

Enterprise JavaBeans™ Specification, v1.1

Please send technical comments to: ejb-spec-comments@sun.com
Please send business comments to: ejb-marketing@sun.com

PRELIMINARY

Public Release 2

Vlada Matena & Mark Hapner



901 San Antonio Road
Palo Alto, CA 94303 USA
650 960-1300 fax: 650 969-9131

Sun Microsystems, Inc.

Copyright 1999 Sun Microsystems, Inc. 901 San Antonio Road, Palo Alto, CA 94303, U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or documentation may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Third party software, including font technology, is copyrighted and licensed from Sun suppliers.

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

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK(R) and Sun(TM) Graphical User Interfaces was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUI's and otherwise comply with Sun's written license agreements.

PostScript is a registered trademark of Adobe Systems, Inc.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a).

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 1999 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, CA 94303, Etats-Unis. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou de sa documentation associée ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a.

Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Sun, Sun Microsystems, le logo Sun, Solaris, Java, et Enterprise JavaBeans, sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays.

Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.



Please
Recycle

Les interfaces d'utilisation graphique OPEN LOOK et Sun (TM) ont été développées par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique, cette licence couvrant aussi les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

Postscript est une marque enregistrée d'Adobe Systems Inc.

LA DOCUMENTATION EST FOURNIE "EN L'ETAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFACON.



Please
Recycle

Table of Contents

Chapter 1	Introduction.....	15
	1.1 Target audience.....	15
	1.2 What is new in EJB 1.1	15
	1.3 Application compatibility and interoperability	16
	1.4 Acknowledgments	17
	1.5 Organization	17
	1.6 Document conventions	18
Chapter 2	Goals	19
	2.1 Overall goals.....	19
	2.2 Goals for Release 1.0.....	20
	2.3 Goals for Release 1.1.....	20
Chapter 3	EJB Roles and Scenarios	21
	3.1 EJB Roles	21
	3.1.1 Enterprise Bean Provider	22
	3.1.2 Application Assembler.....	22
	3.1.3 Deployer.....	22
	3.1.4 EJB Server Provider.....	23
	3.1.5 EJB Container Provider.....	23
	3.1.6 System Administrator	24
	3.2 Scenario: Development, assembly, and deployment	24
Chapter 4	Overview.....	29
	4.1 Enterprise Beans as components	29
	4.1.1 Component characteristics	30
	4.1.2 Flexible component model.....	30
	4.2 Enterprise JavaBeans contracts.....	31
	4.2.1 Client-view contract	31
	4.2.2 Component contract	32
	4.2.3 Ejb-jar file	33
	4.2.4 Contracts summary	33
	4.3 Session and entity objects.....	34
	4.3.1 Session objects	34
	4.3.2 Entity objects.....	35
	4.4 Standard mapping to CORBA protocols	35
Chapter 5	Client View of a Session Bean.....	39
	5.1 Overview	39
	5.2 EJB Container.....	40

	5.2.1	Locating a session bean's home interface	40
	5.2.2	What a container provides	41
5.3		Home interface.....	41
	5.3.1	Creating a session object	42
	5.3.2	Removing a session object	42
5.4		EJBObject	43
5.5		Session object identity	43
5.6		Client view of session object's life cycle.....	44
5.7		Creating and using a session object	45
5.8		Object identity	46
	5.8.1	Stateful session beans.....	46
	5.8.2	Stateless session beans	46
	5.8.3	getPrimaryKey()	47
5.9		Type narrowing	47
Chapter 6		Session Bean Component Contract.....	49
	6.1	Overview.....	49
	6.2	Goals	50
	6.3	A container's management of its working set.....	50
	6.4	Conversational state	51
	6.4.1	Instance passivation and conversational state.....	51
	6.4.2	The effect of transaction rollback on conversational state	53
	6.5	Protocol between a session bean instance and its container	53
	6.5.1	The required <i>SessionBean</i> interface	53
	6.5.2	The <i>SessionContext</i> interface	54
	6.5.3	The optional <i>SessionSynchronization</i> interface	54
	6.5.4	Business method delegation	55
	6.5.5	Session bean's <i>ejbCreate(...)</i> methods	55
	6.5.6	Serializing session bean methods	56
	6.5.7	Transaction context of session bean methods	56
	6.6	STATEFUL Session Bean State Diagram.....	56
	6.6.1	Operations allowed in the methods of a stateful session bean class .	59
	6.6.2	Dealing with exceptions	61
	6.6.3	Missed <i>ejbRemove()</i> calls	61
	6.6.4	Restrictions for transactions	62
	6.7	Object interaction diagrams for a STATEFUL session bean	62
	6.7.1	Notes.....	62
	6.7.2	Creating a session object	63
	6.7.3	Starting a transaction	63
	6.7.4	Committing a transaction	64
	6.7.5	Passivating and activating an instance between transactions	65
	6.7.6	Removing a session object	66
	6.8	Stateless session beans.....	67
	6.8.1	Stateless session bean state diagram	68
	6.8.2	Operations allowed in the methods of a stateless session bean class	69
	6.8.3	Dealing with exceptions	71

6.9	Object interaction diagrams for a STATELESS session bean	71
6.9.1	Client-invoked create().....	71
6.9.2	Business method invocation.....	72
6.9.3	Client-invoked remove()	72
6.9.4	Adding instance to the pool	73
6.10	The responsibilities of the bean provider	74
6.10.1	Classes and interfaces	74
6.10.2	Session bean class	75
6.10.3	ejbCreate methods.....	76
6.10.4	Business methods.....	76
6.10.5	Session bean's remote interface	77
6.10.6	Session bean's home interface	77
6.11	The responsibilities of the container provider	78
6.11.1	Generation of implementation classes	78
6.11.2	Session EJBHome class	78
6.11.3	Session EJBObject class	78
6.11.4	Handle classes	79
6.11.5	EJBMetaData class	79
6.11.6	Non-reentrant instances.....	79
6.11.7	Transaction scoping, security, exceptions	79
Chapter 7	Example Session Scenario	81
7.1	Overview	81
7.2	Inheritance relationship	81
7.2.1	What the session Bean provider is responsible for	83
7.2.2	Classes supplied by container provider.....	83
7.2.3	What the container provider is responsible for	83
Chapter 8	Client View of an Entity.....	85
8.1	Overview	85
8.2	EJB Container.....	86
8.2.1	Locating an entity bean's home interface.....	87
8.2.2	What a container provides.....	87
8.3	Entity bean's home interface	88
8.3.1	create methods.....	89
8.3.2	finder methods.....	90
8.3.3	remove methods	90
8.4	Entity object's life cycle	91
8.5	Primary key and object identity	92
8.6	Entity Bean's remote interface	93
8.7	Entity bean's handle	94
8.8	Entity home handles	95
8.9	Type narrowing of object references	96

Chapter 9	Entity Bean Component Contract	97
	9.1 Concepts	97
	9.1.1 Runtime execution model.....	97
	9.1.2 Granularity of entity beans	99
	9.1.3 Entity persistence (data access protocol)	99
	9.1.3.1 Bean-managed persistence.....	100
	9.1.3.2 Container-managed persistence	101
	9.1.4 Instance life cycle.....	102
	9.1.5 The entity bean component contract	104
	9.1.5.1 Entity bean instance's view:.....	104
	9.1.5.2 Container's view:	107
	9.1.6 Operations allowed in the methods of the entity bean class.....	109
	9.1.7 Caching of entity state and the ejbLoad and ejbStore methods	112
	9.1.7.1 ejbLoad and ejbStore with the NotSupported transaction attribute	113
	9.1.8 Finder method return type	114
	9.1.8.1 Single-object finder.....	114
	9.1.8.2 Multi-object finders.....	114
	9.1.9 Standard application exceptions for Entities	116
	9.1.9.1 CreateException.....	116
	9.1.9.2 DuplicateKeyException	116
	9.1.9.3 FinderException.....	117
	9.1.9.4 ObjectNotFoundException	117
	9.1.9.5 RemoveException	117
	9.1.10 Commit options	118
	9.1.11 Concurrent access from multiple transactions	119
	9.1.12 Non-reentrant and re-entrant instances	120
	9.2 Responsibilities of the Enterprise Bean Provider	121
	9.2.1 Classes and interfaces.....	121
	9.2.2 Enterprise bean class	121
	9.2.3 ejbCreate methods	122
	9.2.4 ejbPostCreate methods	124
	9.2.5 ejbFind methods	124
	9.2.6 Business methods	125
	9.2.7 Entity bean's remote interface.....	125
	9.2.8 Entity bean's home interface	126
	9.2.9 Entity bean's primary key class.....	127
	9.3 The responsibilities of the Container Provider	127
	9.3.1 Generation of implementation classes.....	127
	9.3.2 Entity EJBHome class	128
	9.3.3 Entity EJBObject class	128
	9.3.4 Handle class.....	128
	9.3.5 Home Handle class.....	128
	9.3.6 Meta-data class	129
	9.3.7 Instance's re-entrance.....	129
	9.3.8 Transaction scoping, security, exceptions	129
	9.3.9 Implementation of object references	129
	9.4 Entity beans with container-managed persistence	129
	9.4.1 Container-managed fields.....	130

	9.4.2	ejbCreate, ejbPostCreate	131
	9.4.3	ejbRemove.....	132
	9.4.4	ejbLoad.....	132
	9.4.5	ejbStore	132
	9.4.6	finder methods.....	133
	9.4.7	primary key type	133
	9.4.7.1	Primary key that maps to a single field in the entity bean class	133
	9.4.7.2	Primary key that maps to multiple fields in the entity bean class	133
	9.4.7.3	Special case: Unknown primary key class.....	134
9.5		Object interaction diagrams.....	134
	9.5.1	Notes	135
	9.5.2	Creating an entity object	136
	9.5.3	Passivating and activating an instance in a transaction	138
	9.5.4	Committing a transaction	140
	9.5.5	Starting the next transaction.....	142
	9.5.6	Removing an entity object	145
	9.5.7	Finding an entity object.....	146
	9.5.8	Adding and removing an instance from the pool	147
Chapter 10		Example entity scenario.....	149
	10.1	Overview	149
	10.2	Inheritance relationship	150
	10.2.1	What the entity Bean Provider is responsible for.....	151
	10.2.2	Classes supplied by Container Provider.....	151
	10.2.3	What the container provider is responsible for	151
Chapter 11		Support for Transactions.....	153
	11.1	Overview	153
	11.1.1	Transactions	153
	11.1.2	Transaction model.....	154
	11.1.3	Relationship to JTA and JTS.....	154
	11.2	Sample scenarios	155
	11.2.1	Update of multiple databases	155
	11.2.2	Update of databases via multiple EJB Servers.....	156
	11.2.3	Client-managed demarcation	156
	11.2.4	Container-managed demarcation	157
	11.2.5	Bean-managed demarcation	158
	11.2.6	Interoperability with non-Java clients and servers	158
	11.3	Bean Provider's responsibilities	159
	11.3.1	Bean-managed versus container-managed transaction demarcation.	159
	11.3.1.1	Non-transactional execution	160
	11.3.2	Isolation levels.....	160
	11.3.3	Enterprise beans using bean-managed transaction demarcation.....	161
	11.3.3.1	getRollbackOnly() and setRollbackOnly() method.....	166
	11.3.4	Enterprise beans using container-managed transaction demarcation	167
	11.3.4.1	javax.ejb.SessionSynchronization interface.....	168
	11.3.4.2	javax.ejb.EJBContext.setRollbackOnly() method	168

11.3.4.3	javax.ejb.EJBContext.getRollbackOnly() method	169
11.3.5	Declaration in deployment descriptor	169
11.4	Application Assembler's responsibilities	169
11.4.1	Transaction attributes	169
11.5	Deployer's responsibilities.....	172
11.6	Container Provider responsibilities.....	173
11.6.1	Bean-managed transaction demarcation.....	173
11.6.2	Container-managed transaction demarcation	175
11.6.2.1	NotSupported	175
11.6.2.2	Required	175
11.6.2.3	Supports	176
11.6.2.4	RequiresNew	176
11.6.2.5	Mandatory	177
11.6.2.6	Never	177
11.6.2.7	Transaction attribute summary	177
11.6.2.8	Handling of setRollbackOnly() method.....	178
11.6.2.9	Handling of getRollbackOnly() method	178
11.6.2.10	Handling of getUserTransaction() method	179
11.6.2.11	javax.ejb.SessionSynchronization callbacks	179
11.6.3	Handling of methods that run with "an unspecified transaction context".....	179
11.7	Access from multiple clients in the same transaction context	180
11.7.1	Transaction "diamond" scenario with an entity object.....	180
11.7.2	Container Provider's responsibilities.....	182
11.7.3	Bean Provider's responsibilities	183
11.7.4	Application Assembler and Deployer's responsibilities	184
11.7.5	Transaction diamonds involving session objects.....	184
Chapter 12	Exception handling	187
12.1	Overview and Concepts	187
12.1.1	Application exceptions	187
12.1.2	Goals for exception handling	188
12.2	Bean Provider's responsibilities	188
12.2.1	Application exceptions	188
12.2.2	System exceptions	189
12.2.2.1	javax.ejb.NoSuchEntityException	190
12.3	Container Provider responsibilities.....	190
12.3.1	Exceptions from an enterprise bean's business methods.....	190
12.3.2	Exceptions from container-invoked callbacks.....	193
12.3.3	javax.ejb.NoSuchEntityException.....	193
12.3.4	Non-existing session object.....	193
12.3.5	Exceptions from the management of container-managed transactions.....	194
12.3.6	Release of resources	194
12.3.7	Support for deprecated use of java.rmi.RemoteException.....	194
12.4	Client's view of exceptions.....	195
12.4.1	Application exception.....	195
12.4.2	java.rmi.RemoteException	196
12.4.2.1	javax.transaction.TransactionRolledbackException	197

	12.4.2.2	javax.transaction.TransactionRequiredException.....	197
	12.4.2.3	java.rmi.NoSuchObjectException	197
	12.5	System Administrator's responsibilities	197
	12.6	Differences from EJB 1.0	197
Chapter 13		Support for Distribution	199
	13.1	Overview	199
	13.2	Client-side objects in distributed environment	200
	13.3	Standard distribution protocol	200
Chapter 14		Enterprise bean environment	201
	14.1	Overview	201
	14.2	Enterprise bean's environment as a JNDI naming context	202
	14.2.1	Bean Provider's responsibilities	203
	14.2.1.1	Access to enterprise bean's environment.....	203
	14.2.1.2	Declaration of environment entries.....	204
	14.2.2	Application Assembler's responsibility	207
	14.2.3	Deployer's responsibility	207
	14.2.4	Container Provider responsibility	207
	14.3	EJB references	207
	14.3.1	Bean Provider's responsibilities	208
	14.3.1.1	EJB reference programming interfaces	208
	14.3.1.2	Declaration of EJB references in deployment descriptor ...	208
	14.3.2	Application Assembler's responsibilities	210
	14.3.3	Deployer's responsibility	211
	14.3.4	Container Provider's responsibility	211
	14.4	Resource manager connection factory references	211
	14.4.1	Bean Provider's responsibilities	212
	14.4.1.1	Programming interfaces for resource manager connection factory references	212
	14.4.1.2	Declaration of resource manager connection factory references in deployment descriptor	213
	14.4.1.3	Standard resource manager connection factory types	214
	14.4.2	Deployer's responsibility	215
	14.4.3	Container provider responsibility.....	215
	14.4.4	System Administrator's responsibility	216
	14.5	Deprecated EJBContext.getEnvironment() method	216
	14.6	UserTransaction interface	217
Chapter 15		Security management.....	219
	15.1	Overview	219
	15.2	Bean Provider's responsibilities	220
	15.2.1	Invocation of other enterprise beans	220
	15.2.2	Resource access.....	221
	15.2.3	Access of underlying OS resources	221

15.2.4	Programming style recommendations	221
15.2.5	Programmatic access to caller's security context	221
15.2.5.1	Use of <code>getCallerPrincipal()</code>	223
15.2.5.2	Use of <code>isCallerInRole(String roleName)</code>	224
15.2.5.3	Declaration of security roles referenced from the bean's code	225
15.3	Application Assembler's responsibilities	226
15.3.1	Security roles	227
15.3.2	Method permissions	228
15.3.3	Linking security role references to security roles	232
15.4	Deployer's responsibilities	232
15.4.1	Security domain and principal realm assignment	233
15.4.2	Assignment of security roles	233
15.4.3	Principal delegation	233
15.4.4	Security management of resource access	234
15.4.5	General notes on deployment descriptor processing	234
15.5	EJB Client Responsibilities	234
15.6	EJB Container Provider's responsibilities	235
15.6.1	Deployment tools	235
15.6.2	Security domain(s)	235
15.6.3	Security mechanisms	235
15.6.4	Passing principals on EJB calls	236
15.6.5	Security methods in <code>javax.ejbEJBContext</code>	236
15.6.6	Secure access to resource managers	236
15.6.7	Principal mapping	236
15.6.8	System principal	236
15.6.9	Runtime security enforcement	237
15.6.10	Audit trail	238
15.7	System Administrator's responsibilities	238
15.7.1	Security domain administration	238
15.7.2	Principal mapping	238
15.7.3	Audit trail review	238
Chapter 16	Deployment descriptor	239
16.1	Overview	239
16.2	Bean Provider's responsibilities	240
16.3	Application Assembler's responsibility	242
16.4	Container Provider's responsibilities	244
16.5	Deployment descriptor DTD	244
16.6	Deployment descriptor example	259
Chapter 17	Ejb-jar file	267
17.1	Overview	267
17.2	Deployment descriptor	268
17.3	Class files	268
17.4	ejb-client JAR file	268

	17.5	Deprecated in EJB 1.1	269
	17.5.1	ejb-jar Manifest	269
	17.5.2	Serialized deployment descriptor JavaBeans™ components	269
Chapter 18		Runtime environment.....	271
	18.1	Bean Provider's responsibilities	271
	18.1.1	APIs provided by Container	272
	18.1.2	Programming restrictions	272
	18.2	Container Provider's responsibility	275
	18.2.1	Java 2 based Container	275
	18.2.1.1	Java 2 APIs requirements	275
	18.2.1.2	EJB 1.1 requirements.....	276
	18.2.1.3	JNDI 1.2 requirements	276
	18.2.1.4	JTA 1.0.1 requirements	277
	18.2.1.5	JDBC™ 2.0 extension requirements	277
	18.2.2	JDK™ 1.1 based Container.....	277
	18.2.2.1	JDK 1.1 APIs requirements	277
	18.2.2.2	EJB 1.1 requirements.....	279
	18.2.2.3	JNDI 1.2 requirements	279
	18.2.2.4	JTA 1.0.1 requirements	279
	18.2.2.5	JDBC 2.0 extension requirements	279
	18.2.3	Argument passing semantics	279
Chapter 19		Responsibilities of EJB Roles	281
	19.1	Bean Provider's responsibilities	281
	19.1.1	API requirements	281
	19.1.2	Packaging requirements	281
	19.2	Application Assembler's responsibilities	282
	19.3	EJB Container Provider's responsibilities	282
	19.4	Deployer's responsibilities	282
	19.5	System Administrator's responsibilities	282
	19.6	Client Programmer's responsibilities	282
Chapter 20		Enterprise JavaBeans™ API Reference	283
		package javax.ejb.....	283
		package javax.ejb.deployment	284
Chapter 21		Related documents	285
Appendix A		Features deferred to future releases	287
Appendix B		Frequently asked questions	289
	B.1	Client-demarcated transactions	289

	B.2 Inheritance	290
	B.3 Entities and relationships	291
	B.4 Finder methods for entities with container-managed persistence.....	291
	B.5 JDK 1.1 or Java 2.....	291
	B.6 javax.transaction.UserTransaction versus javax.jts.UserTransaction	291
	B.7 How to obtain database connections.....	292
	B.8 Session beans and primary key.....	292
	B.9 Copying of parameters required for EJB calls within the same JVM	292
Appendix C	Revision History.....	293
	C.1 Changes since Release 0.8.....	293
	C.2 Changes since Release 0.9.....	294
	C.3 Changes since Release 0.95.....	295
	C.4 Changes since 1.0	296
	C.5 Changes since 1.1 Draft 1	297
	C.6 Changes since 1.1 Draft 2.....	297
	C.7 Changes since EJB 1.1 Draft 3	299
	C.8 Changes since EJB 1.1 Public Draft.....	300
	C.9 Changes since EJB 1.1 Public Draft 2.....	301
	C.10 Changes since EJB 1.1 Public Draft 3.....	302
	C.11 Changes since EJB 1.1 Public Release.....	302

List of Figures

Figure 1	Enterprise JavaBeans Contracts	34
Figure 2	Heterogeneous EJB Environment	37
Figure 3	Client View of session beans deployed in a Container.....	41
Figure 4	Lifecycle of a session object.	44
Figure 5	Session Bean Example Objects	45
Figure 6	Lifecycle of a STATEFUL Session bean instance.....	57
Figure 7	OID for Creation of a session object	63
Figure 8	OID for session object at start of a transaction.	64
Figure 9	OID for session object transaction synchronization.....	65
Figure 10	OID for passivation and activation of a session object	66
Figure 11	OID for the removal of a session object	67
Figure 12	Lifecycle of a STATELESS Session bean	69
Figure 13	OID for creation of a STATELESS session object.....	71
Figure 14	OID for invocation of business method on a STATELESS session object.....	72
Figure 15	OID for removal of a STATELESS session object.....	73
Figure 16	OID for Container Adding Instance of a STATELESS session bean to a method-ready pool.....	73
Figure 17	OID for a Container Removing an Instance of a STATELESS Session bean from ready pool	74
Figure 18	Example of Inheritance Relationships Between EJB Classes.....	82
Figure 19	Client view of entity beans deployed in a container	88
Figure 20	Client View of Entity Object Life Cycle	91
Figure 21	Overview of the entity bean runtime execution model.....	98
Figure 22	Client view of underlying data sources accessed through entity bean	100
Figure 23	Life cycle of an entity bean instance.	102
Figure 24	Multiple clients can access the same entity object using multiple instances	119
Figure 25	Multiple clients can access the same entity object using single instance.....	120
Figure 26	OID of Creation of an entity object with bean-managed persistence	136
Figure 27	OID of creation of an entity object with container-managed persistence	137
Figure 28	OID of passivation and reactivation of an entity bean instance with bean-managed persistence	138
Figure 29	OID of passivation and reactivation of an entity bean instance with CMP.....	139
Figure 30	OID of transaction commit protocol for an entity bean instance with bean-managed persistence	140
Figure 31	OID of transaction commit protocol for an entity bean instance with container-managed persistence	141
Figure 32	OID of start of transaction for an entity bean instance with bean-managed persistence	143
Figure 33	OID of start of transaction for an entity bean instance with container-managed persistence	144
Figure 34	OID of removal of an entity bean object with bean-managed persistence.....	145
Figure 35	OID of removal of an entity bean object with container-managed persistence.....	145
Figure 36	OID of execution of a finder method on an entity bean instance with bean-managed persistence	146
Figure 37	OID of execution of a finder method on an entity bean instance with container-managed persistence	147
Figure 38	OID of a container adding an instance to the pool.....	148
Figure 39	OID of a container removing an instance from the pool.....	148

Figure 40	Example of the inheritance relationship between the interfaces and classes:	150
Figure 41	Updates to Simultaneous Databases	155
Figure 42	Updates to Multiple Databases in Same Transaction	156
Figure 43	Updates on Multiple Databases on Multiple Servers	157
Figure 44	Update of Multiple Databases from Non-Transactional Client	158
Figure 45	Interoperating with Non-Java Clients and/or Servers	159
Figure 46	Transaction diamond scenario with entity object	181
Figure 47	Handling of diamonds by a multi-process container	183
Figure 48	Transaction diamond scenario with a session bean	184
Figure 49	Location of EJB Client Stubs.	200

List of Tables

Table 1	EJB Roles in the example scenarios.....	27
Table 2	Operations allowed in the methods of a stateful session bean	60
Table 3	Operations allowed in the methods of a stateless session bean.....	70
Table 4	Operations allowed in the methods of an entity bean	111
Table 5	Summary of commit-time options.....	118
Table 6	Container's actions for methods of beans with bean-managed transaction	174
Table 7	Transaction attribute summary	177
Table 8	Handling of exceptions thrown by a business method of a bean with container-managed transaction demarcation	191
Table 9	Handling of exceptions thrown by a business method of a session with bean-managed transaction demarcation	192
Table 10	Java 2 Security policy for a standard EJB Container	275
Table 11	JDK 1.1 Security manager checks for a standard EJB Container	278

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 1.1

We have tightened the Entity bean specification, and made support for Entity beans mandatory for Container Providers.

The other changes in the EJB 1.1 specification were made to improve the support for the development, application assembly, and deployment of ISV-produced enterprise beans. The specification includes the following primary changes:

- Enhanced support for the enterprise bean's environment. The Bean Provider must specify all the bean's environmental dependencies using entries in a JNDI naming context.
- Added support for Application Assembly in the deployment descriptor.
- Clearly separated the responsibilities of the Bean Provider and Application Assembler.
- Removed the EJB 1.0 deployment descriptor features that describe the Deployer's output. The role of the deployment descriptor is to describe the information that is the *input* to the Deployer, not the Deployer's *output*.

The changes affected mainly Chapters 11, 14, 15, and 16. We minimized the impact on the server vendors who implemented support for EJB 1.0 in their runtime. The only change to the runtime API of the EJB Container is the replacement of the `java.security.Identity` class with the `java.security.Principal` interface, necessitated by changes in JDK 1.2.

We have also added a number of clarifications and corrections to the specification based on the input that we have received from the reviewers.

1.3 Application compatibility and interoperability

EJB 1.1 attempts to provide a high degree of application compatibility for enterprise beans that were written for the EJB 1.0 specification. Principally, the deployment descriptor of EJB 1.0 based enterprise beans must be converted to the EJB 1.1 XML format. However, the EJB 1.0 enterprise bean code does not have to be changed or re-compiled to run in an EJB 1.1 Container, except in the following situations:

- The bean uses the `javax.jts.UserTransaction` interface. The package name of the `javax.jts` interface has changed to `javax.transaction`, and there have been minor changes to the exceptions thrown by the methods of this interface. An enterprise bean that uses the `javax.jts.UserTransaction` interface needs to be modified to use the new name `javax.transaction.UserTransaction`.
- The bean uses the `getCallerIdentity()` or `isCallerInRole(Identity identity)` methods of the `javax.ejb.EJBContext` interface. These methods were deprecated in EJB 1.1 because the class `java.security.Identity` is deprecated in Java 2. While a Container Provider may choose to provide a backward compatible implementation of these two methods, the Container Provider is not required to do so. An enterprise bean written to the EJB 1.0 specification needs to be modified to use the new methods to work in *all* EJB 1.1 Containers.
- The bean is an entity bean with container-managed persistence. The required return value of `ejbPostCreate(...)` is different in EJB 1.1 than in EJB 1.0. An enterprise bean with container-managed persistence written to the EJB 1.0 specification needs to be recompiled to work with all EJB 1.1 compliant Containers.
- The bean is an entity bean whose finders do not define the `FinderException` in the methods' throws clauses. EJB 1.1 requires that all finders define the `FinderException`.
- The bean is an entity bean that uses the `UserTransaction` interface. In EJB 1.1, an entity bean must not use the `UserTransaction` interface.
- The bean uses the `UserTransaction` interface and implements the `SessionSynchronization` interface at the same time. This is disallowed in EJB 1.1.
- The bean violates any of the additional semantic restrictions defined in EJB 1.1 but which were not defined in EJB 1.0.

The client view of an enterprise bean is the same in EJB 1.0 and EJB 1.1. This means that enterprise beans written to EJB 1.1 can seamlessly interoperate with those written to EJB 1.0, and vice versa.

1.4 Acknowledgments

Rick Cattell, Linda DeMichiel, Shel Finkelstein, Graham Hamilton, Li Gong, Rohit Garg, Susan Cheung, Hans Hrasna, Sanjeev Krishnan, Kevin Osborn, Bill Shannon, Anil Vijendran, and Larry Cable have provided invaluable input to the design of Enterprise JavaBeans architecture.

The Enterprise JavaBeans architecture is a broad effort that includes contributions from numerous groups at Sun and at partner companies. The ongoing specification review process has been extremely valuable, and the many comments that we have received helped us to define the specification.

We would also like to thank all the reviewers who sent us feedback during the public review period. Their input helped us to improve the specification.

1.5 Organization

Chapter 2, “Goals” discusses the advantages of 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 Administrators with respect to the Enterprise JavaBeans architecture.

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

Chapters 5 through 7 define Session Beans: Chapter 5 discusses the client view, Chapter 6 presents the Session Bean component contract, and Chapter 7 outlines an example Session Bean scenario.

Chapters 8 through 10 define Entity Beans: Chapter 8 discusses the client view, Chapter 9 presents the Entity Bean component contract, and Chapter 10 outlines an example Entity Bean scenario.

Chapters 11 through 15 discuss transactions, exceptions, distribution, environment, and security.

Chapters 16 and 17 describe the format of the ejb-jar file and its deployment descriptor.

Chapter 18 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 19 summarizes the responsibilities of the individual EJB Roles.

Chapter 20 is the Enterprise JavaBeans API Reference.

Chapter 21 provides a list of related documents.

1.6 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.

Goals

2.1 Overall goals

We have set the following goals for the Enterprise JavaBeans (EJB) architecture:

- *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 make it possible to build distributed applications by combining components developed using tools from different vendors.*
- *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, and 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 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 provide interoperability between enterprise Beans and non-Java programming language applications.*
- *The Enterprise JavaBeans architecture will be compatible with the CORBA protocols.*

2.2 Goals for Release 1.0

In Release 1.0, we focused on the following:

- *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.*

2.3 Goals for Release 1.1

In the EJB 1.1 Release, we focus on the following aspects:

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

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 bean(s). The Bean Provider is responsible for the Java classes that implement the enterprise bean's business methods; the definition of the bean's remote and home 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 the resource managers 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 has inserted the application assembly instruction into the deployment descriptors.

The Application Assembler can also combine enterprise beans with other types of application components (e.g. Java ServerPages™) 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 Beans' remote and home interfaces, he or she does not need to have any knowledge of the enterprise Beans' 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 role, the Deployer uses tools provided by the EJB Container Provider.

The Deployer's output are 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 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 beans' 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 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.

For Entity Beans with container-managed persistence, the entity container is responsible for persistence of the Entity Beans installed in the container. The Container Provider’s tools are used to generate code that moves data between the enterprise Bean’s instance variables and a database or an existing application.

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 Java ServerPages™ (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 ABCPensionPlan) 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 distributed computing. Enterprise beans are components of distributed 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 services information, such as a transaction and security attributes, are separate from the enterprise bean class. This allows the services 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 developer 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 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.

Enterprise beans are intended to be relatively coarse-grained business objects (e.g. purchase order, employee record). Fine-grained objects (e.g. line item on a purchase order, employee's address) should not be modeled as enterprise bean components.

While 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.

A client always uses the same API for object creation, lookup, method invocation, and removal, regardless of how an enterprise bean is implemented or what function it provides to the client.

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 contract

This is a contract between a client and a container. The client-view contract provides 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.

The enterprise bean client view is remotable—both local and remote programs can access an enterprise bean using the same view of the enterprise bean.

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 view of an enterprise bean can also be mapped to non-Java client environments, such as CORBA clients that are not written in the Java™ programming language.

The enterprise Bean Provider and the container provider cooperate to create the enterprise bean's client view. The client view includes:

- Home interface
- Remote interface
- Object identity
- Metadata interface
- Handle

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). The home interface is specified by the Bean Provider; the Container creates a class that implements the home interface. The home interface extends the `javax.ejb.EJBHome` 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 **remote interface**. The remote interface defines the business methods callable by the client. The remote interface is specified by the Bean Provider; the Container creates a class that implements the remote interface. The remote interface extends the `javax.ejb.EJBObject` interface. The `javax.ejb.EJBObject` interface defines the operations that allow the client to access the EJB object's identity and create a persistent handle for the EJB object.

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` interface. The client may also test if two object references refer to the same entity object.

A 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.)

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 provider to delegate the client method invocation to these methods.
- The requirement for the Bean Provider to implement the `ejbCreate`, `ejbPostCreate`, 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 provider to invoke these methods during an EJB object creation, removal, and lookup.
- The requirement for the Bean Provider to define the enterprise bean's home and remote interfaces. The requirement for the Container Provider to provide classes that implement these interfaces.
- For sessions, the requirement for the Bean Provider to implement the Container callbacks defined in the `javax.ejb.SessionBean` interface, and optionally the

[1] In special situations, the primary key type can be specified at deployment time (see subsection 9.4.7.3).

`javax.ejb.SessionSynchronization` interfaces. The requirement for the Container to invoke these callbacks at the appropriate times.

- For entities, 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 Provider to implement persistence for entity beans with container-managed persistence.
- The requirement for the Container Provider to provide the `javax.ejb.SessionContext` interface to session bean instances, and the `javax.ejb.EntityContext` interface to entity 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 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 remote and home 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 remote client running outside of the Container, the client-view API is also applicable to clients that are enterprise Beans deployed in the same Container.

4.3 Session and entity objects

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

- A session object.
- An entity 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.

*Session beans are intended to be stateful. The EJB specification also defines a **stateless Session bean** as a special case of a Session Bean. There are minor differences in the API between stateful (normal) 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.

4.4 Standard mapping to CORBA protocols

To help interoperability for EJB environments that include systems from multiple vendors, we define a standard mapping of the Enterprise JavaBeans client-view contract to the CORBA protocols.

The use of the EJB to CORBA mapping by the EJB Server is not a requirement for EJB 1.1 compliance. A later release of the J2EE platform is likely to require that the J2EE platform vendor implement the EJB to CORBA mapping.

The EJB-to-CORBA mapping covers:

1. Mapping of the EJB remote and home interfaces to RMI-IIOP. This mapping is an identity mapping because every remote and home interface is an RMI-IIOP interface.
2. Propagation of transaction context over IIOP.
3. Propagation of security context over IIOP.
4. Interoperable naming service.

The EJB-to-CORBA mapping not only enables on-the-wire interoperability among multiple vendors' implementations of the EJB Container, but also enables non-Java clients to access server-side applications written as enterprise Beans through standard CORBA APIs.


The EJB-to-CORBA mapping depends on the standard CORBA Object Services protocols for the propagation of the transaction and security context.

The CORBA mapping is defined in an accompanying document [8].

While the EJB-to-CORBA mapping defines the mapping of the EJB application interfaces and transaction interoperability, the mapping must be used in conjunction with other CORBA standards to ensure full "on-the-wire" interoperability. For example, multiple EJB servers must agree on the security protocol to achieve seamless interoperability.

The following figure illustrates a heterogeneous environment that includes systems from five different vendors.

Figure 2 Heterogeneous EJB Environment



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.

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

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

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 the client.

The 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 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.

Multiple enterprise beans can be installed in a container. The container allows the clients 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 an session object is the same, irrespective of the implementation of the session bean and the container.

5.2 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.

5.2.1 Locating a session bean's home interface

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

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

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 enterprise bean.


The lifecycle of the distributed object implementing the home interface (the `EJBHome` 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 home interface object reference to another application. The receiving application can use the home interface in the same way that it would use a home interface object reference obtained via JNDI.

5.2.2 What a container provides

The following diagram illustrates the view that a container provides to clients of session beans.

Figure 3 Client View of session beans deployed in a Container



5.3 Home interface

A Container implements the home interface of the enterprise bean installed in the container. The object that implements a session bean's home interface is called a session EJBHome object. The container makes the session beans' home interfaces available to the client through JNDI.

The 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 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 home interface.

5.3.1 Creating a session object

A home interface defines one or more `create(...)` 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 following example illustrates a home interface that defines a single `create(...)` method:

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

The following example illustrates how a client creates a new session object using a `create(...)` method of the `CartHome` interface:

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

5.3.2 Removing a session object

A client may remove a session object using the `remove()` method on 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.Home.remove(Object primaryKey)` method on a session results in the `javax.ejb.RemoveException`.

5.4 EJBOject

A client never directly accesses instances of the session bean's class. A client always uses the session bean's remote interface to access a session bean's instance. The class that implements the session bean's remote interface is provided by the container; its instances are called session `EJBOjects`.

A session `EJBOject` supports:

- The business logic methods of the object. The session `EJBOject` delegates invocation of a business method to the session bean instance.
- The methods of the `javax.ejb.EJBOject` interface. These methods allow the client to:
 - Get the session object's home interface.
 - Get the session object's handle.
 - 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.EJBOject` interface is provided by the container. They are not delegated to the instances of the session bean class.

5.5 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 `EJBOject.getPrimaryKey()` and `EJBHome.remove(Object primaryKey)` methods result in a `java.rmi.RemoteException` 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 store. 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.)


The client code must use the `javax.rmi.PortableRemoteObject.narrow(...)` method to convert the result of the `getEJBOject()` method invoked on a handle to the remote interface type.

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.

5.6 Client view of session object's life cycle

From a client point of view, the life cycle of a session object is illustrated below

Figure 4 Lifecycle 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 remote interface.

A client that has a reference to a session object 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 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 invocations on a session object that does not exist result in `java.rmi.NoSuchObjectException`.

5.7 Creating and using a session object

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

Figure 5 Session Bean Example Objects



A client creates a `Cart` session object (which provides a shopping service) using a `create(...)` method of the `Cart`'s 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, then using it later to reestablish access to the original `Cart`.

For the following example, we start by looking up the `Cart`'s home interface in JNDI. We then use the 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.create(...);
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)javadoc.rmi.PortableRemoteObject.narrow(
    cartHandle.getEJBObject(), Cart.class);
cart.purchase();
cart.remove();
```

5.8 Object identity

5.8.1 Stateful session beans

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

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

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
    ...
}
```

5.8.2 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)` method 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
    ...
}
```

5.8.3 `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`.

5.9 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 home and remote interface.

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

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 responsibility for managing session bean state.

6.1 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:

- 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.
- 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.

On the other hand, the 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.

6.2 Goals

The goal of the session bean model is to make developing a session bean as simple as developing the same functionality directly in a client.

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.

6.3 A container's management of its working set

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 session bean instance when the instance is **not** in a transaction.

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.

6.4 Conversational state

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 such a session bean must close and open the resources in the `ejbPassivate` and `ejbActivate` notifications.

6.4.1 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^[2].
- A `null`.
- An enterprise bean's remote interface reference, even if the stub class is not serializable.
- An enterprise bean's home interface reference, even if the stub class 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.
- An object that is not directly serializable, but becomes serializable by replacing the references to an enterprise bean's remote and home 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 remote interfaces in the conversational state.

[2] 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.

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 remote 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 remote and home 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 may use, for example, the object replacement technique that is part of the `java.io.ObjectOutputStream` and `java.io.ObjectInputStream` protocol to externalize the remote and home 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^[3]. Declaring the session bean's fields as `transient` is, in general, discouraged.

[3] 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.

6.4.2 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.

6.5 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.

6.5.1 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.

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 session bean instance when it is passivated, most session beans can ignore these notifications. Their purpose is to allow 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.

6.5.2 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 home interface.
- The `getCallerPrincipal` method returns the `java.security.Principal` that identifies the invoker of the bean instance's EJB object.
- 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.

6.5.3 The optional *SessionSynchronization* interface

A 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 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.

6.5.4 Business method delegation

The session bean's remote 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 EJBObject class delegates an invocation of a business method to the matching business method that is implemented in the session bean class.

6.5.5 Session bean's `ejbCreate(...)` methods

A client creates a session bean instance using one of the `create` 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.

The container creates an instance of a session bean in three steps. First, the container calls the bean class' `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 whose signature matches the signature of the `create` method invoked by the client. The input parameters sent from the client are passed to the `ejbCreate` method.

Each session bean class must have at least one `ejbCreate` method. The number and signatures of a session bean's `create` methods are specific to each session bean class.

Since a 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.

6.5.6 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`, `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 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, the container must throw the `java.rmi.RemoteException` to the second client. One implication of this rule is that it is illegal to make a “loopback” call to a session bean instance. An example of a loopback call is when a client calls instance A, instance A calls instance B, and B calls A. The loopback call attempt from B to A would result in the container throwing the `java.rmi.RemoteException` to B.

6.5.7 Transaction context of session bean methods

The implementation of a business method defined in the remote interface 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 11.6.3 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` method 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.

6.6 STATEFUL Session Bean State Diagram

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

Figure 6 Lifecycle 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 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(...)` on the instance and returns the remote reference of the session object 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 11 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` is delivered to 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` to the 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 remote interface to remove the session object, the container issues `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. (This exception is a subclass of `java.rmi.RemoteException`). 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.

Notes:

1. 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.

6.6.1 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, 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.

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> JNDI access to java:comp/env	SessionContext methods: <i>getEJBHome</i> JNDI access to java:comp/env
ejbCreate ejbRemove ejbActivate ejbPassivate	SessionContext methods: <i>getEJBHome</i> , <i>getCallerPrincipal</i> , <i>isCallerInRole</i> , <i>getEJBObject</i> JNDI access to java:comp/env Resource manager access Enterprise bean access	SessionContext methods: <i>getEJBHome</i> , <i>getCallerPrincipal</i> , <i>isCallerInRole</i> , <i>getEJBObject</i> , <i>getUserTransaction</i> UserTransaction methods JNDI access to java:comp/env Resource manager access Enterprise bean access
business method from remote interface	SessionContext methods: <i>getEJBHome</i> , <i>getCallerPrincipal</i> , <i>getRollback-Only</i> , <i>isCallerInRole</i> , <i>setRollback-Only</i> , <i>getEJBObject</i> JNDI access to java:comp/env Resource manager access Enterprise bean access	SessionContext methods: <i>getEJBHome</i> , <i>getCallerPrincipal</i> , <i>isCallerInRole</i> , <i>getEJBObject</i> , <i>getUserTransaction</i> UserTransaction methods JNDI access to java:comp/env Resource manager access Enterprise bean access
afterBegin beforeCompletion	SessionContext methods: <i>getEJBHome</i> , <i>getCallerPrincipal</i> , <i>getRollback-Only</i> , <i>isCallerInRole</i> , <i>setRollback-Only</i> , <i>getEJBObject</i> JNDI access to java:comp/env Resource manager access Enterprise bean access	N/A (a bean with bean-managed transaction demarcation cannot implement the SessionSynchronization interface)
afterCompletion	SessionContext methods: <i>getEJBHome</i> , <i>getCallerPrincipal</i> , <i>isCallerInRole</i> , <i>getEJBObject</i> JNDI access to java:comp/env	

Notes:

- The `ejbCreate`, `ejbRemove`, `ejbPassivate`, and `ejbActivate` methods of a session bean with container-managed transaction demarcation execute with an unspecified transaction context. Refer to Subsection 11.6.3 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` methods is disallowed in the session bean methods in which there is no session object identity established for the instance.
- Invoking the `getCallerPrincipal` and `isCallerInRole` methods is disallowed in the session bean methods for which the Container does not have a client security context.
- 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.

6.6.2 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 12.

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`.

6.6.3 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 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.

6.6.4 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. The container throws the `java.rmi.RemoteException` to the client in such a case.
- 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 remote or home 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.

6.7 Object interaction diagrams for a STATEFUL session bean

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

6.7.1 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.

6.7.2 Creating a session object


The following diagram illustrates the creation of a session object.

Figure 7 OID for Creation of a session object



6.7.3 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.

6.7.4 Committing a transaction

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


Figure 9 OID for session object transaction synchronization



6.7.5 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



6.7.6 Removing a session object

The following diagram illustrates the removal of a session object.

Figure 11 OID for the removal of a session object

6.8 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 serving 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 home interface of a stateless session bean must have one `create` method that takes no arguments and returns the session bean’s remote 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.

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.

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.

6.8.1 Stateless session bean state diagram

When a client calls a method on a stateless session object, 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 Lifecycle of a STATELESS Session bean

The following steps describe the lifecycle of a session bean instance:

- A stateless session bean instance's life starts when the container invokes `newInstance()` 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.
- 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.

6.8.2 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, 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 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> JNDI access to java:comp/env	SessionContext methods: <i>getEJBHome</i> JNDI access to java:comp/env
ejbCreate ejbRemove	SessionContext methods: <i>getEJBHome</i> , <i>getEJBObject</i> JNDI access to java:comp/env	SessionContext methods: <i>getEJBHome</i> , <i>getEJBObject</i> , <i>getUserTransaction</i> UserTransaction methods JNDI access to java:comp/env
business method from remote interface	SessionContext methods: <i>getEJBHome</i> , <i>getCallerPrincipal</i> , <i>getRollback-Only</i> , <i>isCallerInRole</i> , <i>setRollback-Only</i> , <i>getEJBObject</i> JNDI access to java:comp/env Resource manager access Enterprise bean access	SessionContext methods: <i>getEJBHome</i> , <i>getCallerPrincipal</i> , <i>isCallerInRole</i> , <i>getEJBObject</i> , <i>getUserTransaction</i> UserTransaction methods JNDI access to java:comp/env Resource manager access Enterprise bean access

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` method is disallowed in the session bean methods in which there is no session object identity associated with the instance.
- Invoking the `getCallerPrincipal` and `isCallerInRole` methods is disallowed in the session bean methods for which the Container does not have a client security context.

- 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.
- 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.

6.8.3 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 12.

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.


6.9 Object interaction diagrams for a STATELESS session bean

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

6.9.1 Client-invoked `create()`

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


Figure 13 OID for creation of a STATELESS session object



6.9.2 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



6.9.3 Client-invoked *remove()*

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


Figure 15 OID for removal of a STATELESS session object



6.9.4 Adding 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



6.10 The responsibilities of the bean provider

This section describes the responsibilities of session bean provider to ensure that a session bean can be deployed in any EJB Container.

6.10.1 Classes and interfaces

The session bean provider is responsible for providing the following class files:

- Session bean class.
- Session bean's remote interface.
- Session bean's home interface.

6.10.2 Session bean class

The following are the requirements for 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 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 remote interface^[4].
- 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.
- The session bean class may have superclasses and/or superinterfaces. If the session bean has superclasses, then the business methods, the `ejbCreate` methods, the methods of the `SessionBean` interface, 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.

[4] If the session bean class does implement the remote 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 remote interface in the session bean class.

6.10.3 *ejbCreate* methods

The session bean class must define one or more `ejbCreate(...)` methods whose signatures 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 arguments must be legal types for RMI/IIOP.
- 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 is deprecated in EJB 1.1—an EJB 1.1 compliant enterprise bean should throw the `javax.ejb.EJBException` or another `RuntimeException` to indicate non-application exceptions to the Container (see Section 12.2.2).

6.10.4 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.
- 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 is deprecated in EJB 1.1—an EJB 1.1 compliant enterprise bean should throw the `javax.ejb.EJBException` or another `RuntimeException` to indicate non-application exceptions to the Container (see Section 12.2.2).

6.10.5 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 clause 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.

6.10.6 Session bean's home interface

The following are the requirements for the session 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 values must be of valid types for RMI/IIOP, and that their throws clause must include the `java.rmi.RemoteException`.
- The 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 home interface must define one or more `create(...)` methods.
- Each `create` method must be named "**create**", and it must match one of the `ejbCreate` methods defined in the session bean class. The matching `ejbCreate` method must have the same number and types of arguments. (Note that the return type is different.)
- The return type for a `create` method must be the session bean's remote interface type.
- All the exceptions defined in the throws clause of an `ejbCreate` method of the session bean class must be defined in the throws clause of the matching `create` method of the home interface.
- The throws clause must include `javax.ejb.CreateException`.

6.11 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.

6.11.1 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 enterprise 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 home interface (session EJBHome class).
- A class that implements the session bean's remote interface (session EJBObject class).

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.

6.11.2 Session EJBHome class

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

The implementation of each `create(...)` method invokes a matching `ejbCreate(...)` method.

6.11.3 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.

6.11.4 Handle classes

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

6.11.5 EJBMetaData class

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

6.11.6 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 must throw the `java.rmi.RemoteException` to the second request.

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.

6.11.7 Transaction scoping, security, exceptions

The container must follow the rules with respect to transaction scoping, security checking, and exception handling, as described in Chapters 11, 15, and 12, respectively.

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).


7.1 Overview

Wombat Inc. has developed the `CartBean` session Bean. The `CartBean` is deployed in a container provided by the Acme Corporation.

7.2 Inheritance relationship

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

Figure 18 Example of Inheritance Relationships Between EJB Classes



7.2.1 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(...)` methods invoked at an EJB object creation.*
- *Define a home interface (CartHome) for the enterprise Bean. The 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.*

7.2.2 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.

7.2.3 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 home interface. The tools also generate the classes that implement the communication protocol specific artifacts for the 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.

8.1 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.

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

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.

The 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.

A client of an entity object can be another enterprise bean deployed in the same or different Container; or a client can be an arbitrary Java program, such as an application, applet, or servlet. The 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.

Multiple enterprise beans can be deployed in a container. For each entity bean deployed in a container, the container provides a class that implements the entity bean's **home interface**. The home interface allows the client to create, find, and remove entity objects within the enterprise bean's home. 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 the same, irrespective 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.

8.2 EJB Container

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

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 container makes the entity beans' home interfaces (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 API to manipulate the container, and there is no way for a client to tell in which container an enterprise bean is installed.

8.2.1 Locating an entity bean's home interface

A client locates an entity bean's home interface using JNDI. For example, the 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);
```

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.

8.2.2 What a container provides

The following diagram illustrates the view that a container provides to the clients of the entity beans deployed in the container.

Figure 19 Client view of entity beans deployed in a container



8.3 Entity bean's home interface

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

The entity bean's 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.
- 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 meta-data information about the entity bean. The meta-data 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 home interface must extend the `javax.ejb.EJBHome` interface and follow the standard rules for Java programming language remote interfaces.

8.3.1 *create methods*

An entity bean's home interface can define zero or more `create(...)` 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 return type of a `create` method is the entity bean's remote interface.

The throws clause of every `create` method includes the `java.rmi.RemoteException` and the `javax.ejb.CreateException`. It may include additional application-level exceptions.

The following home interface illustrates two 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;
    ...
}
```

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

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

8.3.2 *finder methods*

An entity bean's home interface defines one or more `finder` methods^[5], 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 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 Subsection 9.1.8).

The throws clause of every finder method includes the `java.rmi.RemoteException` and the `javax.ejb.FinderException`.

The home interface of every entity bean includes the `findByPrimaryKey(primaryKey)` method that 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. 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");
```

8.3.3 *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 client result in the `java.rmi.NoSuchObjectException`.


[5] The `findByPrimaryKey(primaryKey)` method is mandatory for all Entity Beans.

8.4 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 interface.)

Figure 20 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 of the home interface (e.g. by inserting a database record), but still may be accessible by a container’s clients 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.

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 home interface.
- Obtain the reference from the entity object’s handle. (see Section 8.7)

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 home interface.
- Pass the reference as a parameter or return value of a remote 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.

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 remote and home interfaces held by clients.

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

8.5 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 Subsection 9.2.9. 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 remote 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.)

A client can test whether two entity object references refer to the same entity object by using the `isIdentical(EJBObject)` 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 {
    acc2 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.

8.6 Entity Bean's remote interface

A client accesses 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 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 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.

8.7 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.

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.
//
ObjectOutputStream 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.
//
ObjectInputStream 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 a 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.

8.8 Entity home handles

The EJB specification allows the client to obtain a handle for the home interface. The client can use the home handle to store a reference to an entity bean’s 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 home interface in the future, but does not know the JNDI name of the home interface.

A handle to a 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 the 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.

8.9 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 home and remote interface.

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

Entity Bean Component Contract

Note: Container support for entity beans is a mandatory feature in the EJB 1.1 release.


The entity bean component contract 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 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.

9.1 Concepts

9.1.1 Runtime execution model

This section describes the runtime model and the classes used in the description of the contract between an entity bean and its container.

Figure 21 Overview of the entity bean runtime execution model

An **enterprise bean instance** is an object whose class was provided by the Bean Provider.

An entity **EJBObject** is an object whose class was generated at deployment time by the Container Provider's tools. The entity EJBObject class implements the entity bean's remote interface. A client never references an entity bean instance directly—a client always references an entity EJBObject whose class is generated by the Container Provider's tools.

An entity **EJBHome** object provides the life cycle operations (create, remove, find) for its entity objects. The class for the entity EJBHome object is generated by the Container Provider's tools at deployment time. The entity EJBHome object implements the entity bean's home interface that was defined by the Bean Provider.

9.1.2 Granularity of entity beans

This section provides guidelines to the Bean Providers for modeling of business objects as entity beans.

In general, an entity bean should represent an independent business object that has an independent identity and lifecycle, and is referenced by multiple enterprise beans and/or clients.

A *dependent object* should not be implemented as an entity bean. Instead, a dependent object is better implemented as a Java class (or several classes) and included as part of the entity bean on which it depends.

A dependent object can be characterized as follows. An object B is a dependent object of an object A, if B is created by A, accessed only by A, and removed by A. This implies, for example, that if B exists when A is being removed, B is automatically removed as well. It also implies that other programs can access the object B only indirectly through object A. In other words, the object A fully manages the lifecycle of the object B.

For example, a purchase order might be implemented as an entity bean, but the individual line items on the purchase order should be implemented as helper classes, not as entity beans. An employee record might be implemented as an entity bean, but the employee address and phone number should be implemented as helper classes, not as entity beans.

The state of an entity object that has dependent objects is often stored in multiple records in multiple database tables.

In addition, the Bean Provider must take into consideration the following factors when making a decision on the granularity of an entity object:

- Every method call to an entity object via the remote and home interface is potentially a remote call. Even if the calling and called entity bean are collocated in the same JVM, the call must go through the container, which must create copies of all the parameters that are passed through the interface by value (i.e. all parameters that do not extend the `java.rmi.Remote` interface). The container is also required to check security and apply the declarative transaction attribute on the inter-component calls. The overhead of an inter-component call will likely be prohibitive for object interactions that are too fine-grained.
- The EJB deployment descriptor does not provide a mechanism for describing object schemas (the relationships among the fine-grained objects, and how fine-grained objects are mapped to the underlying database). If these relationships need to be visible at deployment time, the information describing the relationships must be passed from the Bean Provider to the Deployer through some means outside of the EJB specification.


9.1.3 Entity persistence (data access protocol)

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 allows the entity Bean Provider either 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 (bean-managed persistence), or to delegate the entity bean's persistence to the Container Provider tools used at deployment time (container-managed persistence).

In many cases, the underlying data source may be an existing application rather than a database.

Figure 22 Client view of underlying data sources accessed through entity bean



9.1.3.1 Bean-managed persistence

In the bean-managed persistence case, 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(...)`, `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.

We expect that most enterprise beans 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. This EJB specification does not define the architecture for data access objects or strategies.

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.

If the data access calls are encapsulated in data access components, the data access components may optionally provide deployment interfaces to allow adapting data access to different schemas or even to a different database type. These data access component strategies are beyond the scope of the EJB specification.

9.1.3.2 Container-managed persistence

In the container-managed persistence case, the Bean Provider does not write the database access calls in the entity bean. Instead, the Container Provider's tools generate the database access calls at the entity bean's deployment time (i.e. when the entity bean is installed into a container). The entity Bean Provider must specify in the deployment descriptor the list of the instance fields for which the container provider tools must generate access calls.


The advantage of using container-managed persistence is that the entity bean can be largely independent from the data source in which the entity is stored. The container tools can generate classes that use JDBC or SQLJ to access the entity state in a relational database, or 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.

The disadvantage is that sophisticated tools must be used at deployment time to map the entity bean's fields to a data source. These tools and containers are typically specific to each data source.

The essential difference between an entity with bean-managed persistence and one with container-managed persistence is that in the bean-managed case, the data access components are provided as part of the entity bean, whereas in the container-managed case, the data access components are generated at deployment time by the container tools.

9.1.4 Instance life cycle

Figure 23 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). The instance does **not** move to the ready state during the execution of a finder 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. There are two possible transitions from the pooled to the ready state: through the `ejbCreate(...)` and `ejbPostCreate(...)` methods, or through the `ejbActivate()` method. The container invokes the `ejbCreate(...)` and `ejbPostCreate(...)` 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.
- 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. Invocations of the `ejbLoad()` and `ejbStore()` methods can be arbitrarily mixed with invocations of business methods. 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 two possible transitions from the ready to the pooled state: through the `ejbPassivate()` method and through the `ejbRemove()` method. 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 remote interface, or one of the `remove()` methods on the entity bean's home interface).
- 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 remote interface because the container can use a different entity bean instance to delegate the client’s requests. Exception handling is described further in Chapter 12.
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.

9.1.5 The entity bean component contract

This section specifies the contract between an entity bean and its container. The contract specified here assumes the use of bean-managed persistence. The differences in the contract for container-managed persistence are defined in Section 9.4.

9.1.5.1 Entity bean instance’s view:

The following describes the entity bean instance’s view of the contract:

The entity 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 11.6.3 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(...);`

There are zero^[6] or more `ejbCreate(...)` methods, whose signatures match the signatures of the `create(...)` methods of the entity bean home interface. The container invokes an `ejbCreate(...)` method on an entity bean instance when a client invokes a matching `create(...)` method to create an entity object.

The implementation of the `ejbCreate(...)` 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 must return the primary key for the created entity object.

An `ejbCreate(...)` method executes in the transaction context determined by the transaction attribute of the matching `create(...)` method, as described in subsection 11.6.2.

- `public void ejbPostCreate(...);`

For each `ejbCreate(...)` method, there is a matching `ejbPostCreate(...)` method that has the same input parameters but the return value is `void`. The container invokes the matching `ejbPostCreate(...)` method on an instance after it invokes the `ejbCreate(...)` method with the same arguments. The entity object identity is available during the `ejbPostCreate(...)` method. The instance may, for example, obtain the remote interface of the associated entity object and pass it to another enterprise bean as a method argument.

An `ejbPostCreate(...)` method executes in the same transaction context as the previous `ejbCreate(...)` 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 instance can obtain the identity of the entity object via the `getPrimaryKey()` 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()`.

[6] An entity enterprise Bean has no `ejbCreate(...)` and `ejbPostCreate(...)` methods if it does not define any create methods in its home interface. Such an entity enterprise Bean does not allow the clients to create new EJB objects. The enterprise Bean restricts the clients to accessing entities that were created through direct database inserts.

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()` 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()` 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 in 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 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 invoked on the instance. All business methods invoked between the previous `ejbLoad` or `ejbCreate` 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 11.6.2.

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 9.1.8).

9.1.5.2 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(...);`
`public void ejbPostCreate(...);`

The container invokes these two methods during the creation of an entity object as a result of a client invoking a `create(...)` method on the entity bean's home interface.

The container first invokes the `ejbCreate(...)` method whose signature matches the `create(...)` method invoked by the client. The `ejbCreate(...)` method returns a primary key for the created entity object. The container creates an entity `EJBObject` reference for

the primary key. The container then invokes a matching `ejbPostCreate(...)` 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.

The container must invoke the `ejbCreate(...)` and `ejbPostCreate(...)` methods in the transaction context determined by the transaction attribute of the matching `create(...)` method, as described in subsection 11.6.2.

- `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()` 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.

- `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()` 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()` 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 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 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 11.7).

The container must invoke this method in the same transaction context as the previously invoked `ejbLoad` or `ejbCreate` 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.

After the `ejbFind<METHOD>(...)` method completes, the instance remains in the pooled state. The container may, but is not required to, 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 11.6.2.

If the `ejbFind<METHOD>` method is declared to return a single primary key, the container creates an entity `EJBO` reference for the primary key and returns it to the client. If the `ejbFind<METHOD>` method is declared to return a collection of primary keys, the container creates a collection of entity `EJBO` references for the primary keys returned from `ejbFind<METHOD>`, and returns the collection to the client. (See Subsection 9.1.8 for information on collections.)

9.1.6 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, 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 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> JNDI access to java:comp/env
ejbCreate	EntityContext methods: <i>getEJBHome</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
ejbPostCreate	EntityContext methods: <i>getEJBHome</i> , <i>getCallerPrincipal</i> , <i>getRollbackOnly</i> , <i>isCallerInRole</i> , <i>setRollbackOnly</i> , <i>getEJBObject</i> , <i>getPrimaryKey</i> JNDI access to java:comp/env Resource manager access Enterprise bean access
ejbRemove	EntityContext methods: <i>getEJBHome</i> , <i>getCallerPrincipal</i> , <i>getRollbackOnly</i> , <i>isCallerInRole</i> , <i>setRollbackOnly</i> , <i>getEJBObject</i> , <i>getPrimaryKey</i> JNDI access to java:comp/env Resource manager access Enterprise bean access
ejbFind	EntityContext methods: <i>getEJBHome</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
ejbActivate ejbPassivate	EntityContext methods: <i>getEJBHome</i> , <i>getEJBObject</i> , <i>getPrimaryKey</i> JNDI access to java:comp/env
ejbLoad ejbStore	EntityContext methods: <i>getEJBHome</i> , <i>getCallerPrincipal</i> , <i>getRollbackOnly</i> , <i>isCallerInRole</i> , <i>setRollbackOnly</i> , <i>getEJBObject</i> , <i>getPrimaryKey</i> JNDI access to java:comp/env Resource manager access Enterprise bean access
business method from remote interface	EntityContext methods: <i>getEJBHome</i> , <i>getCallerPrincipal</i> , <i>getRollbackOnly</i> , <i>isCallerInRole</i> , <i>setRollbackOnly</i> , <i>getEJBObject</i> , <i>getPrimaryKey</i> JNDI access to java:comp/env Resource manager access Enterprise bean access

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` and `getPrimaryKey` methods is disallowed in the entity bean methods in which there is no entity object identity associated with the instance.
- Invoking the `getCallerPrincipal` and `isCallerInRole` methods is disallowed in the entity bean methods for which the Container does not have a client security context.
- 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.

9.1.7 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.

9.1.7.1 `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`^[7] transaction attribute is assigned to a remote interface method, the corresponding enterprise bean class method executes with an unspecified transaction context (See Subsection 11.6.3). 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.

[7] This applies also to the `Never` and `Supports` attribute.

9.1.8 Finder method return type

9.1.8.1 Single-object finder

Some finder methods (such as `ejbFindByPrimaryKey`) are designed to return at most one entity object. For these single-object finders, the result type of the `find<METHOD>(...)` method defined in the entity bean's home interface is the entity bean's remote 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.

The following code illustrates the definition of a single-object finder.

```
// Entity's home interface
public AccountHome extends javax.ejb.EJBHome {
    ...
    Account findByPrimaryKey(AccountPrimaryKey primkey)
        throws FinderException, RemoteException;
    ...
}

// Entity's implementation class
public AccountBean implements javax.ejb.EntityBean {
    ...
    public AccountPrimaryKey ejbFindByPrimaryKey(
        AccountPrimaryKey primkey)
        throws FinderException
    {
        ...
    }
    ...
}
```

9.1.8.2 Multi-object finders

Some finder methods are designed to return multiple entity objects. For these multi-object finders, the result type of the `find<METHOD>(...)` method defined in the entity bean's home interface is a *collection* of objects implementing the entity bean's remote interface. 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 JDK™ 1.1 `java.util.Enumeration` interface
- the Java™ 2 `java.util.Collection` interface

A Bean Provider that 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^[8].

[8] The finder will be also compatible with Java 2-based Containers and Clients.

A Bean Provider targeting only containers and clients based on Java 2 can use the `java.util.Collection` interface for the finder's result type.

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.

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 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
    {
        ...
    }
    ...
}
```

The following is an example of a multi-object finder method definition that is compatible only with containers and clients based on Java 2:

```
// Entity's 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
    {
        ...
    }
    ...
}
```

9.1.9 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 bean-managed persistence. The use of the exceptions by entity beans with container-managed persistence is the same, with one additional element: The responsibilities for throwing the exceptions apply to the data access methods generated by the Container Provider's tools.

9.1.9.1 CreateException

From the client's perspective, a `CreateException` (or a subclass of `CreateException`) indicates that an application level error occurred during the `create(...)` 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.)

The Bean Provider throws the `CreateException` (or subclass of `CreateException`) from the `ejbCreate(...)` and `ejbPostCreate(...)` 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(...)` methods are invalid.

The Container treats the `CreateException` as any other application exception. See Section 12.3.

9.1.9.2 DuplicateKeyException

The `DuplicateKeyException` is a subclass of `CreateException`. It is thrown by the `ejbCreate(...)` methods 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.

9.1.9.3 `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 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>(...)` methods 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 12.3.

9.1.9.4 `ObjectNotFoundException`

The `ObjectNotFoundException` is a subclass of `FinderException`. It is thrown by the `ejbFind<METHOD>(...)` methods to indicate that the requested entity object does not exist.

Only single-object finders (see Subsection 9.1.8) may throw this exception. Multi-object finders must not throw this exception.

9.1.9.5 `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 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 12.3.

9.1.10 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 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 5 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 9.5.4 illustrate the three alternative commit options in detail.


9.1.11 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 automatically by the underlying database during the database access calls performed by the business methods; and by the `ejbLoad`, `ejbCreate`, `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 9.5.4 apply to this type of container.


Figure 24 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 9.5.4 applies to this type of container.

Figure 25 Multiple clients can access the same entity object using single instance



9.1.12 Non-reentrant and re-entrant instances

An entity Bean Provider entity 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 the `java.rmi.RemoteException` 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 remote 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.) The container must reject the loopback call and throw the `java.rmi.RemoteException` 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 great 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.

9.2 Responsibilities of the Enterprise Bean Provider

This section describes the responsibilities of an entity Bean Provider to ensure that the entity bean can be deployed in any EJB Container.

The requirements are stated for the provider of an entity bean with bean-managed persistence. The differences for entities with container-managed persistence are defined in Section 9.4.

9.2.1 **Classes and interfaces**

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

- Entity bean class and any dependent classes.
- Entity bean's remote interface
- Entity bean's home interface
- Primary key class

9.2.2 **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 must be defined as `public` and must not be `abstract`.

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 remote interface^[9]. If the class implements the entity bean's remote interface, the class must provide no-op implementations of the methods defined in the `javax.ejb.EJBObject` 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`, `ejbPostCreate`, and `ejbFind<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 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.

9.2.3 *ejbCreate* methods

The entity bean class may define zero or more `ejbCreate(. . .)` methods whose signatures 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 the entity bean's primary key type.

The method argument and return value types must be legal types for RMI-IIOP.

The `throws` clause may define arbitrary application specific exceptions, 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 is deprecated in EJB 1.1—an EJB 1.1 compliant enterprise bean should throw the `javax.ejb.EJBException` or another `java.lang.RuntimeException` to indicate non-application exceptions to the Container (see Section 12.2.2).

[9] If the entity bean class does implement the remote 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 remote interface in the entity bean class.

The entity object created by the `ejbCreate` 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 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.

9.2.4 *ejbPostCreate* methods

For each `ejbCreate(...)` method, the entity bean class must define a matching `ejbPostCreate(...)` method, using the following rules:

The method name must be `ejbPostCreate`.

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.

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 is deprecated in EJB 1.1—an EJB 1.1 compliant enterprise bean should throw the `javax.ejb.EJBException` or another `java.lang.RuntimeException` to indicate non-application exceptions to the Container (see Section 12.2.2).

9.2.5 *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.

The return type of a finder method must be the entity bean’s primary key type, or a collection of primary keys (see Section Subsection 9.1.8).

The `throws` clause may define arbitrary application specific exceptions, including the `javax.ejb.FinderException`.

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).

Compatibility Note: EJB 1.0 allowed the finder methods to throw the `java.rmi.RemoteException` to indicate a non-application exception. This practice is deprecated in EJB 1.1—an EJB 1.1 compliant enterprise bean should throw the `javax.ejb.EJBException` or another `java.lang.RuntimeException` to indicate non-application exceptions to the Container (see Section 12.2.2).

9.2.6 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`.

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 is deprecated in EJB 1.1—an EJB 1.1 compliant enterprise bean should throw the `javax.ejb.EJBException` or another `java.lang.RuntimeException` to indicate non-application exceptions to the Container (see Section 12.2.2).

9.2.7 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.

9.2.8 Entity bean's 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 that their throws clause must include the `java.rmi.RemoteException`.

The 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 home interface must be one of the following:

- A create method.
- A finder method.

Each `create` method must be named "**create**", and it must match one of the `ejbCreate` methods defined in the enterprise Bean class. The matching `ejbCreate` 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 must be the entity bean's remote interface type.

All the exceptions defined in the throws clause of the matching `ejbCreate` and `ejbPostCreate` 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` and `ejbPostCreate` methods)

The throws clause of a `create` 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 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 home interface.

The throws clause of a `finder` method must include the `javax.ejb.FinderException`.

9.2.9 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.

9.3 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 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.

9.3.1 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 home interface (i.e. the entity EJBHome class).
- A class that implements the entity bean's remote interface (i.e. the entity EJBObject 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.

9.3.2 Entity EJBHome class

The entity EJBHome class, which is generated by deployment tools, implements the entity bean's 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 invokes a matching `ejbCreate(...)` method, followed by the matching `ejbPostCreate(...)` method, passing the `create(...)` 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.

9.3.3 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.

9.3.4 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.

9.3.5 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.

9.3.6 Meta-data class

The deployment tools are responsible for implementing the class that provides meta-data information to the client view contract. The class must be a valid RMI-IIOP Value Type, and must implement the `javax.ejb.EJBMetaData` interface.

Because the meta-data class is not entity bean specific, the container may, but is not required to, use a single class for all deployed enterprise beans.

9.3.7 Instance's re-entrance

The container runtime must enforce the rules defined in Section 9.1.12.

9.3.8 Transaction scoping, security, exceptions

The container runtime must follow the rules on transaction scoping, security checking, and exception handling described in Chapters 11, 15, and 12.

9.3.9 Implementation of object references

The container should implement the distribution protocol between the client and the container such that the object references of the 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.

9.4 Entity beans with container-managed persistence

The previous sections described the component contract for entity beans with bean-managed persistence. The contract for an entity bean with container-managed persistence is the same as the contract for an entity bean with bean-managed persistence (as described in the previous sections), except for the differences described in this section.

The deployment descriptor for an entity bean indicates whether the entity bean uses bean-managed persistence or container-managed persistence.

9.4.1 Container-managed fields

An 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 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 following requirements ensure that an entity bean 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 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 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 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 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 instead should use bean-managed persistence.

The provider of entity beans with container-managed persistence must take into account the following limitations of the 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 the container-managed persistence should use bean-managed persistence instead.

9.4.2 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 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 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).

Then container invokes the matching `ejbPostCreate(...)` method on the instance. The instance can discover the primary key by calling `getPrimaryKey()` 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 11.6.2.

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.

9.4.3 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 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 11.6.2.

9.4.4 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).

9.4.5 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.

9.4.6 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 home interface.

Note that the `ejbFind<METHOD>` names and parameter signatures 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 Enterprise JavaBeans architecture does not specify the format of the finder method description.

9.4.7 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 9.2.9.

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.

9.4.7.1 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.

9.4.7.2 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.)

9.4.7.3 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.

9.5 Object interaction diagrams

This section uses object interaction diagrams to illustrate the interactions between an entity bean instance and its container.

9.5.1 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

9.5.2 Creating an entity object

Figure 26 OID of Creation of an entity object with bean-managed persistence





Figure 27 OID of creation of an entity object with container-managed persistence



9.5.3 Passivating and activating an instance in a transaction

Figure 28 OID of passivation and reactivation of an entity bean instance with bean-managed persistence





Figure 29 OID of passivation and reactivation of an entity bean instance with CMP



9.5.4 Committing a transaction

Figure 30 OID of transaction commit protocol for an entity bean instance with bean-managed persistence





Figure 31 OID of transaction commit protocol for an entity bean instance with container-managed persistence



9.5.5 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 32 OID of start of transaction for an entity bean instance with bean-managed persistence





Figure 33 OID of start of transaction for an entity bean instance with container-managed persistence



9.5.6 Removing an entity object

Figure 34 OID of removal of an entity bean object with bean-managed persistence





Figure 35 OID of removal of an entity bean object with container-managed persistence



9.5.7 Finding an entity object

Figure 36 OID of execution of a finder method on an entity bean instance with bean-managed persistence





Figure 37 OID of execution of a finder method on an entity bean instance with container-managed persistence

9.5.8 Adding and removing an instance from the pool

The diagrams in Subsections 9.5.2 through 9.5.7 did not show the sequences between the “does not exist” and “pooled” state (see the diagram in Section 9.1.4).

Figure 38 OID of a container adding an instance to the pool





Figure 39 OID of a container removing an instance from the pool



Example entity scenario

This chapter describes an example development and deployment scenario for an entity bean. 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).


10.1 Overview

Wombat Inc. has developed the `AccountBean` entity bean. The `AccountBean` entity bean is deployed in a container provided by the Acme Corporation.

10.2 Inheritance relationship

Figure 40

Example of the inheritance relationship between the interfaces and classes:



10.2.1 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(...)` and `ejbPostCreate(...)` methods invoked at an entity object creation, and the finder methods (the finders should not have to be implemented if the entity bean uses container-managed persistence).*
- *Define a home interface (AccountHome) for the entity bean. The home interface defines the entity bean's specific create and finder methods. The 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.*

10.2.2 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.*

10.2.3 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 home interface (AccountHome). The tools also generate the classes that implement the communication protocol specific artifacts for the 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.

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.

11.1 Overview

This section provides a brief overview of transactions and illustrates a number of transaction scenarios in EJB.

11.1.1 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^[10] 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 11.6.3 for the description of the "no transaction" case).

Regardless 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, transaction context propagation, and distributed two-phase commit.

11.1.2 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.

11.1.3 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.

[10] 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. These resources are considered to be "managed" by the Container.

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 `javax.transaction.UserTransaction` interface defined in JTA, but it does not require the support for the JTA resource and application server interfaces.

11.2 Sample scenarios


This section describes several scenarios that illustrate the distributed transaction capabilities of the Enterprise JavaBeans architecture.

11.2.1 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. 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. 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 41 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.

11.2.2 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*. *X* updates data in database A, and then calls another enterprise Bean *Y* that is installed in a remote EJB Server. *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 42 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.


11.2.3 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 [10].

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 43 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.


11.2.4 Container-managed demarcation

Whenever a client invokes 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 [11.4.1] 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 remote method on *X*. *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 44 Update of Multiple Databases from Non-Transactional Client



The container automatically commits the transaction at the time *X* returns a reply to the client.

11.2.5 Bean-managed demarcation


A session Bean can use the `javax.transaction.UserTransaction` interface to programmatically demarcate transactions.

11.2.6 Interoperability with non-Java clients and servers

Although the Enterprise JavaBeans architecture focuses on the Java API (and Java programming language) for writing distributed enterprise applications, it is desirable that such applications are also interoperable with non-Java clients and servers.

A container can make it possible for an enterprise Bean to be invoked from a non-Java client. For example, the CORBA mapping of the Enterprise JavaBeans architecture [8] allows any CORBA client to invoke any enterprise Bean object on a CORBA-enabled server using the standard CORBA API.

Figure 45 Interoperating with Non-Java Clients and/or Servers



Providing connectivity to existing server applications is also important. An EJB Server may choose to provide access to existing enterprise applications, such as applications running under CICS on a mainframe. For example, an EJB Server may provide a bridge that makes existing CICS programs accessible to enterprise Beans. The bridge can make the CICS programs visible to the Java programming language-based developer as if the CICS programs were other enterprise Beans installed in some container on the EJB Server.

Note: It is beyond the scope of the Enterprise JavaBeans specification to define the bridging protocols that would enable such interoperability.

11.3 Bean Provider's responsibilities

This section describes the Bean Provider's view of transactions and defines his responsibilities.

11.3.1 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 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. Refer to Table 2 on page 60, Table 3 on page 70, and Table 4 on page 111.

11.3.1.1 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 11.6.3 for how the Container implements this case.

11.3.2 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 level.)
- 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 levels. 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 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.

- For entity beans using 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.

11.3.3 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.

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).

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 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/jdbcDatabase1");
        con1 = ds1.getConnection();

        stmt1 = con1.createStatement();

        // obtain con2 object and set it up for transactions
        ds2 = (javax.sql.DataSource)
            initCtx.lookup("java:comp/env/jdbcDatabase2");
        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 stateful Session Bean that retains a transaction across three client calls, invoked in the following order: *method1*, *method2*, and *method3*.

```
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/jdbcDatabase1");
        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/jdbcDatabase2");
        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
        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/jdbcDatabase");
        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();
    }
    ...
}
```

11.3.3.1 getRollbackOnly() and setRollbackOnly() method

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.

11.3.4 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 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()`.

The enterprise bean's business methods 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();
    }
    ...
}
```

11.3.4.1 `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 6.5.2.

11.3.4.2 `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.

11.3.4.3 `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.

11.3.5 Declaration in deployment descriptor

The Bean Provider of a Session Bean must use the `transaction-type` element to declare whether the Session Bean is of the bean-managed or container-managed transaction demarcation type. (See Chapter 16 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 11.4.1.

11.4 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.

11.4.1 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 an enterprise bean's remote or home interface. The transaction attribute specifies how the Container must manage transactions for a method when a client invokes the business method via the enterprise bean home or remote interface.

The transaction attribute must be specified for the following remote and home interface methods:

- For a session bean, the transaction attributes must be specified for the methods defined in the bean's remote interface and all the direct and indirect superinterfaces of the remote interface, excluding the methods of the `javax.ejb.EJBObject` interface. Transaction attributes must not be specified for the methods of a session bean's home interface.
- For an entity bean, the transaction attributes must be specified for the methods defined in the bean's remote interface and all the direct and indirect superinterfaces of the remote interface, excluding the `getEJBHome`, `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.

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 of the remote and home interfaces 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 11.6.2 for the specification of how the value of the transaction attribute affects the transaction management performed by the Container.

If an enterprise bean implements the `javax.ejb.SessionSynchronization` interface, the Application Assembler can specify only the following values 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 enterprise beans' remote and home interfaces. (See Section 16.5 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 the enterprise beans' remote and home interfaces. 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 remote 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 the remote or home interface of the specified enterprise bean. 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. The method must be one defined in the remote or home interface of the specified enterprise bean. 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 defined in both the remote and home 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>
```

11.5 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.

11.6 Container Provider responsibilities

This section defines the responsibilities of the Container Provider.

Every client method invocation on an enterprise Bean object via the bean's remote and home interface 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 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 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.

11.6.1 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 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 the enterprise bean's remote or home interface, 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 the 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 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^[11].

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.

[11] 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 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.

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 6 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
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, the container invokes the instance with an unspecified transaction context.
- If the client 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.
- 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.
- 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.

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.

11.6.2 Container-managed transaction demarcation

The Container is responsible for providing the transaction demarcation for the enterprise 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 16 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 or remote interface. The Container's responsibilities depend on the value of the transaction attribute.

11.6.2.1 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 11.6.3 for more details of how the Container can implement this case.

11.6.2.2 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.

11.6.2.3 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.

11.6.2.4 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.

11.6.2.5 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.

11.6.2.6 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 calls without a transaction context, the Container performs the same steps as described in the `NotSupported` case.

11.6.2.7 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 7 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

Table 7 Transaction attribute summary

Transaction attribute	Client's transaction	Transaction associated with business method	Transaction associated with resource managers
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 remote 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 11.6.3 for how the Container handles the "none" case in Table 7.

11.6.2.8 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.

11.6.2.9 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.

11.6.2.10 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`.

11.6.2.11 `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.

11.6.3 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`^[12].
- The execution of the `ejbCreate`, `ejbRemove`, `ejbPassivate`, and `ejbActivate` methods of a session bean with container-managed transaction demarcation.

[12] For the `Supports` attribute, the handling described in this section applies only to the case when the client calls without a transaction context.

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.


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.

11.7 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.

11.7.1 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 46 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^[13].

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.

[13] This diamond problem applies only to the case when B and C are in the same transaction.

11.7.2 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.
- 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 47 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.

11.7.3 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.

11.7.4 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.

11.7.5 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 48 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

12.1 Overview and Concepts

12.1.1 Application exceptions

An *application exception* is an exception defined in the throws clause of a method of the enterprise Bean's home and remote interfaces, 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 Subsection 9.1.9). These exceptions are covered by the rules on application exceptions that are defined in this chapter.

12.1.2 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).
- An application exception thrown by an enterprise bean instance should not automatically rollback 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.

12.2 Bean Provider's responsibilities

This section describes the view and responsibilities of the Bean Provider with respect to exception handling.

12.2.1 Application exceptions

The Bean Provider defines the application exceptions in the throws clauses of the methods of the remote and home 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.

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 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 next subsection).

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 Subsection 9.1.9.

Bean Providers may define subclasses of the standard EJB application exceptions and throw instances of the subclasses in the entity bean methods. A subclass will typically provide more information to the client that catches the exception.

12.2.2 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 business method or a container callback method (e.g. `ejbLoad`).

The enterprise bean business method and container callback methods 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^[14], unexpected `RuntimeException`, JVM errors, and so on.

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^[15] 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.

[14] 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`.

[15] A checked exception is one that is not a subclass of `java.lang.RuntimeException`.

The Container catches a non-application exception, logs it (which can result in alerting the System Administrator), and throws the `java.rmi.RemoteException` (or subclass thereof) to the client. 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.

12.2.2.1 `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.

An 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 remote interface.

12.3 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 enterprise bean's business methods.
- Exceptions from container-invoked callbacks on the enterprise bean.
- Exceptions from management of container-managed transaction demarcation.

12.3.1 Exceptions from an enterprise bean's business methods

Business methods are considered to be the methods defined in the enterprise bean's remote and home interface (including all their superinterfaces); and the following methods: `ejbCreate(...)`, `ejbPostCreate(...)`, `ejbRemove()`, and the `ejbFind<METHOD>` methods.

Table 8 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 12.4 describes the client's view of exceptions in detail.)

Table 8 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 <code>Required</code> , <code>Mandatory</code> , and <code>Supports</code> 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 <code>setRollbackOnly()</code>).
	all other exceptions and errors	Log the exception or error [Note B]. Mark the transaction for rollback. Discard instance [Note C]. Throw <code>TransactionRolledBackException</code> to the client.	Receives <code>TransactionRolledBackException</code> . 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 <code>RequiredNew</code> and <code>RequiresNew</code> attributes.	AppException	If the instance called <code>setRollbackOnly()</code> , 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 <code>RemoteException</code> .	Receives <code>RemoteException</code> . If the client executes in a transaction, the client's transaction is not marked for rollback, and client can continue its work.

Table 8 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 <code>NotSupported</code> , <code>Never</code> , and <code>Supports</code> 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.	Receives RemoteException. If the client executes in a transaction, the client's transaction is not marked for rollback, and client can continue its work.

Notes:

- [A] The caller can be another enterprise bean or an arbitrary client program.
- [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.

Table 9 specifies how the Container must handle the exceptions thrown by the business methods for beans with bean-managed transaction demarcation^[16]. 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 12.4 describes the client's view of exceptions in detail.)

Table 9 Handling of exceptions thrown by a business method of a session 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.	Receives RemoteException.

[16] Note that the EJB specification allows only Session beans to use bean-managed transaction demarcation.

12.3.2 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(EntityContext)`, and `unsetEntityContext()` methods of the `EntityBean` interface.
- The `ejbActivate()`, `ejbPassivate()`, and `setSessionContext(SessionContext)` methods of the `SessionBean` interface.
- The `afterBegin()`, `beforeCompletion()` and `afterCompletion(boolean)` 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 instance executed in the client's transaction, the Container should throw the `javax.transaction.TransactionRolledBackException` because it provides more information to the client. (The client knows that it is fruitless to continue the transaction.)

12.3.3 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 12.3.1 and 12.3.2.

To give the client a better indication of the cause of the error, the Container should throw the `java.rmi.NoSuchObjectException` to the client (which is a subclass of `java.rmi.RemoteException`).

12.3.4 Non-existing session object

If a client makes a call to a session object that has been removed, the Container should throw the `java.rmi.NoSuchObjectException` to the client (which is a subclass of `java.rmi.RemoteException`).

12.3.5 Exceptions from the management of container-managed transactions

The container is responsible for starting and committing the container-managed transactions, as described in Subsection 11.6.2. 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 the client.

However, the Container should not throw the `java.rmi.RemoteException` 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` after a number of unsuccessful tries.

12.3.6 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 14.4).

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.

12.3.7 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` is deprecated in EJB 1.1—enterprise beans written for the EJB 1.1 specification should use the `javax.ejb.EJBException` instead.

The EJB 1.1 specification requires 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 home interface still must use the `java.rmi.RemoteException` as required by the EJB specification.

12.4 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 through the enterprise Bean's remote and home interfaces. Both of these interfaces 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.

12.4.1 Application exception

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 causes 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 remote or home interface. The Bean Provider may have specified (although he 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.

12.4.2 java.rmi.RemoteException

The client receives the `java.rmi.RemoteException` 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` 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` 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` 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` to the client. The following subsections describe the several subclasses of the `java.rmi.RemoteException` that may be thrown by the Container to give the client more information.

12.4.2.1 `javax.transaction.TransactionRolledbackException`

The `javax.transaction.TransactionRolledbackException` is a subclass of the `java.rmi.RemoteException`. It is defined in the JTA standard extension.

If a client receives the `javax.transaction.TransactionRolledbackException`, 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.

12.4.2.2 `javax.transaction.TransactionRequiredException`

The `javax.transaction.TransactionRequiredException` is a subclass of the `java.rmi.RemoteException`. It is defined in the JTA standard extension.

The `javax.transaction.TransactionRequiredException` 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.

12.4.2.3 `java.rmi.NoSuchObjectException`

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.

12.5 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.

12.6 Differences from EJB 1.0

The EJB 1.1 specification of exception handling preserved 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 is 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 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 Con-

tainer would not rollback a transaction and discard the instance if the instance threw the `java.rmi.RemoteException`.

- In EJB 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

13.1 Overview

The home and remote interfaces of the enterprise bean's client view are defined as Java™ RMI [3] interfaces. This allows the Container to implement the home and remote interfaces as *distributed objects*. A client using the home and remote interfaces can reside on a different machine than the enterprise bean (location transparency), and the object references of the 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 the legal RMI-IIOP types [7]. This makes it possible for the EJB Container implementors to use RMI-IIOP as the object distribution protocol.

Note: The EJB 1.1 specification does not require Container vendors to use RMI-IIOP. A later release of the J2EE platform is likely to require a J2EE platform implementor to implement the RMI-IIOP protocol for EJB interoperability in heterogeneous server environments.

13.2 Client-side objects in distributed environment

When the RMI-IIOP protocol or similar distribution protocols are used, the client communicates with the enterprise bean using *stubs* for the server-side objects. The stubs implement the home and remote interfaces.

Figure 49 Location of EJB Client Stubs.



The communication stubs used on the client side are artifacts generated at enterprise Bean's deployment time by the EJB Container provider tools. The stubs used on the client are standard if the Container uses RMI-IIOP as the distribution protocol; the stubs are Container-specific otherwise.

13.3 Standard distribution protocol

The standard mapping of the Enterprise JavaBeans architecture to CORBA is defined in [8].

The mapping enables the following interoperability:

- A client using an ORB from one vendor can access enterprise Beans residing on an EJB Server provided by another vendor.
- Enterprise Beans in one EJB Server can access enterprise Beans in another EJB Server.
- A non-Java CORBA client can access any enterprise Bean object.

Enterprise bean environment

This chapter specifies the interfaces for accessing the enterprise bean environment.

14.1 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 14.2 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 14.3 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 14.4 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.

14.2 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. 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.

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.

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.

14.2.1 Bean Provider's responsibilities

This section describes the Bean Provider's view of the enterprise bean's environment, and defines his or her responsibilities.

14.2.1.1 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 > maxExemptions ||
            numberOfExemptions < minExemptions)
            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");
        ...
    }
}
```

14.2.1.2 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 session or entity 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`, `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.

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>
...

```

14.2.2 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.

14.2.3 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.

14.2.4 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.

14.3 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's 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 must be bound to the home of the specified target enterprise bean.

14.3.1 Bean Provider's responsibilities

This subsection describes the Bean Provider's view and responsibilities with respect to EJB references.

14.3.1.1 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 14.3.1.2 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 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 home of another enterprise bean.

14.3.1.2 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` 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` element describes the interface requirements that the referencing enterprise bean has for the referenced enterprise bean. The `ejb-ref` element contains an optional `description` element; and the mandatory `ejb-ref-name`, `ejb-ref-type`, `home`, and `remote` 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` elements specify the expected Java types of the referenced enterprise bean's home and remote interfaces.

An EJB reference is scoped to the session or entity bean whose declaration contains the `ejb-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` 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>
...

```

14.3.2 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 link will be observed by the deployment tools.

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` 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 the same `ejb-jar` file, or in another `ejb-jar` in the same J2EE application unit as the referencing enterprise bean.
- 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 remote 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.

14.3.3 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 remote interfaces of the target enterprise bean must be Java type-compatible with the home and remote interfaces declared in the EJB reference.
- If an EJB reference declaration includes the `ejb-link` element, the Deployer must bind the enterprise bean reference to the home of the enterprise bean specified as the link's target.

14.3.4 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` 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 enterprise 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 enterprise bean.

14.4 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 factories that are configured in the Container. Because these resource 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 factories allow the Container to implement connection pooling and automatic enlistment of the connection with a transaction).

14.4.1 Bean Provider's responsibilities

This subsection describes the Bean Provider's view of locating resource factories and defines his responsibilities.

14.4.1.1 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 14.4.1.2 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.*
- Look up 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 method 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 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();
        ...
    }
}
```

14.4.1.2 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; and the mandatory `res-ref-name`, `res-type`, and `res-auth` elements. The `res-ref-name` element contains the name of the environment entry used in the enterprise bean's code. 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`.

A resource manager connection factory reference is scoped to the session or entity 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 manager connection 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>
    </resource-ref>
    ...
  </session>
</enterprise-beans>
...

```

14.4.1.3 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.QueueConnectionFactory` or the `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.

Note: A future EJB specification will add the “connector” mechanism that will allow an enterprise bean to use the API described in this section to obtain resource objects that provide access to additional back-end systems.

14.4.2 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.

14.4.3 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 set 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.
- 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/pass-

word information for each resource manager connection factory reference declared by the 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.

14.4.4 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.

14.5 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 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 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");
    }
    ...
}
```

14.6 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 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 support for security management.

The deployment aspect of security management has changed significantly since EJB 1.0. These changes were made primarily to support ISV enterprise beans, which are usually written without the knowledge of the target security domain.

15.1 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 method 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 the enterprise beans' home and remote 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 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 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.

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.

15.2 Bean Provider's responsibilities

This section defines the Bean Provider's perspective of the EJB architecture support for security, and defines his responsibilities.

15.2.1 Invocation of other enterprise beans

An enterprise bean business method can invoke another enterprise bean via the other bean's remote or home interface. The EJB architecture provides neither programmatic nor deployment descriptor interfaces for the invoking enterprise bean to control the principal passed to the invoked enterprise bean.

The management of caller principals passed on 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. The default principal management (in the absence of other deployment instructions) is to propagate the caller principal 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.)

15.2.2 Resource access

Section 14.4 defines the protocol for accessing resource managers, including the requirements for security management.

15.2.3 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 15.6.8 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.

15.2.4 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.

15.2.5 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 for which the Container has a client security context, as specified in Table 2 on page 60, Table 3 on page 70, and Table 4 on page 111. If they are invoked when no security context exists, they should throw the `java.lang.IllegalStateException` runtime exception.

The `getCallerIdentity()` and `isCallerInRole(Identity role)` methods are 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 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());  
}
```

15.2.5.1 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.

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 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 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 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).

15.2.5.2 Use of `isCallerInRole(String roleName)`

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 uses the `isCallerInRole(String roleName)` 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 15.3.1), and are assigned to principals or principal groups that exist in the operational environment by the Deployer.

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 caller's
        // who have the security role "payroll"
        if (info.salary != oldInfo.salary &&
            !ejbContext.isCallerInRole("payroll")) {
            throw new SecurityException(...);
        }
        ...
    }
    ...
}
```

15.2.5.3 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.
- Optional: 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.

15.3 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 Applications 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 and remote 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 15.3.3.

15.3.1 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 and remote 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^[17] 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.

[17] 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>

```

15.3.2 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 remote and home interface 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 and remote interfaces of the enterprise beans, including all their superinterfaces (including the methods of the `EJBHome` and `EJBObject` 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.

It is possible that some methods are not assigned to any security roles. This means that none of the security roles defined by the Application Assembler needs access to the methods.

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 remote 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 remote and home 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 remote or home 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 remote or home interface.

The optional `method-intf` element can be used to differentiate methods with the same name and signature that are defined in both the remote and home interfaces.

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>
...
```

15.3.3 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>
...

```

15.4 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.

15.4.1 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.

15.4.2 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 15.3.3 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.)

15.4.3 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 `description` elements of the deployment descriptor, or in a deployment manual).

The default mode is to propagate the caller principal 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). In the absence of instructions from the Application Assembler, the Deployer should configure the enterprise beans such that this “caller propagation” mode is used when one enterprise bean calls another. This ensures that the returned value of `getCallerPrincipal()` will be the same for all the enterprise beans involved in a call chain.

15.4.4 Security management of resource access

The Deployer’s responsibilities with respect to securing resource managers access are defined in subsection 14.4.2.

15.4.5 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.

15.5 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.

15.6 EJB Container Provider's responsibilities

This section describes the responsibilities of the EJB Container and Server Provider.

15.6.1 **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 15.4.

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.

15.6.2 **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.

15.6.3 **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.).*

15.6.4 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 minimal requirement is that the EJB Container must be capable of allowing the Deployer to specify that 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.

15.6.5 Security methods in `javax.ejbEJBContext`

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 this context information during the execution of a business method invoked via the enterprise bean's remote or home interface, as defined in Table 2 on page 60, Table 3 on page 70, and Table 4 on page 111.

The Container must ensure that all enterprise bean method invocations received through the home and remote interface are associated with some principal. The Container must never return a null from the `getCallerPrincipal()` method.

15.6.6 Secure access to resource managers

The EJB Container Provider is responsible for providing secure access to resource managers as described in Subsection 14.4.3.

15.6.7 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.

15.6.8 System principal

The EJB 1.1 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 the 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.

15.6.9 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 operation 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 caller principal 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` to the 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 remote and home 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^[18]. The Container must allow multiple-deployed enterprise beans to co-exist at runtime.

15.6.10 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 remote interfaces, and EJB home interfaces.

15.7 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.

15.7.1 Security domain administration

The System Administrator is responsible for the administration of principals. Security domain administration is beyond the scope of the EJB specification.

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.

15.7.2 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.

15.7.3 Audit trail review

If the EJB Container provides an audit trail facility, the System Administrator is responsible for its management.

[18] The enterprise bean is installed each time using a different JNDI name.

Deployment descriptor

This chapter defines the deployment descriptor that is part of the `ejb-jar` file. Section 16.1 provides an overview of the deployment descriptor. Sections 16.2 through 16.4 describe the information in the deployment descriptor from the perspective of the EJB roles responsible for providing the information. Section 16.5 defines the deployment descriptor's XML DTD. Section 16.7 provides a complete example of a deployment descriptor of an assembled application.

16.1 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.

16.2 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 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 with bean-managed persistence, and *may* (but is not required to) specify the primary key class for an Entity 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-fields` elements.
- **Environment entries.** The Bean Provider must declare all the enterprise bean's environment entries as specified in Subsection 14.2.1.
- **Resource manager connection factory references.** The Bean Provider must declare all the enterprise bean's resource manager connection factory references as specified in Subsection 14.4.1.
- **EJB references.** The Bean Provider must declare all the enterprise bean's references to the homes of other enterprise beans as specified in Subsection 14.3.1.
- **Security role references.** The Bean Provider must declare all the enterprise bean's references to security roles as specified in Subsection 15.2.5.3.

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 16.5. 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">
```

16.3 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 16.2 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 16.5. 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">
```

16.4 Container Provider's responsibilities

The Container provider provides tools that read and import the information contained in the XML deployment descriptor.

16.5 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">
```

We plan to provide an ejb-jar file verifier that can be used by the Bean Provider and Application Assembler Roles to ensure that an ejb-jar is valid. The verifier would check all the requirements for the ejb-jar file and the deployment descriptor stated by this specification.

```

<!--
This is the XML DTD for the EJB 1.1 deployment descriptor.
-->

<!--
The assembly-descriptor element contains application-assembly information.

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 invocations.
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 provide
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 enterprise 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 contains 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 descriptor. 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>

```

```

-->
<!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 inter-
faces.

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>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 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; 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 ses-

sion 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>
```

16.6 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.
        </description>
      </resource-ref>
    </session>
  </enterprise-beans>
</ejb-jar>
```

```

        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>jobTitleHis-
tory</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>
```


Ejb-jar file

The `ejb-jar` file is the standard format for 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).

17.1 Overview

The `ejb-jar` file format is the contract between the Bean Provider and 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. An `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.

17.2 Deployment descriptor

The `ejb-jar` file must contain the deployment descriptor in the format defined in Chapter 16. The deployment descriptor must be stored with the name `META-INF/enterprise-jar.xml` in the `ejb-jar` file.

17.3 Class files

For each enterprise bean, the `ejb-jar` file must include the class files of the following:

- The enterprise bean class.
- The enterprise bean home and remote interface.
- The primary key class if the bean is an entity bean.

The `ejb-jar` file must also contain the class files for all the classes and interfaces that the enterprise bean class, and the remote and home interfaces depend on. This includes their superclasses and superinterfaces, and the classes and interfaces used as method parameters, results, and exceptions.

17.4 ejb-client JAR file

The `ejb-jar` file producer can create an `ejb-client` JAR file for the `ejb-jar` file. The `client-ejb` 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.

The `ejb-client` JAR file is specified in the deployment descriptor of the `ejb-jar` file using the optional `ejb-client-jar` element. The Deployer should ensure that the specified `ejb-client` JAR file is accessible to the client program's class-loader. If no `ejb-client-jar` element is specified, the Deployer should make the entire `ejb-jar` file accessible to the client's class-loader.

The EJB specification does not specify whether the `ejb-jar` file should include by copy or by reference the classes that are in the `ejb-client` JAR. 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, Enterprise Edition Platform specification [10].

17.5 Deprecated in EJB 1.1

This section describes the deployment information that was defined in EJB 1.0, and is deprecated in EJB 1.1.

17.5.1 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. In EJB 1.1, the enterprise beans are identified in the deployment descriptor, so the information in the Manifest is no longer needed.

17.5.2 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 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 Containers.

The chapter also defines the restrictions that the EJB 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.

18.1 Bean Provider's responsibilities

This section describes the view and responsibilities of the Bean Provider.

18.1.1 APIs provided by Container

The EJB Provider can rely on the EJB Container Provider to provide the following APIs:

- JDK 1.1.x or Java 2
- EJB 1.1 Standard Extension
- JDBC 2.0 Standard Extension (support for row sets only)
- JNDI 1.2 Standard Extension
- JTA 1.0.1 Standard Extension (the `UserTransaction` interface only)
- JavaMail 1.1 Standard Extension (for sending mail only)

The Bean Provider must take into consideration that while some Containers will provide JDK 1.1.x APIs, other Containers may provide the Java 2 (i.e. JDK 1.2) APIs. This means that the Bean Providers that want to deploy their enterprise beans in all Containers must restrict the APIs used by the enterprise beans to those that are available in JDK 1.1 and the above listed standard extensions.

18.1.2 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 Container. The restrictions apply to the implementation of the business methods. Section 18.2, 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.

Same reason as above. Synchronization would not work if the EJB Container distributed enterprise bean's instances across multiple JVMs.

- 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()` or `EntityContext.getEJBObject()` instead.

To guarantee portability of the enterprise bean's implementation across all compliant EJB Containers, the Bean Provider should test the enterprise bean using a Container with the security settings defined in Tables 10 and 11. The tables define the minimal functionality that a compliant EJB Container must provide to the enterprise bean instances at runtime.

18.2 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.

18.2.1 Java 2 based Container

A Java™ 2 based EJB Container must make the following APIs available to the enterprise bean instances at runtime:

- Java 2 APIs
- EJB 1.1 APIs
- JNDI 1.2
- JTA 1.0.1, the `UserTransaction` interface only
- JDBC™ 2.0 extension
- JavaMail 1.1, sending mail only

The following subsections describes the requirements in more detail.

18.2.1.1 Java 2 APIs requirements

The Container must provide the full set of Java 2 APIs. The Container is not allowed to subset the Java 2 APIs.

The EJB Container is allowed to make certain Java 2 functionality unavailable to the enterprise bean instances by using the Java 2 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 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 10

Java 2 Security policy for a standard EJB Container

Permission name	EJB Container policy
<code>java.security.AllPermission</code>	deny
<code>java.awt.AWTPermission</code>	deny

Table 10 Java 2 Security policy for a standard EJB Container

Permission name	EJB Container policy
java.io.FilePermission	deny
java.net.NetPermission	deny
java.util.PropertyPermission	deny
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 Java 2.

Some Containers may allow the Deployer to grant more, or fewer, permissions to the enterprise bean instances than specified in Table 10. 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.

18.2.1.2 EJB 1.1 requirements

The container must implement the EJB 1.1 interfaces as defined in this documentation.

18.2.1.3 JNDI 1.2 requirements

At the minimum, the EJB Container must provide a JNDI 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 EJB specification does not require that all the enterprise beans deployed in a Container be presented with the same JNDI name space. However, all the instances of the same enterprise bean must be presented with the same JNDI name space.

18.2.1.4 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 EJB Container is not required to implement the other interfaces defined in the JTA 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.

18.2.1.5 JDBC™ 2.0 extension requirements

The EJB Container must include the JDBC 2.0 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.

18.2.2 JDK™ 1.1 based Container

A JDK 1.1 based EJB Container must make the following APIs available to the enterprise bean instances at runtime:

- JDK 1.1 or higher
- EJB 1.1 APIs
- JNDI 1.2
- JTA 1.0.1, the `UserTransaction` interface only
- JDBC™ 2.0 extension
- JavaMail 1.1, sending mail only

The following subsections describes the requirements in more detail.

18.2.2.1 JDK 1.1 APIs requirements

The Container must provide the full set of JDK 1.1 APIs. The Container is not allowed to subset the JDK 1.1 APIs.

The EJB Container is allowed to make certain JDK 1.1 functionality unavailable to the enterprise bean instances by using the JDK security manager 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 JDK 1.1 security manager checks that the EJB Container must allow to succeed when the check is invoked from an enterprise bean instance.

Table 11 JDK 1.1 Security manager checks for a standard EJB Container

Security manager check	EJB Container's security manager policy
checkAccept(String, int)	throw SecurityException
checkAccess(Thread)	throw SecurityException
checkAccess(ThreadGroup)	throw SecurityException
checkAwtEventQueueAccess()	throw SecurityException
checkConnect(String, int)	allow
checkConnect(String, int, Object)	allow
checkCreateClassLoader()	throw SecurityException
checkDelete(String)	throw SecurityException
checkExec(String)	throw SecurityException
checkExit(int)	throw SecurityException
checkLink(int)	throw SecurityException
checkListen(int)	throw SecurityException
checkMemberAccess(Class, int)	throw SecurityException
checkMulticast(InetAddress)	throw SecurityException
checkMulticast(InetAddress, byte)	throw SecurityException
checkPackageAccess(String)	throw SecurityException
checkPackageDefinition(String)	throw SecurityException
checkPrintJobAccess()	allow
checkPropertiesAccess()	throw SecurityException
checkPropertyAccess(String)	throw SecurityException
checkRead(FileDescriptor)	throw SecurityException
checkRead(String)	throw SecurityException
checkRead(String, Object)	throw SecurityException
checkSecurityAccess(String)	throw SecurityException
checkSetFactory()	throw SecurityException

Table 11 JDK 1.1 Security manager checks for a standard EJB Container

Security manager check	EJB Container's security manager policy
checkSystemClipboardAccess()	throw SecurityException
checkTopLevelWindow(Object)	throw SecurityException
checkWrite(FileDescriptor)	throw SecurityException
checkWrite(String)	throw SecurityException

Some Containers may allow the Deployer to grant more, or fewer, permissions to the enterprise bean instances than specified in Table 10. 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.

18.2.2.2 EJB 1.1 requirements

The container must implement the EJB 1.1 interfaces as defined in this documentation.

18.2.2.3 JNDI 1.2 requirements

Same as defined in Subsection 18.2.1.3.

18.2.2.4 JTA 1.0.1 requirements

Same as defined in Subsection 18.2.1.4.

18.2.2.5 JDBC 2.0 extension requirements

Same as defined in Subsection 18.2.1.5, with the following exception: The EJB Container is not required to provide the support for the RowSet functionality.

This exception was made because the RowSet functionality requires the Java 2 Collections.

18.2.3 Argument passing semantics

The enterprise bean's home and remote interfaces are *remote interfaces* for Java RMI. The Container must ensure the semantics for passing arguments conform to Java RMI. 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.

Responsibilities of EJB Roles

This chapter provides the summary of the responsibilities of each EJB Role.

19.1 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.

19.1.1 API requirements

The enterprise beans must meet all the API requirements defined in the individual chapters of this document.

19.1.2 Packaging requirements

The Bean Provider is responsible for packaging the enterprise beans in an ejb-jar file in the format described in Chapter 17.

The deployment descriptor must include the *structural* information described in Section 16.2.

The deployment descriptor may optionally include any of the *application assembly* information as described in Section 16.3.

19.2 Application Assembler's responsibilities

The requirements for the Application Assembler are in defined in Section 16.3.

19.3 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.

19.4 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.

19.5 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.

19.6 Client Programmer's responsibilities

The EJB client programmer writes applications that access enterprise beans via their home and remote interfaces.

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 EJBMetaData
public interface EJBObject
public interface EnterpriseBean
public interface EntityBean
public interface EntityContext
public interface Handle
public interface HomeHandle
public interface SessionBean
public interface SessionContext
public interface SessionSynchronization
```

Classes:

```
public class CreateException
public class DuplicateKeyException
public class EJBException
public class FinderException
public class ObjectNotFoundException
public class RemoveException
```

package *javax.ejb.deployment*

The `javax.ejb.deployment` package that was defined in the EJB 1.0 specification is 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 compliant Containers.

We intend to a tool which will help convert an EJB 1.0 deployment descriptor to the EJB 1.1 XML-based format. The `javax.ejb.deployment` package will be provided only as part of this tool.

The Javadoc specification of the EJB interface is included in a ZIP file distributed with this document.

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 to IDL Mapping. OMG TC Document TC orbos/98-07-19.
- [8] Enterprise JavaBeans to CORBA Mapping. <http://java.sun.com/products/ejb/docs.html>.
- [9] OMG Object Transaction Service. <http://www.omg.org/corba/sectrans.htm#trans>.
- [10] Java 2 Platform, Enterprise Edition, v1.2 (J2EE). <http://java.sun.com/j2ee>.

Features deferred to future releases

We plan to provide an SPI-level interface for attaching a resource manager (such as a JDBC driver) to the EJB Container as a separate Connector API.

We plan to enhance the support for Entities in the next major release (EJB 2.0). We are looking into the area of use of the UML for the design and analysis of enterprise beans applications.

We plan to provide integration of EJB with JMS as part of EJB 2.0.

Appendix B Frequently asked questions

This Appendix provides the answers to a number of frequently asked questions.

B.1 Client-demarcated transactions

The EJB 1.0 specification did not explain how a client other than another enterprise bean can obtain a the `javax.transaction.UserTransaction` interface.

The Java2, Enterprise Edition specification [10] 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();
...
```

B.2 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 remote interfaces. A component may derive its home and remote interfaces from some "parent" home and remote 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.

B.3 Entities and relationships

The current EJB architecture does not specify how one Entity bean should store an object reference of another Entity bean. The desirable strategy is application-dependent. The enterprise bean (if the bean uses bean-managed persistence) or the Container (if the bean uses container-managed persistence) can use any of the following strategies for maintaining persistently a relationship between entities (the list is not inclusive of all possible strategies):

- Object's primary key. This is applicable if the target object's Home is known and fixed.
- Home name and object's primary key.
- Home object reference and object's primary key.
- Object's handle.

B.4 Finder methods for entities with container-managed persistence

The EJB specification does not provide a *formal* mechanism for the Bean Provider of a bean with container-managed persistence to specify the criteria for the finder methods.

The current mechanism is that Bean Provider describes the finders in a description of the Entity Bean. The current EJB specification does not provide any syntax for describing the finders.

We plan to address this issue in a future release of the specification.

B.5 JDK 1.1 or Java 2

Chapter 18 describes the issue of using JDK 1.1 versus Java 2 in detail.

In summary, the Bean Provider can produce enterprise beans that will run in both JDK 1.1 and Java 2 based Containers. The Container Provider can use either JDK 1.1 or Java 2 as the basis for the implementation of the Container.

B.6 `javax.transaction.UserTransaction` versus `javax.jts.UserTransaction`

The correct spelling is `javax.transaction.UserTransaction`.

The use of `javax.jts.UserTransaction` is deprecated in EJB 1.1.

B.7 How to obtain database connections

Section 14.4 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();
        ...
    }
}
```

B.8 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.

B.9 Copying of parameters required for EJB calls within the same JVM

The enterprise bean's home and remote interfaces are *remote interface* in the Java RMI sense. The Container must ensure the Java RMI argument passing semantics. Non-remote objects must be passed by value.

Specifically, the EJB Container is not allowed to pass local objects by reference on inter-enterprise bean 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.

Revision History

C.1 Changes since Release 0.8

Removed `java.ejb.BeanPermission` from the API. This file was incorrectly included in the 0.8 specification.

Renamed packages to `java.ejb` and `javax.ejb.deployment`. The Enterprise JavaBeans API is packaged as a standard extension, and standard extensions should be prefixed with `javax`. Also renamed `java.jts` to `javax.jts`.

Made clear that a container can support multiple EJB classes. We renamed the `javax.ejb.Container` to `javax.ejb.EJBHome`. Some reviewers pointed out that the use of the term “Container” for the interface that describes the life cycle operations of an EJB class as seen by a client was confusing.

Folded the factory and finder methods into the enterprise bean’s **home interface**. This reduces the number of Java classes per EJB class and the number of round-trips between a client and the container required to create or find an EJB object. It also simplifies the client-view API.

Removed the PINNED mode of a Session Bean. Many reviewers considered this mode to be “dangerous” since it could prevent the container from efficiently managing its memory resources.

Clarified the life cycle of a stateless Session Bean.

Added a chapter with the specification for exception handling.

We have renamed the contract between a component and its container to **component contract**. The previously used term **container contract** confused several reviewers.

Added description of finder methods.

Modified the entity create protocol by breaking the `ejbCreate` method into two: `ejbCreate` and `ejbPostCreate`. This provides a cleaner separation of the discrete steps involved in creating an entity in a database and its associated middle-tier object.

Added more clarification to the description of the entity component protocol.

Added more information about the responsibilities of the enterprise bean provider and container provider.

Renamed `SessionSynchronization.beginTransaction()` to `SessionSynchronization.afterBegin()` to avoid confusion with `UserTransaction.begin()`.

Added the specification of isolation levels for container-managed Entity Beans.

C.2 Changes since Release 0.9

Renamed `javax.ejb.InstanceContext` to `javax.ejb.EJBContext`.

Fixed bugs in the javadoc of the `javax.ejb.EntityContext` interface.

Combined the state diagrams for non-transactional and transactional Session Beans into a single diagram.

Added the definition of the restrictions on using transaction scopes with a Session Bean (a Session Bean can be only in a single transaction at a time).

Allowed the enterprise bean's class to implement the enterprise bean's remote interface. This change was requested by reviewers to facilitate migration of existing Java code to Enterprise JavaBeans.

Removed the `javax.ejb.EJBException` from the specification, and replaced its use by the standard `java.rmi.RemoteException`. This change was necessary because of the previous change that allows the enterprise bean class to implement its remote interface.

Changed some rules regarding exception handling.

Renamed to the `javax.jts.CurrentTransaction` interface to `javax.jts.UserTransaction` to avoid confusion with the `org.omg.CosTransactions.Current` interface. The `javax.jts.UserTransaction` interface defines the subset of operations that are “safe” to use at the application-level, and can be supported by the majority of the transaction managers used by existing platforms.

Added specification for `TX_BEAN_MANAGED` transactions.

Made the isolation levels supplied in the deployment descriptor applicable also to Session Beans and entities with bean-managed persistence.

Renamed the `destroy()` methods to `remove()`. This change was requested by several reviewers who pointed out the potential for name space collisions in their implementations.

Added the create arguments to the `ejbPostCreate` method. This simplifies the programming of an Entity Bean that needs the create arguments in the `ejbPostCreate` method (previously, the bean would have to save these arguments in the `ejbCreate` method).

Added restrictions on the use of per-method deployment attributes.

Added `javax.ejb.EJBMetaData` to the examples, and added the generation of the class that implements this interface as a requirements for the container tools.

Added the `getRollbackOnly` method to the `javax.ejb.EJBContext` interface. This method allows an instance to test if the current transaction has been marked for rollback. The test may help the enterprise bean to avoid fruitless computation after it caught an exception.

We removed the placeholder Appendix for examples. We will provide examples on the Enterprise JavaBeans Web site rather than in this document.

C.3 Changes since Release 0.95

Allowed a container-managed field to be of any Java Serializable type.

Clarified the bean provider responsibilities for the `ejbFind<METHOD>` methods Entity Beans with container-managed persistence.

Added two rules to Subsection xxx on exception handling and transaction management. The new rules are for the `TX_BEAN_MANAGED` beans.

Use the `javax.rmi.PortableRemoteObject.narrow(...)` method to perform the narrow operations after a JNDI lookup in the code samples used in the specification. While some JNDI providers may return from the `lookup(...)` method the exact stub for the home interface making it possible to for the client application to use a Java cast, other providers may return a wider type that requires an explicit narrow to the home interface type. The `javax.rmi.PortableRemoteObject.narrow(...)` method is the standard Java RMI way to perform the explicit narrow operation.

Changed several deployment descriptor method names.

C.4 Changes since 1.0

This section lists the changes since EJB 1.0.

Specified the behavior of `EJBObject.getPrimaryKey()`, `EJBMetaData.getPrimaryKeyClass()`, `EJBHome.remove(Object primaryKey,)` and `isIdentical(Object other)` for Session Beans. As Session Beans do not have client-accessible primary keys, these operations result in exceptions.

Disallowed `TX_BEAN_MANAGED` for Entity Beans.

Disallowed use of `SessionSynchronization` for `TX_BEAN_MANAGED` sessions.

Allowed using `java.lang.String` as a primary key type.

Allowed deferring the specification of the primary key class for entities with container-managed persistence to the deployment time.

Clarified that a matching `ejbPostCreate` is **required** for each `ejbCreate`.

Added requirement for `hashCode` and `equals` for the primary key class.

Deprecated the package `javax.ejb.deployment` by replacing the JavaBeans-based deployment descriptor with an XML-based deployment descriptor.

Improved the information in the deployment descriptor by clearly separating structural information from application assembly information, and by removing support for information that should be supplied by the Deployer rather than by the `ejb-jar` producer (i.e. ISV). The EJB 1.0 deployment descriptor mixed all this information together, making it hard for people to understand the division of responsibility for setting the various values, and it was not clear what values can be changed at application assembly and/or deployment.

Added the requirement for the Bean Provider to specify whether the enterprise bean uses a bean-managed or container-managed transaction.

Added `Never` to the list of possible values of the transaction attributes to allow specification of the case in which an enterprise bean must never be called from a transactional client.

Removed the Appendix describing the `javax.transaction` package. Inclusion of this package in the EJB document is no longer needed because the JTA documentation is publicly available.

Tightened the specification of the responsibilities for transaction management.

Tighten the rules for the runtime environment that the Bean Provider can expect and the EJB Container Provider must provide. See Chapter 18.

C.5 Changes since 1.1 Draft 1

This sections lists the changes since EJB 1.1 Draft 1.

Allow use of the Java 2 `java.util.Collection` interfaces for the result of entity finder methods.

Defining the `FinderException` in the finder methods of the home interface is mandatory now.

Clean up of the exception specification, including minor changes from EJB 1.0 summarized in Section 12.6.

The scope of the EJB specification for managing transaction isolation levels was reduced to sessions with bean-managed transaction demarcation. The current EJB specification does not have any API for managing transaction isolation for beans with container-managed transaction demarcation (note that all Entity beans fall into this category).

Eliminated the `stateless-session` element in the XML DTD. Now the `session` element is used to describe both the stateful and stateless session beans.

Added an optional `description` element to the `method` element. The intention is to allow tools to display the description of the method.

Clarified that the enterprise bean class may have superclasses, and that the business methods and the various container callbacks can be implemented in the enterprise bean class, or in any of its superclasses.

Fixed the example that illustrates the use of handles for session objects. Serialized handles are not guaranteed to be deserializable in a different system, and therefore they cannot be emailed.

Updated the Overview chapter.

Allowed deferring the specification of the primary key class for all entities (not only for those with container-managed persistence as it was the case in Draft 1).

Allow enterprise beans to print. The Container must grant the permission to the enterprise beans to queue printer job.

The `setRollbackOnly()` and `getRollbackOnly()` methods of the `EJBContext` object must not be used by enterprise beans with bean-managed transaction demarcation. There is no need for these beans to use these methods.

C.6 Changes since 1.1 Draft 2

Fix an error in the requirement for how a Container must deal with inter-EJB invocations when both the calling and called bean are in the same JVM. The correct requirement is that the RMI semantics must be ensured, and therefore the Container must not pass non-remote objects by reference.

Clarified the requirements for serialization of the session objects.

Specified that an EJB Compliant Container may always return a null from the deprecated `getCallerIdentity()` method.

Added a section on distributed transaction scenarios involving access to the same entity from multiple clients in the same transaction.

Changed the specification of the return value type of the `ejbCreate(...)` methods for entities with container-managed persistence. The previous specification required that the `ejbCreate` methods are defined as returning void. The new requirement is that the `ejbCreate` methods be defined as returning the primary key class type. The implementation of the `ejbCreate` method should return null. This change is to allow tools, if they wish, to create an entity bean with bean-managed persistence by subclassing an original entity bean with container-managed persistence.

For compatibility with EJB 1.0, added the support for the `java.rmi.RemoteException` to be thrown from the enterprise bean class methods. This is needed to allow an EJB 1.1 Container to support enterprise beans written to the EJB 1.0 specification. The use of the `java.rmi.RemoteException` in the enterprise bean class methods is deprecated, and new applications should throw the `javax.ejb.EJBException` instead.

Removed the deprecated package `javax.ejb.deployment` from the EJB interfaces. The the deprecated package `javax.ejb.deployment` will be distributed only with the deployment descriptor conversion tool.

Updated the examples in the transaction chapter by removing the `setAutoCommit` and `setTransactionIsolation` calls. These calls are not typically done by the enterprise bean.

Added the `<method-intf>` element to allow a method element to differentiate between a method with same signature when defined in both the remote and home interfaces.

Specified the behavior of the `getUserTransaction()`, `setRollbackOnly()`, and `getRollbackOnly()` methods for the cases when the methods are invoked by beans that are not allowed to use these methods. The Container will throw the `java.lang.IllegalException` in these situations.

Specified that `PortableRemoteObject.narrow(...)` must be used by a client to convert the result of `Handle.getEJBObject()` to the remote interface type.

Required portable enterprise bean clients to use the `PortableRemoteObject.narrow(...)`.

Clarified the minimal lifetime for handles.

Clarified that the caller must have **at least one** security role (not **all**) associated with the method permission in order to be allowed to invoke the method.

Support for entities has been made mandatory for the Container Provider.

Added a section to the Exception chapter dealing with the release of resources held by the instance when the instance is being discarded because of a system exception.

Added the `res-auth` element to the deployment descriptor for the Bean Provider to indicate whether the bean code performs an explicit sign-on to a resource manager, or whether the Bean relies on the Container to perform sign-on based on the information supplied by Deployer.

Added `java.io.Serializable` as a superinterface of `javax.ejb.Handle`. The EJB 1.0 spec required that the implementation class implements the `java.io.Serializable` interface, this change expresses the requirement syntactically.

Added the interface `javax.ejb.HomeHandle` to provide support for handles for home objects.

Allowed a Session bean instance to be removed upon a timeout while the instance is in the passivated state.

Add the `javax.ejb.NoSuchEntityException` exception to the API. Added requirements for throwing the `java.rmi.NoSuchObjectException` to the chapter on exceptions.

C.7 Changes since EJB 1.1 Draft 3

Replaced the support for environment properties with the JNDI-based environment entries. The EJB 1.0 style of environment properties access is deprecated in EJB 1.1.

Removed the `finalize()` method from the state diagrams. Specified that an enterprise bean must not define the `finalize()` method in the enterprise bean class. This is because it cannot be guaranteed that the method is called at all in some Container implementations.

Made clear that the result of comparing two object reference using the Java `"=="` operator or the `equals()` method is undefined.

Added Tables 2, 3, and 4 that specify which operations are allowed in the enterprise bean methods.

Clarified what “proper transaction context” means in the Chapter on entities.

Flattened the DTD hierarchy by removing the elements that grouped entries of the same type.

Relaxed the rules for the primary key class. An entity with bean-managed persistence can use any RMI-IIOP Value Type as its primary key type; the primary key type of an entity with container-managed persistence is more constrained.

Added the `isStatelessSession()` method to the `EJBMetaData` interface.

Updated the chapter in distribution to simply reference RMI-IIOP. The original chapter had been written before RMI-IIOP was completed.

C.8 Changes since EJB 1.1 Public Draft

Added the `ejb-client-jar` element to the deployment descriptor to allow the `ejb-jar` file producer to specify a JAR file that contains the classes necessary to access the enterprise beans in the `ejb-jar` file.

The value of the `res-auth` element was changed to `Application` (it was `Bean`) to be consistent with the Java 2, Enterprise Edition platform specification.

Changed the lexical rules for the `env-entry-value` element so that values of the type `String` need not be double-quoted in the deployment descriptor. See subsection 14.2.1.2.

Added the requirement for the Container to provide the `UserTransaction` interface to the enterprise bean instances in the environment JNDI context under the name `java:comp/UserTransaction`. See section 14.6.

Clarified that the container must never return a null from the `getCallerPrincipal()` method.

Allow the stateful session bean's `ejbCreate`, `ejbRemove`, `ejbActivate`, and `ejbPassivate` methods to access resource managers without transaction context. This was allowed in EJB 1.0.

Cleaned up the description of transactions. Removed the confusing term *local transaction*, and created the description of how a Container may deal with the resource manager updates from a method that runs with an unspecified transaction context into Section 11.6.3 (the term *local transaction* was used to refer to the cases that are now covered by this section).

Added an explanation of how `ejbLoad` and `ejbStore` work for entity bean instances that execute with an unspecified transaction context. See subsection 9.1.7.1.

Fixed the argument type of the `isIdentical(EJBObject)` method. The spec showed it incorrectly as `java.lang.Object`. The EJB class files and javadoc have always been correct.

Disallowed the use of `UserTransaction` in the `setSessionContext` method, as specified in Tables 2 and 3.

Specified which methods can be assigned a transaction attribute, and which methods can be included in `method-permission` elements.

Clarified in Subsection 12.3.6 which “resources” the Container is responsible for releasing when an instance is being discarded.

Augmented the restrictions on client's security context in Section 15.5 to cover the case in which requests in the same transaction are received from multiple intermediate objects.

Moved the description of the “transaction diamond” scenario from Subsection 9.1.13 to Subsection 11.7. Described also the transaction diamond scenario for Session Beans.

Clarified the scope of the `env-entry-name`, `ejb-ref-name`, `res-ref-name`, `security-role-ref` and `role-name` elements.

Specified that application exceptions must not be defined as a subclass of `RuntimeException` or `RemoteException`.

Clarified that the container-managed persistence fields must not be defined as `transient`.

Clarified in Subsection 9.2.3 that the entity object created by the `ejbCreate(...)` method must have a unique primary key.

Clarified in Section 6.8 how the Container delegates requests to instances of a stateless Session bean.

Added to Section 18.1.2 the restriction that an enterprise bean must not pass `this` as a method argument or result.

In Section 6.4.1 specified that the Container must be able to preserve to an object reference of the `UserTransaction` interface across passivation. Same for the object references of enterprise beans' home interfaces.

Noted in Chapter 11 that the transaction attributes may be specified either by the Bean Provider or the Application Assembler. The previous text suggested that only the Application Assembler was allowed to specify the transaction attributes.

Made the terminology more consistent throughout the specification. Used the terms *session bean*, *session object*, *session bean instance*, *home interface*, *remote interface*, *session EJBObject*, and *session EJBHome* consistently. Used similar terminology for the entity bean related terms. Note that we turned off the change bars while making this editorial clean up.

Disallowed the use of the `setRollbackOnly` and `getRollbackOnly` method, the use of the `UserTransaction` interface, resource manager access, and enterprise bean access in the `ejbCreate` and `ejbRemove` methods of the stateless session bean (see Table 3 on page 70). The Container does not have a transaction context and client security context during the execution of these two operations of a stateless session bean.

Added support for referencing the deployment descriptor elements through XML IDs. The main reason for this is to make it easier for tools that want to pass additional non-standard deployment information for the enterprise beans. The tools should put the non-standard information into a separate file, and optionally make use the ID mechanism to reference the information in the standard EJB deployment descriptor. Tools are **not** allowed to extend the format of the EJB deployment descriptor.

Made a minor change to the language in Subsection 15.3.2 for the case that a method is not assigned to any security role.

Added a paragraph stating that the Container should implement the object references of the home and remote interfaces of Entity objects such that a client can use the references over a long period of time (Subsection 9.3.9).

C.9 Changes since EJB 1.1 Public Draft 2

We changed the JTA requirements to refer to version JTA 1.0.1.

Added clarifications to subsections 9.1.5.1, 9.1.5.2, and 9.4.3 stating that the container must ensure that the instance's state is synchronized before it invokes `ejbRemove`.

Renamed the `primkey-class` element to `prim-key-class` to be consistent with the rest of the element names.

Clarified in Subsection 12.1.1 that a Bean Provider is allowed to define subclasses of the standard EJB application exceptions.

Added the requirement (in Subsection 9.4.2) for the Container to reset the container-manager fields to the default Java language values prior to invoking `ejbCreate`.

Clarified the reason for allowing the primary key type for CMP entities to be specified at deployment time (Subsection 9.4.7.3).

C.10 Changes since EJB 1.1 Public Draft 3

Made clear in Section 11.2.3 that not all EJB client environments are required to support the `UserTransaction` interface.

Specified the name and URI to be used in the DOCTYPE PUBLIC ID in the EJB XML deployment descriptors.

Corrected Section 5.3.2 to state that a `javax.ejb.RemoveException` be thrown instead of `java.rmi.RemoteException`.

Fixed a few errors in Tables 2, 3, and 4.

Changed the format of the Style 3 method element of the EJB deployment descriptor. See Section 16.5. This change was necessary to disambiguate a Style 3 element for a method with no arguments from a Style 2 element.

C.11 Changes since EJB 1.1 Public Release

Fixed an error in Subsection 15.2.5. The `getCallerPrincipal` and `isCallerInRole` functions, when called in an inappropriate context, must throw the `java.lang.IllegalStateException`, not the `javax.ejb.EJBException`.

Fixed the omission of a requirement in Subsection 9.4.1. The container-managed fields must be declared as `public` fields in the enterprise bean class. This requirement was present in the EJB 1.0 specification, and it was inadvertently left out when the sections on container-managed persistence were reorganized during the EJB 1.1 specification process.

Made it clear in 6.4.1 that an instance is allowed to retain across passivation references to enterprise beans' remote and home interfaces, references to the `SessionContext` object, references to the `java:comp/env` JNDI context and its subcontexts, and references to the `UserTransaction` *anywhere* in the instances conversational state (i.e. not only directly in the fields of the session bean class). For example, it is possible to retain a Collection of remote interfaces in the conversational state.

Changed the version numbers in the DOCTYPE specification in the deployment descriptor from 1.2 to 1.1. The correct DOCTYPE specification is:

```
<!DOCTYPE ejb-jar PUBLIC "-//Sun Microsystems, Inc.//DTD Enterprise  
JavaBeans 1.1//EN" "http://java.sun.com/j2ee/dtds/ejb-jar_1_1.dtd">
```

Clarified in Sections 5.5 and 8.7 that handles are not capabilities in the security sense.

Replaced the term `resource factory reference` with the term `resource manager connection factory reference`. This change makes the terminology consistent with other J2EE specs.

Added a missing entry to Table 11 for `checkLink`.

Clarified in 16.5 that the DTD XML elements' content is case sensitive.

Fixed a few typos in the text.

Index

A

activation, 50, 65
All, 46
APIs
 runtime, 272, 275
Application Assembler, 22
 responsibilities, 282
 transaction attributes, 171
 transaction role, 169
application assembly, 243–244
application exception, 187
application exception See exception

B

Bean Provider, 22
 responsibilities, 281
 responsibility, 74–77
BeanReference
 Interface in package java.beans.enterprise, 285
 Interface in package java.ejb, 285

C

cmp-fields element, 242
commit, 118
Container
 Interface in package java.ejb, 285
Container Provider, 23
 transaction demarcation
 bean managed, 173
 container managed, 175–178
container-transaction element, 171

conversational state, 51
 passivation, 51
 rollback, 53
CORBA mapping, 35
CreateException, 116
CurrentTransaction
 Interface in package java.jts, 293

D

Deployer, 22
 responsibilities, 282
deployment descriptor
 application assembly, 240
 bean structure, 240, 240–242
 DTD, 244–258
 EJB reference, 208
 ejb-link element, 210
 ejb-ref element, 208
 enterprise-beans element, 240
 env-entry element, 204
 environment entry, 204
 res-auth element, 213
 resource-ref element, 213
 role, 240
 transaction attributes, 171
 XML DTD, 244–258
distributed objects, 199
DuplicateKeyException, 116

E

EJB Container Provider, 23
 requirements, 78–79

- responsibilities, 282
- EJB reference, 207
 - Deployer role, 211
 - ejb-link element, 210
 - in deployment descriptor, 208
 - locate home interface, 208
- EJB Role
 - Application Assembler, 22
 - Bean Provider, 22
 - Container Provider, 23
 - Deployer, 22
 - EJB Server Provider, 23
 - System Administrator, 24
- EJB Server Provider, 23
- ejb-class element, 241
- ejb-client JAR file, 268
- EJBContext
 - Interface in package javax.ejb, 284, 285
- EJBHome, 41, 128
 - Interface in package javax.ejb, 285
 - remove method, 42
- ejb-jar file, 33, 242, 267
 - class files, 268
 - deployment descriptor, 268
 - ejb-client JAR file, 268
 - JAR Manifest, 269
- ejb-link element, 210
- EJBMetaData, 129
- ejb-name element, 240
- EJBObject, 39, 43, 128
 - remove method, 42
- ejb-ref element, 208
- ejbRemove, 61
- enterprise bean component
 - characteristics of, 30
- enterprise bean component model, 30
- enterprise bean contract
 - client view, 31
 - CORBA mapping, 35
 - home interface, 31
 - metadata interface, 32
 - object identity, 32
 - remote interface, 32
 - component contract
 - requirements, 32
 - ejb-jar file, 33
- enterprise bean environment
 - JNDI interface, 202
 - InitialContext, 203
- Enterprise Bean Provider, 22
- entity bean
 - allowed method operations, 111
 - bean provider-implemented methods, 104–107
 - business methods, 125
 - class requirements, 121
 - client view of, 85–86
 - commit, 118
 - constructor, 104
 - create method, 89, 122
 - CreateException, 116
 - DuplicateKeyException, 116
 - EJB container, 86
 - ejbActivate, 105
 - ejbCreate, 105, 122
 - ejbFind methods, 124
 - ejbLoad, 106, 112
 - ejbPassivate, 106
 - ejbPostCreate, 105, 124
 - ejbRemove, 106
 - ejbStore, 106, 112
 - exceptions, 116–117
 - find methods, 90, 107, 124
 - return type, 114–115
 - findByPrimaryKey, 90
 - FinderException, 117
 - generated classes, 127
 - getHandle method, 94
 - getPrimaryKey method, 93
 - handle, 94, 128
 - home interface
 - function of, 89
 - requirements, 126
 - home interface handle, 95, 128
 - isIdentical method, 93
 - life cycle, 91–92, 102–104
 - locate home interface, 87
 - methods
 - container view of, 107–109
 - modeling business objects, 99
 - ObjectNotFoundException, 117
 - persistence, 99, 100, 101

- container managed, 129–134
 - primary key, 93, 127
 - reentrancy, 120
 - remote interface, 93, 125
 - remove method, 90
 - RemoveException, 117
 - setEntityContext, 104
 - state, 102
 - state caching, 112
 - transaction demarcation, 160
 - container managed, 167
 - transaction synchronization, 119
 - unsetEntityContext, 105
- entity element, 241
- env-entry element, 204
- environment entry, 204
 - Application Assembler role, 207
 - Deployer role, 207
- environment naming context, 203
- exception
 - application, 187
 - client handling of, 195
 - data integrity, 188
 - defined, 188
 - subclass of, 189
 - client view, 195
 - container handling of, 191
 - container-invoked callbacks, 193
 - containter-managed transaction, 194
 - NoSuchObjectException, 197
 - RemoteException, 194, 196
 - client handling, 196
 - system
 - handling of, 189–190
 - System Administrator, 197
 - transaction commit, 194
 - transaction start, 194
 - TransactionRequiredException, 197
 - TransactionRolledbackException, 197
- F**
- findByPrimaryKey, 90
- FinderException, 117
- G**
- getCallerIdentity, 222

- getCallerPrincipal, 223, 236
- getEnvironment method, 216
- getHandle method, 94
- getPrimaryKey method, 47, 93

H

- home element, 241
- home interface, 31, 40, 67
 - client functionality, 89
 - create method, 89
 - EJB reference to, 207
 - entity bean, 126
 - find methods, 90
 - findByPrimaryKey, 90
 - handle, 95
 - locating, 87
 - remove method, 90

I

Interfaces

- java.beans.enterprise.BeanReference, 285
- java.ejb.BeanReference, 285
- java.ejb.Container, 285
- java.jts.CurrentTransaction, 293
- javax.ejb.EJBContext, 284, 285
- javax.ejb.EJBHome, 285
- javax.ejb.SessionSynchronization, 168
- javax.jts.UserTransaction, 293
- isCallerInRole, 222
- isCallerinRole, 224
- isIdentical method, 46, 93
- isolation level
 - managing in transaction, 160

J

- JAR Manifest file, 269
- Java RMI, 279
- JDBC, 277
- JDK 1.1, 277–279
- JNDI, 276
- JNDI interface, 202
 - InitialContext, 203
- JTA, 277

M

- Mandatory, 170, 177

metadata interface, 32
method-permission element, 229

N

narrow method, 47
Never, 170, 177
NoSuchObjectException, 197
NotSupported, 170, 175

O

object identity, 32
ObjectNotFoundException, 117

P

passivation, 50, 65
 conversational state, 51
 SessionContext interface, 52
 UserTransaction interface, 52
persistence, 99
 bean managed, 100
 entity state caching, 112
 container managed, 101, 129–134
persistence-type element, 241
portability
 programming restrictions, 272–274
primary key, 93, 127
prim-key-class element, 241
principal, 220, 221
 delegation, 233

R

remote element, 241
remote interface, 32, 39, 43
 entity bean, 93, 125
RemoteException, 194, 196
 client handling, 196
RemoveException, 117
Required, 170, 175
RequiresNew, 170, 176
res-auth element, 213
resource
 obtaining connection to, 212
 res-auth element, 213
resource factory, 211
resource factory reference, 211
 resource-ref element, 213

resource-ref element, 213

RMI, 199

role-link element, 232

role-name element, 227

runtime

 APIs, 272, 275

S

security

 audit, 238

 bean provider

 programming recommendations, 221

 client responsibility, 234

 current caller, 223

 deployment descriptor processing, 234

 deployment tools, 235

 EJBContext, 222, 236

 getCallerPrincipal, 222, 223, 236

 isCallerInRole, 222, 224

 mechanism, 235

 principal, 220, 221

 delegation, 233

 passing, 236

 principal realm, 233, 235

 role-link element, 232

 runtime enforcement, 237

 security-role-ref element, 225

security domain, 233, 235

security role, 220, 226, 227

 assigning, 233

 linking, 232

 method permission, 220, 226, 229

 role-name element, 227

security view, 226

security-role element, 227, 233

security-role-ref element, 225

session bean

 access to, 39

 business method requirements, 76

 class requirements, 75

 client operations on, 44

 client view of, 39

 create, 42

 ejbCreate requirements, 76

 ejbRemove call, 61

 exceptions, 61

- getPrimaryKey method, 47
- home interface, 40, 41
- home interface requirements, 77
- identity, 43
- provider responsibility, 74–77
- remote interface, 39, 43
- remote interface requirements, 77
- remove, 42, 66
- requirements, 74–77
- SessionBean interface, 53
- SessionContext interface, 54
- SessionSynchronization interface, 54
- stateful
 - conversational state, 51
 - identity of, 46
 - isIdentical method, 46
 - lifecycle, 57
 - operations in, 59
- stateless, 67–74
 - exceptions, 71
 - home interface, 67
 - identity of, 46
 - isIdentical method, 46
 - lifecycle, 68
 - operations, 70
 - transaction demarcation, 161
 - use of, 67
- transaction context, 56
- transaction demarcation, 160, 161
 - bean managed, 161
 - container managed, 167
- transaction scope, 62
- session bean instance
 - activation, 50, 65
 - characteristics, 49
 - creating, 55
 - diagram of, 63
 - passivation, 50, 65
 - serialization of calls, 56
- session element, 241
- SessionBean interface, 53
- SessionContext interface, 54
 - passivation, 52
- SessionSynchronization interface, 54, 168
 - callbacks, 179
- session-type element, 241

- stateful session bean
 - conversational state, 51
 - lifecycle, 57
 - operations in, 59
- stateless session bean. See session bean
- Supports, 170, 176
- System Administrator, 24
 - responsibilities, 282

T

- transaction
 - attributes, 154
 - definition, 169
 - deployment descriptor, 171
 - Mandatory, 177
 - Never, 177
 - NotSupported, 175
 - Required, 175
 - RequiresNew, 176
 - Supports, 176
 - values, 170
- bean managed, 154, 159–169
 - container responsibilities, 173
- committing, 64
- container managed, 154, 159–169
 - container responsibilities, 175–178
 - getRollbackOnly method, 169, 178
 - getUserTransaction method, 179
 - SessionSynchronization callbacks, 179
 - setRollbackOnly method, 168, 178
- isolation level, 160
- JTA, 155
- JTS, 155
- multiple client access, 180–185
- nested, 154
- SessionSynchronization interface, 168
- starting, 63
- synchronizing, 119
- unspecified transaction context, 179
- UserTransaction interface, 154

- transaction context
 - session bean, 56
- transaction scope
 - session bean, 62
- TransactionRequiredException, 197

- TransactionRolledbackException, 197
- transaction-type element, 241
- trans-attribute element, 171
- type narrowing, 47, 96

U

- UserTransaction
 - Interface in package javax.jts, 293
- UserTransaction interface, 154, 161, 217
 - passivation, 52