examples:

```
<ejb-link>EmployeeRecord</ejb-link>
<ejb-link>../products/product.jar#ProductEJB</ejb-link>
```

-->

<!ELEMENT ejb-link (#PCDATA)>

<!--

The ejb-local-ref element is used for the declaration of a reference to an enterprise bean's local home. The declaration consists of an optional description; the EJB reference name used in the code of the enterprise bean that is referencing the enterprise bean; the expected type of the referenced enterprise bean; the expected local home and local interfaces of the referenced enterprise bean; and optional ejb-link information, used to specify the referenced enterprise bean.

Used in: entity, session, message-driven

-->

<!ELEMENT ejb-local-ref (description?, ejb-ref-name, ejb-ref-type, local-home, local, ejb-link?)>

<!--

The ejb-name element specifies an enterprise bean's name. This name is assigned by the ejb-jar file producer to name the enterprise bean in the ejb-jar file's deployment descriptor. The name must be unique among the names of the enterprise beans in the same ejb-jar file.

There is no architected relationship between the ejb-name in the deployment descriptor and the JNDI name that the Deployer will assign to the enterprise bean's home.

The name for an entity bean with cmp-version 2.x must conform to the lexical rules for an NMTOKEN. The name for an entity bean with cmp-version 2.x must not be a reserved literal in EJB QL.

Used in: entity, message-driven, session, method, and relationship-role-source.

Example:

```
<ejb-name>EmployeeService</ejb-name>
```

-->

<!ELEMENT ejb-name (#PCDATA)>

<!--

The ejb-ql element contains the EJB QL query string that defines a finder or select query. This element is defined within the scope of a query element whose contents specify the finder or the select method that uses the query. The content must be a valid EJB QL query string for the entity bean for which the query is specified.

The ejb-ql element must be specified for all queries that are expressible in EJB QL.

Used in: query

Example:


```

<query>
  <query-method>
    <method-name>ejbSelectPendingLineitems</method-name>
    <method-params/>
  </query-method>
  <ejb-ql>SELECT OBJECT(l)
    FROM LineItems l
    WHERE l.shipped <> TRUE
  </ejb-ql>
</query>

```

```

-->
<!--ELEMENT ejb-ql (#PCDATA)>

```

<!--
The ejb-ref element is used for the declaration of a reference to an enterprise bean's home. The declaration consists of an optional description; the EJB reference name used in the code of the enterprise bean that is referencing the enterprise bean; the expected type of the referenced enterprise bean; the expected home and remote interfaces of the referenced enterprise bean; and optional ejb-link information, used to specify the referenced enterprise bean.

Used in: entity, message-driven, and session

```

-->
<!--ELEMENT ejb-ref (description?, ejb-ref-name, ejb-ref-type, home,
  remote, ejb-link?)>

```

<!--
The ejb-ref-name element contains the name of an EJB reference. The EJB reference is an entry in the enterprise bean's environment and is relative to the java:comp/env context. The name must be unique within the enterprise bean.

It is recommended that name is prefixed with "ejb/".

Used in: ejb-ref, ejb-local-ref

Example:

```

  <ejb-ref-name>ejb/Payroll</ejb-ref-name>
-->
<!--ELEMENT ejb-ref-name (#PCDATA)>

```

<!--
The ejb-ref-type element contains the expected type of the referenced enterprise bean.

The ejb-ref-type element must be one of the following:

```

  <ejb-ref-type>Entity</ejb-ref-type>
  <ejb-ref-type>Session</ejb-ref-type>

```

Used in: ejb-ref, ejb-local-ref

```

-->
<!--ELEMENT ejb-ref-type (#PCDATA)>

```

<!--
The ejb-relation element describes a relationship between two entity beans with container-managed persistence. An ejb-relation element contains a description; an optional ejb-relation-name element; and

exactly two relationship role declarations, defined by the `ejb-relationship-role` elements. The name of the relationship, if specified, is unique within the `ejb-jar` file.

Used in: `relationships`

-->

```
<!--ELEMENT ejb-relation (description?, ejb-relation-name?,
    ejb-relationship-role, ejb-relationship-role)>
```

<!--

The `ejb-relation-name` element provides a unique name for a relationship.

Used in: `ejb-relation`

-->

```
<!--ELEMENT ejb-relation-name (#PCDATA)>
```

<!--

The `ejb-relationship-role` element describes a role within a relationship. There are two roles in each relationship.

The `ejb-relationship-role` element contains an optional description; an optional name for the relationship role; a specification of the multiplicity of the role; an optional specification of cascade-delete functionality for the role; the role source; and a declaration of the `cmr-field`, if any, by means of which the other side of the relationship is accessed from the perspective of the role source.

The multiplicity and `relationship-role-source` element are mandatory.

The `relationship-role-source` element designates an entity bean by means of an `ejb-name` element. For bidirectional relationships, both roles of a relationship must declare a `relationship-role-source` element that specifies a `cmr-field` in terms of which the relationship is accessed. The lack of a `cmr-field` element in an `ejb-relationship-role` specifies that the relationship is unidirectional in navigability and that entity bean that participates in the relationship is "not aware" of the relationship.

Used in: `ejb-relation`

Example:

```
<ejb-relation>
  <ejb-relation-name>Product-LineItem</ejb-relation-name>
  <ejb-relationship-role>
    <ejb-relationship-role-name>
      product-has-lineitems
    </ejb-relationship-role-name>
    <multiplicity>One</multiplicity>
    <relationship-role-source>
      <ejb-name>ProductEJB</ejb-name>
    </relationship-role-source>
  </ejb-relationship-role>
-->
<!--ELEMENT ejb-relationship-role (description?,
    ejb-relationship-role-name?, multiplicity, cascade-delete?,
    relationship-role-source, cmr-field?)>
```

```
<!--
The ejb-relationship-role-name element defines a name for a role that
is unique within an ejb-relation. Different relationships can use the
same name for a role.
```

```
Used in: ejb-relationship-role
```

```
-->
```

```
<!ELEMENT ejb-relationship-role-name (#PCDATA)>
```

```
<!--
```

```
The enterprise-beans element contains the declarations of one or more
enterprise beans.
```

```
-->
```

```
<!ELEMENT enterprise-beans (session | entity | message-driven)+>
```

```
<!--
```

```
The entity element declares an entity bean. The declaration consists
of: an optional description; an optional display name; an optional
small icon file name; optional large icon file name; a unique name
assigned to the enterprise bean in the deployment descriptor; the
names of the entity bean's home and remote interfaces, if any; the
names of the entity bean's local home and local interfaces, if any;
the entity bean's implementation class; the entity bean's persistence
management type; the entity bean's primary key class name; an indica-
tion of the entity bean's reentrancy; an optional specification of the
entity bean's cmp-version; an optional specification of the entity
bean's abstract schema name; an optional list of container-managed
fields; an optional specification of the primary key field; an
optional declaration of the bean's environment entries; an optional
declaration of the bean's EJB references; an optional declaration of
the bean's local EJB references; an optional declaration of the secu-
rity role references; an optional declaration of the security iden-
tity to be used for the execution of the bean's methods; an optional
declaration of the bean's resource manager connection factory refer-
ences; an optional declaration of the bean's resource environment
references; an optional set of query declarations for finder and
select methods for an entity bean with cmp-version 2.x.
```

The optional abstract-schema-name element must be specified for an entity bean with container-managed persistence and cmp-version 2.x.

The optional primkey-field may be present in the descriptor if the entity's persistence-type is Container.

The optional cmp-version element may be present in the descriptor if the entity's persistence-type is Container. If the persistence-type is Container and the cmp-version element is not specified, its value defaults to 2.x.

The optional home and remote elements must be specified if the entity bean cmp-version is 1.x.

The optional home and remote elements must be specified if the entity bean has a remote home and remote interface.

The optional local-home and local elements must be specified if the entity bean has a local home and local interface.

Either both the local-home and the local elements or both the home and

the remote elements must be specified.

The optional query elements must be present if the persistence-type is Container and the cmp-version is 2.x and query methods other than findByPrimaryKey have been defined for the entity bean.

The other elements that are optional are "optional" in the sense that they are omitted if the lists represented by them are empty.

At least one cmp-field element must be present in the descriptor if the entity's persistence-type is Container and the cmp-version is 1.x, and none must not be present if the entity's persistence-type is Bean.

Used in: enterprise-beans

-->

```
<!--ELEMENT entity (description?, display-name?, small-icon?,
    large-icon?, ejb-name, home?, remote?, local-home?,
    local?, ejb-class, persistence-type, prim-key-class,
    reentrant, cmp-version?, abstract-schema-name?,
    cmp-field*, primkey-field?, env-entry*, ejb-ref*,
    ejb-local-ref*, security-role-ref*, security-identity?,
    resource-ref*, resource-env-ref*, query*)>
```

<!--

The env-entry element contains the declaration of an enterprise bean's environment entry. The declaration consists of an optional description, the name of the environment entry, and an optional value. If a value is not specified, one must be supplied during deployment.

Used in: entity, message-driven, and session

-->

```
<!--ELEMENT env-entry (description?, env-entry-name, env-entry-type,
    env-entry-value?)>
```

<!--

The env-entry-name element contains the name of an enterprise bean's environment entry. The name is a JNDI name relative to the java:comp/env context. The name must be unique within an enterprise bean.

Used in: env-entry

Example:

```
<env-entry-name>minAmount</env-entry-name>
```

-->

```
<!--ELEMENT env-entry-name (#PCDATA)>
```

<!--

The env-entry-type element contains the fully-qualified Java type of the environment entry value that is expected by the enterprise bean's code.

The following are the legal values of env-entry-type: java.lang.Boolean, java.lang.String, java.lang.Character, java.lang.Integer, java.lang.Double, java.lang.Byte, java.lang.Short, java.lang.Long, and java.lang.Float.

Used in: env-entry

Example:

```
<env-entry-type>java.lang.Boolean</env-entry-type>
-->
<!ELEMENT env-entry-type (#PCDATA)>
```

<!--

The env-entry-value element contains the value of an enterprise bean's environment entry. The value must be a String that is valid for the constructor of the specified type that takes a single String parameter, or for java.lang.Character, a single character.

Used in: env-entry

Example:

```
<env-entry-value>100.00</env-entry-value>
-->
<!ELEMENT env-entry-value (#PCDATA)>
```

<!--

The exclude-list element defines one or more methods which the Assembler marks to be uncallable. If the method permission relation contains methods that are in the exclude list, the Deployer should consider those methods to be uncallable.

Used in: assembly-descriptor

-->

```
<!ELEMENT exclude-list (description?, method+)>
```

<!--

The field-name element specifies the name of a container-managed field.

The name of the cmp-field of an entity bean with cmp-version 2.x must begin with a lowercase letter. This field is accessed by methods whose names consists of the name of the field specified by field-name in which the first letter is uppercased, prefixed by "get" or "set".

The name of the cmp-field of an entity bean with cmp-version 1.x must denote a public field of the enterprise bean class or one of its superclasses.

Used in: cmp-field

Example:

```
<field-name>firstName</field-name>
-->
<!ELEMENT field-name (#PCDATA)>
```

<!--

The home element contains the fully-qualified name of the enterprise bean's home interface.

Used in: ejb-ref, entity, and session

Example:

```
<home>com.aardvark.payroll.PayrollHome</home>
-->
```

```
<!--ELEMENT home (#PCDATA)>
```

```
<!--
```

The large-icon element contains the name of a file containing a large (32 x 32) icon image. The file name is relative path within the ejb-jar file.

The image must be either in the JPEG or GIF format.
The icon can be used by tools.

Example:

```
<large-icon>employee-service-icon32x32.jpg</large-icon>
```

```
-->
```

```
<!--ELEMENT large-icon (#PCDATA)>
```

```
<!--
```

The local element contains the fully-qualified name of the enterprise bean's local interface.

Used in: entity, session, ejb-local-ref

```
-->
```

```
<!--ELEMENT local (#PCDATA)>
```

```
<!--
```

The local-home element contains the fully-qualified name of the enterprise bean's local home interface.

Used in: entity, session, ejb-local-ref

```
-->
```

```
<!--ELEMENT local-home (#PCDATA)>
```

```
<!--
```

The message-driven element declares a message-driven bean. The declaration consists of: an optional description; optional display name; optional small icon file name; optional large icon file name; a name assigned to the enterprise bean in the deployment descriptor; the message-driven bean's implementation class; the message-driven bean's transaction management type; an optional declaration of the message-driven bean's message selector; an optional declaration of the acknowledgment mode for the message-driven bean if bean-managed transaction demarcation is used; an optional declaration of the intended destination type of the message-driven bean; an optional declaration of the bean's environment entries; an optional declaration of the bean's EJB references; an optional declaration of the bean's local EJB references; an optional declaration of the security identity to be used for the execution of the bean's methods; an optional declaration of the bean's resource manager connection factory references; and an optional declaration of the bean's resource environment references.

Used in: enterprise-beans

```
-->
```

```
<!--ELEMENT message-driven (description?, display-name?, small-icon?,
    large-icon?, ejb-name, ejb-class, transaction-type,
    message-selector?, acknowledge-mode?,
    message-driven-destination?, env-entry*, ejb-ref*,
    ejb-local-ref*, security-identity?, resource-ref*,
    resource-env-ref*)>
```

```
<!--
The message-driven-destination element provides advice to the
Deployer as to whether a message-driven bean is intended for a Queue
or a Topic. The declaration consists of: the type of the mes-
sage-driven bean's intended destination and an optional declaration
of whether a durable or non-durable subscription should be used if the
destination-type is javax.jms.Topic.
```

```
Used in: message-driven
```

```
-->
```

```
<!ELEMENT message-driven-destination (destination-type,
subscription-durability?)>
```

```
<!--
```

```
The message-selector element is used to specify the JMS message selec-
tor to be used in determining which messages a message-driven bean is
to receive.
```

```
Example:
```

```
<message-selector>JMSType = 'car' AND color = 'blue' AND weight >
2500</message-selector>
```

```
Used in: message-driven
```

```
-->
```

```
<!ELEMENT message-selector (#PCDATA)>
```

```
<!--
```

```
The method element is used to denote a method of an enterprise bean's
home or component interface, or, in the case of a message-driven bean,
the bean's onMessage method, or a set of such methods. The ejb-name
element must be the name of one of the enterprise beans declared in
the deployment descriptor; the optional method-intf element allows to
distinguish between a method with the same signature that is multiply
defined across the home and component interfaces; the method-name
element specifies the method name; and the optional method-params
elements identify a single method among multiple methods with an over-
loaded method name.
```

There are three possible styles of the method element syntax:

1.

```
<method>
    <ejb-name>EJBNAME</ejb-name>
    <method-name>*</method-name>
</method>
```

This style is used to refer to all the methods of the specified enterprise bean's home and component interfaces.

2.

```
<method>
    <ejb-name>EJBNAME</ejb-name>
    <method-name>METHOD</method-name>
</method>>
```

This style is used to refer to the specified method of the specified enterprise bean. If there are multiple methods with the same overloaded name, the element of this style refers to all the methods with the overloaded name.

```

3.  <method>
    <ejb-name>EJBNAME</ejb-name>
    <method-name>METHOD</method-name>
    <method-params>
      <method-param>PARAM-1</method-param>
      <method-param>PARAM-2</method-param>
      ...
      <method-param>PARAM-n</method-param>
    </method-params>
  </method>

```

This style is used to refer to a single method within a set of methods with an overloaded name. PARAM-1 through PARAM-n are the fully-qualified Java types of the method's input parameters (if the method has no input arguments, the method-params element contains no method-param elements). Arrays are specified by the array element's type, followed by one or more pair of square brackets (e.g. int[][]). If there are multiple methods with the same overloaded name, this style refers to all of the overloaded methods.

Used in: method-permission, container-transaction, exclude-list.

Examples:

Style 1: The following method element refers to all the methods of the EmployeeService bean's home and component interfaces:

```

<method>
  <ejb-name>EmployeeService</ejb-name>
  <method-name>*</method-name>
</method>

```

Style 2: The following method element refers to all the *create* methods of the EmployeeService bean's home interface(s):

```

<method>
  <ejb-name>EmployeeService</ejb-name>
  <method-name>create</method-name>
</method>

```

Style 3: The following method element refers to the *create(String firstName, String LastName)* method of the EmployeeService bean's home interface(s).

```

<method>
  <ejb-name>EmployeeService</ejb-name>
  <method-name>create</method-name>
  <method-params>
    <method-param>java.lang.String</method-param>
    <method-param>java.lang.String</method-param>
  </method-params>
</method>

```

The following example illustrates a Style 3 element with more complex parameter types. The method

*foobar(char s, int i, int[] iar, mypackage.MyClass mycl,
mypackage.MyClass[][] myclaar)*
would be specified as:

```
<method>
  <ejb-name>EmployeeService</ejb-name>
  <method-name>foobar</method-name>
  <method-params>
    <method-param>char</method-param>
    <method-param>int</method-param>
    <method-param>int[]</method-param>
    <method-param>mypackage.MyClass</method-param>
    <method-param>mypackage.MyClass[][]</method-param>
  </method-params>
</method>
```

The optional `method-intf` element can be used when it becomes necessary to differentiate between a method that is multiply defined across the enterprise bean's home and component interfaces the same name and signature.

For example, the method element

```
<method>
  <ejb-name>EmployeeService</ejb-name>
  <method-intf>Remote</method-intf>
  <method-name>create</method-name>
  <method-params>
    <method-param>java.lang.String</method-param>
    <method-param>java.lang.String</method-param>
  </method-params>
</method>
```

can be used to differentiate the `create(String, String)` method defined in the remote interface from the `create(String, String)` method defined in the remote home interface, which would be defined as

```
<method>
  <ejb-name>EmployeeService</ejb-name>
  <method-intf>Home</method-intf>
  <method-name>create</method-name>
  <method-params>
    <method-param>java.lang.String</method-param>
    <method-param>java.lang.String</method-param>
  </method-params>
</method>
```

and the `create` method that is defined in the local home interface, which would be defined as

```
<method>
  <ejb-name>EmployeeService</ejb-name>
  <method-intf>LocalHome</method-intf>
  <method-name>create</method-name>
  <method-params>
    <method-param>java.lang.String</method-param>
    <method-param>java.lang.String</method-param>
  </method-params>
```

```
</method>
```

The method-intf element can be used with all three Styles of the method element usage. For example, the following method element example could be used to refer to all the methods of the EmployeeService bean's remote home interface.

```
<method>
  <ejb-name>EmployeeService</ejb-name>
  <method-intf>Home</method-intf>
  <method-name>*</method-name>
</method>
```

```
-->
```

```
<!--ELEMENT method (description?, ejb-name, method-intf?, method-name,
method-params?)>
```

```
<!--
```

The method-intf element allows a method element to differentiate between the methods with the same name and signature that are multiply defined across the home and component interfaces (e.g., in both an enterprise bean's remote and local interfaces, or in both an enterprise bean's home and remote interfaces, etc.)

The method-intf element must be one of the following:

```
<method-intf>Home</method-intf>
<method-intf>Remote</method-intf>
<method-intf>LocalHome</method-intf>
<method-intf>Local</method-intf>
```

Used in: method

```
-->
```

```
<!--ELEMENT method-intf (#PCDATA)>
```

```
<!--
```

The method-name element contains a name of an enterprise bean method, or the asterisk (*) character. The asterisk is used when the element denotes all the methods of an enterprise bean's component and home interfaces.

Used in: method, query-method

```
-->
```

```
<!--ELEMENT method-name (#PCDATA)>
```

```
<!--
```

The method-param element contains the fully-qualified Java type name of a method parameter.

Used in: method-params

```
-->
```

```
<!--ELEMENT method-param (#PCDATA)>
```

```
<!--
```

The method-params element contains a list of the fully-qualified Java type names of the method parameters.

Used in: method, query-method

```
-->
```

```
<!--ELEMENT method-params (method-param*)>
```

```
<!--
The method-permission element specifies that one or more security
roles are allowed to invoke one or more enterprise bean methods. The
method-permission element consists of an optional description, a list
of security role names or an indicator to specify that the methods are
not to be checked for authorization, and a list of method elements.
```

The security roles used in the method-permission element must be defined in the security-role elements of the deployment descriptor, and the methods must be methods defined in the enterprise bean's component and/or home interfaces.

Used in: assembly-descriptor

```
-->
```

```
<!ELEMENT method-permission (description?, (role-name+|unchecked),
method+)>
```

```
<!--
```

The multiplicity element describes the multiplicity of the role that participates in a relation.

The multiplicity element must be one of the two following:

```
<multiplicity>One</multiplicity>
<multiplicity>Many</multiplicity>
```

Used in: ejb-relationship-role

```
-->
```

```
<!ELEMENT multiplicity (#PCDATA)>
```

```
<!--
```

The persistence-type element specifies an entity bean's persistence management type.

The persistence-type element must be one of the two following:

```
<persistence-type>Bean</persistence-type>
<persistence-type>Container</persistence-type>
```

Used in: entity

```
-->
```

```
<!ELEMENT persistence-type (#PCDATA)>
```

```
<!--
```

The prim-key-class element contains the fully-qualified name of an entity bean's primary key class.

If the definition of the primary key class is deferred to deployment time, the prim-key-class element should specify java.lang.Object.

Used in: entity

Examples:

```
<prim-key-class>java.lang.String</prim-key-class>
<prim-key-class>com.wombat.empl.EmployeeID</prim-key-class>
<prim-key-class>java.lang.Object</prim-key-class>
```

```
-->
```

```
<!ELEMENT prim-key-class (#PCDATA)>
```

<!--

The `primkey-field` element is used to specify the name of the primary key field for an entity with container-managed persistence.

The `primkey-field` must be one of the fields declared in the `cmp-field` element, and the type of the field must be the same as the primary key type.

The `primkey-field` element is not used if the primary key maps to multiple container-managed fields (i.e. the key is a compound key). In this case, the fields of the primary key class must be public, and their names must correspond to the field names of the entity bean class that comprise the key.

Used in: entity

Example:

```
<primkey-field>EmployeeId</primkey-field>
-->
<!ELEMENT primkey-field (#PCDATA)>
```

<!--

The `query` element is used to specify a finder or select query. It contains an optional description of the query; the specification of the finder or select method it is used by; an optional specification of the result type mapping, if the query is for a select method and entity objects are returned; and the EJB QL query string that defines the query.

Queries that are expressible in EJB QL must use the `ejb-ql` element to specify the query. If a query is not expressible in EJB QL, the `description` element should be used to describe the semantics of the query and the `ejb-ql` element should be empty.

The `result-type-mapping` element is optional. It can only be present if the `query-method` element specifies a select method that returns entity objects. The default value for the `result-type-mapping` element is `Local`.

Used in: entity

```
-->
<!ELEMENT query (description?, query-method, result-type-mapping?,
ejb-ql)>
```

<!--

The `query-method` element is used to specify the method for a finder or select query.

The `method-name` element specifies the name of a finder or select method in the entity bean's implementation class.

Each `method-param` must be defined for a `query-method` using the `method-params` element.

Used in: query

Example:

```
<query>
  <description>Method finds large orders</description>
```

```

        <query-method>
          <method-name>findLargeOrders</method-name>
          <method-params></method-params>
        </query-method>
        <ejb-ql>SELECT OBJECT(o) FROM Order o WHERE o.amount >
1000</ejb-ql>
      </query>

-->
<!--ELEMENT query-method (method-name, method-params)-->

<!--
The reentrant element specifies whether an entity bean is reentrant or
not.

The reentrant element must be one of the two following:
    <reentrant>True</reentrant>
    <reentrant>False</reentrant>

Used in: entity
-->
<!--ELEMENT reentrant (#PCDATA)-->

<!--
The relationships element describes the relationships in which entity
beans with container-managed persistence participate. The relation-
ships element contains an optional description; and a list of
ejb-relation elements, which specify the container-managed relation-
ships.

Used in: ejb-jar
-->
<!--ELEMENT relationships (description?, ejb-relation+)-->

<!--
The relationship-role-source element designates the source of a role
that participates in a relationship. A relationship-role-source ele-
ment uniquely identifies an entity bean.

Used in: ejb-relationship-role
-->
<!--ELEMENT relationship-role-source (description?, ejb-name)-->

<!--
The remote element contains the fully-qualified name of the enter-
prise bean's remote interface.

Used in: ejb-ref, entity, and session

Example:
    <remote>com.wombat.empl.EmployeeService</remote>
-->
<!--ELEMENT remote (#PCDATA)-->

<!--
The res-auth element specifies whether the enterprise bean code signs
on programmatically to the resource manager, or whether the Container
will sign on to the resource manager on behalf of the enterprise bean.
In the latter case, the Container uses information that is supplied by

```

the Deployer.

The value of this element must be one of the two following:

```
<res-auth>Application</res-auth>
<res-auth>Container</res-auth>
```

```
-->
```

```
<!--ELEMENT res-auth (#PCDATA)-->
```

```
<!--
```

The res-ref-name element specifies the name of a resource manager connection factory reference. The name is a JNDI name relative to the java:comp/env context. The name must be unique within an enterprise bean.

Used in: resource-ref

```
-->
```

```
<!--ELEMENT res-ref-name (#PCDATA)-->
```

```
<!--
```

The res-sharing-scope element specifies whether connections obtained through the given resource manager connection factory reference can be shared. The value of this element, if specified, must be one of the two following:

```
<res-sharing-scope>Shareable</res-sharing-scope>
<res-sharing-scope>Unshareable</res-sharing-scope>
```

The default value is Shareable.

Used in: resource-ref

```
-->
```

```
<!--ELEMENT res-sharing-scope (#PCDATA)-->
```

```
<!--
```

The res-type element specifies the type of the data source. The type is specified by the Java language class or interface expected to be implemented by the data source.

Used in: resource-ref

```
-->
```

```
<!--ELEMENT res-type (#PCDATA)-->
```

```
<!--
```

The resource-env-ref element contains a declaration of an enterprise bean's reference to an administered object associated with a resource in the enterprise bean's environment. It consists of an optional description, the resource environment reference name, and an indication of the resource environment reference type expected by the enterprise bean code.

Used in: entity, message-driven and session

Examples:

```
<resource-env-ref>
  <resource-env-ref-name>jms/StockQueue
  </resource-env-ref-name>
  <resource-env-ref-type>javax.jms.Queue
  </resource-env-ref-type>
</resource-env-ref>
```

```
-->
```

```
<!ELEMENT resource-env-ref (description?, resource-env-ref-name,
resource-env-ref-type)>
```

```
<!--
```

The resource-env-ref-name element specifies the name of a resource environment reference; its value is the environment entry name used in the enterprise bean code. The name is a JNDI name relative to the java:comp/env context and must be unique within an enterprise bean.

```
Used in: resource-env-ref
```

```
-->
```

```
<!ELEMENT resource-env-ref-name (#PCDATA)>
```

```
<!--
```

The resource-env-ref-type element specifies the type of a resource environment reference. It is the fully qualified name of a Java language class or interface

```
Used in: resource-env-ref
```

```
-->
```

```
<!ELEMENT resource-env-ref-type (#PCDATA)>
```

```
<!--
```

The resource-ref element contains a declaration of enterprise bean's reference to an external resource. It consists of an optional description, the resource manager connection factory reference name, the indication of the resource manager connection factory type expected by the enterprise bean code, the type of authentication (Application or Container), and an optional specification of the shareability of connections obtained from the resource (Shareable or Unshareable).

```
Used in: entity, message-driven, and session
```

```
Example:
```

```
    <resource-ref>
      <res-ref-name>jdbc/EmployeeAppDB</res-ref-name>
      <res-type>javax.sql.DataSource</res-type>
      <res-auth>Container</res-auth>
      <res-sharing-scope>Shareable</res-sharing-scope>
    </resource-ref>
```

```
-->
```

```
<!ELEMENT resource-ref (description?, res-ref-name, res-type,
res-auth, res-sharing-scope?)>
```

```
<!--
```

The result-type-mapping element is used in the query element to specify whether an abstract schema type returned by a query for a select method is to be mapped to an EJBLocalObject or EJBObject type.

The value of the result-type-mapping element, if specified, must be one of the two following:

```
<result-type-mapping>Remote</result-type-mapping>
<result-type-mapping>Local</result-type-mapping>
```

The default value is Local.

```
Used in: query
```

```
-->
```

```

<!--
The role-link element is a reference to a defined security role. The
role-link element must contain the name of one of the security roles
defined in the security-role elements.

Used in: security-role-ref
-->
<!--
The role-name element contains the name of a security role.

The name must conform to the lexical rules for an NMTOKEN.

Used in: method-permission, run-as, security-role, and secu-
rity-role-ref
-->
<!--
The run-as element specifies the run-as identity to be used for the
execution of the methods of an enterprise bean. It contains an
optional description, and the name of a security role.

Used in: security-identity
-->
<!--
The security-identity element specifies whether the caller's security
identity is to be used for the execution of the methods of the enter-
prise bean or whether a specific run-as identity is to be used. It
contains an optional description and a specification of the security
identity to be used.

Used in: session, entity, message-driven
-->
<!--
The security-role element contains the definition of a security role.
The definition consists of an optional description of the security
role, and the security role name.

Used in: assembly-descriptor

Example:
    <security-role>
      <description>
        This role includes all employees who are authorized
        to access the employee service application.
      </description>
      <role-name>employee</role-name>
    </security-role>
-->

```


<!--

The security-role-ref element contains the declaration of a security role reference in the enterprise bean's code. The declaration consists of an optional description, the security role name used in the code, and an optional link to a defined security role.

The value of the role-name element must be the String used as the parameter to the EJBContext.isCallerInRole(String roleName) method.

The value of the role-link element must be the name of one of the security roles defined in the security-role elements.

Used in: entity and session

-->

<!ELEMENT security-role-ref (description?, role-name, role-link?)>

<!--

The session-type element describes whether the session bean is a stateful session or stateless session.

The session-type element must be one of the two following:

```
<session-type>Stateful</session-type>
<session-type>Stateless</session-type>
```

-->

<!ELEMENT session-type (#PCDATA)>

<!--

The session element declares an session bean. The declaration consists of: an optional description; an optional display name; an optional small icon file name; an optional large icon file name; a name assigned to the enterprise bean in the deployment description; the names of the session bean's remote home and remote interfaces, if any; the names of the session bean's local home and local interfaces, if any; the session bean's implementation class; the session bean's state management type; the session bean's transaction management type; an optional declaration of the bean's environment entries; an optional declaration of the bean's EJB references; an optional declaration of the bean's local EJB references; an optional declaration of the security role references; an optional declaration of the security identity to be used for the execution of the bean's methods; an optional declaration of the bean's resource manager connection factory references; and an optional declaration of the bean's resource environment references.

The elements that are optional are "optional" in the sense that they are omitted when if lists represented by them are empty.

Either both the local-home and the local elements or both the home and the remote elements must be specified for the session bean.

Used in: enterprise-beans

-->

<!ELEMENT session (description?, display-name?, small-icon?, large-icon?, ejb-name, home?, remote?, local-home?, local?, ejb-class, session-type, transaction-type, env-entry*, ejb-ref*, ejb-local-ref*, security-role-ref*, security-identity?, resource-ref*, resource-env-ref*)>

<!--
 The small-icon element contains the name of a file containing a small
 (16 x 16) icon image. The file name is relative path within the
 ejb-jar file.

The image must be either in the JPEG or GIF format.
 The icon can be used by tools.

Example:

```
<small-icon>employee-service-icon16x16.jpg</small-icon>
```

-->

<!ELEMENT small-icon (#PCDATA)>

<!--

The subscription-durability element specifies whether a JMS topic
 subscription is intended to be durable or nondurable.

The subscription-durability element must be one of the two following:

```
<subscription-durability>Durable</subscription-durability>
<subscription-durability>NonDurable</subscription-durability>
```

Used in: message-driven-destination

-->

<!ELEMENT subscription-durability (#PCDATA)>

<!--

The transaction-type element specifies an enterprise bean's transac-
 tion management type.

The transaction-type element must be one of the two following:

```
<transaction-type>Bean</transaction-type>
<transaction-type>Container</transaction-type>
```

Used in: session and message-driven

-->

<!ELEMENT transaction-type (#PCDATA)>

<!--

The trans-attribute element specifies how the container must manage
 the transaction boundaries when delegating a method invocation to an
 enterprise bean's business method.

The value of trans-attribute must be one of the following:

```
<trans-attribute>NotSupported</trans-attribute>
<trans-attribute>Supports</trans-attribute>
<trans-attribute>Required</trans-attribute>
<trans-attribute>RequiresNew</trans-attribute>
<trans-attribute>Mandatory</trans-attribute>
<trans-attribute>Never</trans-attribute>
```

Used in: container-transaction

-->

<!ELEMENT trans-attribute (#PCDATA)>

<!--

The unchecked element specifies that a method is not checked for
 authorization by the container prior to invocation of the method.

Used in: method-permission

-->

<!ELEMENT unchecked EMPTY>

<!--

The use-caller-identity element specifies that the caller's security identity be used as the security identity for the execution of the enterprise bean's methods.

Used in: security-identity

-->

<!ELEMENT use-caller-identity EMPTY>

<!--

The ID mechanism is to allow tools that produce additional deployment information (i.e., information beyond the standard EJB deployment descriptor information) to store the non-standard information in a separate file, and easily refer from these tools-specific files to the information in the standard deployment descriptor.

The EJB architecture does not allow the tools to add the non-standard information into the EJB deployment descriptor.

-->

<!ATTLIST abstract-schema-name id ID #IMPLIED>

<!ATTLIST acknowledge-mode id ID #IMPLIED>

<!ATTLIST assembly-descriptor id ID #IMPLIED>

<!ATTLIST cascade-delete id ID #IMPLIED>

<!ATTLIST cmp-field id ID #IMPLIED>

<!ATTLIST cmp-version id ID #IMPLIED>

<!ATTLIST cmr-field id ID #IMPLIED>

<!ATTLIST cmr-field-name id ID #IMPLIED>

<!ATTLIST cmr-field-type id ID #IMPLIED>

<!ATTLIST container-transaction id ID #IMPLIED>

<!ATTLIST description id ID #IMPLIED>

<!ATTLIST destination-type id ID #IMPLIED>

<!ATTLIST display-name id ID #IMPLIED>

<!ATTLIST ejb-class id ID #IMPLIED>

<!ATTLIST ejb-client-jar id ID #IMPLIED>

<!ATTLIST ejb-jar id ID #IMPLIED>

<!ATTLIST ejb-link id ID #IMPLIED>

<!ATTLIST ejb-local-ref id ID #IMPLIED>

<!ATTLIST ejb-name id ID #IMPLIED>

<!ATTLIST ejb-ql id ID #IMPLIED>

<!ATTLIST ejb-ref id ID #IMPLIED>

<!ATTLIST ejb-ref-name id ID #IMPLIED>

<!ATTLIST ejb-ref-type id ID #IMPLIED>

<!ATTLIST ejb-relation id ID #IMPLIED>

<!ATTLIST ejb-relation-name id ID #IMPLIED>

<!ATTLIST ejb-relationship-role id ID #IMPLIED>

<!ATTLIST ejb-relationship-role-name id ID #IMPLIED>

<!ATTLIST enterprise-beans id ID #IMPLIED>

<!ATTLIST entity id ID #IMPLIED>

<!ATTLIST env-entry id ID #IMPLIED>

<!ATTLIST env-entry-name id ID #IMPLIED>

<!ATTLIST env-entry-type id ID #IMPLIED>

<!ATTLIST env-entry-value id ID #IMPLIED>

<!ATTLIST exclude-list id ID #IMPLIED>

<!ATTLIST field-name id ID #IMPLIED>

<!ATTLIST home id ID #IMPLIED>

```
<!ATTLIST large-icon id ID #IMPLIED>
<!ATTLIST local id ID #IMPLIED>
<!ATTLIST local-home id ID #IMPLIED>
<!ATTLIST message-driven id ID #IMPLIED>
<!ATTLIST message-driven-destination id ID #IMPLIED>
<!ATTLIST message-selector id ID #IMPLIED>
<!ATTLIST method id ID #IMPLIED>
<!ATTLIST method-intf id ID #IMPLIED>
<!ATTLIST method-name id ID #IMPLIED>
<!ATTLIST method-param id ID #IMPLIED>
<!ATTLIST method-params id ID #IMPLIED>
<!ATTLIST method-permission id ID #IMPLIED>
<!ATTLIST multiplicity id ID #IMPLIED>
<!ATTLIST persistence-type id ID #IMPLIED>
<!ATTLIST prim-key-class id ID #IMPLIED>
<!ATTLIST primkey-field id ID #IMPLIED>
<!ATTLIST query id ID #IMPLIED>
<!ATTLIST query-method id ID #IMPLIED>
<!ATTLIST reentrant id ID #IMPLIED>
<!ATTLIST relationships id ID #IMPLIED>
<!ATTLIST relationship-role-source id ID #IMPLIED>
<!ATTLIST remote id ID #IMPLIED>
<!ATTLIST res-auth id ID #IMPLIED>
<!ATTLIST res-ref-name id ID #IMPLIED>
<!ATTLIST res-sharing-scope id ID #IMPLIED>
<!ATTLIST res-type id ID #IMPLIED>
<!ATTLIST resource-env-ref id ID #IMPLIED>
<!ATTLIST resource-env-ref-name id ID #IMPLIED>
<!ATTLIST resource-env-ref-type id ID #IMPLIED>
<!ATTLIST resource-ref id ID #IMPLIED>
<!ATTLIST result-type-mapping id ID #IMPLIED>
<!ATTLIST role-link id ID #IMPLIED>
<!ATTLIST role-name id ID #IMPLIED>
<!ATTLIST run-as id ID #IMPLIED>
<!ATTLIST security-identity id ID #IMPLIED>
<!ATTLIST security-role id ID #IMPLIED>
<!ATTLIST security-role-ref id ID #IMPLIED>
<!ATTLIST session-type id ID #IMPLIED>
<!ATTLIST session id ID #IMPLIED>
<!ATTLIST small-icon id ID #IMPLIED>
<!ATTLIST subscription-durability id ID #IMPLIED>
<!ATTLIST transaction-type id ID #IMPLIED>
<!ATTLIST trans-attribute id ID #IMPLIED>
<!ATTLIST unchecked id ID #IMPLIED>
<!ATTLIST use-caller-identity id ID #IMPLIED>
```

EJB 1.1 Deployment Descriptor

This appendix defines the EJB 1.1 deployment descriptor. All EJB 2.1 and EJB 2.0 compliant implementations must support EJB 1.1 deployment descriptors. Section C.1 provides an overview of the deployment descriptor. Sections C.2 through C.4 describe the information in the deployment descriptor from the perspective of the EJB roles responsible for providing the information. Section C.5 defines the deployment descriptor's XML DTD. Section C.6 provides a complete example of a deployment descriptor of an assembled application.

C.1 Overview

The deployment descriptor is part of the contract between the ejb-jar file producer and consumer. This contract covers both the passing of enterprise beans from the Bean Provider to Application Assembler, and from the Application Assembler to the Deployer.

An ejb-jar file produced by the Bean Provider contains one or more enterprise beans and typically does not contain application assembly instructions. An ejb-jar file produced by an Application Assembler contains one or more enterprise beans, plus application assembly information describing how the enterprise beans are combined into a single application deployment unit.

The J2EE specification defines how enterprise beans and other application components contained in multiple ejb-jar files can be assembled into an application.

The role of the deployment descriptor is to capture the declarative information (i.e information that is not included directly in the enterprise beans' code) that is intended for the consumer of the ejb-jar file.

There are two basic kinds of information in the deployment descriptor:

- *Enterprise beans' structural information.* Structural information describes the structure of an enterprise bean and declares an enterprise bean's external dependencies. Providing structural information in the deployment descriptor is mandatory for the ejb-jar file producer. The structural information cannot, in general, be changed because doing so could break the enterprise bean's function.
- *Application assembly information.* Application assembly information describes how the enterprise bean (or beans) in the ejb-jar file is composed into a larger application deployment unit. Providing assembly information in the deployment descriptor is optional for the ejb-jar file producer. Assembly level information can be changed without breaking the enterprise bean's function, although doing so may alter the behavior of an assembled application.

C.2 Bean Provider's Responsibilities

The Bean Provider is responsible for providing the structural information for each enterprise bean in the deployment descriptor.

The Bean Provider must use the `enterprise-beans` element to list all the enterprise beans in the ejb-jar file.

The Bean Provider must provide the following information for each enterprise bean:

- **Enterprise bean's name.** The Bean Provider must assign a logical name to each enterprise bean in the ejb-jar file. There is no architected relationship between this name, and the JNDI API name that the Deployer will assign to the enterprise bean. The Bean Provider specifies the enterprise bean's name in the `ejb-name` element.
- **Enterprise bean's class.** The Bean Provider must specify the fully-qualified name of the Java class that implements the enterprise bean's business methods. The Bean Provider specifies the enterprise bean's class name in the `ejb-class` element.
- **Enterprise bean's home interfaces.** The Bean Provider must specify the fully-qualified name of the enterprise bean's home interface in the `home` element.
- **Enterprise bean's remote interfaces.** The Bean Provider must specify the fully-qualified name of the enterprise bean's remote interface in the `remote` element.
- **Enterprise bean's type.** The enterprise beans types are session and entity. The Bean Provider must use the appropriate `session` or `entity` element to declare the enterprise bean's structural information.
- **Re-entrancy indication.** The Bean Provider must specify whether an entity bean is re-entrant or not. Session beans are never re-entrant.

- **Session bean's state management type.** If the enterprise bean is a session bean, the Bean Provider must use the `session-type` element to declare whether the session bean is stateful or stateless.
- **Session bean's transaction demarcation type.** If the enterprise bean is a session bean, the Bean Provider must use the `transaction-type` element to declare whether transaction demarcation is performed by the enterprise bean or by the container.
- **Entity bean's persistence management.** If the enterprise bean is an entity bean, the Bean Provider must use the `persistence-type` element to declare whether persistence management is performed by the enterprise bean or by the container.
- **Entity bean's primary key class.** If the enterprise bean is an entity bean, the Bean Provider specifies the fully-qualified name of the entity bean's primary key class in the `prim-key-class` element. The Bean Provider *must* specify the primary key class for an entity bean with bean-managed persistence, and *may* (but is not required to) specify the primary key class for an entity bean with container-managed persistence.
- **Container-managed fields.** If the enterprise bean is an entity bean with container-managed persistence, the Bean Provider must specify the container-managed fields using the `cmp-field` elements.
- **Environment entries.** The Bean Provider must declare all the enterprise bean's environment entries as specified in Subsection 20.2.1.
- **Resource manager connection factory references.** The Bean Provider must declare all the enterprise bean's resource manager connection factory references as specified in Subsection 20.5.1.
- **EJB references.** The Bean Provider must declare all the enterprise bean's references to the homes of other enterprise beans as specified in Subsection 20.3.1.
- **Security role references.** The Bean Provider must declare all the enterprise bean's references to security roles as specified in Subsection 21.2.5.3.

The deployment descriptor produced by the Bean Provider must be well formed in the XML sense, and valid with respect to the DTD in Section C.5. The content of the deployment descriptor must conform to the semantics rules specified in the DTD comments and elsewhere in this specification. The deployment descriptor must refer to the DTD using the following statement:

```
<!DOCTYPE ejb-jar PUBLIC "-//Sun Microsystems, Inc.//DTD Enterprise  
JavaBeans 1.1//EN" "http://java.sun.com/j2ee/dtds/ejb-jar_1_1.dtd">
```

C.3 Application Assembler's Responsibility

The Application Assembler assembles enterprise beans into a single deployment unit. The Application Assembler's input is one or more ejb-jar files provided by one or more Bean Providers, and the output is also one or more ejb-jar files. The Application Assembler can combine multiple input ejb-jar files into a single output ejb-jar file, or split an input ejb-jar file into multiple output ejb-jar files. Each output ejb-jar file is either a deployment unit intended for the Deployer, or a partially assembled application that is intended for another Application Assembler.

The Bean Provider and Application Assembler may be the same person or organization. In such a case, the person or organization performs the responsibilities described both in this and the previous sections.

The Application Assembler may modify the following information that was specified by the Bean Provider:

- **Enterprise bean's name.** The Application Assembler may change the enterprise bean's name defined in the `ejb-name` element.
- **Values of environment entries.** The Application Assembler may change existing and/or define new values of environment properties.
- **Description fields.** The Application Assembler may change existing or create new `description` elements.

The Application Assembler must not, in general, modify any other information listed in Section C.2 that was provided in the input ejb-jar file.

In addition, the Application Assembler may, but is not required to, specify any of the following *application assembly* information:

- **Binding of enterprise bean references.** The Application Assembler may link an enterprise bean reference to another enterprise bean in the ejb-jar file. The Application Assembler creates the link by adding the `ejb-link` element to the referencing bean.
- **Security roles.** The Application Assembler may define one or more security roles. The security roles define the *recommended* security roles for the clients of the enterprise beans. The Application Assembler defines the security roles using the `security-role` elements.
- **Method permissions.** The Application Assembler may define method permissions. Method permission is a binary relation between the security roles and the methods of the remote and home interfaces of the enterprise beans. The Application Assembler defines method permissions using the `method-permission` elements.
- **Linking of security role references.** If the Application Assembler defines security roles in the deployment descriptor, the Application Assembler must link the security role references declared by the Bean Provider to the security roles. The Application Assembler defines these links using the `role-link` element.

- **Transaction attributes.** The Application Assembler may define the value of the transaction attributes for the methods of the remote and home interfaces of the enterprise beans that require container-managed transaction demarcation. All entity beans and the session beans declared by the Bean Provider as transaction-type Container require container-managed transaction demarcation. The Application Assembler uses the `container-transaction` elements to declare the transaction attributes.

If an input `ejb-jar` file contains application assembly information, the Application Assembler is allowed to change the application assembly information supplied in the input `ejb-jar` file. (This could happen when the input `ejb-jar` file was produced by another Application Assembler.)

The deployment descriptor produced by the Bean Provider must be well formed in the XML sense, and valid with respect to the DTD in Section C.5. The content of the deployment descriptor must conform to the semantic rules specified in the DTD comments and elsewhere in this specification. The deployment descriptor must refer to the DTD using the following statement:

```
<!DOCTYPE ejb-jar PUBLIC "-//Sun Microsystems, Inc.//DTD Enterprise  
JavaBeans 1.1//EN" "http://java.sun.com/j2ee/dtds/ejb-jar_1_1.dtd">
```

C.4 Container Provider's Responsibilities

The Container Provider provides tools that read and import the information contained in the XML deployment descriptor.

C.5 Deployment Descriptor DTD

This section defines the XML DTD for the EJB 1.1 deployment descriptor. The comments in the DTD specify additional requirements for the syntax and semantics that cannot be easily expressed by the DTD mechanism.

The content of the XML elements is in general case sensitive. This means, for example, that

```
<reentrant>True</reentrant>
```

must be used, rather than:

```
<reentrant>true</reentrant>.
```

All valid ejb-jar deployment descriptors must contain the following DOCTYPE declaration:

```
<!DOCTYPE ejb-jar PUBLIC "-//Sun Microsystems, Inc.//DTD Enterprise
JavaBeans 1.1//EN" "http://java.sun.com/j2ee/dtds/ejb-jar_1_1.dtd">
```

```
<!--
This is the XML DTD for the EJB 1.1 deployment descriptor.
-->
```

```
<!--
The assembly-descriptor element contains application-assembly infor-
mation.
```

The application-assembly information consists of the following parts: the definition of security roles, the definition of method permissions, and the definition of transaction attributes for enterprise beans with container-managed transaction demarcation.

All the parts are optional in the sense that they are omitted if the lists represented by them are empty.

Providing an assembly-descriptor in the deployment descriptor is optional for the ejb-jar file producer.

```
Used in: ejb-jar
-->
<!--ELEMENT assembly-descriptor (security-role*, method-permission*,
        container-transaction*)-->
```

```
<!--
The cmp-field element describes a container-managed field. The field
element includes an optional description of the field, and the name of
the field.
```

```
Used in: entity
-->
<!--ELEMENT cmp-field (description?, field-name)-->
```

```
<!--
The container-transaction element specifies how the container must
manage transaction scopes for the enterprise bean's method invoca-
tions. The element consists of an optional description, a list of
method elements, and a transaction attribute. The transaction
attribute is to be applied to all the specified methods.
```

```
Used in: assembly-descriptor
-->
<!--ELEMENT container-transaction (description?, method+,
        trans-attribute)-->
```

```
<!--
The description element is used by the ejb-jar file producer to pro-
vide text describing the parent element.
```

The description element should include any information that the ejb-jar file producer wants to provide to the consumer of the ejb-jar

file (i.e. to the Deployer). Typically, the tools used by the ejb-jar file consumer will display the description when processing the parent element.

Used in: cmp-field, container-transaction, ejb-jar, entity, env-entry, ejb-ref, method, method-permission, resource-ref, security-role, security-role-ref, and session.

```
-->
<!--ELEMENT description (#PCDATA)-->
```

```
<!--
The display-name element contains a short name that is intended to be
display by tools.
```

Used in: ejb-jar, session, and entity

Example:

```
<!--
  <display-name>Employee Self Service</display-name>
-->
<!--ELEMENT display-name (#PCDATA)-->
```

```
<!--
The ejb-class element contains the fully-qualified name of the enter-
prise bean's class.
```

Used in: entity and session

Example:

```
<!--
  <ejb-class>com.wombat.empl.EmployeeServiceBean</ejb-class>
-->
<!--ELEMENT ejb-class (#PCDATA)-->
```

```
<!--
The optional ejb-client-jar element specifies a JAR file that con-
tains the class files necessary for a client program to access the
enterprise beans in the ejb-jar file. The Deployer should make the
ejb-client JAR file accessible to the client's class-loader.
```

Used in: ejb-jar

Example:

```
<!--
  <ejb-client-jar>employee_service_client.jar</ejb-client-jar>
-->
<!--ELEMENT ejb-client-jar (#PCDATA)-->
```

```
<!--
The ejb-jar element is the root element of the EJB deployment descrip-
tor. It contains an optional description of the ejb-jar file, optional
display name, optional small icon file name, optional large icon file
name, mandatory structural information about all included enterprise
beans, optional application-assembly descriptor, and an optional name
of an ejb-client-jar file for the ejb-jar.
```

```
-->
<!--ELEMENT ejb-jar (description?, display-name?, small-icon?,
  large-icon?, enterprise-beans, assembly-descriptor?,
  ejb-client-jar?)-->
```

```
<!--
The ejb-link element is used in the ejb-ref element to specify that an
```

EJB reference is linked to another enterprise bean in the ejb-jar file.

The value of the ejb-link element must be the ejb-name of an enterprise bean in the same ejb-jar file, or in another ejb-jar file in the same J2EE application unit.

Used in: ejb-ref

Example:

```
<ejb-link>EmployeeRecord</ejb-link>
-->
<!--ELEMENT ejb-link (#PCDATA)>
```

<!--

The ejb-name element specifies an enterprise bean's name. This name is assigned by the ejb-jar file producer to name the enterprise bean in the ejb-jar file's deployment descriptor. The name must be unique among the names of the enterprise beans in the same ejb-jar file.

The enterprise bean code does not depend on the name; therefore the name can be changed during the application-assembly process without breaking the enterprise bean's function.

There is no architected relationship between the ejb-name in the deployment descriptor and the JNDI name that the Deployer will assign to the enterprise bean's home.

The name must conform to the lexical rules for an NMTOKEN.

Used in: entity, method, and session

Example:

```
<ejb-name>EmployeeService</ejb-name>
-->
<!--ELEMENT ejb-name (#PCDATA)>
```

<!--

The ejb-ref element is used for the declaration of a reference to another enterprise bean's home. The declaration consists of an optional description; the EJB reference name used in the code of the referencing enterprise bean; the expected type of the referenced enterprise bean; the expected home and remote interfaces of the referenced enterprise bean; and an optional ejb-link information.

The optional ejb-link element is used to specify the referenced enterprise bean. It is used typically in ejb-jar files that contain an assembled application.

Used in: entity and session

```
-->
<!--ELEMENT ejb-ref (description?, ejb-ref-name, ejb-ref-type, home,
remote, ejb-link?)>
```

<!--

The ejb-ref-name element contains the name of an EJB reference. The EJB reference is an entry in the enterprise bean's environment.

It is recommended that name is prefixed with "ejb/".

Used in: ejb-ref

Example:

```
<ejb-ref-name>ejb/Payroll</ejb-ref-name>
```

```
-->
```

```
<!--ELEMENT ejb-ref-name (#PCDATA)>
```

```
<!--
```

The ejb-ref-type element contains the expected type of the referenced enterprise bean.

The ejb-ref-type element must be one of the following:

```
<ejb-ref-type>Entity</ejb-ref-type>
```

```
<ejb-ref-type>Session</ejb-ref-type>
```

Used in: ejb-ref

```
-->
```

```
<!--ELEMENT ejb-ref-type (#PCDATA)>
```

```
<!--
```

The enterprise-beans element contains the declarations of one or more enterprise beans.

```
-->
```

```
<!--ELEMENT enterprise-beans (session | entity)+>
```

```
<!--
```

The entity element declares an entity bean. The declaration consists of: an optional description; optional display name; optional small icon file name; optional large icon file name; a name assigned to the enterprise bean in the deployment descriptor; the names of the entity bean's home and remote interfaces; the entity bean's implementation class; the entity bean's persistence management type; the entity bean's primary key class name; an indication of the entity bean's reentrancy; an optional list of container-managed fields; an optional specification of the primary key field; an optional declaration of the bean's environment entries; an optional declaration of the bean's EJB references; an optional declaration of the security role references; and an optional declaration of the bean's resource manager connection factory references.

The optional primkey-field may be present in the descriptor if the entity's persistency-type is Container.

The other elements that are optional are "optional" in the sense that they are omitted if the lists represented by them are empty.

At least one cmp-field element must be present in the descriptor if the entity's persistency-type is Container, and none must not be present if the entity's persistence-type is Bean.

Used in: enterprise-beans

```
-->
```

```
<!--ELEMENT entity (description?, display-name?, small-icon?,
    large-icon?, ejb-name, home, remote, ejb-class,
    persistence-type, prim-key-class, reentrant,
    cmp-field*, primkey-field?, env-entry*,
    ejb-ref*, security-role-ref*, resource-ref*)>
```

<!--
 The env-entry element contains the declaration of an enterprise bean's environment entries. The declaration consists of an optional description, the name of the environment entry, and an optional value.

Used in: entity and session

-->

<!ELEMENT env-entry (description?, env-entry-name, env-entry-type, env-entry-value?)>

<!--

The env-entry-name element contains the name of an enterprise bean's environment entry.

Used in: env-entry

Example:

`<env-entry-name>minAmount</env-entry-name>`

-->

<!ELEMENT env-entry-name (#PCDATA)>

<!--

The env-entry-type element contains the fully-qualified Java type of the environment entry value that is expected by the enterprise bean's code.

The following are the legal values of env-entry-type: java.lang.Boolean, java.lang.String, java.lang.Integer, java.lang.Double, java.lang.Byte, java.lang.Short, java.lang.Long, and java.lang.Float.

Used in: env-entry

Example:

`<env-entry-type>java.lang.Boolean</env-entry-type>`

-->

<!ELEMENT env-entry-type (#PCDATA)>

<!--

The env-entry-value element contains the value of an enterprise bean's environment entry.

Used in: env-entry

Example:

`<env-entry-value>100.00</env-entry-value>`

-->

<!ELEMENT env-entry-value (#PCDATA)>

<!--

The field-name element specifies the name of a container-managed field. The name must be a public field of the enterprise bean class or one of its superclasses.

Used in: cmp-field

Example:

`<field-name>firstName</field-name>`

-->

<!ELEMENT field-name (#PCDATA)>

<!--

The home element contains the fully-qualified name of the enterprise bean's home interface.

Used in: ejb-ref, entity, and session

Example:

```
<home>com.aardvark.payroll.PayrollHome</home>
```

-->

<!ELEMENT home (#PCDATA)>

<!--

The large-icon element contains the name of a file containing a large (32 x 32) icon image. The file name is relative path within the ejb-jar file.

The image must be either in the JPEG or GIF format, and the file name must end with the suffix ".jpg" or ".gif" respectively. The icon can be used by tools.

Example:

```
<large-icon>employee-service-icon32x32.jpg</large-icon>
```

-->

<!ELEMENT large-icon (#PCDATA)>

<!--

The method element is used to denote a method of an enterprise bean's home or remote interface, or a set of methods. The ejb-name element must be the name of one of the enterprise beans in declared in the deployment descriptor; the optional method-intf element allows to distinguish between a method with the same signature that is defined in both the home and remote interface; the method-name element specifies the method name; and the optional method-params elements identify a single method among multiple methods with an overloaded method name.

There are three possible styles of the method element syntax:

1.

```
<method>
  <ejb-name>EJBNAME</ejb-name>
  <method-name>*</method-name>
</method>
```

This style is used to refer to all the methods of the specified enterprise bean's home and remote interfaces.

2.

```
<method>
  <ejb-name>EJBNAME</ejb-name>
  <method-name>METHOD</method-name>
</method>>
```

This style is used to refer to the specified method of the specified enterprise bean. If there are multiple methods with the same overloaded name, the element of this style refers to all the methods with the overloaded name.

```

3.  <method>
    <ejb-name>EJBNAME</ejb-name>
    <method-name>METHOD</method-name>
    <method-params>
      <method-param>PARAM-1</method-param>
      <method-param>PARAM-2</method-param>
      ...
      <method-param>PARAM-n</method-param>
    </method-params>
  </method>

```

This style is used to refer to a single method within a set of methods with an overloaded name. PARAM-1 through PARAM-n are the fully-qualified Java types of the method's input parameters (if the method has no input arguments, the method-params element contains no method-param elements). Arrays are specified by the array element's type, followed by one or more pair of square brackets (e.g. `int[][]`).

Used in: method-permission and container-transaction

Examples:

Style 1: The following method element refers to all the methods of the `EmployeeService` bean's home and remote interfaces:

```

<method>
  <ejb-name>EmployeeService</ejb-name>
  <method-name>*</method-name>
</method>

```

Style 2: The following method element refers to all the `create` methods of the `EmployeeService` bean's home interface:

```

<method>
  <ejb-name>EmployeeService</ejb-name>
  <method-name>create</method-name>
</method>

```

Style 3: The following method element refers to the `create(String firstName, String LastName)` method of the `EmployeeService` bean's home interface.

```

<method>
  <ejb-name>EmployeeService</ejb-name>
  <method-name>create</method-name>
  <method-params>
    <method-param>java.lang.String</method-param>
    <method-param>java.lang.String</method-param>
  </method-params>
</method>

```

The following example illustrates a Style 3 element with more complex parameter types. The method `foobar(char s, int i, int[] iar, mypackage.MyClass mycl,`

mypackage.MyClass[][] myclaar)
would be specified as:

```
<method>
  <ejb-name>EmployeeService</ejb-name>
  <method-name>foobar</method-name>
  <method-params>
    <method-param>char</method-param>
    <method-param>int</method-param>
    <method-param>int[]</method-param>
    <method-param>mypackage.MyClass</method-param>
    <method-param>mypackage.MyClass[][]</method-param>
  </method-params>
</method>
```

The optional method-intf element can be used when it becomes necessary to differentiate between a method defined in the home interface and a method with the same name and signature that is defined in the remote interface.

For example, the method element

```
<method>
  <ejb-name>EmployeeService</ejb-name>
  <method-intf>Remote</method-intf>
  <method-name>create</method-name>
  <method-params>
    <method-param>java.lang.String</method-param>
    <method-param>java.lang.String</method-param>
  </method-params>
</method>
```

can be used to differentiate the create(String, String) method defined in the remote interface from the create(String, String) method defined in the home interface, which would be defined as

```
<method>
  <ejb-name>EmployeeService</ejb-name>
  <method-intf>Home</method-intf>
  <method-name>create</method-name>
  <method-params>
    <method-param>java.lang.String</method-param>
    <method-param>java.lang.String</method-param>
  </method-params>
</method>
```

The method-intf element can be used with all three Styles of the method element usage. For example, the following method element example could be used to refer to all the methods of the EmployeeService bean's home interface.

```
<method>
  <ejb-name>EmployeeService</ejb-name>
  <method-intf>Home</method-intf>
  <method-name>*</method-name>
</method>
```

-->

```
<!--ELEMENT method (description?, ejb-name, method-intf?, method-name,
method-params?)>
```

```
<!--
```

The method-intf element allows a method element to differentiate between the methods with the same name and signature that are defined in both the remote and home interfaces.

The method-intf element must be one of the following:

```
<method-intf>Home</method-intf>
<method-intf>Remote</method-intf>
```

Used in: method

```
-->
```

```
<!--ELEMENT method-intf (#PCDATA)>
```

```
<!--
```

The method-name element contains a name of an enterprise bean method, or the asterisk (*) character. The asterisk is used when the element denotes all the methods of an enterprise bean's remote and home interfaces.

Used in: method

```
-->
```

```
<!--ELEMENT method-name (#PCDATA)>
```

```
<!--
```

The method-param element contains the fully-qualified Java type name of a method parameter.

Used in: method-params

```
-->
```

```
<!--ELEMENT method-param (#PCDATA)>
```

```
<!--
```

The method-params element contains a list of the fully-qualified Java type names of the method parameters.

Used in: method

```
-->
```

```
<!--ELEMENT method-params (method-param*)>
```

```
<!--
```

The method-permission element specifies that one or more security roles are allowed to invoke one or more enterprise bean methods. The method-permission element consists of an optional description, a list of security role names, and a list of method elements.

The security roles used in the method-permission element must be defined in the security-role element of the deployment descriptor, and the methods must be methods defined in the enterprise bean's remote and/or home interfaces.

Used in: assembly-descriptor

```
-->
```

```
<!--ELEMENT method-permission (description?, role-name+, method+)>
```

```
<!--
```

The persistence-type element specifies an entity bean's persistence

management type.

The persistence-type element must be one of the two following:

```
<persistence-type>Bean</persistence-type>
<persistence-type>Container</persistence-type>
```

Used in: entity

-->

<!ELEMENT persistence-type (#PCDATA)>

<!--

The prim-key-class element contains the fully-qualified name of an entity bean's primary key class.

If the definition of the primary key class is deferred to deployment time, the prim-key-class element should specify java.lang.Object.

Used in: entity

Examples:

```
<prim-key-class>java.lang.String</prim-key-class>
<prim-key-class>com.wombat.empl.EmployeeID</prim-key-class>
<prim-key-class>java.lang.Object</prim-key-class>
```

-->

<!ELEMENT prim-key-class (#PCDATA)>

<!--

The primkey-field element is used to specify the name of the primary key field for an entity with container-managed persistence.

The primkey-field must be one of the fields declared in the cmp-field element, and the type of the field must be the same as the primary key type.

The primkey-field element is not used if the primary key maps to multiple container-managed fields (i.e., the key is a compound key). In this case, the fields of the primary key class must be public, and their names must correspond to the field names of the entity bean class that comprise the key.

Used in: entity

Example:

```
<primkey-field>EmployeeId</primkey-field>
```

-->

<!ELEMENT primkey-field (#PCDATA)>

<!--

The reentrant element specifies whether an entity bean is reentrant or not.

The reentrant element must be one of the two following:

```
<reentrant>True</reentrant>
<reentrant>False</reentrant>
```

Used in: entity

-->

<!ELEMENT reentrant (#PCDATA)>

<!--
The remote element contains the fully-qualified name of the enterprise bean's remote interface.

Used in: ejb-ref, entity, and session

Example:

```
<remote>com.wombat.empl.EmployeeService</remote>
-->
<!--ELEMENT remote (#PCDATA)>
```

<!--
The res-auth element specifies whether the enterprise bean code signs on programmatically to the resource manager, or whether the Container will sign on to the resource manager on behalf of the bean. In the latter case, the Container uses information that is supplied by the Deployer.

The value of this element must be one of the two following:

```
<res-auth>Application</res-auth>
<res-auth>Container</res-auth>
-->
<!--ELEMENT res-auth (#PCDATA)>
```

<!--
The res-ref-name element specifies the name of a resource manager connection factory reference.

Used in: resource-ref

```
-->
<!--ELEMENT res-ref-name (#PCDATA)>
```

<!--
The res-type element specifies the type of the data source. The type is specified by the Java interface (or class) expected to be implemented by the data source.

Used in: resource-ref

```
-->
<!--ELEMENT res-type (#PCDATA)>
```

<!--
The resource-ref element contains a declaration of enterprise bean's reference to an external resource. It consists of an optional description, the resource manager connection factory reference name, the indication of the resource manager connection factory type expected by the enterprise bean code, and the type of authentication (bean or container).

Used in: entity and session

Example:

```
<resource-ref>
  <res-ref-name>jdbc/EmployeeAppDB</res-ref-name>
  <res-type>javax.sql.DataSource</res-type>
  <res-auth>Container</res-auth>
</resource-ref>
-->
```

```
<!--ELEMENT resource-ref (description?, res-ref-name, res-type,
res-auth)>
```

```
<!--
The role-link element is used to link a security role reference to a
defined security role. The role-link element must contain the name of
one of the security roles defined in the security-role elements.
```

```
Used in: security-role-ref
```

```
-->
```

```
<!--ELEMENT role-link (#PCDATA)>
```

```
<!--
```

```
The role-name element contains the name of a security role.
```

```
The name must conform to the lexical rules for an NMTOKEN.
```

```
Used in: method-permission, security-role, and security-role-ref
```

```
-->
```

```
<!--ELEMENT role-name (#PCDATA)>
```

```
<!--
```

```
The security-role element contains the definition of a security role.
The definition consists of an optional description of the security
role, and the security role name.
```

```
Used in: assembly-descriptor
```

```
Example:
```

```
    <security-role>
      <description>
        This role includes all employees who are authorized
        to access the employee service application.
      </description>
      <role-name>employee</role-name>
    </security-role>
```

```
-->
```

```
<!--ELEMENT security-role (description?, role-name)>
```

```
<!--
```

```
The security-role-ref element contains the declaration of a security
role reference in the enterprise bean's code. The declaration con-
sists of an optional description, the security role name used in the
code, and an optional link to a defined security role.
```

```
The value of the role-name element must be the String used as the
parameter to the EJBContext.isCallerInRole(String roleName) method.
```

```
The value of the role-link element must be the name of one of the
security roles defined in the security-role elements.
```

```
Used in: entity and session
```

```
-->
```

```
<!--ELEMENT security-role-ref (description?, role-name, role-link?)>
```

```
<!--
```

```
The session-type element describes whether the session bean is a
stateful session, or stateless session.
```

The session-type element must be one of the two following:

```
<session-type>Stateful</session-type>
<session-type>Stateless</session-type>
```

```
-->
```

```
<!--ELEMENT session-type (#PCDATA)-->
```

```
<!--
```

The session element declares an session bean. The declaration consists of: an optional description; optional display name; optional small icon file name; optional large icon file name; a name assigned to the enterprise bean in the deployment description; the names of the session bean's home and remote interfaces; the session bean's implementation class; the session bean's state management type; the session bean's transaction management type; an optional declaration of the bean's environment entries; an optional declaration of the bean's EJB references; an optional declaration of the security role references; and an optional declaration of the bean's resource manager connection factory references.

The elements that are optional are "optional" in the sense that they are omitted when if lists represented by them are empty.

Used in: enterprise-beans

```
-->
```

```
<!--ELEMENT session (description?, display-name?, small-icon?,
    large-icon?, ejb-name, home, remote, ejb-class,
    session-type, transaction-type, env-entry*,
    ejb-ref*, security-role-ref*, resource-ref*)-->
```

```
<!--
```

The small-icon element contains the name of a file containing a small (16 x 16) icon image. The file name is relative path within the ejb-jar file.

The image must be either in the JPEG or GIF format, and the file name must end with the suffix ".jpg" or ".gif" respectively.

The icon can be used by tools.

Example:

```
<small-icon>employee-service-icon16x16.jpg</small-icon>
```

```
-->
```

```
<!--ELEMENT small-icon (#PCDATA)-->
```

```
<!--
```

The transaction-type element specifies an enterprise bean's transaction management type.

The transaction-type element must be one of the two following:

```
<transaction-type>Bean</transaction-type>
<transaction-type>Container</transaction-type>
```

Used in: session

```
-->
```

```
<!--ELEMENT transaction-type (#PCDATA)-->
```

```
<!--
```

The trans-attribute element specifies how the container must manage

the transaction boundaries when delegating a method invocation to an enterprise bean's business method.

The value of trans-attribute must be one of the following:

```
<trans-attribute>NotSupported</trans-attribute>
<trans-attribute>Supports</trans-attribute>
<trans-attribute>Required</trans-attribute>
<trans-attribute>RequiresNew</trans-attribute>
<trans-attribute>Mandatory</trans-attribute>
<trans-attribute>Never</trans-attribute>
```

Used in: container-transaction

-->

<!ELEMENT trans-attribute (#PCDATA)>

<!--

The ID mechanism is to allow tools that produce additional deployment information (i.e information beyond the standard EJB deployment descriptor information) to store the non-standard information in a separate file, and easily refer from these tools-specific files to the information in the standard deployment descriptor.

The EJB architecture does not allow the tools to add the non-standard information into the EJB deployment descriptor.

-->

```
<!ATTLIST assembly-descriptor id ID #IMPLIED>
<!ATTLIST cmp-field id ID #IMPLIED>
<!ATTLIST container-transaction id ID #IMPLIED>
<!ATTLIST description id ID #IMPLIED>
<!ATTLIST display-name id ID #IMPLIED>
<!ATTLIST ejb-class id ID #IMPLIED>
<!ATTLIST ejb-client-jar id ID #IMPLIED>
<!ATTLIST ejb-jar id ID #IMPLIED>
<!ATTLIST ejb-link id ID #IMPLIED>
<!ATTLIST ejb-name id ID #IMPLIED>
<!ATTLIST ejb-ref id ID #IMPLIED>
<!ATTLIST ejb-ref-name id ID #IMPLIED>
<!ATTLIST ejb-ref-type id ID #IMPLIED>
<!ATTLIST enterprise-beans id ID #IMPLIED>
<!ATTLIST entity id ID #IMPLIED>
<!ATTLIST env-entry id ID #IMPLIED>
<!ATTLIST env-entry-name id ID #IMPLIED>
<!ATTLIST env-entry-type id ID #IMPLIED>
<!ATTLIST env-entry-value id ID #IMPLIED>
<!ATTLIST field-name id ID #IMPLIED>
<!ATTLIST home id ID #IMPLIED>
<!ATTLIST large-icon id ID #IMPLIED>
<!ATTLIST method id ID #IMPLIED>
<!ATTLIST method-intf id ID #IMPLIED>
<!ATTLIST method-name id ID #IMPLIED>
<!ATTLIST method-param id ID #IMPLIED>
<!ATTLIST method-params id ID #IMPLIED>
<!ATTLIST method-permission id ID #IMPLIED>
<!ATTLIST persistence-type id ID #IMPLIED>
<!ATTLIST prim-key-class id ID #IMPLIED>
<!ATTLIST primkey-field id ID #IMPLIED>
<!ATTLIST reentrant id ID #IMPLIED>
<!ATTLIST remote id ID #IMPLIED>
<!ATTLIST res-auth id ID #IMPLIED>
```

```
<!--ATTLIST res-ref-name id ID #IMPLIED>
<!--ATTLIST res-type id ID #IMPLIED>
<!--ATTLIST resource-ref id ID #IMPLIED>
<!--ATTLIST role-link id ID #IMPLIED>
<!--ATTLIST role-name id ID #IMPLIED>
<!--ATTLIST security-role id ID #IMPLIED>
<!--ATTLIST security-role-ref id ID #IMPLIED>
<!--ATTLIST session-type id ID #IMPLIED>
<!--ATTLIST session id ID #IMPLIED>
<!--ATTLIST small-icon id ID #IMPLIED>
<!--ATTLIST transaction-type id ID #IMPLIED>
<!--ATTLIST trans-attribute id ID #IMPLIED>
```

C.6 Deployment Descriptor Example

The following example illustrates a sample deployment descriptor for the ejb-jar containing the Wombat's assembled application described in Section 3.2.

Note: The text in the <description> elements has been formatted by adding whitespace to appear properly indented in this document. In a real deployment descriptor document, the <description> elements would likely contain no extra whitespace characters.

```
<!DOCTYPE ejb-jar PUBLIC "-//Sun Microsystems, Inc.//DTD Enterprise
JavaBeans 1.1/EN" "http://java.sun.com/j2ee/dtds/ejb-jar_1_1.dtd">
<ejb-jar>
  <description>
    This ejb-jar file contains assembled enterprise beans that are
    part of employee self-service application.
  </description>

  <enterprise-beans>
    <session>
      <description>
        The EmployeeService session bean implements a session
        between an employee and the employee self-service
        application.
      </description>
      <ejb-name>EmployeeService</ejb-name>
      <home>com.wombat.empl.EmployeeServiceHome</home>
      <remote>com.wombat.empl.EmployeeService</remote>
      <ejb-class>com.wombat.empl.EmployeeServiceBean</ejb-class>
      <session-type>Stateful</session-type>
      <transaction-type>Bean</transaction-type>

      <env-entry>
        <env-entry-name>envvar1</env-entry-name>
        <env-entry-type>String</env-entry-type>
        <env-entry-value>some value</env-entry-value>
      </env-entry>

      <ejb-ref>
        <ejb-ref-name>ejb/EmplRecords</ejb-ref-name>
        <ejb-ref-type>Entity</ejb-ref-type>
        <home>com.wombat.empl.EmployeeRecordHome</home>
        <remote>com.wombat.empl.EmployeeRecord</remote>
        <ejb-link>EmployeeRecord</ejb-link>
      </ejb-ref>

      <ejb-ref>
        <ejb-ref-name>ejb/Payroll</ejb-ref-name>
        <ejb-ref-type>Entity</ejb-ref-type>
        <home>com.aardvark.payroll.PayrollHome</home>
        <remote>com.aardvark.payroll.Payroll</remote>
        <ejb-link>AardvarkPayroll</ejb-link>
      </ejb-ref>

      <ejb-ref>
        <ejb-ref-name>ejb/PensionPlan</ejb-ref-name>
        <ejb-ref-type>Session</ejb-ref-type>
        <home>com.wombat.empl.PensionPlanHome</home>
        <remote>com.wombat.empl.PensionPlan</remote>
      </ejb-ref>

      <resource-ref>
        <description>
          This is a reference to a JDBC database.
```

```

        EmployeeService keeps a log of all
        transactions performed through the
        EmployeeService bean for auditing
        purposes.
    </description>
    <res-ref-name>jdbc/EmployeeAppDB</res-ref-name>
    <res-type>javax.sql.DataSource</res-type>
    <res-auth>Container</res-auth>
  </resource-ref>
</session>

<session>
  <description>
    The EmployeeServiceAdmin session bean implements
    the session used by the application's administrator.
  </description>

  <ejb-name>EmployeeServiceAdmin</ejb-name>
  <home>com.wombat.empl.EmployeeServiceAdminHome</home>
  <remote>com.wombat.empl.EmployeeServiceAdmin</remote>
  <ejb-class>com.wombat.empl.EmployeeServiceAdmin-
Bean</ejb-class>
  <session-type>Stateful</session-type>
  <transaction-type>Bean</transaction-type>

  <resource-ref>
    <description>
      This is a reference to a JDBC database.
      EmployeeService keeps a log of all
      transactions performed through the
      EmployeeService bean for auditing
      purposes.
    </description>
    <res-ref-name>jdbc/EmployeeAppDB</res-ref-name>
    <res-type>javax.sql.DataSource</res-type>
    <res-auth>Container</res-auth>
  </resource-ref>
</session>

<entity>
  <description>
    The EmployeeRecord entity bean encapsulates access
    to the employee records. The deployer will use
    container-managed persistence to integrate the
    entity bean with the back-end system managing
    the employee records.
  </description>

  <ejb-name>EmployeeRecord</ejb-name>
  <home>com.wombat.empl.EmployeeRecordHome</home>
  <remote>com.wombat.empl.EmployeeRecord</remote>
  <ejb-class>com.wombat.empl.EmployeeRecordBean</ejb-class>
  <persistence-type>Container</persistence-type>
  <prim-key-class>com.wombat.empl.EmployeeID</prim-key-class>
  <reentrant>True</reentrant>

  <cmp-field><field-name>employeeID</field-name></cmp-field>
  <cmp-field><field-name>firstName</field-name></cmp-field>
  <cmp-field><field-name>lastName</field-name></cmp-field>

```

```

        <cmp-field><field-name>address1</field-name></cmp-field>
        <cmp-field><field-name>address2</field-name></cmp-field>
        <cmp-field><field-name>city</field-name></cmp-field>
        <cmp-field><field-name>state</field-name></cmp-field>
        <cmp-field><field-name>zip</field-name></cmp-field>
        <cmp-field><field-name>homePhone</field-name></cmp-field>
        <cmp-field><field-name>jobTitle</field-name></cmp-field>
        <cmp-field><field-name>managerID</field-name></cmp-field>
        <cmp-field><field-name>jobTitleHistory</field-name></cmp-field>
    </entity>

    <entity>
        <description>
            The Payroll entity bean encapsulates access
            to the payroll system. The deployer will use
            container-managed persistence to integrate the
            entity bean with the back-end system managing
            payroll information.
        </description>

        <ejb-name>AardvarkPayroll</ejb-name>
        <home>com.aardvark.payroll.PayrollHome</home>
        <remote>com.aardvark.payroll.Payroll</remote>
        <ejb-class>com.aardvark.payroll.PayrollBean</ejb-class>
        <persistence-type>Bean</persistence-type>
        <prim-key-class>com.aardvark.payroll.Account-
            tID</prim-key-class>
        <reentrant>False</reentrant>

        <security-role-ref>
            <role-name>payroll-org</role-name>
            <role-link>payroll-department</role-link>
        </security-role-ref>
    </entity>
</enterprise-beans>

<assembly-descriptor>
    <security-role>
        <description>
            This role includes the employees of the
            enterprise who are allowed to access the
            employee self-service application. This role
            is allowed only to access his/her own
            information.
        </description>
        <role-name>employee</role-name>
    </security-role>

    <security-role>
        <description>
            This role includes the employees of the human
            resources department. The role is allowed to
            view and update all employee records.
        </description>
        <role-name>hr-department</role-name>
    </security-role>

    <security-role>

```

```

    <description>
      This role includes the employees of the payroll
      department. The role is allowed to view and
      update the payroll entry for any employee.
    </description>
    <role-name>payroll-department</role-name>
  </security-role>

  <security-role>
    <description>
      This role should be assigned to the personnel
      authorized to perform administrative functions
      for the employee self-service application.
      This role does not have direct access to
      sensitive employee and payroll information.
    </description>
    <role-name>admin</role-name>
  </security-role>

  <method-permission>
    <role-name>employee</role-name>
    <method>
      <ejb-name>EmployeeService</ejb-name>
      <method-name>*</method-name>
    </method>
  </method-permission>

  <method-permission>
    <role-name>employee</role-name>
    <method>
      <ejb-name>EmployeeRecord</ejb-name>
      <method-name>findByPrimaryKey</method-name>
    </method>
    <method>
      <ejb-name>EmployeeRecord</ejb-name>
      <method-name>getDetail</method-name>
    </method>
    <method>
      <ejb-name>EmployeeRecord</ejb-name>
      <method-name>updateDetail</method-name>
    </method>
  </method-permission>

  <method-permission>
    <role-name>employee</role-name>
    <method>
      <ejb-name>AardvarkPayroll</ejb-name>
      <method-name>findByPrimaryKey</method-name>
    </method>
    <method>
      <ejb-name>AardvarkPayroll</ejb-name>
      <method-name>getEmployeeInfo</method-name>
    </method>
    <method>
      <ejb-name>AardvarkPayroll</ejb-name>
      <method-name>updateEmployeeInfo</method-name>
    </method>
  </method-permission>

```

```
<method-permission>
  <role-name>admin</role-name>
  <method>
    <ejb-name>EmployeeServiceAdmin</ejb-name>
    <method-name>*</method-name>
  </method>
</method-permission>

<method-permission>
  <role-name>hr-department</role-name>
  <method>
    <ejb-name>EmployeeRecord</ejb-name>
    <method-name>create</method-name>
  </method>
  <method>
    <ejb-name>EmployeeRecord</ejb-name>
    <method-name>remove</method-name>
  </method>
  <method>
    <ejb-name>EmployeeRecord</ejb-name>
    <method-name>changeManager</method-name>
  </method>
  <method>
    <ejb-name>EmployeeRecord</ejb-name>
    <method-name>changeJobTitle</method-name>
  </method>
  <method>
    <ejb-name>EmployeeRecord</ejb-name>
    <method-name>findByPrimaryKey</method-name>
  </method>
  <method>
    <ejb-name>EmployeeRecord</ejb-name>
    <method-name>getDetail</method-name>
  </method>
  <method>
    <ejb-name>EmployeeRecord</ejb-name>
    <method-name>updateDetail</method-name>
  </method>
</method-permission>

<method-permission>
  <role-name>payroll-department</role-name>
  <method>
    <ejb-name>AardvarkPayroll</ejb-name>
    <method-name>findByPrimaryKey</method-name>
  </method>
  <method>
    <ejb-name>AardvarkPayroll</ejb-name>
    <method-name>getEmployeeInfo</method-name>
  </method>
  <method>
    <ejb-name>AardvarkPayroll</ejb-name>
    <method-name>updateEmployeeInfo</method-name>
  </method>
  <method>
    <ejb-name>AardvarkPayroll</ejb-name>
    <method-name>updateSalary</method-name>
  </method>
</method-permission>
```

```
<container-transaction>
  <method>
    <ejb-name>EmployeeRecord</ejb-name>
    <method-name>*</method-name>
  </method>
  <trans-attribute>Required</trans-attribute>
</container-transaction>

<container-transaction>
  <method>
    <ejb-name>AardvarkPayroll</ejb-name>
    <method-name>*</method-name>
  </method>
  <trans-attribute>Required</trans-attribute>
</container-transaction>
</assembly-descriptor>
</ejb-jar>
```

Frequently Asked Questions

This Appendix provides the answers to a number of frequently asked questions.

D.1 Client-Demarcated Transactions

The Java 2, Enterprise Edition specification [9] defines how a client can obtain the `javax.transaction.UserTransaction` interface using JNDI.

The following is an example of how a Java application can obtain the `javax.transaction.UserTransaction` interface.

```
...
Context ctx = new InitialContext();
UserTransaction utx =
    (UserTransaction)ctx.lookup("java:comp/UserTransaction");

//
// Perform calls to enterprise beans in a transaction.
//
utx.begin();
... call one or more enterprise beans
utx.commit();
...
```

D.2 Container-Managed Persistence

EJB 2.0 supports both method-based container-managed persistence and EJB 1.1 field-based container-managed persistence. The EJB 2.0 specification recommends that the new EJB 2.0 method-based mechanism for container-managed persistence be used for new work because of the added functionality that it provides.

The use of both EJB 2.0 method-based and EJB 1.1 field-based container-managed persistence entity beans can be combined in the same EJB 2.0 application. The beans that are written to the EJB 1.1 container-managed persistence API, however, must be indicated as such in the EJB 2.0 deployment descriptor.

EJB 2.0 containers must support the EJB 1.1 mechanism for container-managed persistence. EJB 1.1 entity beans and deployment descriptors are supported by EJB 2.0. The EJB 1.1 container-managed persistence mechanism has not been deprecated.

D.3 Inheritance

The current EJB specification does not specify the concept of *component inheritance*. There are complex issues that would have to be addressed in order to define component inheritance (for example, the issue of how the primary key of the derived class relates to the primary key of the parent class, and how component inheritance affects the parent component's persistence).

However, the Bean Provider can take advantage of the Java language support for inheritance as follows:

- *Interface inheritance.* It is possible to use the Java language interface inheritance mechanism for inheritance of the home and component interfaces. A component may derive its home and component interfaces from some "parent" home and component interfaces; the component then can be used anywhere where a component with the parent interfaces is expected. This is a Java language feature, and its use is transparent to the EJB Container.
- *Implementation class inheritance.* It is possible to take advantage of the Java class implementation inheritance mechanism for the enterprise bean class. For example, the class `CheckingAccountBean` class can extend the `AccountBean` class to inherit the implementation of the business methods.

D.4 How to Obtain Database Connections

Section 20.5 specifies how an enterprise bean should obtain connections to resources such as JDBC connections. The connection acquisition protocol uses resource manager connection factory references that are part of the enterprise bean's environment.

The following is an example of how an enterprise bean obtains a JDBC connection:

```
public class EmployeeServiceBean implements SessionBean {
    EJBContext ejbContext;

    public void changePhoneNumber(...) {
        ...

        // obtain the initial JNDI context
        Context initCtx = new InitialContext();

        // perform JNDI lookup to obtain resource manager
        // connection factory
        javax.sql.DataSource ds = (javax.sql.DataSource)
            initCtx.lookup("java:comp/env/jdbc/EmployeeAppDB");

        // Invoke factory to obtain a connection. The security
        // principal is not given, and therefore
        // it will be configured by the Deployer.
        java.sql.Connection con = ds.getConnection();
        ...
    }
}
```

D.5 Session Beans and Primary Key

The EJB 1.1 specification specifies the Container's behavior for the cases when a client attempts to access the primary key of a session object. In summary, the Container must throw an exception on a client's attempt to access the primary key of a session object.

D.6 Copying of Parameters Required for Remote EJB Calls Within the Same JVM

The enterprise bean's remote home and remote interfaces are *remote interfaces* in the Java RMI sense. The Container must ensure the Java RMI argument passing semantics. Non-remote objects must be passed by value across these interfaces.

Specifically, the EJB Container is not allowed to pass objects by reference on inter-enterprise bean invocations of methods on the enterprise bean's remote home and remote interface when the calling and called enterprise beans are collocated in the same JVM. Doing so could result in the multiple beans sharing the state of a Java object, which would break the enterprise bean's semantics.

EJB 2.0 entity beans and session beans support local and local home interfaces. The arguments of the methods of these interfaces are passed by reference, and the local clients of such beans must be collocated in the same JVM. Such beans must be coded to assume that the state of any Java object that is passed as an argument or result is shared by caller and callee.

Revision History

This appendix lists the significant changes that have been made during the development of the EJB 2.1 specification.

E.1 Version 0.1

Created document from EJB 2.0 Final Release.

Extended message-driven bean component type to provide support for other messaging types.

E.2 Version 0.2

Added ORDER BY clause to EJB QL.

Added AVG, MAX, MIN, SUM, and COUNT functions to EJB QL.

Loosened the restrictions on EJB QL equality and comparison operations in EJB 2.0.

Changes to EJB QL BNF to distinguish between path expressions leading to cmp-fields and path expressions leading to cmr-fields.

Added MOD function.

Clarified requirements for returning null values from queries and finder and select methods.

E.3 Version 0.3

Added web service endpoint interface to stateless session beans to support invocations over SOAP, HTTP from web service clients.

Added description of web service client view.

Added service-ref construct to allow enterprise beans to declare dependencies on and access external web services.

Modifications to chapters on Transactions, Exceptions, Environment, Deployment descriptor to include stateless session beans with web service endpoint interfaces.

Change result type of EJB QL COUNT operator to long.

Moved EJB 2.0 deployment descriptor chapter to appendix.

E.4 Version 0.4

Added EJB Timer Service.

Modifications to chapters on enterprise beans, as well as chapters on Transactions, Security, Exceptions to reflect addition of timer service.

Updated security chapter to support security role specifications for web service endpoint interfaces.

Updates to reflect JMS 1.1 API additions.

Message-driven bean deployment descriptor changes to integrate with Connector 1.5 activation-config requirements.

Updated Runtime chapter to reflect J2EE 1.4 APIs.

E.5 Version 0.5

Clarified section on HandleDelegate to explicitly indicate use of standard IIOP abstract interface format.

Extended use of Timer Service to stateless session beans.

Extended use of transaction attributes to ejbTimeout method.

Allowed use of input parameters and numeric literals in EJB QL IN expressions.

Allowed use of input parameters in EJB QL IS [NOT] NULL expressions.

Clarified that path expression navigation corresponds to “inner join” semantics with regard to the handling of null cmr-field values.

E.6 Version 0.6: Community Review Draft

Minor clarifications for timers and message-driven beans.

E.7 Version 0.7

Eliminated unnecessary nonterminals from EJB QL BNF: `single_value_designator`, `scalar_expression`, `conditional_test`.

Added `getMessageContext` method to `SessionContext`. Only stateless session beans with a web service endpoint interface can call this method. It returns the `javax.xml.rpc.handler.MessageContext` interface.

Added message destination references and links to allow the Application Assembler to specify the messaging flow within an application.

Minor editorial improvements.

E.8 Version 0.8: Public Draft

Converted EJB 2.1 deployment descriptor DTD to XML Schema.

E.9 Version 0.9: Proposed Final Draft

Removed restriction that a message listener method must have only a single argument.

Removed restriction that the message listener interface must specify only a single method.

Removed restriction that the methods of a message listener interface must not throw application exceptions.

Clarified that the enterprise bean class must be a top level class.

Clarified the requirements for shareability of connections in conformance with [12].

Extended use of Timer Service to message-driven beans.

Added requirement for support of JAX-RPC message handlers.

Added extensibility elements to XML Schema for EJB 2.1 deployment descriptor.

Minor editorial changes.

E.10 Proposed Final Draft 2

Removed deployment-extension elements from deployment descriptor.

Added warning about use of untrimmed strings in primary key fields.

Added note on required support for JAX-RPC Holder classes.

Clarified that the container is not required to call `ejbStore` before `ejbFindByPrimary` key for container-managed persistence entity beans.

Added note clarifying that the assignment rules for collection-valued cmr-fields pertain to entity objects contained in non-managed source collections as well.

EJB QL query methods using the aggregate operators AVG, SUM, MAX, MIN, COUNT must be single-object select methods. Relaxed the rules for the result types of these methods.

Clarified that the container does not eliminate duplicates from an EJB QL query unless `SELECT DISTINCT` is specified or the query method result type is `java.util.Set`.

Clarified that 1 designates the first position in a string for EJB QL start and length functions.

Clarification of semantics of commit options for entity beans with container-managed persistence.

Editorial sweep.

E.11 Final Release

Moved `j2ee:java-typeType` definition from `ejb-jar_2_1.xsd` to `j2ee_1_4.xsd`.

Clarified requirement for the container to return a representation for the unauthenticated identity from `getCallerPrincipal` when the security identity of the caller has not been established.

Allowed use-caller-identity to be specified for message-driven beans and beans that implement the `TimedObject` interface.

Relaxed restrictions on invocation of `getCallerPrincipal` and `isCallerInRole`.