



J2EE™ Connector Architecture Specification

Version 1.5

Sun Microsystems, Inc.
www.sun.com

Final Release
November 2003

Submit comments about this document to: j2ee-connector-comments@sun.com.

J2EE™ Connector Architecture Specification ("Specification")

Version: 1.5

Status: FCS

Release: November 24, 2003

Copyright 2003 Sun Microsystems, Inc.

4150 Network Circle, Santa Clara, California 95054, U.S.A

All rights reserved.

NOTICE; LIMITED LICENSE GRANTS

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

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

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

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

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

TRADEMARKS

No right, title, or interest in or to any trademarks, service marks, or trade names of Sun or Sun's licensors is granted hereunder. Sun, Sun Microsystems, the Sun logo, Java, the Java Coffee Cup logo, and J2EE are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

DISCLAIMER OF WARRANTIES

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

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

LIMITATION OF LIABILITY

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



Please
Recycle



Adobe PostScript

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

RESTRICTED RIGHTS LEGEND

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

REPORT

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

(LFI#136187/Form ID#011801)

Contents

1. Introduction 1-1

- 1.1 Overview 1-1
- 1.2 Scope 1-2
- 1.3 Target Audience 1-3
- 1.4 JDBC and Connector Architecture 1-4
- 1.5 Organization 1-4
- 1.6 Document Convention 1-5
- 1.7 Connector Architecture Expert Group for Version 1.5 (JSR-112) 1-5
- 1.8 Acknowledgements for Version 1.5 1-6
- 1.9 Connector Architecture Expert Group for Version 1.0 (JSR-16) 1-7
- 1.10 Acknowledgements for Version 1.0 1-8

2. Overview 2-1

- 2.1 Definitions 2-1
 - 2.1.1 Enterprise Information System (EIS) 2-1
 - 2.1.2 Connector Architecture 2-2
 - 2.1.3 EIS Resource 2-2
 - 2.1.4 Resource Manager (RM) 2-2
 - 2.1.5 Managed Environment 2-3
 - 2.1.6 Non-Managed Environment 2-3

2.1.7	Connection	2-3
2.1.8	Application Component	2-3
2.1.9	Container	2-3
2.2	Rationale	2-4
2.2.1	System Contracts	2-4
2.2.2	Common Client Interface	2-5
2.3	Goals	2-6
3.	The Connector Architecture	3-1
3.1	System Contracts	3-2
3.2	Client API	3-4
3.3	Requirements	3-5
3.4	Non-Managed Environment	3-5
4.	Roles and Scenarios	4-1
4.1	Roles	4-1
4.1.1	Resource Adapter Provider	4-1
4.1.2	Application Server Vendor	4-2
4.1.3	Container Provider	4-2
4.1.4	Application Component Provider	4-3
4.1.5	Enterprise Tools Vendors	4-3
4.1.6	Application Assembler	4-4
4.1.7	Deployer	4-4
4.1.8	System Administrator	4-5
4.2	Scenario: Integrated Purchase Order System	4-5
4.3	Scenario: Business Integration	4-8
5.	Lifecycle Management	5-1
5.1	Overview	5-1
5.2	Goals	5-2

5.3	Lifecycle Management Model	5-3
5.3.1	ResourceAdapter JavaBean and Bootstrapping a Resource Adapter Instance	5-5
5.3.2	ManagedConnectionFactory JavaBean and Outbound Communication	5-8
5.3.3	ActivationSpec JavaBean and Inbound Communication	5-9
5.3.4	Resource Adapter Shutdown Procedure	5-10
5.3.4.1	Phase One	5-10
5.3.4.2	Phase Two	5-11
5.3.5	Requirements	5-11
5.3.6	Resource Adapter Implementation Guidelines	5-13
5.3.7	JavaBean Configuration and Deployment	5-13
5.3.7.1	ResourceAdapter JavaBean instance Configuration	5-14
5.3.7.2	Resource Adapter Deployment	5-14
5.3.7.3	ManagedConnectionFactory JavaBean Instance Configuration	5-14
5.3.7.4	ActivationSpec JavaBean instance Configuration	5-14
5.3.7.5	Resource Adapter Implementation Guidelines	5-15
5.3.8	Lifecycle Management In A Non-Managed Environment	5-15
5.3.9	A Sample Resource Adapter Implementation	5-16
6.	Connection Management	6-1
6.1	Overview	6-1
6.2	Goals	6-1
6.3	Architecture: Connection Management	6-2
6.3.1	Overview: Managed Application Scenario	6-2
6.4	Application Programming Model	6-6
6.4.1	Managed Application Scenario	6-6
6.4.2	Non-managed Application Scenario	6-7

6.5	Interface/Class Specification	6-7
6.5.1	ConnectionFactory and Connection	6-10
6.5.1.1	Requirements	6-11
6.5.1.2	ConnectionRequestInfo	6-12
6.5.1.3	Additional Requirements	6-13
6.5.2	ConnectionManager	6-14
6.5.2.1	Interface	6-14
6.5.2.2	Requirements	6-15
6.5.3	ManagedConnectionFactory	6-16
6.5.3.1	Interface	6-17
6.5.3.2	Requirements	6-18
6.5.3.3	Connection Pool Implementation	6-19
6.5.3.4	Detecting Invalid Connections	6-20
6.5.3.5	Requirement for XA Recovery	6-20
6.5.4	ManagedConnection	6-20
6.5.4.1	Interface	6-22
6.5.4.2	Connection Sharing and Multiple Connection Handles	6-23
6.5.4.3	Connection Matching Contract	6-23
6.5.4.4	Cleanup of ManagedConnection	6-24
6.5.4.5	Requirements	6-25
6.5.5	ManagedConnectionMetaData	6-25
6.5.5.1	Interface	6-25
6.5.5.2	Requirements	6-26
6.5.6	ConnectionEventListener	6-26
6.5.6.1	Interface	6-26
6.5.7	ConnectionEvent	6-27
6.6	Error Logging and Tracing	6-28
6.6.1	ManagedConnectionFactory	6-28

6.6.2	ManagedConnection	6-29
6.7	Object Diagram	6-29
6.8	Illustrative Scenarios	6-32
6.8.1	Scenario: Connection Pool Management	6-32
6.8.2	Scenario: Connection Matching	6-37
6.8.3	Scenario: Connection Event Notifications and Connection Close	6-40
6.8.3.1	Connection Cleanup	6-41
6.8.3.2	Connection Destroy	6-41
6.9	Architecture: Non-managed Environment	6-44
6.9.1	Scenario: Programmatic Access to ConnectionFactory	6-45
6.9.2	Scenario: Connection Creation in Non-managed Application Scenario	6-47
6.10	Requirements	6-49
6.10.1	Resource Adapter	6-49
6.10.2	Application Server	6-50
7.	Transaction Management	7-1
7.1	Overview	7-2
7.2	Transaction Management Scenarios	7-3
7.2.1	Transactions across multiple Resource Managers	7-4
7.2.2	Local Transaction Management	7-5
7.3	Transaction Management Contract	7-6
7.3.1	Interface: ManagedConnection	7-9
7.3.2	Interface: XAResource	7-11
7.3.2.1	Implementation	7-11
7.3.3	Interface: LocalTransaction	7-12
7.4	Relationship to JTA and JTS	7-12
7.4.1	JTA Interfaces	7-12
7.5	Object Diagram	7-13

7.6	XAResource-based Transaction Contract	7-16
7.6.1	Scenarios Supported	7-17
7.6.2	Resource Adapter Requirements	7-18
7.6.2.1	General	7-19
7.6.2.2	One-phase Commit	7-19
7.6.2.3	Two-phase Commit	7-19
7.6.2.4	Transaction Association and Calling Protocol	7-20
7.6.2.5	Unilateral Roll-back	7-20
7.6.2.6	Read-Only Optimization	7-21
7.6.2.7	XID Support	7-21
7.6.2.8	Support for Failure Recovery	7-21
7.6.3	Transaction Manager Requirements	7-21
7.6.3.1	Interfaces	7-22
7.6.3.2	XID requirements	7-22
7.6.3.3	One-phase Commit Optimization	7-22
7.6.3.4	Implementation Options	7-22
7.6.4	Scenario: Transactional Setup for a ManagedConnection	7-22
7.6.5	Scenario: Connection Close and JTA Transactional Cleanup	7-26
7.6.6	OID: Transaction Completion	7-29
7.7	Local Transaction Management Contract	7-31
7.7.1	Interface: Local Transaction	7-31
7.7.2	Interface: ConnectionEventListener	7-31
7.7.2.1	Requirements	7-32
7.8	Scenarios: Local Transaction Management	7-33
7.8.1	Local Transaction Cleanup	7-33
7.8.2	Component Termination	7-34
7.8.3	Transaction Interleaving	7-34
7.8.3.1	Scenario	7-34

7.9	Connection Sharing	7-35
7.9.1	Sharing Violation Detection	7-36
7.9.1.1	Scenario 1	7-37
7.9.1.2	Scenario 2	7-37
7.10	Transaction Scenarios	7-37
7.10.1	Requirements	7-37
7.10.2	Illustrative Scenarios	7-39
7.10.3	Scenario: Local Transaction	7-40
7.11	Connection Association	7-45
7.11.1	Scenario	7-45
7.11.2	Connection Association	7-46
7.11.3	Requirements	7-47
7.12	Local Transaction Optimization	7-47
7.12.1	Requirements	7-48
7.13	Requirements	7-48
7.13.1	Resource Adapter	7-48
7.13.1.1	Auto Commit	7-49
7.13.2	Application Server	7-49
7.14	Connection Optimizations	7-50
7.14.1	Lazy Connection Association Optimization	7-50
7.14.1.1	API Additions	7-55
7.14.2	Lazy Transaction Enlistment Optimization	7-55
7.14.2.1	API Additions	7-56
8.	Security Architecture	8-1
8.1	Overview	8-1
8.2	Goals	8-1
8.3	Terminology	8-2
8.4	Application Security Model	8-3

- 8.4.1 Scenario: Container-Managed Sign-on 8-4
- 8.4.2 Scenario: Component-Managed Sign-on 8-5
- 8.5 EIS Sign-on 8-5
 - 8.5.1 Authentication Mechanism 8-6
 - 8.5.2 Resource Principal 8-6
 - 8.5.3 Authorization Model 8-7
 - 8.5.4 Secure Association 8-8
- 8.6 Roles and Responsibilities 8-9
 - 8.6.1 Application Component Provider 8-9
 - 8.6.2 Deployer 8-9
 - 8.6.3 Application Server 8-10
 - 8.6.4 EIS Vendor 8-11
 - 8.6.5 Resource Adapter Provider 8-11
 - 8.6.6 System Administrator 8-11

9. Security Contract 9-1

- 9.1 Security Contract 9-1
 - 9.1.1 Interfaces and Classes 9-1
 - 9.1.2 Subject 9-2
 - 9.1.3 Resource Principal 9-2
 - 9.1.4 GenericCredential 9-3
 - 9.1.4.1 Interface 9-4
 - 9.1.4.2 Implementation 9-4
 - 9.1.5 GSSCredential 9-5
 - 9.1.5.1 Implementation 9-5
 - 9.1.6 PasswordCredential 9-5
 - 9.1.7 ConnectionManager 9-7
 - 9.1.8 ManagedConnectionFactory 9-9
 - 9.1.8.1 Contract for the Application Server 9-10

9.1.8.2	Contract for Resource Adapter	9-12
9.1.9	ManagedConnection	9-14
9.2	Requirements	9-15
9.2.1	Resource Adapter	9-15
9.2.2	Application Server	9-15
10.	Work Management	10-1
10.1	Overview	10-1
10.2	Goals	10-2
10.3	Work Management Model	10-2
10.3.1	Requirements	10-4
10.3.2	Work Interface	10-9
10.3.3	WorkManager Interface	10-10
10.3.3.1	Work Submit	10-12
10.3.3.2	Work Accepted	10-13
10.3.3.3	Work Rejected	10-13
10.3.3.4	Work Started	10-14
10.3.3.5	Work Completed	10-14
10.3.3.6	Requirements	10-14
10.3.4	WorkListener Interface and WorkEvent Class	10-18
10.3.4.1	Requirements	10-20
10.3.5	ExecutionContext Class	10-20
10.3.6	Resource Adapter Thread Usage Recommendations	10-23
10.3.7	Periodic Execution of Work Instances	10-23
10.3.8	Illustration: Using a Work Instance to Listen on Multiple Network Endpoints	10-25
10.3.9	Work Management in a Non-Managed Environment	10-25
11.	Inbound Communicaton	11-1
11.1	Overview	11-1

11.2 An Illustrative Use Case 11-2

12. Message Inflow 12-1

12.1 Overview 12-1

12.2 Goals 12-3

12.3 Message Inflow Model 12-4

12.4 Endpoint Deployment 12-9

12.4.1 Message Endpoint 12-10

12.4.2 Resource Adapter 12-11

12.4.2.1 List of Supported Message Listener Types 12-11

12.4.2.2 ActivationSpec JavaBean 12-11

12.4.2.3 Administered Objects 12-12

12.4.2.4 Configuring Administered Objects 12-13

12.4.3 Endpoint Deployer 12-13

12.4.4 Application Server 12-14

12.4.5 Message Provider 12-16

12.4.6 Endpoint Deployment Steps 12-18

12.4.7 Requirements 12-18

12.4.8 Structure of a Message Listener Interface 12-19

12.4.9 Multiple Endpoint Activations With Similar Activation Configuration 12-20

12.4.9.1 Requirements 12-20

12.5 Message Delivery 12-22

12.5.1 Sample Resource Adapter Code To Illustrate Message Delivery 12-24

12.5.1.1 Requirements 12-26

12.5.2 Message Redelivery Upon Crash Recovery 12-26

12.5.3 Durable Message Delivery Setup 12-27

12.5.4 Concurrent Delivery of Messages 12-28

12.5.4.1 Requirements 12-28

12.5.5	Delivery Semantics and Acknowledgement	12-28
12.5.6	Transacted Delivery (Using Container-Managed Transaction)	12-29
12.5.7	Non-transacted Delivery	12-31
12.5.8	Transacted Delivery Using an Imported Transaction	12-32
12.5.9	Requirements	12-33
12.6	Endpoint Undeployment	12-33
12.7	Java Message Service (JMS) Use Case	12-38
12.7.0.1	A Sample JMS Resource Adapter Deployment Descriptor	12-39
12.7.0.2	A Sample JMS ActivationSpec Implementation	12-44
12.7.0.3	A Sample EJB 2.0 Message-driven Bean Deployment Descriptor	12-44
12.7.0.4	A Sample EJB 2.1 Message-driven Bean Deployment Descriptor	12-45
12.7.0.5	A Sample EJB 2.1 EJB Deployment Descriptor	12-47
12.7.1	Message-driven Bean Asynchronously Receiving Messages	12-47
12.7.1.1	Message-Driven Bean Deployment	12-47
12.7.1.2	Message Delivery	12-48
12.7.1.3	Message-Driven Bean Undeployment	12-48
12.7.2	EJB Using JMS API to Send and Synchronously Receive Messages Via a JMS Resource Adapter	12-49
12.7.2.1	Using JMS API to Send Messages	12-50
12.7.2.2	J2EE Component Using JMS API to Synchronously Receive Messages	12-50
12.8	A Non-JMS Use Case	12-51
12.8.1	Resource Adapter Deployment Descriptor	12-51
12.8.2	Resource Adapter Deployment	12-53
12.8.3	Message-Driven Bean Asynchronously Receiving Notifications From an EIS	12-54

12.8.3.1	The Message-Driven Bean Deployment Descriptor	12-54
12.8.4	Message-Driven Bean and Resource Adapter Activation	12-55
12.8.5	Message Delivery	12-55
13.	EJB Invocation	13-1
13.1	Overview	13-1
13.2	EJB Invocation Model	13-2
13.3	An Illustrative Use Case	13-3
13.3.0.1	Message-driven Bean Dispatcher Pattern	13-5
14.	Transaction Inflow	14-1
14.1	Overview	14-1
14.2	Goals	14-2
14.3	Use Case Scenario	14-2
14.4	Transaction Inflow Model	14-4
14.4.1	Processing of Transactional Calls	14-4
14.4.2	Transaction Completion Processing	14-5
14.4.3	Crash Recovery Processing	14-8
14.4.4	Requirements	14-12
14.4.5	Non-Requirements	14-13
14.4.6	Recommendations	14-13
14.4.7	Transaction Inflow in a Non-managed Environment	14-13
15.	Common Client Interface	15-1
15.1	Overview	15-2
15.2	Goals	15-2
15.3	Scenarios	15-3
15.3.1	Enterprise Application Integration Framework	15-3
15.3.2	Metadata Repository and API	15-4

15.3.3	Enterprise Application Development Tool	15-4
15.4	Common Client Interface	15-5
15.4.1	Requirements	15-6
15.5	Connection Interfaces	15-9
15.5.1	ConnectionFactory	15-9
15.5.1.1	Requirements	15-10
15.5.2	ConnectionSpec	15-10
15.5.3	Connection	15-11
15.5.3.1	Auto Commit	15-13
15.6	Interaction Interfaces	15-13
15.6.1	Interaction	15-13
15.6.2	InteractionSpec	15-14
15.6.2.1	Standard Properties	15-15
15.6.2.2	ResultSet Properties	15-16
15.6.2.3	Additional Properties	15-16
15.6.2.4	Implementation	15-17
15.6.2.5	Administered Object	15-17
15.6.2.6	Illustrative Scenario	15-17
15.6.3	LocalTransaction	15-18
15.6.3.1	Requirements	15-18
15.7	Basic Metadata Interfaces	15-18
15.7.1	ConnectionMetaData	15-19
15.7.1.1	Implementation	15-19
15.7.2	ResourceAdapterMetaData	15-19
15.8	Service Endpoint Message Listener Interface	15-21
15.9	Exception Interfaces	15-21
15.9.1	ResourceException	15-21
15.9.2	ResourceWarning	15-22

15.10	Record	15-22
15.10.1	Component-view Contract	15-24
15.10.1.1	Type Mapping	15-25
15.10.1.2	Record Interface	15-26
15.10.1.3	MappedRecord and IndexedRecord Interfaces	15-28
15.10.1.4	RecordFactory	15-28
15.10.2	Interaction and Record	15-29
15.10.3	Resource Adapter-view Contract	15-29
15.10.3.1	Streamable Interface	15-30
15.11	ResultSet	15-31
15.11.1	ResultSet Interface	15-33
15.11.1.1	Type Mapping	15-33
15.11.1.2	ResultSet Types	15-33
15.11.1.3	Scrolling	15-34
15.11.1.4	Concurrency Types	15-34
15.11.1.5	Updatability	15-35
15.11.1.6	Persistence of Java Objects	15-35
15.11.1.7	Support for SQL Types	15-35
15.11.1.8	Support for Customized SQL Type Mapping	15-36
15.11.2	ResultSetMetaData	15-36
15.11.3	ResultSetInfo	15-36
15.12	Code Samples	15-38
15.12.1	Connection	15-38
15.12.2	InteractionSpec	15-39
15.12.3	Mapped Record	15-39
15.12.4	ResultSet	15-40
15.12.5	Custom Record	15-41

16. API Requirements 16-1

16.1	Requirements of the Application Server	16-1
16.2	Requirements of the Resource adapter	16-1
16.3	JavaBean Requirements	16-2
16.4	Equality Constraints	16-3
17.	Packaging Requirements	17-1
17.1	Overview	17-1
17.2	Packaging	17-4
	17.2.0.1 Resource Adapter Archive	17-4
	17.2.0.2 RAR Contents	17-5
	17.2.0.3 Sample Directory Structure	17-5
	17.2.0.4 Requirements	17-5
17.3	Deployment	17-6
	17.3.1 Resource Adapter Provider	17-6
	17.3.2 Deployer	17-9
	17.3.2.1 Stand-Alone Resource Adapter Module	17-9
	17.3.2.2 Resource Adapter Module with J2EE Application	17-9
	17.3.2.3 Configuration	17-10
	17.3.2.4 Security Configuration	17-10
17.4	Interfaces/Classes	17-11
	17.4.1 ResourceAdapter	17-11
	17.4.1.1 Requirements	17-11
	17.4.2 ManagedConnectionFactory	17-11
	17.4.2.1 Requirements	17-12
	17.4.3 Properties Conventions	17-12
	17.4.4 Standard Properties	17-12
17.5	JNDI Configuration and Lookup	17-13
	17.5.1 Responsibilities	17-14
	17.5.1.1 Deployer	17-14

17.5.1.2	Resource Adapter	17-14
17.5.1.3	Application Server	17-15
17.5.2	Scenario: Serializable	17-15
17.5.3	Scenario: Referenceable	17-17
17.5.3.1	ObjectFactory Implementation	17-17
17.5.3.2	Deployment	17-19
17.5.3.3	Scenario: Connection Factory Lookup	17-19
17.5.4	Requirements	17-23
17.6	Resource Adapter XML Schema Definition	17-23
18.	Runtime Environment	18-1
18.1	Programming APIs	18-1
18.2	Security Permissions	18-2
18.3	Requirements	18-5
18.3.1	Example	18-6
18.4	Privileged Code	18-7
18.4.1	Example	18-7
19.	Exceptions	19-1
19.1	ResourceException	19-1
19.2	System Exceptions	19-2
19.2.1	Exception Hierarchy	19-2
19.3	Work Exceptions	19-4
19.4	Additional Exceptions	19-5
20.	Projected Items	20-1
A.	Previous Version DTDs	A-1
A.1	J2EE Connector Architecture 1.0 Resource Adapter XML DTD	A-1
B.	Activation Configuration for Message Inflow to JMS Endpoints	B-1

B.1	Introduction	B-1
B.2	JMS ActivationSpec JavaBean	B-1
B.2.1	JMS ActivationSpec JavaBean Properties	B-2
B.2.1.1	destination	B-2
B.2.1.2	destinationType	B-2
B.2.1.3	messageSelector	B-3
B.2.1.4	acknowledgeMode	B-3
B.2.1.5	subscriptionDurability	B-3
B.2.1.6	clientId	B-4
B.2.1.7	subscriptionName	B-4
B.2.2	JMS ActivationSpec JavaBean Property Values	B-4
B.3	JMS Endpoint with EJB 2.1 Activation Configuration Elements	B-5
B.4	JMS Endpoint with EJB 2.0 Deployment Descriptor Elements	B-6
C.	Caching Manager	C-1
C.1	Overview	C-1
C.2	Synchronization contract	C-2
C.2.1	Interface	C-3
C.2.2	Implementation	C-3
D.	Security Scenarios	D-1
D.1	EStore Application	D-1
D.1.1	Scenario	D-2
D.1.2	Security Environment	D-2
D.1.3	Deployment	D-3
D.2	Employee Self Service Application	D-4
D.2.1	Architecture	D-5
D.2.2	Security Environment	D-5
D.2.3	Deployment	D-6

D.2.4	Scenario	D-7
D.3	Integrated Purchasing Application	D-7
D.3.1	Architecture	D-7
D.3.2	Security Environment	D-8
D.3.3	Deployment	D-9
E.	JAAS Based Security Architecture	E-1
E.1	Java Authentication and Authorization Service (JAAS)	E-1
E.2	Requirements	E-2
E.3	Security Architecture	E-3
E.3.1	JAAS Modules	E-5
E.3.2	Illustrative Examples: JAAS Module	E-5
E.3.2.1	Principal Mapping Module	E-5
E.3.2.2	Credential Mapping Module	E-6
E.3.2.3	Kerberos Module	E-6
E.4	Security Configuration	E-7
E.4.1	JAAS Configuration	E-7
E.5	Scenarios	E-8
E.5.1	Scenario: Resource Adapter Managed Authentication	E-8
E.5.2	Scenario: Kerberos and Principal Delegation	E-10
E.5.3	Scenario: GSS-API	E-12
E.5.4	Scenario: Kerberos Authentication After Principal Mapping	E-13
E.5.5	Scenario: EIS-Specific Authentication	E-14
F.	Related Documents	F-1
G.	Change History for Version 1.0	G-1
G.1	Version 0.9	G-1
G.2	Version 1.0 - Public Draft 1	G-2
G.3	Version 1.0 - Public Draft 2	G-2

G.4 Version 1.0 - Proposed Final Draft 2 G-3

G.5 Version 1.0 - Final Release G-4

H. Change History for Version 1.5 H-1

H.1 Public Review Draft H-1

H.2 Proposed Final Draft H-1

H.3 Proposed Final Draft 2 H-2

H.4 Final Release H-3

Tables

TABLE 1-1	Connector 1.5 Expert Group Members	1–5
TABLE 1-2	Connector 1.0 Expert Group Members	1–7
TABLE 7-1	Transaction Management Scenarios	7–17
TABLE 12-1	Application Server Behavior For Transacted and Non-transacted Message Delivery	12–33
TABLE 15-1	Standard Properties for ConnectionSpec	15–11
TABLE 16-1	Resource Adapter API Requirements	16–2
TABLE 17-1	Description of RAR File Contents	17–5
TABLE 17-2	Standard Properties of the Connector Architecture	17–12
TABLE 18-1	Default Security Permission Set	18–3
TABLE 18-2	Methods and Security Permissions required	18–6
TABLE B-1	JMS ActivationSpec JavaBean Property Values	B–4
TABLE B-2	Activation Configuration Elements (EJB 2.1) and ActivationSpec JavaBean Properties	B–5
TABLE B-3	EJB 2.0 Deployment Descriptor Elements and JMS ActivationSpec JavaBean Properties	B–6

Figures

FIGURE 2-1	System Level Pluggability between Application Servers and EISs	2–5
FIGURE 3-1	Overview of the Connector Architecture	3–2
FIGURE 4-1	Illustration of a Scenario Based On the Connector Architecture	4–6
FIGURE 4-2	Connector Architecture Usage In Business Integration Scenario	4–9
FIGURE 5-1	Lifecycle Management Contract (interfaces)	5–3
FIGURE 5-2	Lifecycle Management (object diagram)	5–4
FIGURE 5-3	Resource Adapter Instance (composition)	5–7
FIGURE 5-4	Resource Adapter Lifecycle (state diagram)	5–12
FIGURE 5-5	Lifecycle Management Model (sequence diagram)	5–17
FIGURE 6-1	Architecture Diagram: Managed Application scenario	6–5
FIGURE 6-2	Class Diagram: Connection Management Architecture	6–9
FIGURE 6-3	ConnectionManager and Application Server specific services	6–16
FIGURE 6-4	Object Diagram: Connection Management architecture	6–31
FIGURE 6-5	OID: Connection Pool Management with new Connection Creation	6–36
FIGURE 6-6	OID: Connection Pool Management with Connection Matching	6–39
FIGURE 6-7	OID: Connection Event Notification	6–43
FIGURE 6-8	Architecture Diagram: Non-Managed application scenario	6–45
FIGURE 6-9	OID: Connection Creation in a Non-managed Application Scenario	6–48
FIGURE 7-1	Transaction Management Contract	7–2
FIGURE 7-2	Scenario: Transactions Across Multiple Resource Managers	7–4

FIGURE 7-3	Scenario: Local Transaction on a Single Resource Manager	7-6
FIGURE 7-4	Architecture Diagram: Transaction Management	7-8
FIGURE 7-5	ManagedConnection Interface for Transaction Management	7-10
FIGURE 7-6	Object Diagram: Transaction Management	7-15
FIGURE 7-7	OID: Transactional Setup For Newly Created ManagedConnection Instances	7-25
FIGURE 7-8	OID: Connection Close and Transactional Cleanup	7-28
FIGURE 7-9	OID: Transaction Completion	7-30
FIGURE 7-10	Scenario to illustrate Local Transaction Management	7-40
FIGURE 7-11	OID: Connection Sharing across Component instances	7-43
FIGURE 7-12	Connection Sharing Scenario	7-45
FIGURE 7-13	State Diagram of Application-Level Connection Handle	7-46
FIGURE 7-14	Connection Acquisition Processing	7-52
FIGURE 7-15	Connection Re-association Processing	7-53
FIGURE 7-16	State Diagram of a Dissociatable Application-level Connection Handle	7-54
FIGURE 9-1	Security Contract	9-8
FIGURE 9-2	Security Contract: Subject Interface and its Containment Hierarchy	9-10
FIGURE 10-1	Work Management Contract (object diagram)	10-5
FIGURE 10-2	Work Management Contract (interfaces)	10-6
FIGURE 10-3	Work Processing Stages and their Outcomes	10-13
FIGURE 10-4	Blocking Durations of Various Work Submissions	10-16
FIGURE 10-5	Work Submission - Blocking Behavior (sequence diagram)	10-17
FIGURE 10-6	Work Submission - Callback Mechanism (sequence diagram)	10-22
FIGURE 11-1	Inbound Communication Model	11-1
FIGURE 11-2	Inbound Communication Model (an illustrative use case)	11-3
FIGURE 12-1	Message Inflow Contract	12-2
FIGURE 12-2	Message Inflow Contract (object diagram)	12-4
FIGURE 12-3	Message Inflow Contract (interfaces)	12-5
FIGURE 12-4	Endpoint (Message-Driven Bean) Deployment (Actors)	12-9
FIGURE 12-5	Endpoint (message-driven bean) Deployment Steps	12-17
FIGURE 12-6	Endpoint Deployment (sequence diagram)	12-21

FIGURE 12-7	Transacted Message Delivery: Option A (sequence diagram)	12–35
FIGURE 12-8	Transacted Message Delivery: Option B (sequence diagram)	12–36
FIGURE 12-9	Non-transacted Message Delivery (sequence diagram)	12–37
FIGURE 12-10	Endpoint Undeployment (sequence diagram)	12–38
FIGURE 13-1	EJB Invocation Model	13–2
FIGURE 14-1	Transaction Inflow Contract	14–2
FIGURE 14-2	Transaction Inflow Contract (object diagram)	14–3
FIGURE 14-3	Transactional Calls and Transaction Completion Flow (sequence diagram)	14–7
FIGURE 14-4	Crash Recovery Flows When Application Server Crashes (sequence diagram)	14–10
FIGURE 14-5	Crash Recovery Flows When EIS Crashes (sequence diagram)	14–11
FIGURE 15-1	Common Client Interface	15–2
FIGURE 15-2	Scenario: EAI Framework	15–4
FIGURE 15-3	Scenario: Enterprise Application Development Tool	15–5
FIGURE 15-4	Class Diagram: Common Client Interface	15–8
FIGURE 15-5	Record at Development-time and Runtime	15–23
FIGURE 15-6	Component-view Contract	15–26
FIGURE 15-7	Streamable Interface	15–30
FIGURE 15-8	ResultSet Interface	15–32
FIGURE 17-1	Packaging and Deployment Lifecycle of a Resource adapter	17–1
FIGURE 17-2	Deployment of a Resource Adapter Module	17–3
FIGURE 17-3	OID: Lookup of a ConnectionFactory Instance from JNDI	17–22
FIGURE C-1	Synchronization Contract between Caching Manager and Application Server	C–2
FIGURE D-1	Illustrative Architecture of an Estore Application	D–2
FIGURE D-2	Resource Principal for Estore Application Scenario	D–4
FIGURE D-3	Illustrative Architecture of an Employee Self-service Application	D–6
FIGURE D-4	Principal Mapping	D–6
FIGURE D-5	Illustrative Architecture of an Integrated Purchasing Application	D–8
FIGURE D-6	Principal Mapping	D–9
FIGURE E-1	Security Architecture.	E–4
FIGURE E-2	Resource Adapter-Managed Authentication	E–9

FIGURE E-3	Kerberos Authentication with Principal Delegation	E-10
FIGURE E-4	GSS-API use by Resource Adapter	E-12
FIGURE E-5	Kerberos Authentication After Principal Mapping	E-13
FIGURE E-6	Authentication Through EIS-Specific JAAS Module	E-14

Code Samples

CODE EXAMPLE 5-1	Sample Resource Adapter	5-16
CODE EXAMPLE 10-1	javax.resource.spi.work	10-7
CODE EXAMPLE 12-1	javax.resource.spi	12-6
CODE EXAMPLE 12-2	javax.resource.spi.endpoint	12-6
CODE EXAMPLE 12-3	Message-Driven Bean Deployment Descriptor	12-10
CODE EXAMPLE 12-4	Message Delivery In a Resource Adapter	12-24
CODE EXAMPLE 12-5	Sample JMS Resource Adapter Deployment Descriptor	12-39
CODE EXAMPLE 12-6	Message-Driven Bean Deployment Descriptor	12-45
CODE EXAMPLE 12-7	Sample EJB 2.1 Deployment Descriptor	12-47
CODE EXAMPLE 12-8	Sending Messages Using the JMS API	12-50
CODE EXAMPLE 12-9	Synchronously Receiving Messages in a J2EE Component	12-50
CODE EXAMPLE 12-10	Deployment Descriptor for a Resource Adapter	12-52
CODE EXAMPLE 12-11	Deployment Descriptor for a Message-Driven Bean	12-54
CODE EXAMPLE 13-1	A Message-Driven Bean Implementation	13-3
CODE EXAMPLE 17-1	Schema Definition for the Deployment Descriptor for a Resource Adapter	17-23
CODE EXAMPLE A-1	Connector Architecture 1.0 Resource Adapter DTD	A-1

Introduction

The Java™ 2 Platform, Enterprise Edition (J2EE™ platform) provides containers for client applications, web components based on servlets and Java Server Pages™ (JSP™), and Enterprise JavaBeans™ (EJB™) components. These containers provide deployment and runtime support for application components. They provide a federated view of the services provided by the underlying application server for the application components.

Containers can run on existing systems; for example, web servers for the web containers; application servers, TP monitors, and database systems for EJB containers. This enables enterprises to leverage both the advantages of their existing systems and those of J2EE. Enterprises can write, or rewrite, new applications using J2EE capabilities and can also encapsulate parts of existing applications in Enterprise JavaBeans, Java Server Pages or servlets.

Enterprise applications access functions and data associated with applications running on Enterprise Information Systems (EIS). Application servers extend their containers and support connectivity to heterogeneous EISs. Enterprise tools and Enterprise Application Integration (EAI) vendors add value by providing tools and frameworks to simplify the EIS integration task.

For enterprise application integration, bi-directional connectivity between enterprise applications and EIS is essential. The J2EE Connector architecture defines standard contracts that allow bi-directional connectivity between enterprise applications and EISs. It also formalizes the relationships, interactions, and the packaging of the integration layer, thus enabling enterprise application integration.

1.1 Overview

The J2EE Connector architecture defines a standard architecture for connecting the J2EE platform to heterogeneous EISs. Examples of EISs include Enterprise Resource Planning (ERP), mainframe transaction processing (TP), and database systems.

The connector architecture defines a set of scalable, secure, and transactional mechanisms that enable the integration of EISs with application servers¹ and enterprise applications.

The connector architecture also defines a Common Client Interface (CCI) for EIS access. The CCI defines a client API for interacting with heterogeneous EISs.

The connector architecture enables an EIS vendor to provide a standard resource adapter for its EIS. A resource adapter is a system-level software driver that is used by a Java application to connect to an EIS. The resource adapter plugs into an application server and provides connectivity between the EIS, the application server, and the enterprise application. The resource adapter serves as a protocol adapter that allows any arbitrary EIS communication protocol to be used for connectivity.

An application server vendor extends its system once to support the connector architecture and is then assured of seamless connectivity to multiple EISs. Likewise, an EIS vendor provides one standard resource adapter which has the capability to plug in to any application server that supports the connector architecture.

1.2 Scope

Version 1.0 of the connector architecture defines:

- A standard set of system-level contracts between an application server and EIS. These contracts focus on the important system-level aspects of integration: connection management, transaction management, and security.
- A Common Client Interface (CCI) that defines a client API for interacting with multiple EISs.
- A standard deployment and packaging protocol for resource adapters.

Refer to section 2.2.2 for the rationale behind the Common Client Interface.

Version 1.5 of the connector architecture defines:

- Lifecycle management contract: A contract between an application server and a resource adapter that allows an application server to manage the lifecycle of a resource adapter. This contract provides a mechanism for the application server to bootstrap a resource adapter instance during its deployment or application server startup, and to notify the resource adapter instance during its undeployment or during an orderly shutdown of the application server.
- Work management contract: A contract between an application server and a resource adapter that allows a resource adapter to do work (monitor network endpoints, call application components, etc.) by submitting Work instances to an

1. Application server is a generic term used in this document to refer to a middle-tier component server that is compliant with the Java 2 Platform, Enterprise Edition.

application server for execution. The application server dispatches threads to execute submitted Work instances. This allows a resource adapter to avoid creating or managing threads directly, and allows an application server to efficiently pool threads and have more control over its runtime environment. The resource adapter can control the security context and transaction context with which Work instances are executed.

- **Transaction inflow contract:** A contract between an application server and a resource adapter that allows a resource adapter to propagate an imported transaction to an application server. This contract also allows a resource adapter to transmit transaction completion and crash recovery calls initiated by an EIS, and ensures that the ACID (Atomicity, Consistency, Isolation and Durability) properties of the imported transaction are preserved.
- **Message inflow contract:** A standard, generic contract between an application server and a resource adapter that allows a resource adapter to asynchronously deliver messages to message endpoints residing in the application server independent of the specific messaging style, messaging semantics, and messaging infrastructure used to deliver messages. This contract also serves as the standard message provider pluggability contract that allows a wide range of message providers (Java Message Service (JMS), Java API for XML Messaging (JAXM), etc.) to be plugged into any J2EE compatible application server via a resource adapter.
- Describes the packaging model for different types of resource adapters (outbound only, inbound only, or both).

1.3 Target Audience

The target audience for this specification includes:

- EIS vendors and resource adapter providers
- Messaging system vendors
- Application server vendors and container providers
- Enterprise application developers and system integrators
- Enterprise tool and EAI vendors

The system-level contracts between an application server and an EIS are targeted towards EIS vendors (or resource adapter providers, if the two roles are different) and application server vendors. The CCI is targeted primarily towards enterprise tools and EAI vendors.

1.4 JDBC and Connector Architecture

The JDBC™ API defines a standard Java API for accessing relational databases. JDBC provides an API for sending SQL statements to a database and processing the tabular data returned by the database.

The connector architecture is a standard architecture for integrating J2EE applications with EISs that are not relational databases. Each of these EISs currently provides a native function call API for identifying a function to call, specifying its input data, and processing its output data. The goal of the Common Client Interface (CCI) is to provide an EIS independent API for coding these EIS function calls.

The CCI is targeted at EIS development tools and other sophisticated users of EISs. The CCI provides a way to minimize the EIS specific code required by such tools. Most J2EE developers will access EISs using these tools rather than using CCI directly.

It is expected that many J2EE applications will combine relational database access using JDBC with EIS access using EIS access tools based on CCI.

The connector architecture defines a standard SPI (Service Provider Interface) for integrating the transaction, security, and connection management facilities of an application server with those of a transactional resource manager. The JDBC 3.0 specification (Section 3., “JDBC API specification, version 3.0” on page F-1) specifies the relationship of JDBC to the SPI specified in the connector architecture.

1.5 Organization

This document begins by describing the rationale and goals for creating a standard architecture to integrate an application server with multiple heterogeneous EISs. It then describes the key concepts relevant to the connector architecture. These sections provide an overall picture of the architecture.

This document then describes typical scenarios for using the connector architecture. This chapter introduces the various roles and responsibilities involved in the development and deployment of enterprise applications that integrate with multiple EISs.

After these descriptive sections, this document focuses on the prescriptive aspects of the connector architecture.

1.6 Document Convention

A regular Palatino font is used for describing the connector architecture.

An *italic* font is used for paragraphs that contain descriptive notes providing clarifications.

A regular `Courier` font is used for Java source code, class, interface and method names.

Note that the scenarios described in this document are illustrative in scope. The intent of the scenarios is not to specify a prescriptive way of implementing a particular contract.

The requirements section occurring in various chapters of this document highlight only the salient requirements, but does not contain all the requirements. So, this entire document must be used as a requirements specification.

This document uses the EJB component model to describe some scenarios. The EJB specification (Section 1., “Enterprise JavaBeans (EJB™) specification, version 2.1:” on page F-1) provides the latest details of the EJB component model.

1.7 Connector Architecture Expert Group for Version 1.5 (JSR-112)

Refer to the URL <http://www.jcp.org/en/jsr/detail?id=112> for details on JSR-112. The following table lists the members of the Connector Expert Group.

TABLE 1-1 Connector 1.5 Expert Group Members

Company	Expert Member
BEA	Jim Gish, Mitch Upton, Tom Mitchell
Bull S.A.	Michael Giroux
Ericsson Infotech AB	Peter Kristiansson
Fujitsu Limited	Ivar Alexander
Hewlett-Packard	John Speidel
IBM	Michael Beisiegel, Piotr Przybylski
Bahwan Cybertek Technologies Inc.	Vijay Sundhar

TABLE 1-1 Connector 1.5 Expert Group Members

Company	Expert Member
Individual Expert	Charlton Barreto
IONA Technologies PLC	Gary Tully
MicroFocus	Stephen Gennard
NEON Systems, Inc.	Yongtao You
Oracle	Anthony Lai
SAP AG Walldorf	Nikolai Tankov
Siemens AG	Ulf Lange
Silverstream Software	Roman Kishchenko
Sonic Software	Jaime Meritt
Sun Microsystems, Inc.	Ram Jeyaraman (Specification Lead)
Sybase	Mark DeLaFranier
TIBCO Software Inc.	Jon Dart, Anuradha M. Sastri
Unisys	Albert DeNigris
WebMethods Corp	Bruce Tran, Bruno Kurtic

1.8 Acknowledgements for Version 1.5

This specification is the work of many people. Bill Shannon, Mark Hapner, Kenneth Saks, Rahul Sharma, and Kate Stout contributed to the overall design. Umit Yalcinalp designed the Connector 1.5 XML Schema Definition.

Venkat Amirisetty, Qingqing Ouyang, Binod P G, Balaji Raghunathan, Srikanth Padakandla and Aditya Gore from the J2EE Implementation team, Gursharan Singh from J2EE Compatibility Test Suite team, provided useful feedback on the various drafts of the specification. Ian Evans provided excellent editorial assistance.

Vijay Sarathy and Vivek Khandelwal did a great job evangelizing, strategizing, and guiding the overall direction of this specification. Anita Jindal, Jennifer Douglas, David Heisser, Bonnie Kellett, and Peter Walker provided management support.

Sanjeev Krishnan, Max Mortazavi, Tony Ng, Linda DeMichiel, George Tharakan, Colleen Evans (Sonic Software), Hemanth Puttaswamy, Sheetal Vartak and Shivaram Mysore provided useful comments on the specification.

And of course, this specification was formed and molded based on conversations with and review feedback from the expert group members.

1.9 Connector Architecture Expert Group for Version 1.0 (JSR-16)

Refer to the URL <http://www.jcp.org/en/jsr/detail?id=16> for details on JSR-16. The following are part of the expert group and have made invaluable contributions to the Connector architecture specification:

TABLE 1-2 Connector 1.0 Expert Group Members

Company	Expert Member
BEA	Pete Homan
Fujitsu	Yoshi Otagiri, Ivar Alexander
IBM	Tom Freund, Michael Beisiegel
Inline	Jack Greenfield
Inprise	Charlton Barreto
IPlanet	Tony Pan, Pavan Bhatnagar
Motorola	Guy Bieber
Oracle	Dan Coyle
SAP	Marek Barwicki
Sun	Rahul Sharma (Specification Lead) Fred H. Carter
Sybase	Rajini Balay, K. Swaminathan
Tibco	Jon Dart
Unisys	Lester Lee

1.10 Acknowledgements for Version 1.0

Shel Finkelstein, Mark Hapner, Vlada Matena, Tony Ng, Bill Shannon, and Sekhar Vajjhala (all from Sun Microsystems) have provided invaluable technical input and guidance to the Connector architecture specification. Jean Zeng and Pong Ching also provided useful input to the specification.

Rick Cattell, Shel Finkelstein, Bonnie Kellett, and Jeff Jackson have provided huge support to the specification lead in the management of the Connectors expert group.

Tony Ng is leading the effort of providing a reference implementation for the Connector architecture as part of J2EE 1.3 platform. Liz Blair has worked on providing the Compatibility Test Suite (CTS) plan for the Connector architecture.

Beth Stearns was a great help in doing an editorial review of this document.

Overview

This chapter introduces key concepts that are required to understand the connector architecture. It lays down a reference framework to facilitate a formal specification of the connector architecture in the subsequent chapters of this document.

2.1 Definitions

2.1.1 Enterprise Information System (EIS)

An EIS provides the information infrastructure for an enterprise. An EIS offers a set of services to its clients. These services are exposed to clients as local and/or remote interfaces. Examples of an EIS include:

- ERP system
- Mainframe transaction processing system
- Legacy database system

There are two aspects of an EIS:

- System level services—for example, SAP RFC, CICS ECI
- An application specific interface—for example, the table schema and specific stored procedures of a database, the specific CICS TP program

2.1.2 Connector Architecture

An architecture for integrating J2EE servers with EISs. There are two parts to this architecture: an EIS vendor-provided resource adapter and an application server that allows this resource adapter to be plugged in. This architecture defines a set of contracts (such as transactions, security, connection management) that a resource adapter has to support to plug in to an application server.

These contracts support bi-directional communication (outbound and inbound) between an application server and an EIS via a resource adapter. That is, the application server may use the resource adapter for outbound communication to the EIS, and it may also use the resource adapter for inbound communication from the EIS.

2.1.3 EIS Resource

An EIS resource provides EIS-specific functionality to its clients. Examples are:

- A record or set of records in a database system
- A business object in an ERP system
- A transaction program in a transaction processing system

2.1.4 Resource Manager (RM)

A resource manager manages a set of shared EIS resources. A client requests access to a resource manager to use its managed resources. A transactional resource manager can participate in transactions that are externally controlled and coordinated by a transaction manager.

In the context of the connector architecture, a client of a resource manager can either be a middle-tier application server or a client-tier application. A resource manager is typically in a different address space or on a different machine from the client that accesses it.

This document refers to an EIS as a resource manager when it is mentioned in the context of transaction management. Examples of resource managers are a database system, a mainframe TP system, and an ERP system.

2.1.5 Managed Environment

A managed environment defines an operational environment for a J2EE-based, multi-tier, web-enabled application that accesses EISs. The application consists of one or more application components—EJBs, JSPs, servlets—which are deployed on containers. These containers can be one of the following:

- Web containers that host JSPs, servlets, and static HTML pages
- EJB containers that host EJB components
- Application client containers that host standalone application clients

2.1.6 Non-Managed Environment

A non-managed environment defines an operational environment for a two-tier application. An application client directly uses a resource adapter to access the EIS, which defines the second tier of a two-tier application.

2.1.7 Connection

A connection provides connectivity to a resource manager. It enables an application client to connect to a resource manager, perform transactions, and access services provided by that resource manager. A connection can be either transactional or non-transactional. Examples include a database connection and an SAP R/3 connection. A connection to a resource manager may be used by a client for bi-directional communication, depending on the capabilities of the resource manager.

2.1.8 Application Component

An application component can be a server-side component, such as an EJB, JSP, or servlet, that is deployed, managed, and executed on an application server. It can also be a component executed on the web-client tier but made available to the web-client by an application server. Examples of the latter type of application component include a Java applet, and a DHTML page.

2.1.9 Container

A container is a part of an application server that provides deployment and runtime support for application components. It provides a federated view of the services provided by the underlying application server for the application components. For

more details on different types of standard containers, refer to the EJB (Section 1., “Enterprise JavaBeans (EJB™) specification, version 2.1:” on page F-1), JSP, and servlet specifications.

2.2 Rationale

The following section describes the rationale behind the connector architecture.

2.2.1 System Contracts

A standard architecture is needed to integrate various EISs with an application server. Without a standard, EIS vendors and application server vendors may have to use vendor-specific architectures to provide EIS integration.

The connector architecture provides a Java solution to the problem of bi-directional connectivity between the multitude of application servers and EISs. By using the connector architecture, it is no longer necessary for EIS vendors to customize their product for each application server. An application server vendor who conforms to the connector architecture also does not need to add custom code whenever it wants to extend its application server to support connectivity to a new EIS.

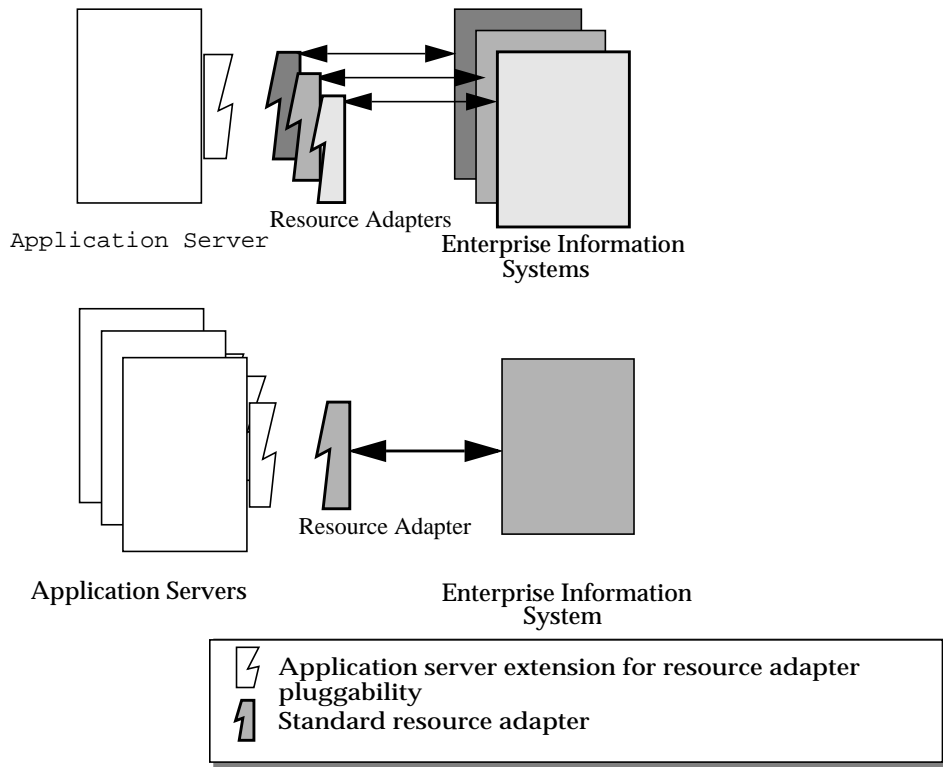
The connector architecture enables an EIS vendor to provide a standard resource adapter for its EIS. The resource adapter plugs into an application server and provides the underlying infrastructure for the integration between an EIS and the application server.

An application server vendor extends its system only once to support the connector architecture and is then assured of connectivity to multiple EISs. Likewise, an EIS vendor provides one standard resource adapter and it has the capability to plug in to any application server that supports the connector architecture.

FIGURE 2-1 on page 5 shows that a standard EIS resource adapter can plug into multiple application servers. Similarly, multiple resource adapters for different EISs can plug into an application server. This system-level pluggability is made possible through the connector architecture.

If there are m application servers and n EISs, the connector architecture reduces the scope of the integration problem from an $m \times n$ problem to an $m + n$ problem.

FIGURE 2-1 System Level Pluggability between Application Servers and EISs



2.2.2 Common Client Interface

An enterprise tools vendor provides tools that lead to a simple application programming model for EIS access, thereby reducing the effort required in EIS integration. An EAI vendor provides a framework that supports integration across multiple EISs. Both types of vendors need to integrate across heterogeneous EISs.

Each EIS typically has a client API that is specific to the EIS. Examples of EIS client APIs are RFC for SAP R/3 and ECI for CICS.

An enterprise tools vendor adapts different client APIs for target EISs to a common client API. The adapted API is typically specific to a tools vendor and supports an application programming model common across all EISs. Adapting the API requires significant effort on the part of a tools vendor. In this case, the $m \times n$ integration problem applies to tools vendors.

The connector architecture provides a solution for the $m \times n$ integration problem for tools and EAI vendors. The architecture specifies a standard Common Client Interface (CCI) that supports a common client API across heterogeneous EISs.

All EIS resource adapters that support CCI are capable of being plugged into enterprise tools and EAI frameworks in a standard way. A tools vendor need not do any API adoption; the vendor can focus on providing its added value of simplifying EIS integration.

The CCI drastically reduces the effort and learning requirements for tools vendor by narrowing the scope of a $m \times n$ problem to a $m + n$ problem if there are m tools and n EISs.

2.3 Goals

The connector architecture has been designed with the following goals:

- Simplify the development of scalable, secure, and transactional resource adapters for a wide range of EISs—ERP systems, database systems, mainframe-based transaction processing systems.
- Be sufficiently general to cover a wide range of heterogeneous EISs. The sufficient generality of the architecture ensures that there are various implementation choices for different resource adapters; each choice is based on the characteristics and mechanisms of an underlying EIS.
- Be not tied to a specific application server implementation, but applicable to all J2EE platform compliant application servers from multiple vendors.
- Provide a standard client API for enterprise tools and EAI vendors. The standard API will be common across heterogeneous EISs.
- Express itself in a manner that allows an organization to unambiguously determine whether or not an implementation is compatible.
- Be simple to understand and easy to follow, regardless of whether one is designing a resource adapter for a particular EIS or developing/deploying application components that need to access multiple EISs. This simplicity means the architecture introduces only a few new concepts, and places minimal implementation requirements so that it can be leveraged across different integration scenarios and environments.
- Define contracts and responsibilities for various roles that provide pieces for standard bi-directional connectivity to an EIS. This enables a standard resource adapter from a EIS vendor to be pluggable across multiple application servers.

- Enable an enterprise application programmer in a non-managed application environment to directly use the resource adapter to access the underlying EIS. This is in addition to managed access to an EIS, with the resource adapter deployed in the middle-tier application server.

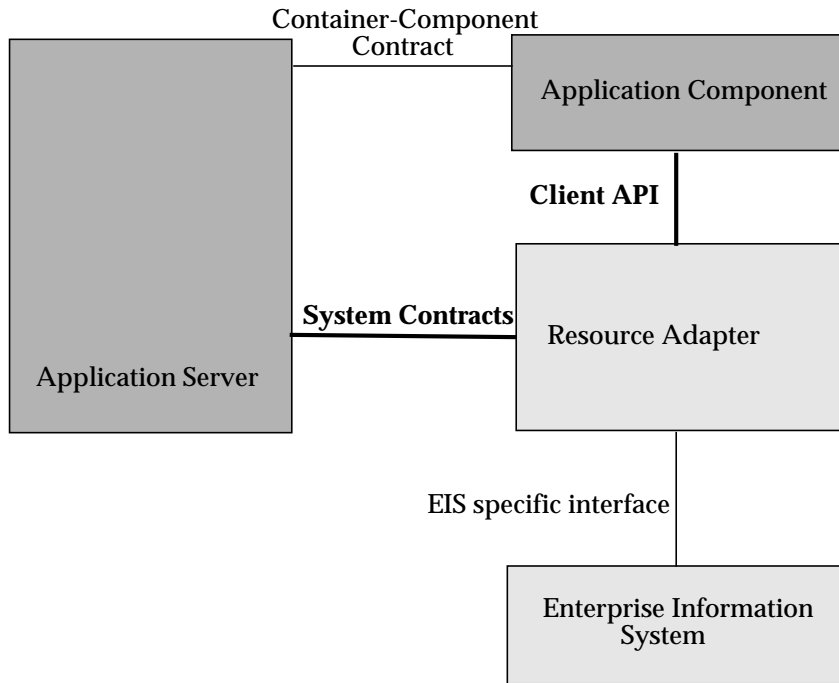
The Connector Architecture

The following chapter gives an overview of the connector architecture.

Multiple resource adapters—that is, one resource adapter per type of EIS—are pluggable into an application server. This capability enables application components deployed on the application server to access the underlying EISs.

An application server and an EIS collaborate to keep all system-level mechanisms—transactions, security, and connection management—transparent from the application components. As a result, an application component provider focuses on the development of business and presentation logic for its application components and need not get involved in the system-level issues related to EIS integration. This leads to an easier and faster cycle for the development of scalable, secure, and transactional enterprise applications that require connectivity with multiple EISs.

FIGURE 3-1 Overview of the Connector Architecture



3.1 System Contracts

To achieve a standard system-level pluggability between application servers and EISs, the connector architecture defines a standard set of system-level contracts between an application server and an EIS. The EIS side of these system-level contracts are implemented in a resource adapter.

A resource adapter is specific to an underlying EIS. It is a system-level software driver that is used by an application server or an application client to connect to an EIS.

A resource adapter plugs into an application server. The resource adapter and application server collaborate to provide the underlying mechanisms—transactions, security, connection pooling, and dispatch to application components.

A resource adapter is used within the address space of the application server. Examples of resource adapters are:

- A JDBC driver to connect to a relational database (as specified in the JDBC (See Related Documents, page 1, 3.) specification)
- A resource adapter to connect to an ERP system
- A resource adapter to connect to a TP system
- A resource adapter to plug-in a messaging system

A resource adapter may provide different types of connectivity between an application and an EIS.

- **Outbound communication:** The resource adapter allows an application to connect to an EIS system and perform work. All communication is initiated by the application. In this case, the resource adapter serves as a passive library for connecting to an EIS, and executes in the context of the application threads.
- **Inbound communication:** The resource adapter allows an EIS to call application components and perform work. All communication is initiated by the EIS. The resource adapter may request threads from the application server or create its own threads.
- **Bi-directional communication:** The resource adapter supports both outbound and inbound communication.

The connector architecture defines the following set of standard contracts between an application server and EIS:

- A connection management contract that enables an application server to pool connections to an underlying EIS, and enables application components to connect to an EIS. This leads to a scalable application environment that can support a large number of clients requiring access to EISs.
- A transaction management contract between the transaction manager and an EIS that supports transactional access to EIS resource managers. This contract enables an application server to use a transaction manager to manage transactions across multiple resource managers. This contract also supports transactions that are managed internal to an EIS resource manager without the necessity of involving an external transaction manager.
- A security contract that enables secure access to an EIS. This contract provides support for a secure application environment that reduces security threats to the EIS and protects valuable information resources managed by the EIS.
- A lifecycle management contract that allows an application server to manage the lifecycle of a resource adapter. This contract provides a mechanism for the application server to bootstrap a resource adapter instance during its deployment or application server startup, and to notify the resource adapter instance during its undeployment or during an orderly shutdown of the application server.
- A work management contract that allows a resource adapter to do work (monitor network endpoints, call application components, etc.) by submitting `Work` instances to an application server for execution. The application server dispatches threads to execute submitted `Work` instances. This allows a resource adapter to

avoid creating or managing threads directly, and allows an application server to efficiently pool threads and have more control over its runtime environment. The resource adapter can control the security context and transaction context with which `Work` instances are executed.

- A transaction inflow contract that allows a resource adapter to propagate an imported transaction to an application server. This contract also allows a resource adapter to transmit transaction completion and crash recovery calls initiated by an EIS, and ensures that the ACID properties of the imported transaction are preserved.
- A message inflow contract that allows a resource adapter to asynchronously deliver messages to message endpoints residing in the application server independent of the specific messaging style, messaging semantics, and messaging infrastructure used to deliver messages. This contract also serves as the standard message provider pluggability contract that allows a wide range of message providers (Java Message Service (JMS), Java API for XML Messaging (JAXM), etc.) to be plugged into any J2EE compatible application server via a resource adapter.

FIGURE 3-1 on page 2 does not illustrate any contracts that are internal to an application server implementation. The specific mechanisms and contracts within an application server are outside the scope of the connector architecture specification. This specification focuses on the system-level contracts between the application server and the EIS.

In FIGURE 3-1, the application server and resource adapter are shown as separate entities. This is done to illustrate that there is a logical separation of the respective roles and responsibilities defined for the support of the system level contracts. However, this separation does not imply a physical separation, as in an application server and a resource adapter running in separate processes.

3.2 Client API

The client API used by application components for EIS access may be defined as:

- The standard Common Client Interface (CCI) as specified in the chapter “Common Client Interface” on page 15-1.
- A client API specific to the type of a resource adapter and its underlying EIS. An example of such an EIS specific client API is JDBC for relational databases.

The Common Client Interface (CCI) defines a common client API for accessing EISs. The CCI is targeted towards Enterprise Application Integration (EAI) and enterprise tools vendors.

3.3 Requirements

The connector architecture requires that the connector architecture-compliant resource adapter and the application server support the system contracts. Detailed requirements for each system contract are specified in later chapters.

The connector architecture recommends, though it does not mandate, that a resource adapter support CCI as the client API. The recommendation enables the connector architecture to provide a solution for the $m \times n$ integration problem for application development tools and EAI vendors.

The connector architecture allows a resource adapter with an EIS-specific client API to support system contracts and to be capable of standard connector architecture-based pluggability into an application server.

3.4 Non-Managed Environment

The connector architecture supports access to EISs from non-managed application clients; for example, Java applications and applets.

In a non-managed two-tier application environment, an application client directly uses a resource adapter library. A resource adapter, in this case, exposes its low-level transactions and security APIs to its clients. An application client has to take responsibility for managing security and transactions (and rely on connection pooling if done by the resource adapter internally) by using the low-level APIs exposed by the resource adapter. This model is similar to the way a two-tier JDBC application client accesses a database system in a non-managed environment.

Roles and Scenarios

This chapter describes a set of roles specific to the connector architecture. The goal of this chapter is to specify contracts that ensure that the output of each role is compatible with the input of the other role. Later chapters specify a detailed set of responsibilities for each role, relative to the system-level contracts.

4.1 Roles

The following section describes the roles and responsibilities specific to the connector architecture.

4.1.1 Resource Adapter Provider

The resource adapter provider is an expert in the technology related to an EIS and is responsible for providing a resource adapter for an EIS. Since this role is highly EIS specific, an EIS vendor typically provides the resource adapter for its system.

A third party vendor (who is not an EIS vendor) may also provide an EIS resource adapter and its associated set of application development tools. Such a provider typically specializes in writing resource adapters and related tools for a large number of EISs.

4.1.2 Application Server Vendor

The application server vendor provides an implementation of a J2EE-compliant application server that provides support for component based enterprise applications. A typical application server vendor is an OS vendor, middleware vendor, or database vendor. The role of an application server vendor is typically the same as that of a container provider.

The J2EE platform specification (See Related Documents, page 1, 8.) specifies requirements for a J2EE platform provider.

4.1.3 Container Provider

The container provider is responsible for providing a container implementation for a specific type of application component. For example, the container provider may provide a container for EJB components. Each type of application component—EJB, servlet, JSP, applet—has its own set of responsibilities for its container provider. The respective specifications outline these responsibilities.

A container implementation typically provides the following functionality:

- It provides deployed application components with transaction and security management, distribution of clients, scalable management of resources, and other services that are generally required as part of a managed server platform.
- It provides application components with connectivity to an EIS by transparently managing security, resources, and transactions using the system-level contracts with the EIS-specific resource adapter.
- It insulates application components from the specifics of the underlying system-level mechanisms by supporting a simple, standard contract with the application component. Refer to the Enterprise JavaBeans specification (See Related Documents, page 1, 1.) for more details on the EJB component contract.

The expertise of the container provider is system-level programming, with its focus on the development of a scalable, secure, and transaction-enabled container.

The container provider is also responsible for providing deployment tools necessary for the deployment of application components and resource adapters. It is also required to provide runtime support for the deployed application components.

The container provider typically provides tools that allow the system administrator to monitor and manage a container and application components during runtime.

4.1.4 Application Component Provider

In the context of the connector architecture, the application component provider produces an application component that accesses one or more EISs to provide its application functionality.

The application component provider is an application domain expert. In the case of application components targeted towards integration with multiple EISs, various business tasks and entities are implemented based on access to EIS data and functions.

The application component provider typically programs against easy-to-use Java abstractions produced by application development tools. These Java abstractions are based on the Common Client interface (CCI).

The application component provider is not required to be an expert at system level programming. The application component provider does not program transactions, security, concurrency, or distribution, but relies on a container to provide these services transparently.

The application component provider is responsible for specifying structural information for an application component and its external dependencies. This information includes, for example, the name and type of the connection factories, and security information.

The output of an application component provider is a Java archive (JAR) file that contains the application components and any additional Java classes required to connect to EISs.

4.1.5 Enterprise Tools Vendors

The application component provider relies on tools to simplify application development and EIS integration. Since programming client access to EIS data and functions is a complex application development task, an application development tool reduces the effort and complexity involved in this task.

Enterprise tools serve different roles in the application development process, as follows:

- Data and function mining tool—enables application component providers to look at the scope and structure of data and functions existing in an EIS.
- Analysis and design tool—enables application component providers to design an application in terms of EIS data and functions.
- Code generation tool—generates Java classes for accessing EIS data and functions. A mapping tool that bridges across two different programming models (object to relational or vice-versa) falls into this category of tools.

- Application composition tool—enables application component providers to compose application components from Java classes generated by a code generation tool. This type of tool typically uses the JavaBeans component model to enhance the ease of programming and composition.
- Deployment tool—used by application component providers and deployers to set transaction, security, and other deployment time requirements.

A number of these tools may be integrated together to form an end-to-end application development environment.

In addition, various tools and middleware vendors offer EAI frameworks that simplify integration across heterogeneous EISs.

4.1.6 Application Assembler

The application assembler combines various application components into a larger set of deployable units. The input of the application assembler is one or more JAR files produced by an application component provider and the output is one or more JAR files with a deployment descriptor.

The application assembler is typically a domain expert who assembles application components to produce an enterprise application. To achieve this goal, the application assembler takes application components, possibly from multiple application component providers, and assembles these components.

4.1.7 Deployer

The deployer takes one or more deployable units of application components, produced by the application assembler or component provider, and deploys the application components in a target operational environment. An operational environment is comprised of an application server and multiple connected EISs.

The deployer is responsible for resolving all external dependencies declared by the application component provider. For example, the deployer ensures that all connection factories used by the application components are present in an operational environment. To perform its role, the deployer typically uses the application server-provided deployment tools.

The deployer is also responsible for the deployment of resource adapters. Since an operational environment may include multiple EISs, the role of the deployer is more intensive and complex than that in a non-EIS scenario. The deployer has to understand security, transaction, and connection management-related aspects of multiple EISs that are configured in an operational environment.

4.1.8 System Administrator

The system administrator is responsible for the configuration and administration of a complete enterprise infrastructure that includes multiple containers and EISs.

In an operational environment that has multiple EISs, the deployer should manage the operational environment by working closely with the system administrators of respective EISs. This enables the deployer to resolve deployment issues while deploying application components and resource adapters in a target operational environment.

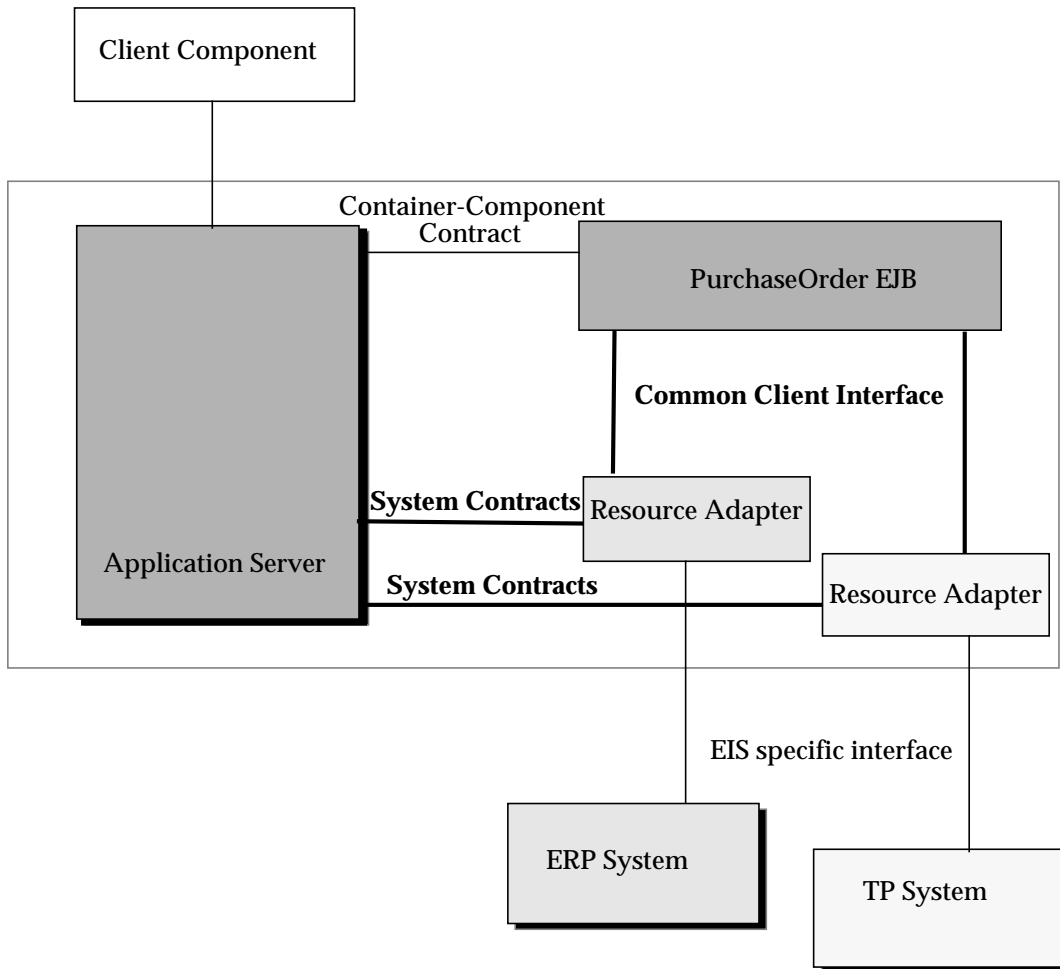
This chapter introduced the roles involved in the connector architecture. The later chapters specify responsibilities for each role in more detail.

4.2 Scenario: Integrated Purchase Order System

This section describes a scenario that illustrates the use of the connector architecture. The following description is kept at a high level. Specific scenarios related to transaction management, security, connection management, and inbound communications are described in subsequent chapters.

The following diagram shows the different pieces that comprise this scenario:

FIGURE 4-1 Illustration of a Scenario Based On the Connector Architecture



ERP Software Inc. is an enterprise system vendor that provides an enterprise resource planning (ERP) system. ERP Software wants to integrate its ERP system with various application servers. It achieves this goal by providing a standard resource adapter for its ERP system. The resource adapter for ERP systems supports the standard inbound communication, transaction, connection management and security contracts. The resource adapter also supports the Common Client Interface (CCI) as its client API.

TPSoft Inc. is another enterprise system vendor that provides a transaction processing (TP) system. TPSOft has also developed a standard resource adapter for its TP system. The resource adapter library supports CCI as part of its implementation.

AppServer Inc. is a system vendor that has an application server product which supports the development and deployment of component-based enterprise applications. This application server product has an EJB container that provides deployment and runtime support for EJB components. The application server supports the system-level contracts that enable a resource adapter, which also supports these contracts, to plug into the application server and provide bi-directional connectivity to the underlying EIS. The EJB container insulates EJB components from the communication, transaction, security, and connection management mechanisms required for connecting to the EIS.

Manufacturer Corp. is a big manufacturing firm that uses a purchase order processing system based on the ERP system for its business processes. Recently, Manufacturer has acquired a firm that uses TPSOft's TP system for its purchase order processing. Manufacturer aims to integrate these two systems together into a single integrated purchase order system. It wants a scalable, multi-user, secure, transaction-enabled integrated purchase order system that is not tied to a specific computing platform. Manufacturer plans to deploy the middle-tier of this system on the application server from AppServer Inc.

The MIS department of Manufacturer develops a PurchaseOrder EJB that provides an integrated view of the two underlying purchase order systems. While developing PurchaseOrder EJB, the bean provider does not program the transactions, security, connection management or inbound communication mechanisms required for connectivity to the ERP and TP systems; it relies on the EJB container and application server to provide these services.

The bean provider uses an application programming model based on the CCI to access the business objects and function modules for purchase order processing in the ERP system. The bean provider uses a similar application programming model based on the CCI to access the purchase order processing programs in the TP system.

The MIS department of Manufacturer assembles an integrated web-based purchase order application using PurchaseOrder EJB with other types of application components, such as JSPs and servlets.

The MIS department installs and configures the application server, ERP, and TP system as part of its operational environment. It then deploys the integrated purchase order application on this operational environment. As part of the deployment, the MIS department configures the operational environment based on the deployment requirements for the various application components that have been assembled into the integrated enterprise application.

After deploying and successfully testing the integrated purchase order system, the MIS department makes the system available for other departments to use.

4.3 Scenario: Business Integration

This scenario illustrates the use of the connector architecture in a business integration scenario.

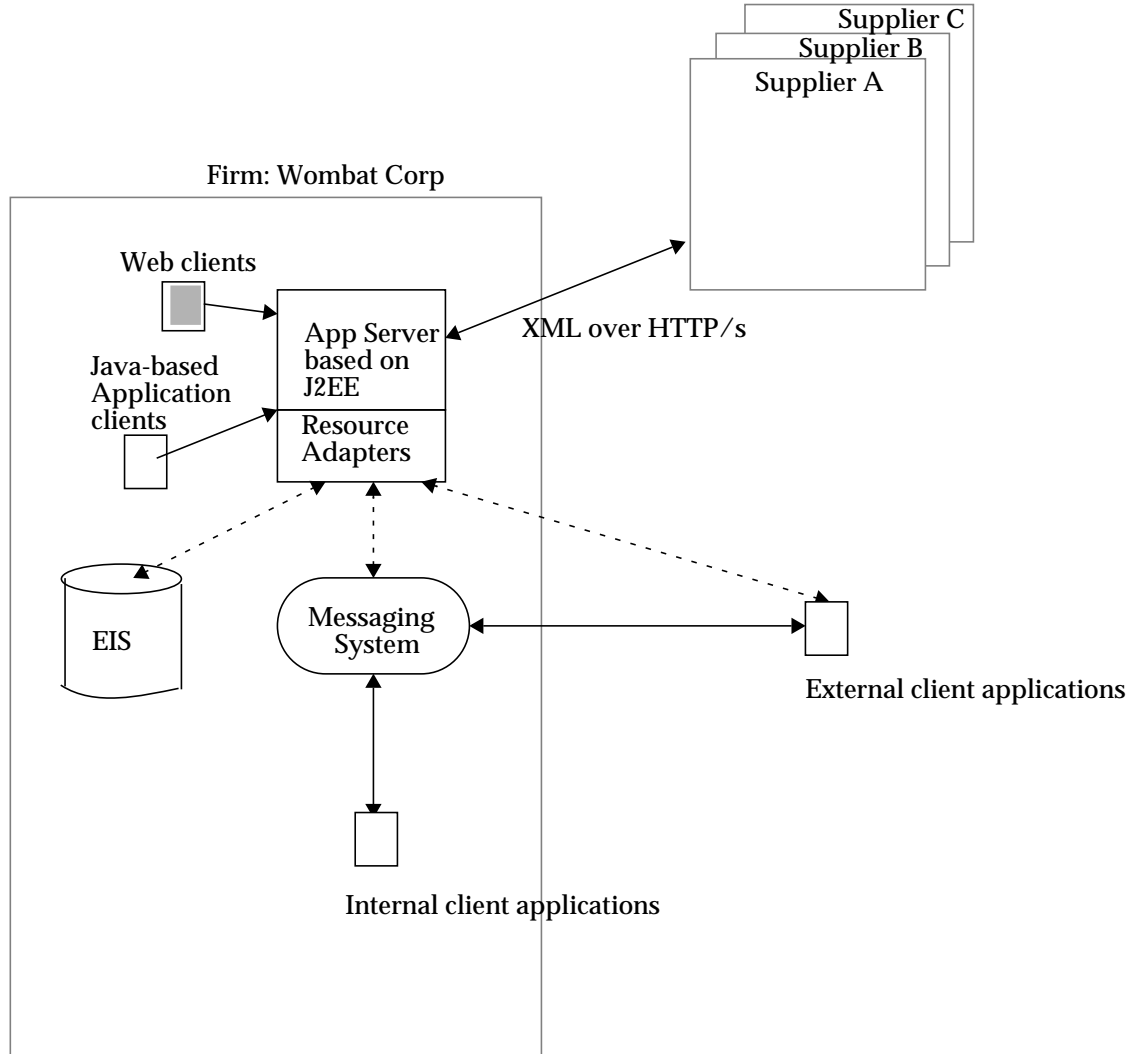
Wombat Corp. is a manufacturing firm that aims to adopt an e-business strategy. Wombat has huge existing investments in its EIS systems. The EISs include ERP systems, mainframe transaction processing systems, and message providers.

Wombat needs to interact with its various partners. In order to do this, it needs support for different interaction mechanisms. It also needs a mechanism to involve all its EIS systems in the interaction. Further, it needs an application server to host its business applications which participate in the various interactions.

Wombat buys a J2EE based application server from EComm, Inc. to host its business applications which interact with its EISs and its various partners. The application server supports the connector architecture contracts which make it possible to use appropriate resource adapters to drive interactions with its partners and its EISs.

The connector architecture enables Wombat to integrate its existing infrastructure with the application server. Wombat buys off-the-shelf resource adapters for its existing set of EISs and to support interactions with its partners and uses them to integrate its business applications (deployed on the application server).

FIGURE 4-2 Connector Architecture Usage In Business Integration Scenario



Lifecycle Management

This chapter specifies a contract between an application server and a resource adapter that allows an application server to manage the lifecycle of a resource adapter. This contract provides a mechanism for the application server to bootstrap a resource adapter instance during its deployment or application server startup, and to notify the resource adapter instance during its undeployment or during an orderly shutdown of the application server.

5.1 Overview

A resource adapter is a system component which is deployed in an application server. When a resource adapter is deployed, or during application server startup, an application server needs to bootstrap an instance of the resource adapter in its address space. When a resource adapter is undeployed, or during application server shutdown, the application server needs a mechanism to notify the resource adapter instance to stop functioning so that it can be safely unloaded.

The lifecycle management contract provides such a mechanism for an application server to manage the lifecycle of a resource adapter instance. This allows an application server to bootstrap a resource adapter instance during resource adapter deployment or application server startup and also to expose some of its useful facilities to the resource adapter instance. It also provides a mechanism to notify the resource adapter instance while it is undeployed or during an orderly shutdown of the application server.

5.2 Goals

- Provide a mechanism for an application server to manage the lifecycle of a resource adapter instance.

5.3 Lifecycle Management Model

FIGURE 5-1 Lifecycle Management Contract (interfaces)

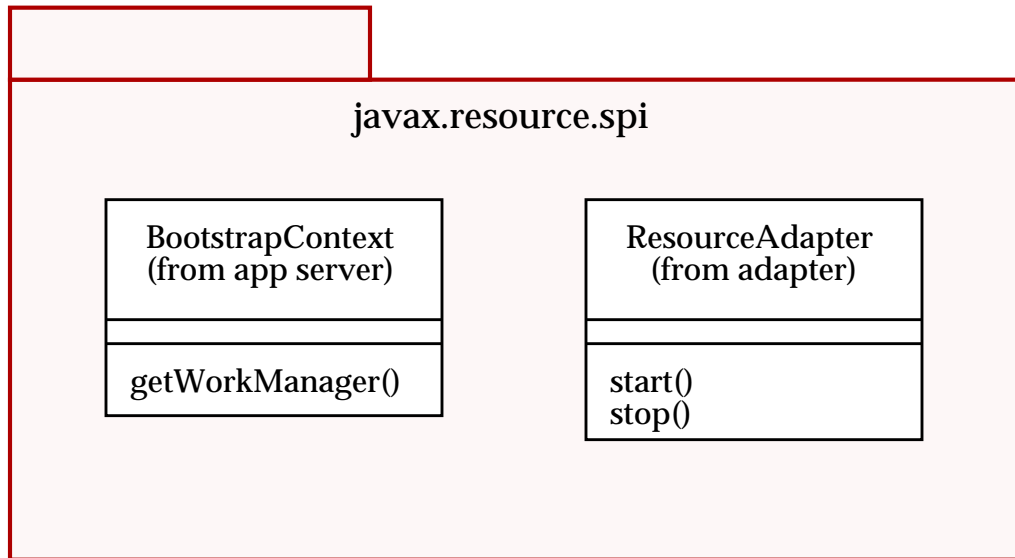
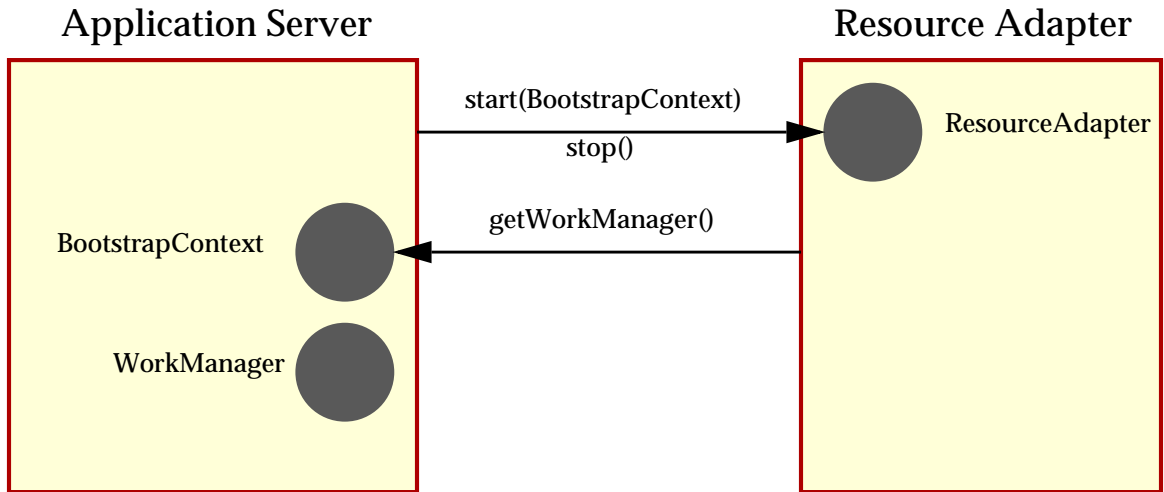


FIGURE 5-2 Lifecycle Management (object diagram)



```
package javax.resource.spi;

import javax.resource.spi.work.WorkManager;

public interface ResourceAdapter {
    void start(BootstrapContext) // startup notification
        throws ResourceAdapterInternalException;
    void stop(); // shutdown notification
    ... // other operations
}

public interface BootstrapContext {
    WorkManager getWorkManager();
    ... // other operations
}
```

An application server implements the `BootstrapContext` and `WorkManager` interfaces. A resource adapter implements the `ResourceAdapter` interface.

5.3.1 ResourceAdapter JavaBean and Bootstrapping a Resource Adapter Instance

The implementation class name of the `ResourceAdapter` interface is specified in the resource adapter deployment descriptor. The `ResourceAdapter` class must be a JavaBean. Refer to Section 16.3, “JavaBean Requirements” on page 16-2. During resource adapter deployment, the resource adapter deployer creates a `ResourceAdapter` JavaBean and configures it with the appropriate properties.

When a resource adapter is deployed, or during application server startup, an application server bootstraps an instance of the resource adapter in its address space. In order to bootstrap a resource adapter instance, the application server must use the configured `ResourceAdapter` JavaBean and call its `start` method. The `start` method call is a startup notification from the application server, and this method is called by an application server thread.

During the `start` method call the `ResourceAdapter` JavaBean is responsible for initializing the resource adapter instance. This may involve creating resource adapter instance specific objects, creating threads (refer to Chapter 10, “Work Management”), and setting up network endpoints. A `ResourceAdapter` JavaBean represents exactly one functional resource adapter unit or instance. The application server must instantiate exactly one `ResourceAdapter` JavaBean per functional resource adapter instance. The application server must create at least one functional resource adapter instance per resource adapter deployment¹.

The application server is allowed to have multiple instances of a `ResourceAdapter` JavaBean active simultaneously, in the same JVM, provided the instances are not equal. Their equality is determined using the `equals` method, and therefore, the `ResourceAdapter` JavaBean is required to implement the `equals` method.

During the `start` method call, an application server must provide a `BootstrapContext` instance containing references to some of the application server facilities (for example, `WorkManager`) for use by the resource adapter instance. The application server facilities exposed via the `BootstrapContext` instance may be used by the resource adapter instance during its lifetime.

During the `start` method call, the resource adapter instance initializes itself, and may use the `WorkManager` to submit `Work` instances for execution (see Chapter 10, “Work Management”). The `start` method call should return in a timely manner, and should avoid blocking calls, such as use of `doWork` method call on the `WorkManager` instance. The application server may throw a `WorkRejectedException` in response to any or all `doWork` method calls on the

1. An application server may create more than one functional resource adapter instance per resource adapter deployment, in order to create replicas of a single functional resource adapter instance on multiple Virtual Machines (VMs). In general, however, there should be just one functional resource adapter instance per deployment.

`WorkManager` instance, in order to enforce that a `start` method call does not block. Resource adapter implementations are strongly recommended to use `startWork` and `scheduleWork` methods on the `WorkManager`, instead of the `doWork` method.

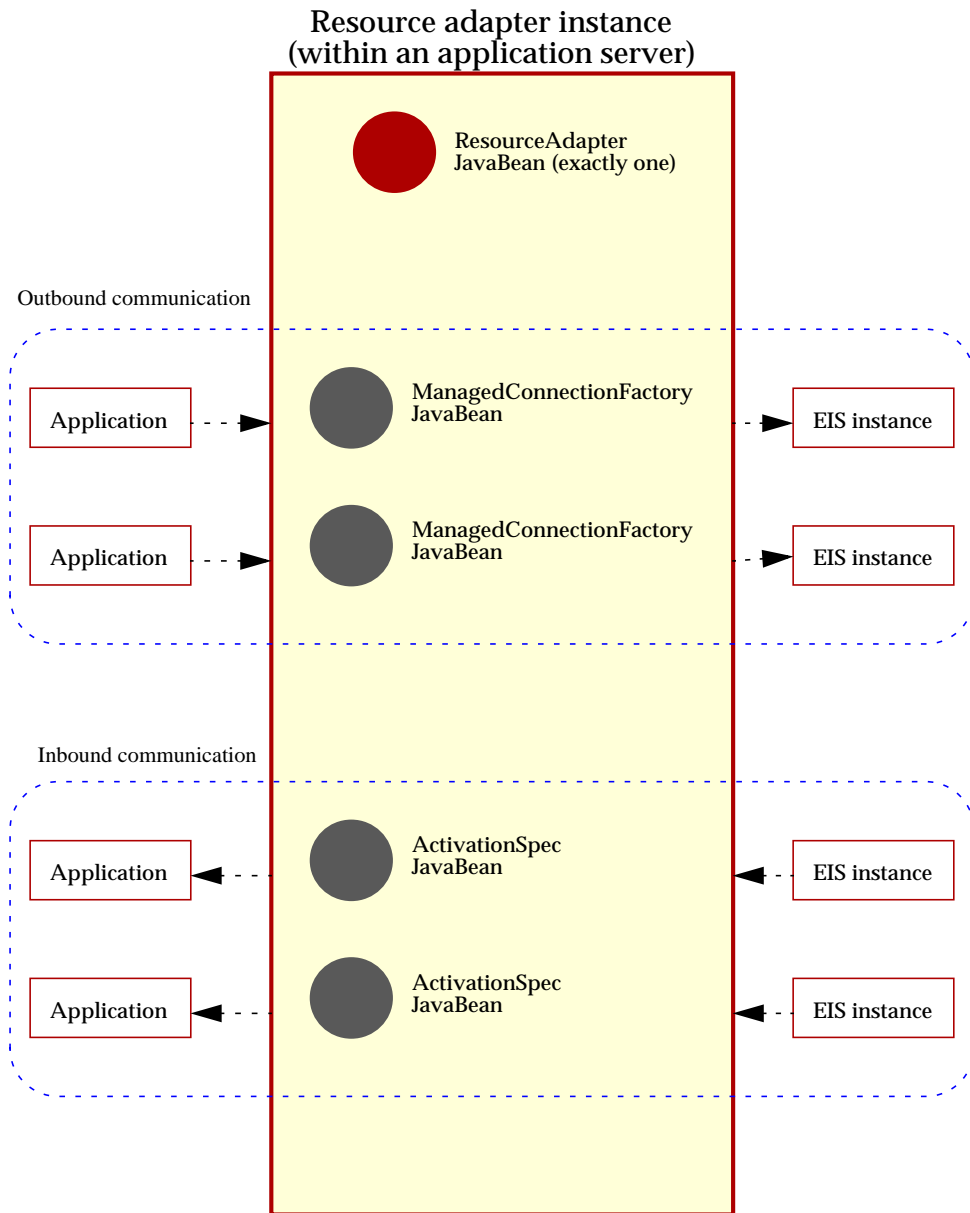
Any exception thrown during the `start` method call indicates an error condition, and the attempt by the application server to create a resource adapter instance fails. A future version of the specification may add a two-phase startup procedure.

A resource adapter instance at runtime may contain several objects that may be created and discarded during its lifetime. Such objects include `ManagedConnectionFactory` `JavaBean` (refer to Chapter 6, “Connection Management”), `ActivationSpec` `JavaBean` (refer to Chapter 12, “Message Inflow”), various connection objects, resource adapter private objects, and other resource adapter specific objects that are exposed to applications.

The `ResourceAdapter` `JavaBean` represents a resource adapter instance and contains the configuration information pertaining to that resource adapter instance. This configuration information may also be used as global defaults for `ManagedConnectionFactory` and `ActivationSpec` `JavaBeans`. That is, when `ManagedConnectionFactory` or `ActivationSpec` `JavaBeans` are created they may inherit the global defaults (`ResourceAdapter` `JavaBean` configuration information), which make it easier to configure them.

A resource adapter instance may provide bi-directional connectivity to multiple EIS instances. A `ManagedConnectionFactory` `JavaBean` can be used to provide outbound connectivity to a single EIS instance. An `ActivationSpec` `JavaBean` can be used to provide inbound connectivity from an EIS instance. A resource adapter instance may contain several such `ManagedConnectionFactory` and `ActivationSpec` `JavaBeans`. FIGURE 5-3 on page 7 describes the association between a resource adapter instance and its various `ManagedConnectionFactory` and `ActivationSpec` `JavaBeans`.

FIGURE 5-3 Resource Adapter Instance (composition)



5.3.2 ManagedConnectionFactory JavaBean and Outbound Communication

A `ManagedConnectionFactory` JavaBean represents outbound connectivity information to an EIS instance from an application via a specific resource adapter instance. This contains the configuration information pertaining to outbound connectivity to an EIS instance. Refer to Chapter 6, “Connection Management” for more details on the `ManagedConnectionFactory` JavaBean.

When a `ManagedConnectionFactory` JavaBean is created, it may inherit the `ResourceAdapter` JavaBean (which represents the resource adapter instance) configuration information, and overrides specific global defaults, if any, and may add other configuration information specific to outbound connectivity.

That is, in the case of outbound communication, the outbound connectivity configuration is a union of `ResourceAdapter` JavaBean and `ManagedConnectionFactory` JavaBean configuration, with the intersecting configuration properties based on the `ManagedConnectionFactory` JavaBean settings.

Outbound communication is initiated by an application and the communication occurs in the context of an application thread, even though resource adapter threads may be involved in the interaction. Note, a resource adapter may use the work management contract (refer to Chapter 10, “Work Management”) to request threads to do work.

```
import javax.resource.spi.ResourceAdapterAssociation;
import javax.resource.spi.ManagedConnectionFactory;

public class ManagedConnectionFactoryImpl
    implements ManagedConnectionFactory,
    ResourceAdapterAssociation {

    ResourceAdapter getResourceAdapter();
    void setResourceAdapter(ResourceAdapter) throws
    ResourceException;
    ... // other methods
}
```

The `ResourceAdapterAssociation` interface specifies the methods to associate a `ManagedConnectionFactory` JavaBean with a `ResourceAdapter` JavaBean.

Prior to using a `ManagedConnectionFactory` JavaBean, the application server must create an association between the `ManagedConnectionFactory` JavaBean and a `ResourceAdapter` JavaBean, by calling the `setResourceAdapter` method on

the `ManagedConnectionFactory` `JavaBean`. A successful association is established only when the `setResourceAdapter` method on the `ManagedConnectionFactory` `JavaBean` returns without throwing an exception.

The `setResourceAdapter` method on the `ManagedConnectionFactory` `JavaBean` must be called exactly once; that is, the association must not change during the lifetime of a `ManagedConnectionFactory` `JavaBean`.

5.3.3 ActivationSpec JavaBean and Inbound Communication

An `ActivationSpec` `JavaBean` represents inbound connectivity information from an EIS instance to an application via a specific resource adapter instance. This contains the configuration information pertaining to inbound connectivity from an EIS instance. Refer to Chapter 12, “Message Inflow” for more details on the `ActivationSpec` `JavaBean`.

When an `ActivationSpec` `JavaBean` is created, it may inherit the `ResourceAdapter` `JavaBean` (which represents the resource adapter instance) configuration information, and overrides specific global defaults, if any, and may add other configuration information specific to inbound connectivity.

That is, in the case of inbound communication, the inbound connectivity configuration is a union of `ResourceAdapter` `JavaBean` and `ActivationSpec` `JavaBean` configuration, with the intersecting configuration properties based on the `ActivationSpec` `JavaBean` settings.

Inbound communication is initiated by an EIS instance and the communication occurs in the context of a resource adapter thread. There are no application threads involved. Note, a resource adapter may use the work management contract (refer to Chapter 10, “Work Management”) to request threads to do work.

```
import javax.resource.spi.ActivationSpec;

// ActivationSpec interface extends ResourceAdapterAssociation
interface.

public class ActivationSpecImpl implements ActivationSpec {
    ResourceAdapter getResourceAdapter();
    void setResourceAdapter(ResourceAdapter) throws
ResourceException;
    ... // other methods
}
```

The `ResourceAdapterAssociation` interface specifies the methods to associate an `ActivationSpec` `JavaBean` with a `ResourceAdapter` `JavaBean`.

Prior to using an `ActivationSpec` `JavaBean`, the application server must create an association between the `ActivationSpec` `JavaBean` and a `ResourceAdapter` `JavaBean`, by calling the `setResourceAdapter` method on the `ActivationSpec` `JavaBean`. A successful association is established only when the `setResourceAdapter` method on the `ActivationSpec` `JavaBean` returns without throwing an exception.

The `setResourceAdapter` method on the `ActivationSpec` `JavaBean` must be called exactly once; that is, the association must not change during the lifetime of an `ActivationSpec` `JavaBean`.

5.3.4 Resource Adapter Shutdown Procedure

The following are some likely situations during which an application server would shutdown a resource adapter instance:

- The application server is being shutdown.
- The resource adapter is being undeployed.

Irrespective of what causes a resource adapter instance to be shutdown, the application server must use the following two phases to shutdown a resource adapter instance.

5.3.4.1 Phase One

Before calling the `stop` method on the `ResourceAdapter` `JavaBean`, the application server must ensure that all dependant applications using the specific resource adapter instance are stopped. This includes deactivating all message endpoints receiving messages via the specific resource adapter. Note, however, since dependant applications typically cannot be stopped until they are undeployed, the application server may have to delay stopping the resource adapter instance, until all such dependant applications are undeployed.

Completion of phase one guarantees that application threads will not use the resource adapter instance, even though the resource adapter instance specific objects may still be in the memory heap. This ensures that all application activities including transactional activities are completed.

Thus, phase one ensures that even if a resource adapter instance does not properly shutdown during phase two, the resource adapter instance is practically unusable.

5.3.4.2 Phase Two

- The application server calls the `stop` method on the `ResourceAdapter` `JavaBean` to notify the resource adapter instance to stop functioning so that it can be safely unloaded. This is a graceful shutdown notification from the application server, and this method is called by an application server thread.

The `ResourceAdapter` `JavaBean` is responsible for performing an orderly shutdown of the resource adapter instance during the `stop` method call. This may involve closing network endpoints, relinquishing threads, releasing all active `Work` instances, allowing resource adapter internal in-flight transactions to complete if they are already in the process of doing a commit, and flushing any cached data to the EIS.

The resource adapter instance is considered fully functional until the application server calls the `stop` method on the `ResourceAdapter` `JavaBean`.

Any unchecked exception thrown by the `stop` method call does not alter the processing of the application server shutdown or resource adapter undeployment that caused the `stop` method call. The application server may log the exception information for error reporting purposes.

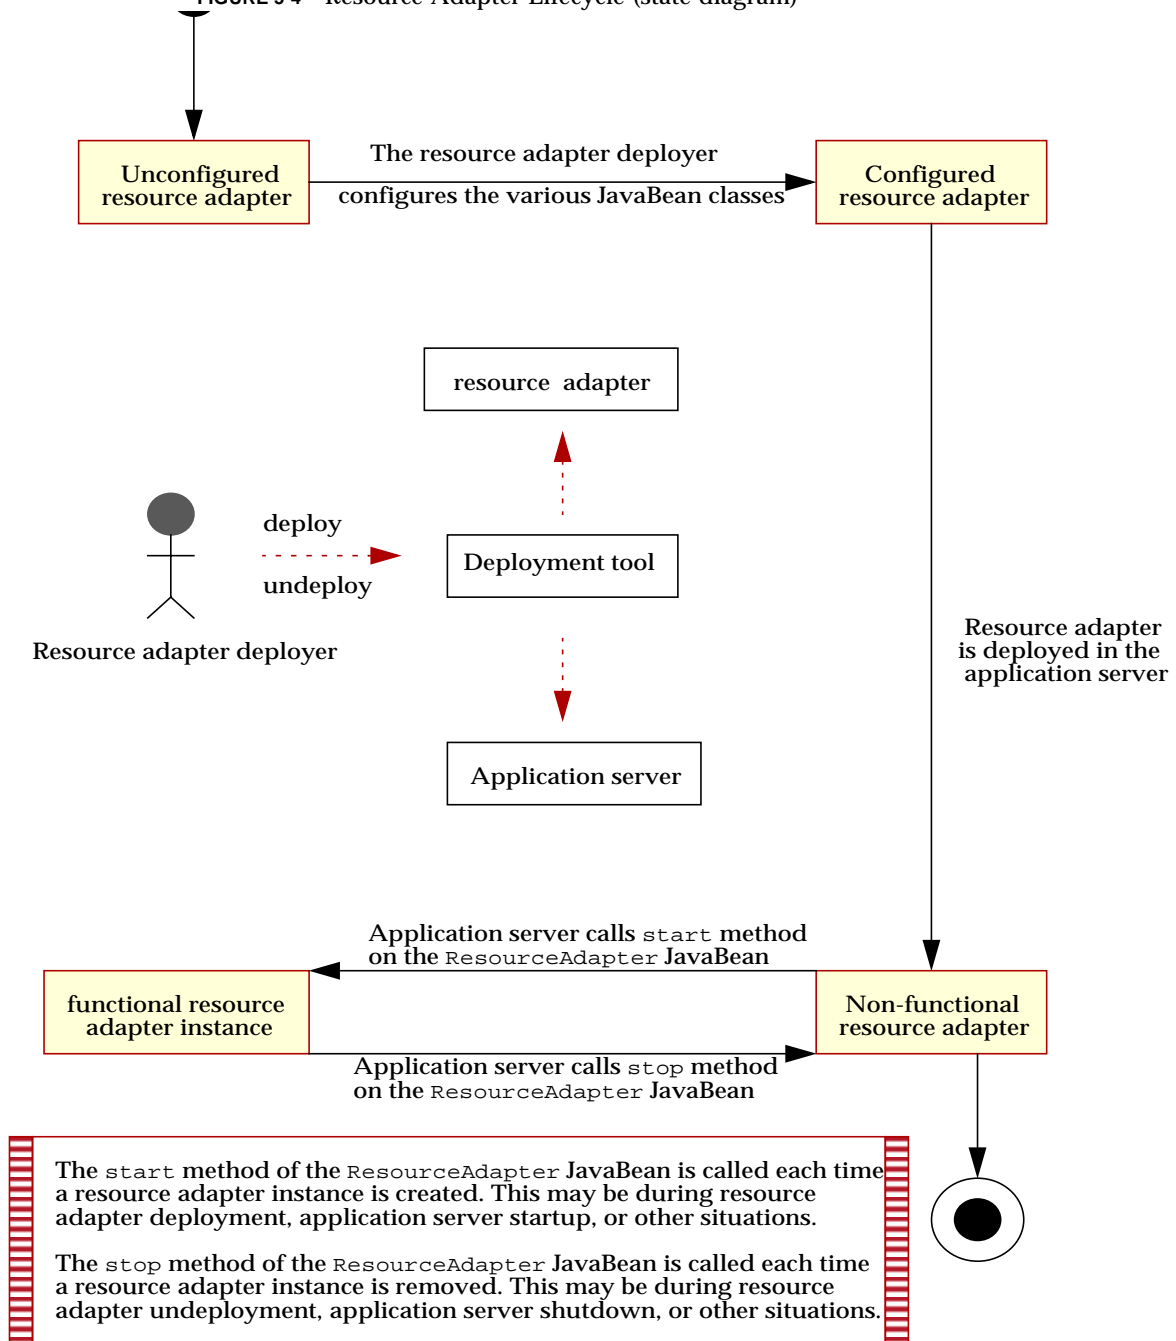
Note, it is possible for a resource adapter instance to become non-functional during its lifetime even before the `stop` method is called, due to EIS failure or other reasons. In such cases, the resource adapter instance should throw exceptions to indicate the failure condition, when it is accessed by an application (during outbound communication) or the application server.

A future version of the specification may add a forced shutdown method in addition to the current graceful `stop` method.

5.3.5 Requirements

- The application server must use a new `ResourceAdapter` `JavaBean` for managing the lifecycle of each resource adapter instance and must discard the `ResourceAdapter` `JavaBean` after its `stop` method has been called. That is, the application server must not reuse the same `ResourceAdapter` `JavaBean` object to manage multiple instances of a resource adapter, since the `ResourceAdapter` `JavaBean` object may contain resource adapter instance specific state information.
- The application server must call the `start` method on the `ResourceAdapter` `JavaBean` (in order to create a functional resource adapter instance), before accessing other methods on the `ResourceAdapter` `JavaBean` instance or before using other objects that belong to the same resource adapter instance.
- The application server thread which calls the `start` and the `stop` method on the `ResourceAdapter` `JavaBean` executes in an unspecified context. However, the application server thread must have at least the same level of security permissions as that of the resource adapter instance.

FIGURE 5-4 Resource Adapter Lifecycle (state diagram)



5.3.6 Resource Adapter Implementation Guidelines

The `ResourceAdapter` `JavaBean` should be treated as a central authority or registry for resource adapter instance specific information, and it should have access to the overall state of the resource adapter instance (network endpoints, etc.). This helps in the manageability of the resource adapter instance, and in performing an orderly shutdown.

Some conventions to follow:

- Any resource adapter specific object (for example, `ManagedConnectionFactory` `JavaBean`, `ActivationSpec` `JavaBean`, or others) which creates network endpoints should register them with the `ResourceAdapter` `JavaBean`.
- The resource adapter threads should periodically scan the `ResourceAdapter` `JavaBean` state and behave accordingly. It is desirable that such threads avoid boundless blocking on I/O calls, and instead use a bounded blocking duration. This helps in resource adapter shutdown, and also potentially avoids deadlock situations during shutdown.

The above conventions enable a `ResourceAdapter` `JavaBean` to effectively manage the resource adapter instance and to perform an orderly shutdown of the resource adapter instance.

5.3.7 JavaBean Configuration and Deployment

There is at most one `ResourceAdapter` `JavaBean` instance per resource adapter instance. But there can be many `ManagedConnectionFactory`, `ActivationSpec` or administered object instances (Section 12.4.2.3, “Administered Objects” on page 12-12) per resource adapter instance.

The `ResourceAdapter` `JavaBean` instance is created and configured during resource adapter deployment. The `ManagedConnectionFactory`, `ActivationSpec` and administered object instances are created and configured during the lifetime of a resource adapter instance.

At runtime, the resource adapter internally uses a union of the configured `ResourceAdapter` and `ManagedConnectionFactory` `JavaBean` properties, to represent outbound communication configuration.

Similarly, at runtime, the resource adapter internally uses a union of the configured `ResourceAdapter` and `ActivationSpec` `JavaBean` properties, to represent inbound communication configuration.

5.3.7.1 ResourceAdapter JavaBean instance Configuration

- Create a `ResourceAdapter` JavaBean instance. This will initialize the instance with the defaults specified via the JavaBean mechanism.
- Apply the `ResourceAdapter` class configuration properties specified in the resource adapter deployment descriptor, on the `ResourceAdapter` instance. This may override some of the default values specified via the JavaBean mechanism.
- The `ResourceAdapter` deployer may further override the values of the `ResourceAdapter` instance before deployment.

5.3.7.2 Resource Adapter Deployment

The `ResourceAdapter` instance property values may be stored separately and reused later while configuring `ManagedConnectionFactory`, `ActivationSpec`, or administered object instances.

5.3.7.3 ManagedConnectionFactory JavaBean Instance Configuration

- Create a `ManagedConnectionFactory` JavaBean instance. This will initialize the instance with the defaults specified via the JavaBean mechanism.
- Apply the `ResourceAdapter` instance property values, that were stored earlier, on the `ManagedConnectionFactory` instance. Note, that the `ManagedConnectionFactory` JavaBean may have none, some or all of the properties of the `ResourceAdapter` JavaBean.
- Apply the `ManagedConnectionFactory` class configuration properties specified in the resource adapter deployment descriptor, on the `ManagedConnectionFactory` instance.
- The `ManagedConnectionFactory` deployer may further override the values of the `ManagedConnectionFactory` instance before deployment.

At runtime, the resource adapter internally uses a union of the configured `ResourceAdapter` and `ManagedConnectionFactory` JavaBean properties, to represent outbound communication configuration. Note, the `ManagedConnectionFactory` instance and the `ResourceAdapter` instance may have intersecting property names. In such a situation, the values specified in the `ManagedConnectionFactory` instance takes precedence.

5.3.7.4 ActivationSpec JavaBean instance Configuration

- Create an `ActivationSpec` JavaBean instance. This will initialize the instance with the defaults specified via the JavaBean mechanism.
- Apply the `ResourceAdapter` instance property values, that were stored earlier, on the `ActivationSpec` instance. Note, that the `ActivationSpec` JavaBean may have none, some, or all of the properties of the `ResourceAdapter` JavaBean.

- Apply the `ActivationSpec` class configuration properties specified in the application deployment descriptor, on the `ActivationSpec` instance.
- The `ActivationSpec` deployer may further override the values of the `ActivationSpec` instance before deployment.

At runtime, the resource adapter internally uses a union of the configured `ResourceAdapter` and `ActivationSpec` `JavaBean` properties, to represent inbound communication configuration. Note, the `ActivationSpec` instance and the `ResourceAdapter` instance may have intersecting property names. In such a situation, the values specified in the `ActivationSpec` instance takes precedence.

5.3.7.5 Resource Adapter Implementation Guidelines

A resource adapter implementation may choose to use common properties, that is, a `ManagedConnectionFactory` or an `ActivationSpec` `JavaBean`, may contain some or all of the properties of the `ResourceAdapter` `JavaBean`. The choice is up to the resource adapter implementation.

In general, there is no need for common properties, since these various objects are associated at runtime with the `ResourceAdapter` `JavaBean`. However, there may be situations, for example, a `ManagedConnectionFactory` `JavaBean` may need to override the `ResourceAdapter` `JavaBean` values in order to successfully connect to a different EIS. In such a scenario, providing common properties between the `ResourceAdapter` and `ManagedConnectionFactory` `JavaBeans`, allows the `ManagedConnectionFactory` deployer to override the `ResourceAdapter` property values and configure the `ManagedConnectionFactory` appropriately.

5.3.8 Lifecycle Management In A Non-Managed Environment

Although the lifecycle management contract is primarily intended for a managed environment, it may still be used in a non-managed environment provided that the application that bootstraps a resource adapter instance is capable of managing its lifecycle.

5.3.9 A Sample Resource Adapter Implementation

CODE EXAMPLE 5-1 Sample Resource Adapter

```
package com.xyz.adapter;

import javax.resource.spi.ResourceAdapter;
import javax.resource.spi.BootstrapContext;
import javax.resource.spi.work.*;

public class MyResourceAdapterImpl implements ResourceAdapter {

    void start(BootstrapContext serverCtx) {
        // 1. setup network endpoints
        ...

        // 2. get WorkManager reference
        WorkManager wm = serverCtx.getWorkManager();

        // 3. provide Work objects to WorkManager
        for (i = 0; i < 10; i++) {
            Work work = new MyWork(...);
            try {
                wm.startWork(work);
            } catch (WorkException we) { // handle the exception }
        }
    }

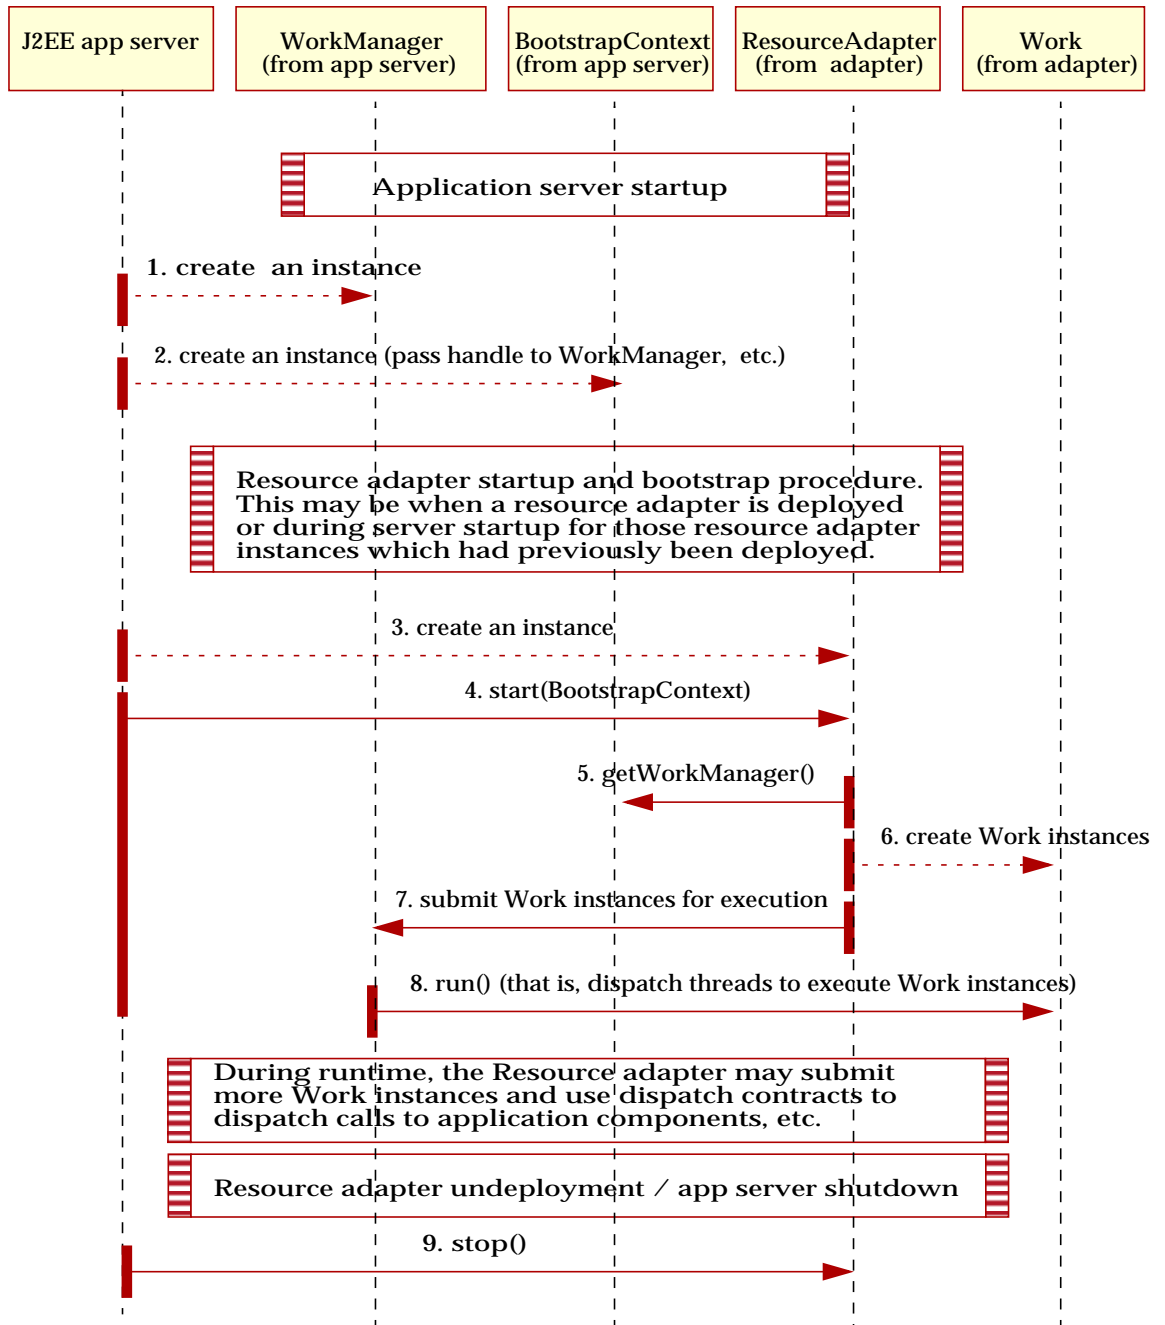
    void stop() { // release Work instances, do cleanup and return.}
}

public class MyWork implements Work {

    void release() {
        // set a flag to hint the Work instance to complete.
        // Note, the calling thread is different from
        // the active thread in which this instance is executing.
    }

    void run() {
        // do work (call application components, monitor
        // network ports, etc.).
    }
}
```


FIGURE 5-5 Lifecycle Management Model (sequence diagram)



Connection Management

This chapter specifies the connection management contract between an application server and a resource adapter. It introduces the concepts and mechanisms relevant to this contract, and delineates the responsibilities of the roles of the resource adapter provider and application server vendor in terms of their system-level support for the connection management contract. To complete the description of the connection management contract, this chapter also refers to the responsibilities of the application component provider and deployer. The chapter includes scenarios to illustrate the connection management contract.

6.1 Overview

An application component uses a connection factory to access a connection instance, which the component then uses to connect to the underlying EIS. A resource adapter acts as a factory of connections. Examples of connections include database connections, JMS (Java Message Service) connections, and SAP R/3 connections.

Connection pooling manages connections that are expensive to create and destroy. Connection pooling of expensive connections leads to better scalability and performance in an operational environment. The connection management contract provides support for connection pooling.

6.2 Goals

The connection management contract has been designed with the following goals:

- To provide a consistent application programming model for connection acquisition for both managed and non-managed (two-tier) applications.

- To enable a resource adapter to provide a connection factory and connection interfaces based on the CCI specific to the type of resource adapter and EIS. This enables JDBC drivers to be aligned with the connector architecture with minimum impact on the existing JDBC APIs.
- To provide a generic mechanism by which an application server can provide different services—transactions, security, advanced pooling, error tracing/logging—for its configured set of resource adapters.
- To provide support for connection pooling.

The goal of the connector architecture is to enable efficient, scalable, and extensible connection pooling mechanisms, not to specify a mechanism or implementation for connection pooling. The goal is accomplished by defining a standard contract for connection management with the providers of connections—that is, resource adapters. An application server should use the connection management contract to implement a connection pooling mechanism in its own implementation-specific way.

6.3 Architecture: Connection Management

The connection management contract specifies an architected contract between an application server and a resource adapter. This connection management contract is shown with bold flow lines in FIGURE 6-1 on page 5. It includes the set of interfaces shown in the architecture diagram.

6.3.1 Overview: Managed Application Scenario

The application server uses the deployment descriptor mechanism (specified in section Section 17.5.4, “Requirements” on page 17-23) to configure the resource adapter in the operational environment.

The resource adapter provides connection and connection factory interfaces. A connection factory acts as a factory for EIS connections. For example, `javax.sql.DataSource` and `java.sql.Connection` interfaces are JDBC-based interfaces for connecting to a relational database.

The CCI (specified in Chapter 15, “Common Client Interface”) defines `javax.resource.cci.ConnectionFactory` and `javax.resource.cci.Connection` as interfaces for a connection factory and a connection, respectively.

The application component does a lookup of a connection factory in the Java Naming and Directory Interface (JNDI) name space. It uses the connection factory to get a connection to the underlying EIS. The connection factory instance delegates the connection creation request to the `ConnectionManager` instance.

The `ConnectionManager` enables the application server to provide different quality of services in the managed application scenario. These quality of services include transaction management, security, error logging and tracing, and connection pool management. The application server provides these services in its own implementation-specific way. The connector architecture does not specify how the application server implements these services.

The `ConnectionManager` instance, on receiving a connection creation request from the connection factory, does a lookup in the connection pool provided by the application server. If there is no connection in the pool that can satisfy the connection request, the application server uses the `ManagedConnectionFactory` interface (implemented by the resource adapter) to create a new physical connection to the underlying EIS. If the application server finds a matching connection in the pool, it uses the matching `ManagedConnection` instance to satisfy the connection request.

If a new `ManagedConnection` instance is created, the application server adds the new `ManagedConnection` instance to the connection pool.

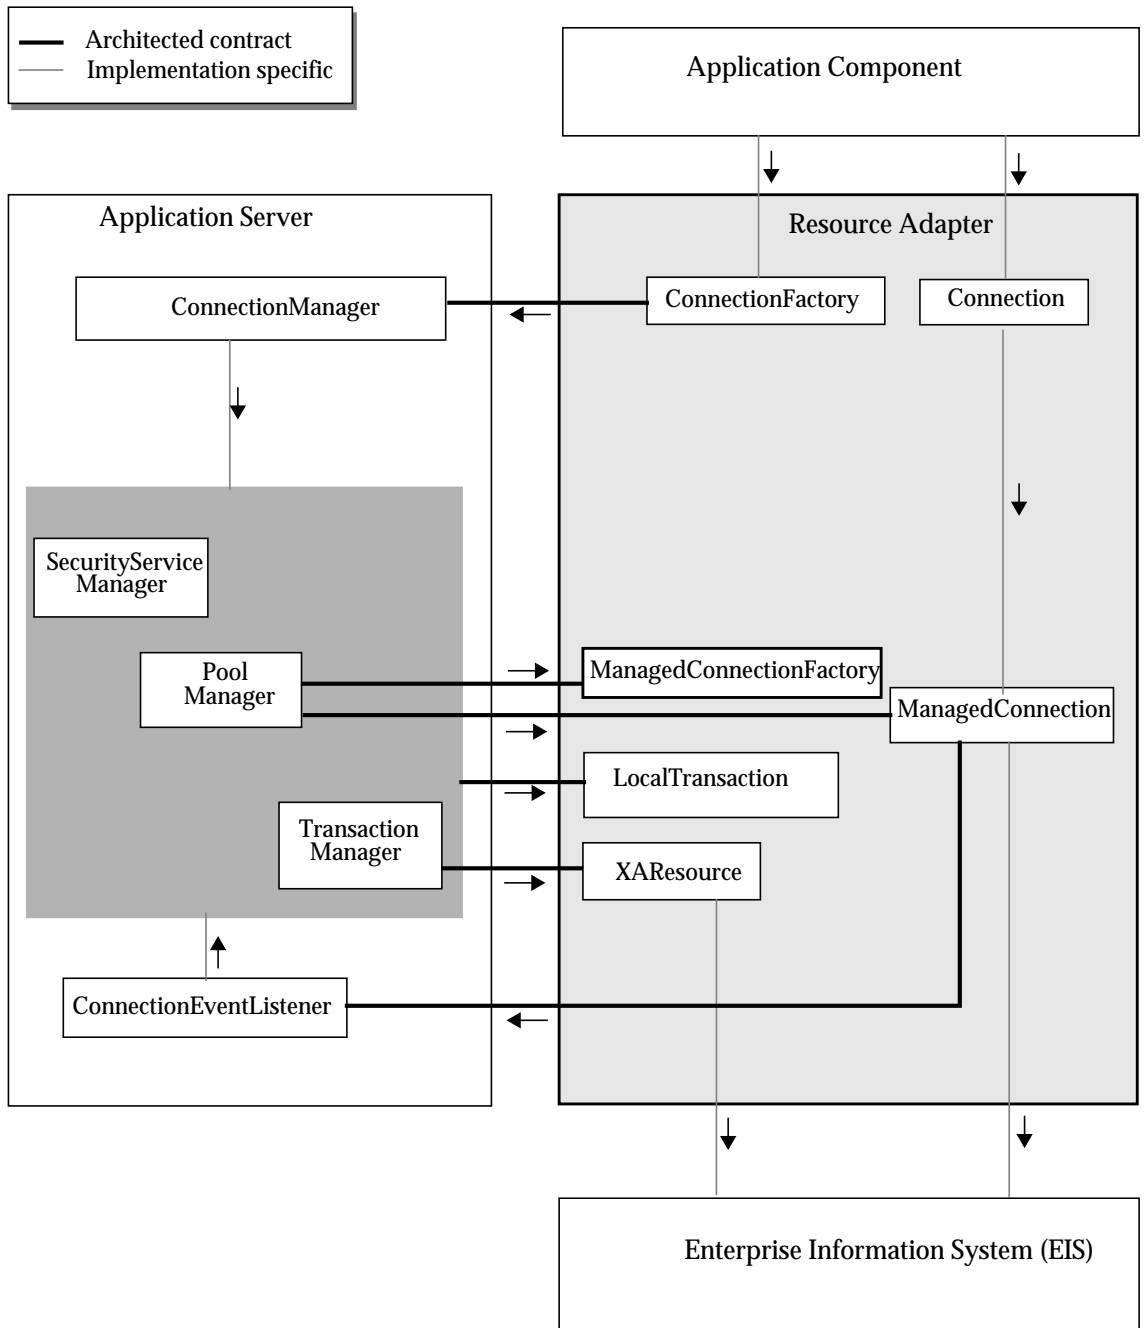
The application server registers a `ConnectionEventListener` with the `ManagedConnection` instance. This listener enables the application server to get event notifications related to the state of the `ManagedConnection` instance. The application server uses these notifications to manage connection pooling, manage transactions, cleanup connections, and handle any error conditions.

The application server uses the `ManagedConnection` instance to get a connection instance that acts as an application-level handle to the underlying physical connection. An instance of type `javax.resource.cci.Connection` is an example of such a connection handle. An application component uses the connection handle to access EIS resources.

The resource adapter implements the `XAResource` interface to provide support for transaction management. The resource adapter also implements the `LocalTransaction` interface so that the application server can manage transactions internal to a resource manager. The chapter on transaction management describes

this transaction management contract between the application server (and its transaction manager) and the resource adapter (and its underlying resource manager).

FIGURE 6-1 Architecture Diagram: Managed Application scenario



6.4 Application Programming Model

The application programming model for getting an EIS connection is similar across both managed (application server based) and non-managed scenarios. The following sections explain a typical application programming model scenario.

6.4.1 Managed Application Scenario

The following steps are involved in a managed scenario:

- The application assembler or component provider specifies connection factory requirements for an application component using a deployment descriptor mechanism. For example, a bean provider specifies the following elements in the deployment descriptor for a connection factory reference. Note that the connection factory reference is part of the deployment descriptor for EJB components and not the resource adapter. Refer EJB specification (“Enterprise JavaBeans (EJB™) specification, version 2.1:” on page F-1) for details on the deployment mechanism for EJB components:
 - `res-ref-name: eis/MyEIS`
 - `res-type: javax.resource.cci.ConnectionFactory`
 - `res-auth: Application or Container`
- During resource adapter deployment, the deployer sets the configuration information (example: server name, port number) for the resource adapter. The application server uses a configured resource adapter to create physical connections to the underlying EIS. Refer to Chapter 16, “API Requirements” for details on packaging and deployment of a resource adapter.
- The application component looks up a connection factory instance in the component’s environment using the JNDI interface.

```
// obtain the initial JNDI Naming context
Context initctx = new InitialContext();

// perform JNDI lookup to obtain the connection factory
javax.resource.cci.ConnectionFactory cxf =
    (javax.resource.cci.ConnectionFactory)
        initctx.lookup("java:comp/env/eis/MyEIS");
```

The JNDI name passed in the method `NamingContext.lookup` is the same as that specified in the `res-ref-name` element of the deployment descriptor. The JNDI lookup results in a connection factory instance of type `java.resource.cci.ConnectionFactory` as specified in the `res-type` element.

- The application component invokes the `getConnection` method on the connection factory to get an EIS connection. The returned connection instance represents an application-level handle to an underlying physical connection.

An application component obtains multiple connections by calling the method `getConnection` on the connection factory multiple times.

```
javax.resource.cci.Connection cx = cxf.getConnection();
```

- The application component uses the returned connection to access the underlying EIS via the resource adapter. Chapter 15, “Common Client Interface” specifies in detail the application programming model for EIS access.

Note – The JNDI context of an accessing application is available to a resource adapter via the application thread that uses its connection object. The resource adapter may use the JNDI context to access other resources.

- After the component finishes with the connection, it closes the connection using the `close` method on the `Connection` interface.

```
cx.close();
```

- If an application component fails to close an allocated connection after its use, that connection is considered an unused connection. The application server manages the cleanup of unused connections. When a container terminates a component instance, the container cleans up all connections used by that component instance. Refer section 6.5.4 and scenario 6.8.3 for details on the cleanup of connections.

6.4.2 Non-managed Application Scenario


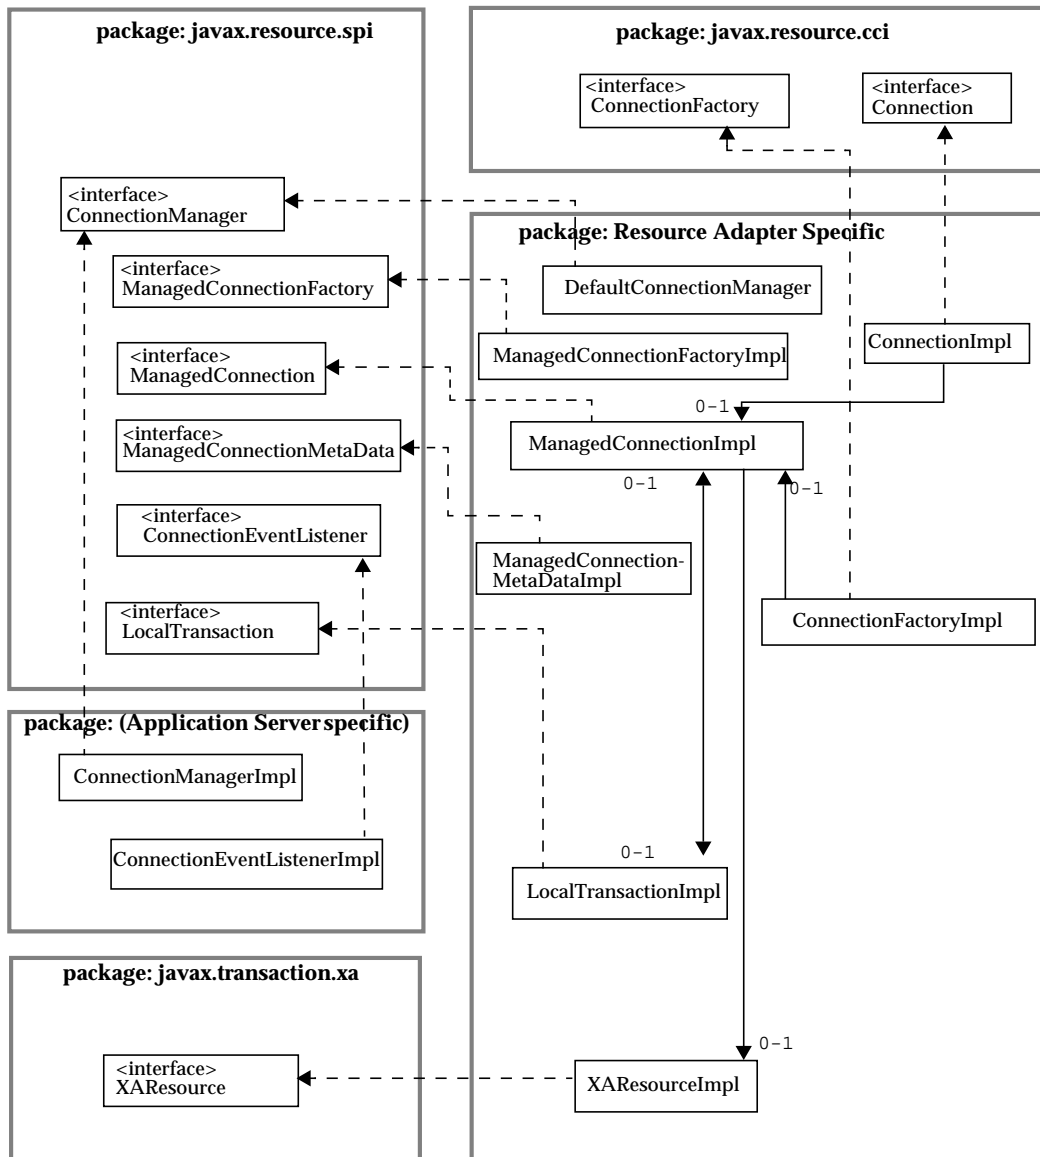
In a non-managed application scenario, the application developer follows a similar programming model to the managed application scenario. The non-managed case involves looking up of a connection factory instance, getting an EIS connection, using the connection for EIS access, and finally closing the connection.

6.5 Interface/Class Specification

This section specifies the Java classes and interfaces defined as part of the connection management contract. For a complete specification of these classes and interfaces, refer to the API documentation distributed with this document.

FIGURE 6-2 shows the class hierarchy for the connection management contract. The diagram also illustrates the responsibilities for the definition of an interface and its implementation:

FIGURE 6-2 Class Diagram: Connection Management Architecture



6.5.1 ConnectionFactory and Connection¹

A connection factory provides an interface to get a connection to an EIS instance. A connection provides connectivity to an underlying EIS.

One goal of the connector architecture is to support a consistent application programming model across both CCI and EIS specific client APIs. To achieve this goal, the connector architecture recommends a design pattern (specified as an interface template) for both the connection factory and connection interfaces.

The CCI connection factory and connection interfaces (defined in the package `javax.resource.cci`) are based on the above design pattern. Refer to Section 15.5, “Connection Interfaces” on page 15-9 for details on the CCI connection factory and connection interfaces. The following code sample shows the CCI interfaces:

```
public interface javax.resource.cci.ConnectionFactory
    extends java.io.Serializable, javax.resource.Referenceable {

    public javax.resource.cci.Connection getConnection()
        throws javax.resource.ResourceException;
    ...
}

public interface javax.resource.cci.Connection {
    public void close() throws javax.resource.ResourceException;
    ...
}
```

1. In this document, the term ‘physical connection’ refers to a `ManagedConnection` instance, while the term ‘connection handle’ refers to an application-level connection handle. When the distinction between ‘physical connection’ and ‘connection handle’ is not important, the term ‘connection’ is used to refer to an EIS connection.

An example of a non-CCI interface is a resource adapter that uses the package `com.myeis` for its EIS specific interfaces, as follows:

```
public interface com.myeis.ConnectionFactory
    extends java.io.Serializable, javax.resource.Referenceable {

    public com.myeis.Connection getConnection()
        throws com.myeis.ResourceException;

    ...
}

public interface com.myeis.Connection {
    public void close() throws com.myeis.ResourceException;
    ...
}
```

The JDBC interfaces—`javax.sql.DataSource`, `java.sql.Connection`—are examples of non-CCI connection factory and connection interfaces.

Note that the methods defined on a non-CCI interface are not required to throw a `ResourceException`. The exception can be specific to a resource adapter, for example: `java.sql.SQLException` for JDBC (Section 3., “JDBC API specification, version 3.0” on page F-1) interfaces.

The following are additional guidelines for the recommended interface template:

- A resource adapter is allowed to add additional `getConnection` methods to its definition of a connection factory interface. These additional methods are specific to a resource adapter and its EIS. For example, CCI defines a variant of the `getConnection` method that takes `java.resource.cci.ConnectionSpec` as a parameter.
- A resource adapter should only introduce additional `getConnection` methods if it requires additional flexibility (beyond that offered by the default `getConnection` method) in the connection request invocations.
- A connection interface must provide a `close` method to close the connection. The behavior of such an application-level connection closure is described in the OID FIGURE 6-7 on page 43.

The above design pattern leads to a consistent application programming model for connection creation and connection closing.

6.5.1.1 Requirements

A resource adapter must provide implementations for both the connection factory and connection interfaces.

In the connector architecture, a resource adapter provides an implementation of the connection factory interface in both managed and non-managed scenarios. This differs from the JDBC (Section 3., “JDBC API specification, version 3.0” on page F-1) architecture.

In the JDBC architecture, an application server provides the implementation of `javax.sql.DataSource` interface. Using a similar design approach for the connector architecture would have required an application server to provide implementations of various connection factory interfaces defined by different resource adapters. Since the connection factory interface may be defined as specific to an EIS, the application server may find it difficult to provide implementations of connection factory interfaces without any code generation.

The connection factory implementation class delegates the `getConnection` method invocation from an application component to the associated `ConnectionManager` instance. The `ConnectionManager` instance is associated with a connection factory instance at its instantiation [refer to the OID shown in FIGURE 17-3 on page 22].

Note that the connection factory implementation class must call the `ConnectionManager.allocateConnection` method in the same thread context in which the application component had called the `getConnection` method.

The connection factory implementation class is responsible for taking connection request information and passing it in a form required by the `ConnectionManager.allocateConnection` method.

```
public interface javax.resource.spi.ConnectionManager
    extends java.io.Serializable {

    public Object allocateConnection(
        ManagedConnectionFactory mcf,
        ConnectionRequestInfo cxRequestInfo)
        throws ResourceException;
}

public interface javax.resource.spi.ConnectionRequestInfo {
    public boolean equals(Object other);
    public int hashCode();
}
```

6.5.1.2 ConnectionRequestInfo

The `ConnectionRequestInfo` parameter to the `ConnectionManager.allocateConnection` method enables a resource adapter to pass its own request-specific data structure across the connection request flow.

A resource adapter extends the `ConnectionRequestInfo` interface to support its own data structure for the connection request.

This is typically used to allow a resource adapter to handle application component-specified per-connection request properties (for example, `client ID` and `language`). The application server passes these properties to the `createManagedConnection` and `matchManagedConnections` method calls on the `ManagedConnectionFactory`. These properties remain opaque to the application server during the connection request flow.

It is important to note that the properties passed through the `ConnectionRequestInfo` instance should be client-specific (for example, `user name`, `password`, `language`) and not related to the configuration of a target EIS instance (for example, `port number`, `server name`).

The `ManagedConnectionFactory` instance is configured with properties required for the creation of a connection to a specific EIS instance. Note that a configured `ManagedConnectionFactory` instance must have the complete set of properties that are needed for the creation of the physical connections. This enables the container to manage connection request without requiring an application component to pass any explicit connection parameters. Configured properties on a `ManagedConnectionFactory` can be overridden through `ConnectionRequestInfo` in cases when a component provides client-specific properties in the `getConnection` method invocation. Refer to Section 17.4.1, “ResourceAdapter” on page 17-11 for details on the configuration of a `ManagedConnectionFactory`.

When the `ConnectionRequestInfo` reaches the `createManagedConnection` or `matchManagedConnections` methods on the `ManagedConnectionFactory` instance, the resource adapter uses this additional per-request information to create and match connections.

A resource adapter must implement the `equals` and `hashCode` methods defined in the `ConnectionRequestInfo` interface. The equality must be defined in the complete set of properties for the `ConnectionRequestInfo` instance. An application server can use these methods to structure its connection pool in an implementation-specific way. Since `ConnectionRequestInfo` represents a resource adapter specific data structure, the conditions for equality are defined and implemented by a resource adapter.

6.5.1.3 Additional Requirements

A resource adapter implementation is not required to support the mechanism for passing resource adapter-specific connection request information. It can choose to pass `null` for `ConnectionRequestInfo` in the `allocateConnection` invocation.

An implementation class for a connection factory interface must implement `java.io.Serializable`. This enables a connection factory instance to be stored in the JNDI naming environment. A connection factory implementation class must implement the interface `javax.resource.Referenceable`. Note that the `javax.resource.Referenceable` interface extends the `javax.naming.Referenceable` interface. Refer to section Section 17.5.3, “Scenario: Referenceable” on page 17-17 for details on the JNDI reference mechanism.

A connection implementation class implements its methods in a resource adapter implementation-specific way. It must use a `javax.resource.spi.ManagedConnection` instance as its underlying physical connection.

6.5.2 ConnectionManager

The `javax.resource.spi.ConnectionManager` interface provides a hook for a resource adapter to pass a connection request to an application server. An application server provides different quality of service as part of its handling of the connection request.

6.5.2.1 Interface

The connection management contract defines a standard interface for the `ConnectionManager` as follows:

```
public interface javax.resource.spi.ConnectionManager
    extends java.io.Serializable {

    public Object allocateConnection(
        ManagedConnectionFactory mcf,
        ConnectionRequestInfo cxRequestInfo)
        throws ResourceException;

}
```

The method `allocateConnection` is called by a resource adapter’s connection factory instance so that the instance can delegate a connection request to the `ConnectionManager` instance.

The `ConnectionRequestInfo` parameter represents information specific to a resource adapter to handle the connection request.

6.5.2.2 Requirements

An application server must provide an implementation of the `ConnectionManager` interface. This implementation is not specific to any particular resource adapter or connection factory interface.

The `ConnectionManager` implementation delegates to the internal mechanisms of an application server to provide various services: security, connection pool management, transaction management, and error logging and tracing.

An application server should implement these services in a generic manner, independent of any resource adapter and EIS-specific mechanisms. The connector architecture does not specify how an application server implements these services; the implementation is specific to each application server.

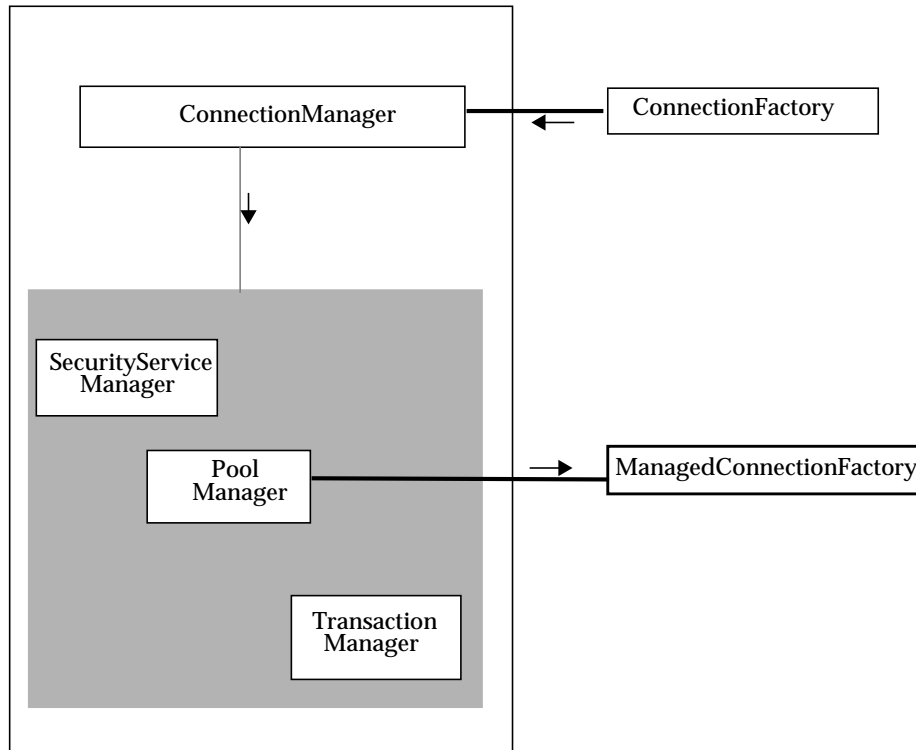
After an application server hooks-in its services, the connection request is delegated to a `ManagedConnectionFactory` instance either for the creation of a new physical connection or for the matching of an already existing physical connection.

An implementation class for the `ConnectionManager` interface must implement the `java.io.Serializable` interface.

A resource adapter must provide a default implementation of the `javax.resource.spi.ConnectionManager` interface. The implementation class comes into play when a resource adapter is used in a non-managed two-tier application scenario. In an application server-managed environment, the resource adapter must not use the default `ConnectionManager` implementation class. A default implementation of `ConnectionManager` enables the resource adapter to provide services specific to itself. These services can include connection pooling, error logging and tracing, and security management. The default `ConnectionManager` delegates to the `ManagedConnectionFactory` the creation of physical connections to the underlying EIS.

An implementation of the `ConnectionManager` interface may only be provided by a resource adapter, for the purpose described in this section, or by an application server that fully meets the requirements of this specification.

FIGURE 6-3 `ConnectionManager` and Application Server specific services



6.5.3 ManagedConnectionFactory

A `javax.resource.spi.ManagedConnectionFactory` instance is a factory of both `ManagedConnection` and connection factory instances. This interface supports connection pooling by defining methods for matching and creating connections.

6.5.3.1 Interface

The following code extract shows the interface specification for the `ManagedConnectionFactory`.

```
public interface javax.resource.spi.ManagedConnectionFactory
    extends java.io.Serializable {

    public Object createConnectionFactory(
        ConnectionManager connectionManager)
        throws ResourceException;

    public Object createConnectionFactory()
        throws ResourceException;

    public ManagedConnection createManagedConnection(
        javax.security.auth.Subject subject,
        ConnectionRequestInfo cxRequestInfo)
        throws ResourceException;

    public ManagedConnection matchManagedConnections(
        java.util.Set connectionSet,
        javax.security.auth.Subject subject,
        ConnectionRequestInfo cxRequestInfo)
        throws ResourceException;

    public boolean equals(Object other);
    public int hashCode();

}
```

The method `createConnectionFactory` creates a connection factory instance. For CCI, the connection factory instance is of the type `javax.resource.cci.ConnectionFactory`. The connection factory instance is initialized with the `ConnectionManager` instance provided by the application server.

When the `createConnectionFactory` method takes no arguments, `ManagedConnectionFactory` provides a default `ConnectionManager` instance. This occurs in a non-managed application scenario.

The method `createManagedConnection` creates a new physical connection to the underlying EIS instance. The `ManagedConnectionFactory` instance uses the security information (passed as a `Subject` instance) and an optional `ConnectionRequestInfo` instance to create this new physical connection (refer to Chapter 9, “Security Contract” for more details).

A created `ManagedConnection` instance typically maintains internal information about the security context (under which the connection has been created) and any connection-specific parameters (for example, the socket connection).

The `matchManagedConnections` method enables the application server to use resource adapter-specific criteria for matching a `ManagedConnection` instance to service a connection request. The application server finds a candidate set of `ManagedConnection` instances from its connection pool based on application server-specific criteria, and passes this candidate set to the `matchManagedConnections` method. If the application server implements connection pooling, it must use the `matchManagedConnections` method to choose a suitable connection.

The `matchManagedConnections` method matches a candidate set of connections using criteria known internally to the resource adapter. The criteria used for matching connections is specific to a resource adapter and is not specified by the connector architecture.

A `ManagedConnection` instance has specific internal state information based on its security context and physical connection. The `ManagedConnectionFactory` implementation compares this information for each `ManagedConnection` instance in the candidate set against the information passed in through the `matchManagedConnections` method and the configuration of this `ManagedConnectionFactory` instance. The `ManagedConnectionFactory` uses the results of this comparison to choose the `ManagedConnection` instance that can best satisfy the current connection request.

If the resource adapter cannot find an acceptable `ManagedConnection` instance, it returns a `null` value. In this case, the application server requests the resource adapter to create a new connection instance.

If the resource adapter does not support connection matching, it must throw a `NotSupportedException` when `matchManagedConnections` method is invoked. This allows an application server to avoid pooling connections obtained from that resource adapter.

6.5.3.2 Requirements

A resource adapter must provide an implementation of the `ManagedConnectionFactory` interface.

It is required that the `ManagedConnectionFactory` implementation class extend the implementation of the `hashCode` and `equals` methods defined in `java.lang.Object`. These two methods are used by an application server to structure its connection pool in an implementation-specific way. The `equals` and

`hashCode` method implementation should be based on a complete set of configuration properties that make a `ManagedConnectionFactory` instance unique and specific to an EIS instance.

An implementation class for `ManagedConnectionFactory` interface must be a `JavaBean`. Refer to Section 16.3, “`JavaBean` Requirements” on page 16-2.

6.5.3.3 Connection Pool Implementation

The connector architecture does not specify how an application server implements connection pooling. However, it recommends that an application server should structure its connection pool such that it uses the connection creation and matching facility in an efficient manner and does not cause resource starvation.

The following paragraphs provide non-prescriptive guidelines for the connection pool implementation by an application server.

An application server may partition its pool on a per `ManagedConnectionFactory` instance (and thereby on a per EIS instance) basis. An application server may choose to guarantee, in an implementation specific way, that it will always partition connection pools with at least per `ManagedConnectionFactory` instance granularity.

The per-`ManagedConnectionFactory` instance pool may be further partitioned based on the transaction or security context or any client-specific parameters (as associated with the `ConnectionRequestInfo`). When an application server calls the matching facility, it is recommended that the application server narrow down the candidate set of `ManagedConnection` instances to a reasonable limit, and achieves matching efficiently. For example, an application server may pass only those `ManagedConnection` instances to the `matchManagedConnections` method that are associated with the target `ManagedConnectionFactory` instance (and thereby a specific target EIS instance).

An application server may use additional parameters for its search and matching criteria used in its connection pool management. These parameters may be EIS- or application server- specific. The `equals` and `hashCode` methods defined in both `ManagedConnectionFactory` and `ConnectionRequestInfo` facilitate connection pool management and structuring by an application server.

6.5.3.4 Detecting Invalid Connections

```
import java.util.Set;

interface ValidatingManagedConnectionFactory {
    Set getInvalidConnections(Set connectionSet) throws
    ResourceException;
}
```

This interface may be implemented by a `ManagedConnectionFactory` instance that supports the ability to validate `ManagedConnection` objects. The `getInvalidConnections` method returns a set of invalid `ManagedConnection` objects chosen from a specified set of `ManagedConnection` objects.

This optional functionality may be used by the application server to prune invalid `ManagedConnection` objects from its connection pool periodically.

6.5.3.5 Requirement for XA Recovery

The `ManagedConnectionFactory` implementation for a transaction authority (XA) protocol capable resource adapter (refer to Chapter 7, “Transaction Management” for more details on transactions) must support the `createManagedConnection` method that takes a `Subject` and a null for the parameter `ConnectionRequestInfo`. This enables the application server to get an `XAResource` instance using `ManagedConnection.getXAResource` and then call the `XAResource.recover` method. Note that the application server uses this special case only to get to the `XAResource` instance for the underlying resource manager.

The reason for this requirement is that the application server may not have a valid `ConnectionRequestInfo` instance when it needs to get the `ManagedConnection` instance to initiate recovery. Refer to Section 9.1.8, “`ManagedConnectionFactory`” on page 9-9 for additional details on the `ManagedConnectionFactory.createManagedConnection` method.

6.5.4 ManagedConnection

A `javax.resource.spi.ManagedConnection` instance represents a physical connection to an underlying EIS.

Note – The connector architecture allows one or more `ManagedConnection` instances to be multiplexed over a single physical pipe to an EIS. However, for simplicity, this specification describes a `ManagedConnection` instance as being mapped 1-1 to a physical connection.

The creation of a `ManagedConnection` instance typically results in the allocation of EIS and resource adapter resources (for example, memory and network sockets) for each physical connection. Since these resources can be costly and scarce, an application server pools `ManagedConnection` instances in a managed environment.

Connection pooling improves the scalability of an application environment. An application server uses the `ManagedConnectionFactory` and `ManagedConnection` interfaces to implement connection pool management.

An application server also uses the transaction management-related methods (`getXAResource` and `getLocalTransaction`) on the `ManagedConnection` interface to manage transactions. These methods are discussed in more detail in Chapter 7, “Transaction Management”.

The `ManagedConnection` interface also provides methods to support error logging and tracing in a managed environment.

6.5.4.1 Interface

The connection management contract defines the following interface for a `ManagedConnection`. The following code extract shows only the methods that are used for connection pool management. The remaining methods are introduced in other parts of the specification.

```
public interface javax.resource.spi.ManagedConnection {
    public Object getConnection(
        javax.security.auth.Subject subject,
        ConnectionRequestInfo cxRequestInfo)
        throws ResourceException;
    public void destroy() throws ResourceException;
    public void cleanup() throws ResourceException;

    // Methods for Connection and transaction event notifications
    public void addConnectionEventListener(
        ConnectionEventListener listener);
    public void removeConnectionEventListener(
        ConnectionEventListener listener);

    public ManagedConnectionMetaData getMetaData()
        throws ResourceException;

    // Additional methods - specified in the other sections
    ...
}
```

The `getConnection` method creates a new application-level connection handle. A connection handle is tied to an underlying physical connection represented by a `ManagedConnection` instance. For CCI, the connection handle created by a `ManagedConnection` instance is of the type `javax.resource.cci.Connection`. A connection handle is tied to its `ManagedConnection` instance in a resource adapter implementation-specific way.

A `ManagedConnection` instance may use the `getConnection` method to change the state of the physical connection based on the `Subject` and `ConnectionRequestInfo` arguments. For example, a resource adapter can re-authenticate a physical connection to the underlying EIS when the application server calls the `getConnection` method. Section 9.1.9, “`ManagedConnection`” on page 9-14 specifies re-authentication requirements in more detail.

The method `addConnectionEventListener` allows a connection event listener to register with a `ManagedConnection` instance. The `ManagedConnection` instance notifies connection close/error and local transaction-related events to its registered set of listeners.

The `removeConnectionEventListener` method removes a registered `ConnectionEventListener` instance from a `ManagedConnection` instance.

The method `getMetaData` returns the metadata information (represented by the `ManagedConnectionMetaData` interface) for a `ManagedConnection` and the connected EIS instance.

6.5.4.2 Connection Sharing and Multiple Connection Handles

To support connection sharing, the application server can call `getConnection` multiple times on a `ManagedConnection` instance. In this case, a call to the method `ManagedConnection.getConnection` does not invalidate any previously created connection handles. Multiple connection handles can exist concurrently for a single `ManagedConnection` instance. This design supports the connection sharing mechanism. Refer to Section 7.9, “Connection Sharing” on page 7-35 for more details.

Because multiple connection handles to a single `ManagedConnection` can exist concurrently, a resource adapter implementation may:

- Provide thread-safe semantics for a `ManagedConnection` implementation to support concurrent access to a `ManagedConnection` instance from multiple connection handles. It is strongly recommended that resource adapters provide support for concurrent access to a `ManagedConnection` instance from multiple connection handles. This may be required in a future release of the specification.
- Ensure that there is at most one connection handle associated actively with a `ManagedConnection` instance. The active connection handle is the only connection using the `ManagedConnection` instance until an application-level `close` is called on this connection handle. For example, a `ManagedConnection.getConnection` method implementation associates a newly created connection handle as the active connection handle. Any operations on the `ManagedConnection` from any previously created connection handles should result in an application level exception. An example application level exception extends the `javax.resource.ResourceException` interface and is specific to a resource adapter. A scenario illustrating this implementation is shown in the Section 7.10.3, “Scenario: Local Transaction” on page 7-40.

6.5.4.3 Connection Matching Contract

The application server invokes the `ManagedConnectionFactory.matchManagedConnections` method (implemented by a resource adapter) to find a matching `ManagedConnection` for servicing a connection request. The application server passes a candidate set of `ManagedConnection` instances to the `matchManagedConnections` method.

The application server should use the connection matching contract for `ManagedConnection` instances that have no existing connection handles. A candidate set passed to the `matchManagedConnections` method should not have any `ManagedConnection` instances with existing connection handles.

There is no requirement that the `matchManagedConnections` implementation be capable of performing a match across a candidate set that includes `ManagedConnection` instances with existing connection handles. Note that a resource adapter can return a successful match with the requirement that the `ManagedConnection.getConnection` method will later change the state of the matched `ManagedConnection`. To avoid any unexpected matching behavior, the application server should not pass a `ManagedConnection` instance with existing connection handles to the `matchManagedConnections` method as part of a candidate set.

A connection request can lead to the creation of additional connection handles for a `ManagedConnection` instance that already has one or more existing connection handles. In this case, the application server should take the responsibility of checking whether or not the chosen `ManagedConnection` instance can service such a request. Refer to Section 7.9, “Connection Sharing” on page 7-35 for details.

6.5.4.4 Cleanup of `ManagedConnection`

A resource adapter typically allocates system resources (outside a JVM) for a `ManagedConnection` instance. Additionally, a `ManagedConnection` instance can have state specific to a client, such as security context, data/function access structures, and result set from a query.

The method `ManagedConnection.cleanup` initiates a cleanup of any client-specific state maintained by a `ManagedConnection` instance. The cleanup must invalidate all connection handles created using the `ManagedConnection` instance. Any attempt by an application component to use the associated connection handle after cleanup of the underlying `ManagedConnection` should result in an exception.

The container always drives the cleanup of a `ManagedConnection` instance. The container keeps track of created connection handles in an implementation specific mechanism. It invokes `ManagedConnection.cleanup` when it has to invalidate all connection handles associated with this `ManagedConnection` instance and put the `ManagedConnection` instance back in to the pool. This may be called after the end of a connection sharing scope or when the last associated connection handle is closed for a `ManagedConnection` instance.

The invocation of the `ManagedConnection.cleanup` method on an already cleaned-up connection should not throw an exception.

The cleanup of a `ManagedConnection` instance resets its client-specific state and prepares the connection to be put back into a connection pool. The `cleanup` method should not cause the resource adapter to close the physical pipe and reclaim system resources associated with the physical connection.

An application server should explicitly call `ManagedConnection.destroy` to destroy a physical connection. An application server should destroy a physical connection to manage the size of its connection pool and to reclaim system resources.

A resource adapter should destroy all allocated system resources for this `ManagedConnection` instance when the method `destroy` is called.

6.5.4.5 Requirements

A resource adapter must provide an implementation of the `ManagedConnection` interface.

6.5.5 ManagedConnectionMetaData

The method `ManagedConnection.getMetaData` returns a `javax.resource.spi.ManagedConnectionMetaData` instance. The `ManagedConnectionMetaData` provides information about a `ManagedConnection` and the connected EIS instance. This information is only available to the caller of this method if a valid physical connection exists for an EIS instance.

6.5.5.1 Interface

The `ManagedConnectionMetaData` interface provides the following information about an EIS instance:

- Product name of the EIS instance
- Product version of the EIS instance
- Maximum number of concurrent connections from different processes that an EIS instance can support
- User name for this connection, as known to the EIS instance

The method `getUserName` returns the user name known to the underlying EIS instance for an active connection. The name corresponds to the resource principal under whose security context the connection to the EIS instance has been established.

6.5.5.2 Requirements

A resource adapter must provide an implementation of the `ManagedConnectionMetaData` interface. An instance of this implementation class should be returned from the `ManagedConnection.getMetaData` method.

6.5.6 ConnectionEventListener

The connector architecture provides an event callback mechanism that enables an application server to receive notifications from a `ManagedConnection` instance. An application server uses these event notifications to manage its connection pool, to clean up invalid or terminated connections, and to manage local transactions. Chapter 7, “Transaction Management” discusses local transaction-related event notifications in more detail.

An application server implements the `javax.resource.spi.ConnectionEventListener` interface. It uses the `ManagedConnection.addConnectionEventListener` method to register a connection listener with a `ManagedConnection` instance.

6.5.6.1 Interface

The following code extract specifies the `ConnectionEventListener` interface:

```
public interface javax.resource.spi.ConnectionEventListener {
    public void connectionClosed(ConnectionEvent event);
    public void connectionErrorOccurred(ConnectionEvent event);

    // Local Transaction Management related events
    public void localTransactionStarted(ConnectionEvent event);
    public void localTransactionCommitted(ConnectionEvent event);
    public void localTransactionRolledback(ConnectionEvent
event);
}
```

A `ManagedConnection` instance calls the `ConnectionEventListener.connectionClosed` method to notify its registered set of listeners when an application component closes a connection handle. The application server uses this connection close event to make a decision on whether or not to put the `ManagedConnection` instance back into the connection pool.

The `ManagedConnection` instance calls the `ConnectionEventListener.connectionErrorOccurred` method to notify its registered listeners of the occurrence of a physical connection-related error. The event notification happens just before a resource adapter throws an exception to the application component using the connection handle.

The `connectionErrorOccurred` method indicates that the associated `ManagedConnection` instance is now invalid and unusable. The application server handles the connection error event notification by initiating application server-specific cleanup (for example, removing `ManagedConnection` instance from the connection pool) and then calling `ManagedConnection.destroy` method to destroy the physical connection.

A `ManagedConnection` instance also notifies its registered listeners for transaction-related events by calling the following methods—`localTransactionStarted`, `localTransactionCommitted`, and `localTransactionRolledback`. An application server uses these notifications to manage local transactions. See Section 7.7, “Local Transaction Management Contract” on page 7-31 for details on the local transaction management.

The processing of event notifications by the registered event listeners may be synchronous or asynchronous. That is, a listener may process an event notification immediately (as part of the notification method call) or it may defer event processing to a later in time. The resource adapter must not assume the processing of event notifications by its listeners to be synchronous or asynchronous.

6.5.7 ConnectionEvent

A `javax.resource.spi.ConnectionEvent` class provides information about the source of a connection-related event. A `ConnectionEvent` instance contains the following information:

- Type of the connection event
- `ManagedConnection` instance that has generated the connection event. A `ManagedConnection` instance is returned from the `ConnectionEvent.getSource` method.
- Connection handle associated with the `ManagedConnection` instance; required for the `CONNECTION_CLOSED` event and optional for the other event types.
- Optionally, an exception indicating a connection related error. Refer to Section 19.2, “System Exceptions” on page 19-2 for details on the system exception. Note that the exception is used for the `CONNECTION_ERROR_OCCURRED` notification.

This class defines the following types of event notifications:

- `CONNECTION_CLOSED`
- `LOCAL_TRANSACTION_STARTED`

- LOCAL_TRANSACTION_COMMITTED
- LOCAL_TRANSACTION_ROLLEDBACK
- CONNECTION_ERROR_OCCURRED

6.6 Error Logging and Tracing

The connector architecture provides basic support for error logging and tracing in both managed and non-managed environments. This support enables an application server to detect errors related to a resource adapter and its EIS, and to use error information for debugging.

6.6.1 ManagedConnectionFactory

The `javax.resource.spi.ManagedConnectionFactory` interface defines the following methods for error logging and tracing:

```
public interface javax.resource.spi.ManagedConnectionFactory
    extends java.io.Serializable {

    public void setLogWriter(java.io.PrintWriter out)
        throws ResourceException;

    public java.io.PrintWriter getLogWriter()
        throws ResourceException;

    ...

}
```

The log writer is a character output stream to which all logging and tracing messages for a `ManagedConnectionFactory` instance are printed.

A character output stream can be registered with a `ManagedConnectionFactory` instance using the `setLogWriter` method. A `ManagedConnectionFactory` implementation uses this character output stream to output error log and trace information.

An application server manages the association of a log writer with a `ManagedConnectionFactory`. When a `ManagedConnectionFactory` instance is created, the log writer is initially `null` and logging is disabled. Associating a log writer with a `ManagedConnectionFactory` instance enables logging and tracing for the `ManagedConnectionFactory` instance.

An application server administrator primarily uses the error and trace information printed on a log writer by a `ManagedConnectionFactory` instance. This information is typically system-level in nature (for example, information related to connection pooling and transactions) rather than of direct interest to application developers.

6.6.2 ManagedConnection

The `javax.resource.spi.ManagedConnection` interface defines the following methods to support error logging and tracing specific to a physical connection.

```
public interface javax.resource.spi.ManagedConnection {
    public void setLogWriter(java.io.PrintWriter out)
        throws ResourceException;
    public java.io.PrintWriter getLogWriter()
        throws ResourceException;
    ...
}
```

A newly created `ManagedConnection` instance gets the default log writer from the `ManagedConnectionFactory` instance that creates the `ManagedConnection` instance. The default log writer can be overridden by an application server using the `ManagedConnection.setLogWriter` method. The setting of the log writer on a `ManagedConnection` enables an application server to manage error logging and tracing specific to the physical connection represented by a `ManagedConnection` instance.

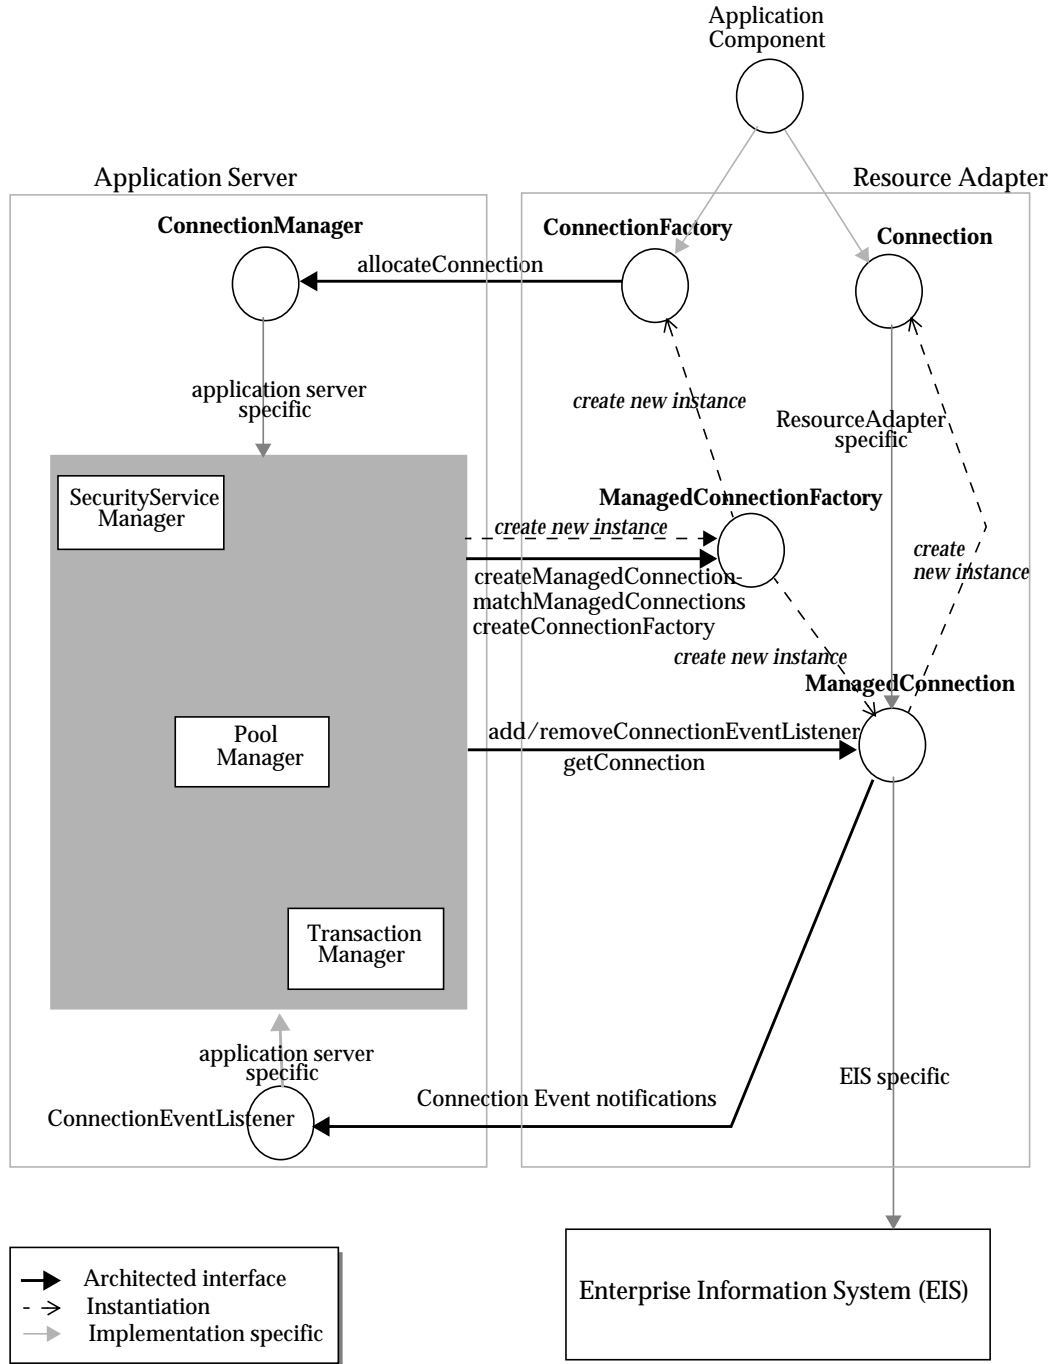
An application server can optionally disassociate the log writer from a `ManagedConnection` instance when this connection instance is put back into the connection pool by using `setLogWriter` and passing `null`.

6.7 Object Diagram

FIGURE 6-4 shows the object diagram for the connection management architecture. It shows invocations across the various object instances that correspond to the architected interfaces in the connection management contract, as opposed to those instances specific to implementations of the application server and the resource adapter.

To keep the diagram simple, it does not show the transaction management contract-related interfaces (`XAResource` and `LocalTransaction`) and invocations.

FIGURE 6-4 Object Diagram: Connection Management architecture



6.8 Illustrative Scenarios

The following section uses sequence diagrams to illustrate various interactions between the object instances involved in the connection management contract.

Some sequence diagrams include a box labeled “Application Server”. This box refers to various modules and classes internal to an application server. These modules and classes communicate through contracts that are application server implementation specific.

In this section, the CCI interfaces—`javax.resource.cci.ConnectionFactory` and `javax.resource.cci.Connection`—represent connection factory and connection interfaces respectively.

The description of these sequence diagrams does not include transaction-related details. These are covered in Chapter 7, “Transaction Management.”

6.8.1 Scenario: Connection Pool Management

The following object interactions are involved in the scenario shown in FIGURE 6-5 on page 36:

- The application component calls the `getConnection` method on the `javax.resource.cci.ConnectionFactory` instance (returned from the JNDI lookup) to get a connection to the underlying EIS instance. Refer to Section 17.5, “JNDI Configuration and Lookup” on page 17-13 for details on the JNDI configuration and lookup.
- The `ConnectionFactory` instance initially handles the connection request from the application component in a resource adapter-specific way. It then delegates the connection request to the associated `ConnectionManager` instance. The `ConnectionManager` instance has been associated with the `ConnectionFactory` instance when the `ConnectionFactory` was instantiated.

The `ConnectionFactory` instance receives all connection request information passed through the `getConnection` method and, in turn, passes it in a form required by the method `ConnectionManager.allocateConnection`. The `ConnectionRequestInfo` parameter to the `allocateConnection` method enables a `ConnectionFactory` implementation class to pass on client-specific connection request information. This information is opaque to an application server and is used subsequently by a resource adapter to do connection matching and creation.

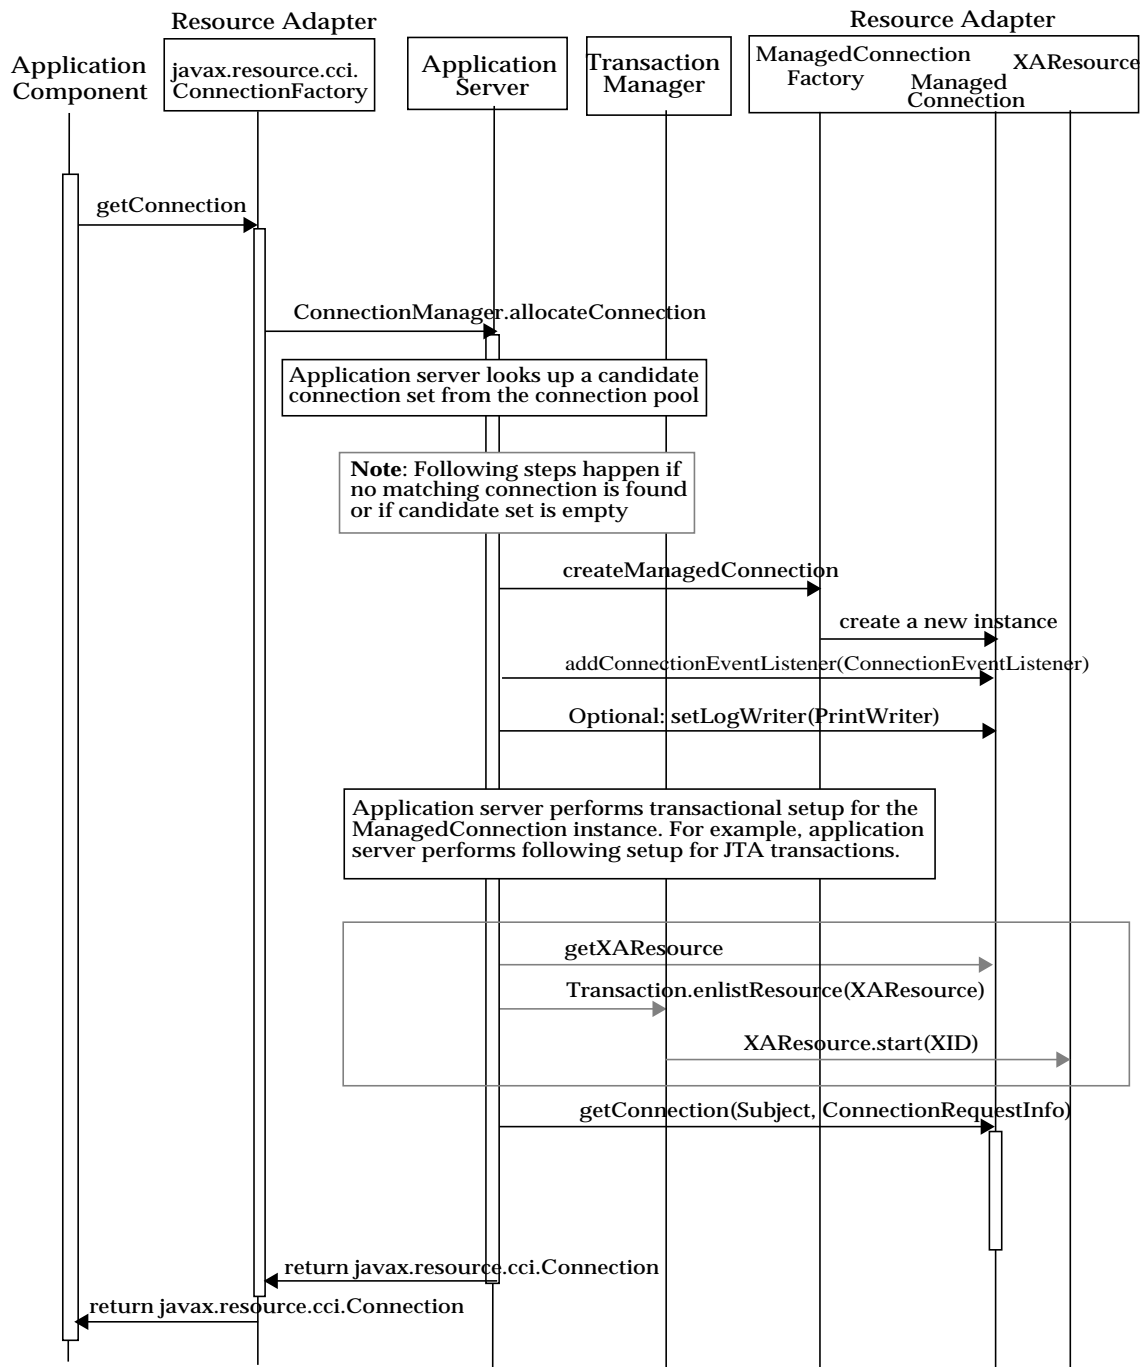
- The `ConnectionManager` instance (provided by the application server) handles the `allocateConnection` request by interacting with the application server-specific connection pool manager. The interaction between a `ConnectionManager` instance and pool manager is internal and specific to an application server.
- The application server finds a candidate set of `ManagedConnection` instances from its connection pool. The candidate set includes all `ManagedConnection` instances that the application server considers suitable for handling the current connection allocation request. The application server finds the candidate set using its own implementation-specific structuring and lookup criteria for the connection pool. Refer to Section 6.5.3 “`ManagedConnectionFactory`” for guidelines of connection pool implementation by an application.
- If the application server finds no matching `ManagedConnection` instance that can best handle this connection allocation request, or if the candidate set is empty, the application server calls the `ManagedConnectionFactory.createManagedConnection` method to create a new physical connection to the underlying EIS instance. The application server passes necessary security information (as JAAS Subject) as part of this method invocation. For details on the security contract, refer to the *Security Management* chapter. It can also pass the `ConnectionRequestInfo` information to the resource adapter. The connection request information has been associated with the connection allocation request by the resource adapter and is used during connection creation.
- The `ManagedConnectionFactory` instance creates a new physical connection to the underlying EIS to handle the `createManagedConnection` method. This new physical connection is represented by a `ManagedConnection` instance. The `ManagedConnectionFactory` uses the security information (passed as a Subject instance), `ConnectionRequestInfo`, and its default set of configured properties (port number, server name) to create a new `ManagedConnection` instance. Refer to Chapter 9, “Security Contract” for more details on the `createManagedConnection` method.
- The `ManagedConnectionFactory` instance initializes the created `ManagedConnection` instance and returns it to the application server.
- The application server registers a `ConnectionEventListener` instance with the `ManagedConnection` instance, enabling it to receive notifications for events on this connection. The application server uses these event notifications to manage connection pooling and transactions.
- The `ManagedConnection` instance obtains its log writer (for error logging and tracing support) from the `ManagedConnectionFactory` instance that created this connection. However, an application server can set a new log writer with a `ManagedConnection` instance to do additional error logging and tracing at the level of a `ManagedConnection`.
- The application server does the necessary transactional setup for the `ManagedConnection` instance. Chapter 7, “Transaction Management” explains this step in more detail.

- Next, the application server calls `ManagedConnection.getConnection` method to get an application level connection handle of type `javax.resource.cci.Connection`. A `ManagedConnection` instance uses the `Subject` and `ConnectionRequestInfo` parameters to the `getConnection` method to change the state of the `ManagedConnection`.

Calling the `getConnection` method does not necessarily create a new physical connection to the EIS instance. Calling `getConnection` produces a temporary connection handle that is used by an application component to access the underlying physical connection. The actual underlying physical connection is represented by a `ManagedConnection` instance.

- The application server returns the connection handle to the resource adapter. The resource adapter then passes the connection handle to the application component that initiated the connection request.

FIGURE 6-5 OID: Connection Pool Management with new Connection Creation



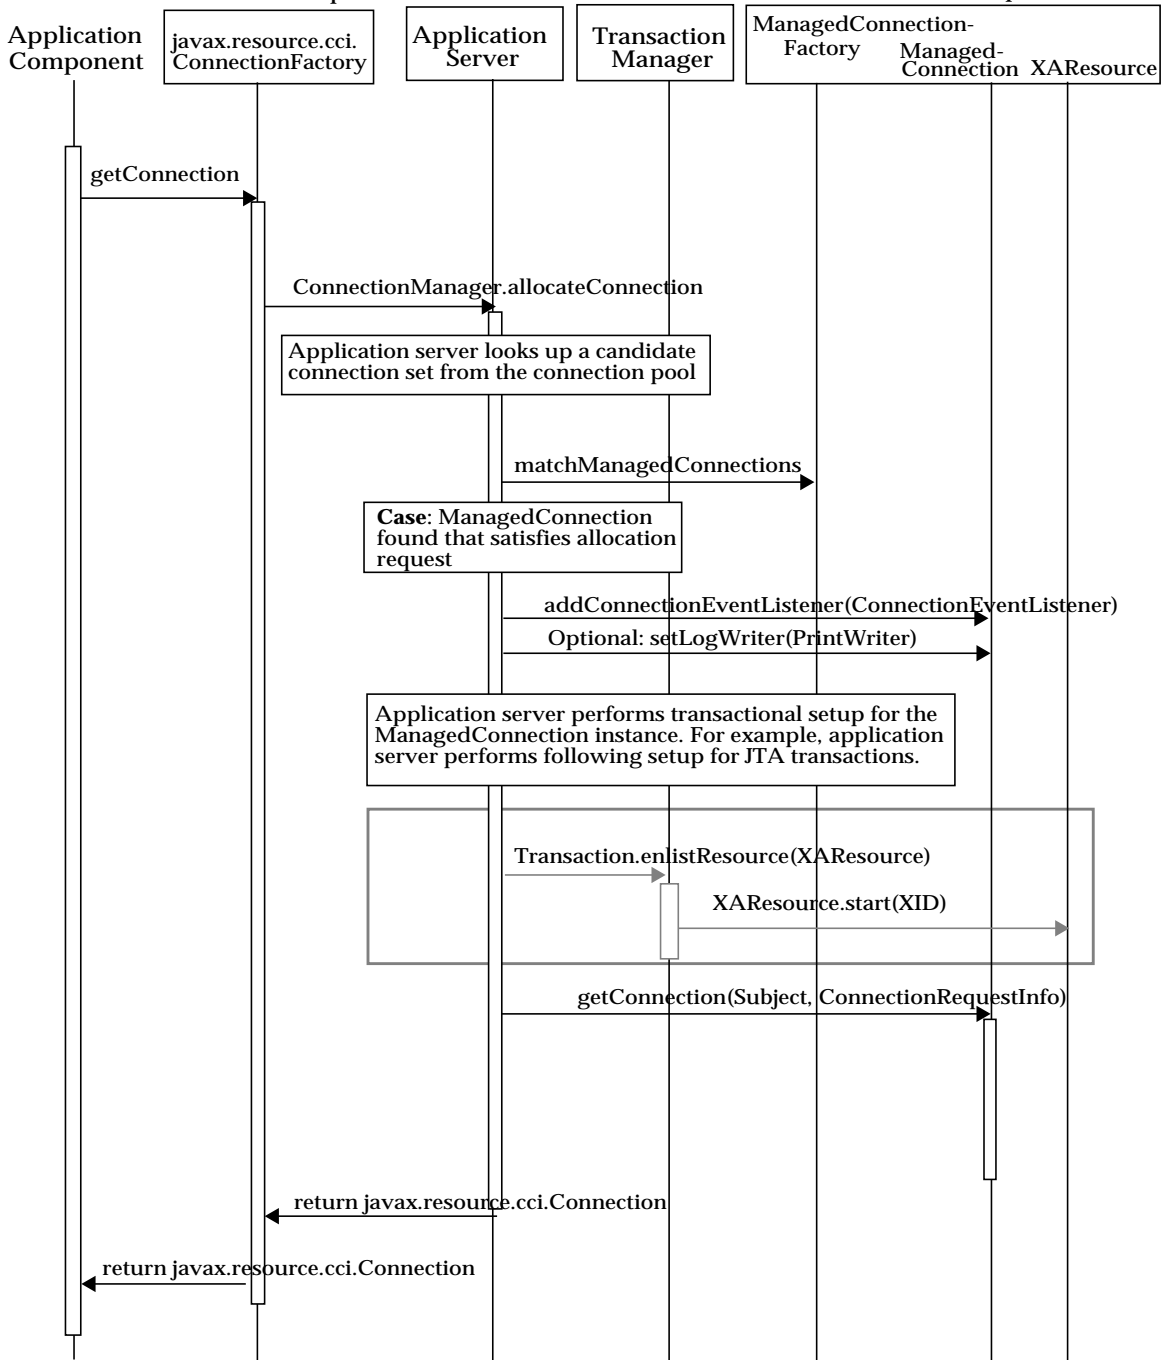
6.8.2 Scenario: Connection Matching

FIGURE 6-6 on page 39 shows the object interactions for a connection matching scenario—that is, a scenario in which the application server finds a non-empty candidate connection set and calls the resource adapter to do matching on the candidate set. The following steps are involved in this scenario:

- The application server handles the connection allocation request by creating a candidate set of `ManagedConnection` instances from the connection pool. The candidate set includes the `ManagedConnection` instances that the application server considers suitable for handling the current connection allocation request. The application server finds this candidate set using its own implementation-specific structuring and lookup criteria for the connection pool. Refer to Section 6.5.3 “`ManagedConnectionFactory`” for guidelines on connection pool implementation by an application.
- The application server calls the `ManagedConnectionFactory.matchManagedConnections` method to enable the resource adapter to do the connection matching. It passes the candidate connection set, security information (as a `Subject` instance associated with the current connection request), and any `ConnectionRequestInfo`.
- The `ManagedConnectionFactory` instance matches the candidate set of connections using the criteria known internally to the resource adapter. The `matchManagedConnections` method returns a `ManagedConnection` instance that the resource adapter considers to be an acceptable match for the current connection allocation request.
- The application server can set a new log writer with the `ManagedConnection` instance to do error logging and tracing at the level of the `ManagedConnection`.
- The application server does the necessary transactional setup for the `ManagedConnection` instance. Chapter 7, “Transaction Management” explains this step in more detail.
- The application server calls the `ManagedConnection.getConnection` method to get a new application level connection handle.
- The `ManagedConnection.getConnection` method implementation uses the `Subject` parameter and any `ConnectionRequestInfo` to set the state of the `ManagedConnection` instance based on the current connection allocation request. Refer to Section 9.1.9, “`ManagedConnection`” on page 9-14 for details if a resource adapter implements support for re-authentication of a `ManagedConnection` instance.

- The application server returns the connection handle to the resource adapter. The resource adapter then passes the connection handle to the application component that initiated the connection request.

FIGURE 6-6 OID: Connection Pool Management with Connection Matching
Resource Adapter Resource Adapter



6.8.3 Scenario: Connection Event Notifications and Connection Close

For each `ManagedConnection` instance in the pool, the application server registers a `ConnectionEventListener` instance to receive close and error events on the connection. This scenario explains how the connection event callback mechanism enables an application server to manage connection pooling.

The scenario involves the following steps (see FIGURE 6-7 on page 43) when an application component initiates a connection close:

- The application component releases an allocated connection handle using the `close` method on the `javax.resource.cci.Connection` instance. The `Connection` instance delegates the `close` method to the associated `ManagedConnection` instance. The delegation happens through an association between `ManagedConnection` instance and the corresponding connection handle `Connection` instance. The mechanism by which this association is achieved is specific to the implementation of a resource adapter.
- The connection management contract places a requirement that a `ManagedConnection` instance must not alter the state of a physical connection while handling the connection close.
- The `ManagedConnection` instance notifies all its registered listeners of the application's connection close request using the `ConnectionEventListener.connectionClosed` method. It passes a `ConnectionEvent` instance with the event type set to `CONNECTION_CLOSED`.
- On receiving the connection close event notification, the application server performs the transaction management-related cleanup of the `ManagedConnection` instance. Refer to FIGURE 6-7 on page 43 for details on the cleanup of a `ManagedConnection` instance participating in a JTA transaction.
- The application server also uses the connection close event notification to manage its connection pool. On receiving the connection close notification, the application server calls the `ManagedConnection.cleanup` method (depending on whether the `ManagedConnection` is shared and the presence of other active connection handles) to perform cleanup on the `ManagedConnection` instance that raised the connection close event. The application server-initiated cleanup of a `ManagedConnection` instance prepares this `ManagedConnection` instance to be reused for subsequent connection requests. See Section 7.9, "Connection Sharing" on page 7-35 for a discussion of connection sharing and its implications on `ManagedConnection` cleanup.
- After initiating the necessary cleanup for the `ManagedConnection` instance, the application server puts the `ManagedConnection` instance back into the connection pool. The application server should be able to use this available `ManagedConnection` instance to handle future connection allocation requests from application components.

6.8.3.1 Connection Cleanup

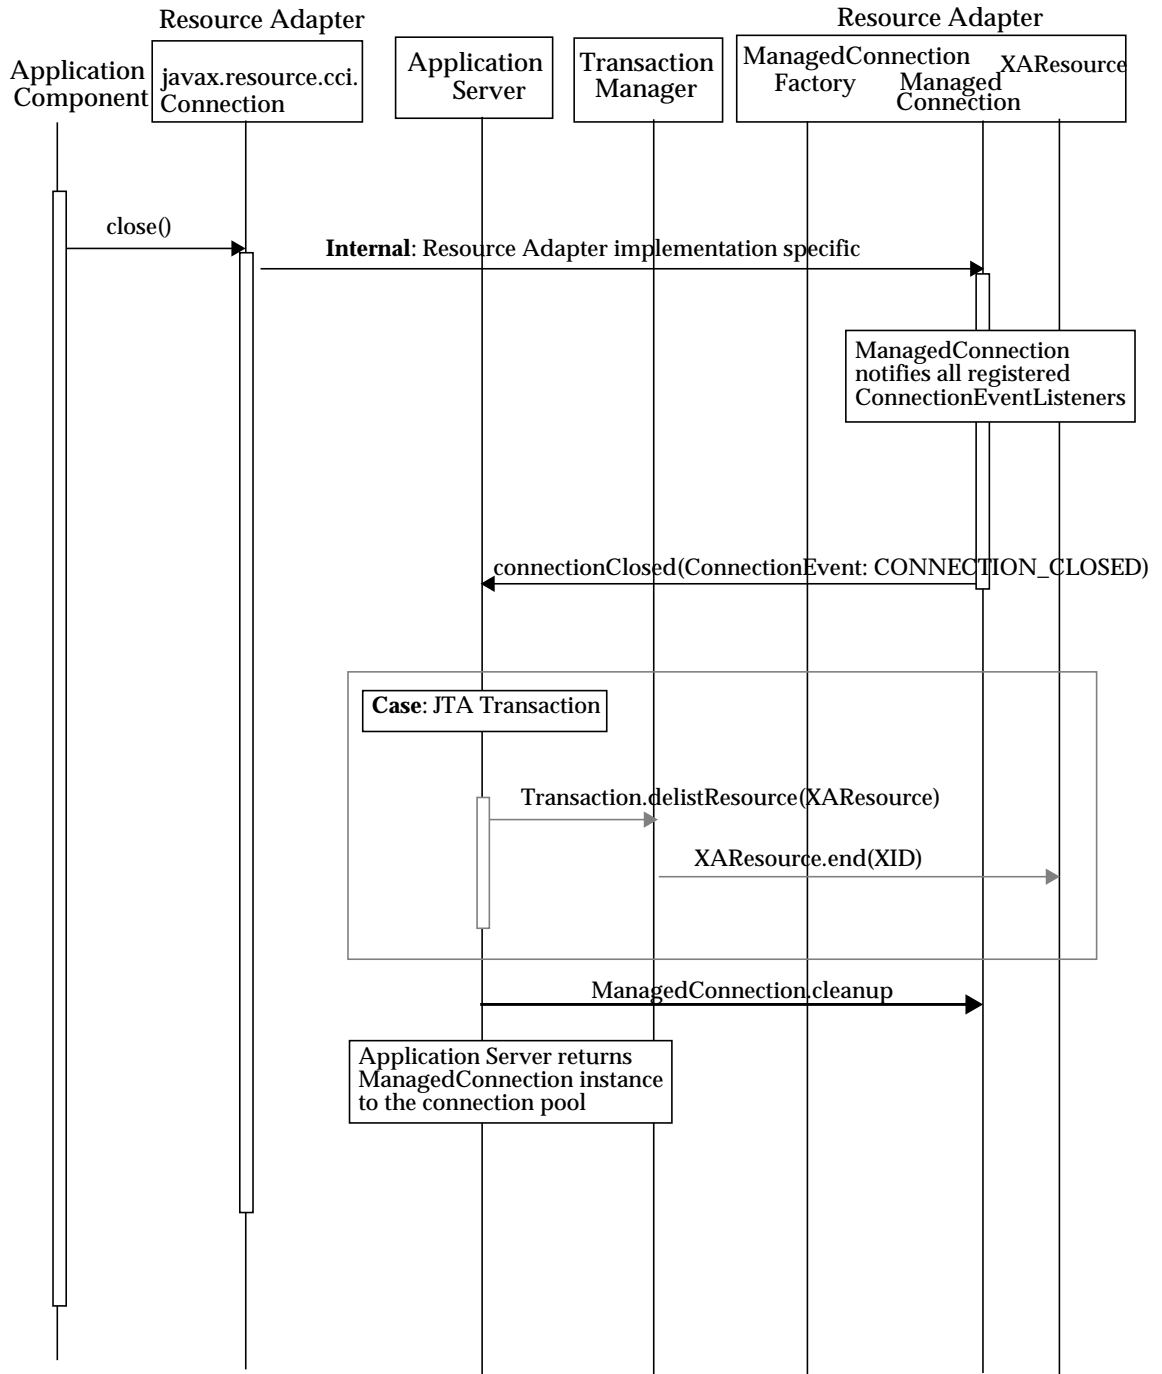
The application server can also initiate cleanup of a `ManagedConnection` instance when the container terminates the application component instance that has the corresponding connection handle. The application server should call `ManagedConnection.cleanup` to initiate the connection cleanup. After the cleanup, the application server puts the `ManagedConnection` instance into the pool to serve future allocation requests.

6.8.3.2 Connection Destroy

To manage the size of the connection pool, the application server can call `ManagedConnection.destroy` method to destroy a `ManagedConnection`. A `ManagedConnection` instance handles this method call by closing the physical connection to the EIS instance and releasing all system resources held by this instance.

The application server also calls `ManagedConnection.destroy` when it receives a connection error event notification that signals a fatal error on the physical connection.

FIGURE 6-7 OID: Connection Event Notification



6.9 Architecture: Non-managed Environment

The connection management contract enables a resource adapter to be used in a two-tier application directly from an application client.

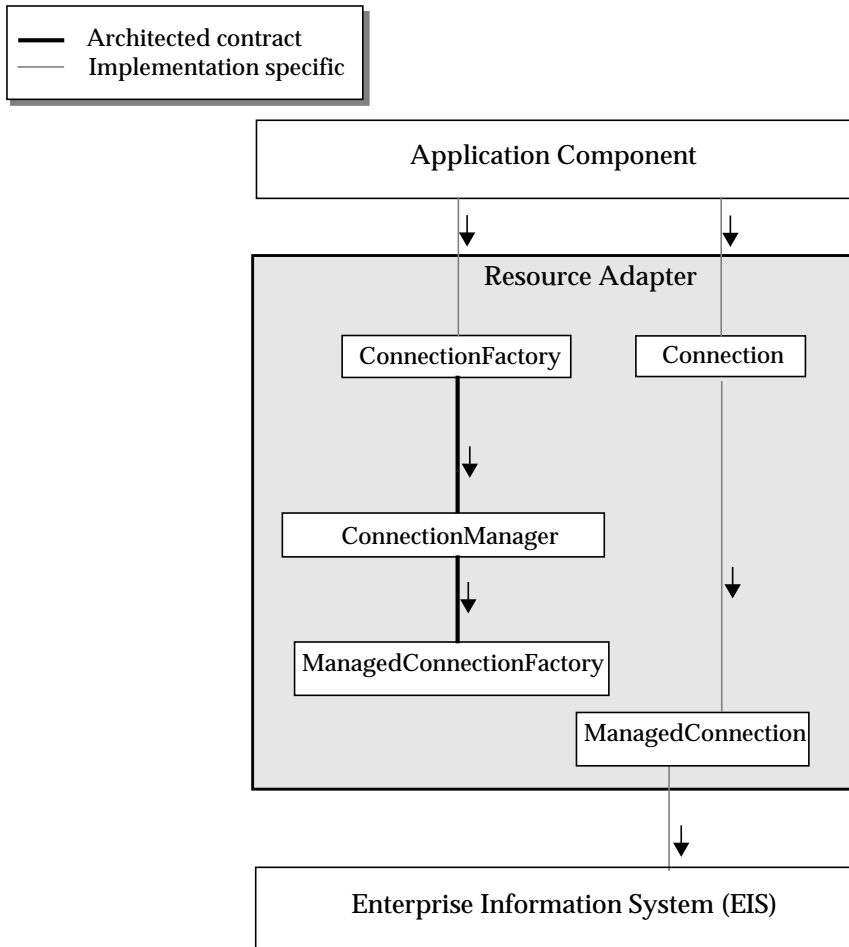
In a non-managed application scenario, the `ConnectionManager` implementation class may be provided either by a resource adapter (as a default `ConnectionManager` implementation) or by application developers. Note that a default implementation of the `ConnectionManager` should be defined for a resource adapter (in terms of the functionality provided and third-party components added) only at development time.

The default `ConnectionManager` instance interposes on the connection request and delegates the request to the `ManagedConnectionFactory` instance. The `ManagedConnectionFactory` creates a physical connection (represented by a `ManagedConnection` instance) to the underlying EIS. The `ConnectionManager` gets a connection handle (of type `javax.resource.cci.Connection` for CCI) from the `ManagedConnection` and returns it to the connection factory. The connection factory returns the connection handle to the application.

A resource adapter supports interactions (shown as light shaded lines in FIGURE 6-8) between its internal objects in an implementation-specific way. For example, a resource adapter can use the connection event listening mechanism as part of its

ManagedConnection implementation for connection management. However, the resource adapter is not required to use the connection event mechanism to drive its internal interactions.

FIGURE 6-8 Architecture Diagram: Non-Managed application scenario



6.9.1 Scenario: Programmatic Access to ConnectionFactory

To maintain the consistency of the application programming model across both managed and non-managed environments, application code should use the `JNDI` namespace to look-up a connection factory instance.

The following code extract shows how an application client accesses a connection factory instance in a non-managed environment. The code extract does not show the use of JNDI. It is used as an example to illustrate the use of `ManagedConnectionFactory` and `ConnectionFactory` interfaces in the application code. Refer to section Section 17.5, “JNDI Configuration and Lookup” on page 17-13 for details on JNDI configuration and lookup.

```
// Application Client Code
// Create an instance of the ManagedConnectionFactory
// implementation class passing in initialization parameters
// (if any) for this instance
com.myeis.ManagedConnectionFactoryImpl mcf =
    new com.myeis.ManagedConnectionFactoryImpl(...);

// Set properties on the ManagedConnectionFactory instance
// Note: Properties are defined on the implementation class
// and not on the javax.resource.spi.ManagedConnectionFactory
// interface
mcf.setServerName(...);
mcf.setPortNumber(...);

// set remaining properties
...

// Get access to connection factory. The ConnectionFactory
instance
// gets initialized with the default ConnectionManager provided
// by the resource adapter
javax.resource.cci.ConnectionFactory cxf =
    (javax.resource.cci.ConnectionFactory)
    mcf.createConnectionFactory();

// Get a connection using the ConnectionFactory instance
javax.resource.cci.Connection cx = cxf.getConnection(...);

// use connection to access the underlying EIS instance
...

// Close the connection
cx.close();
```

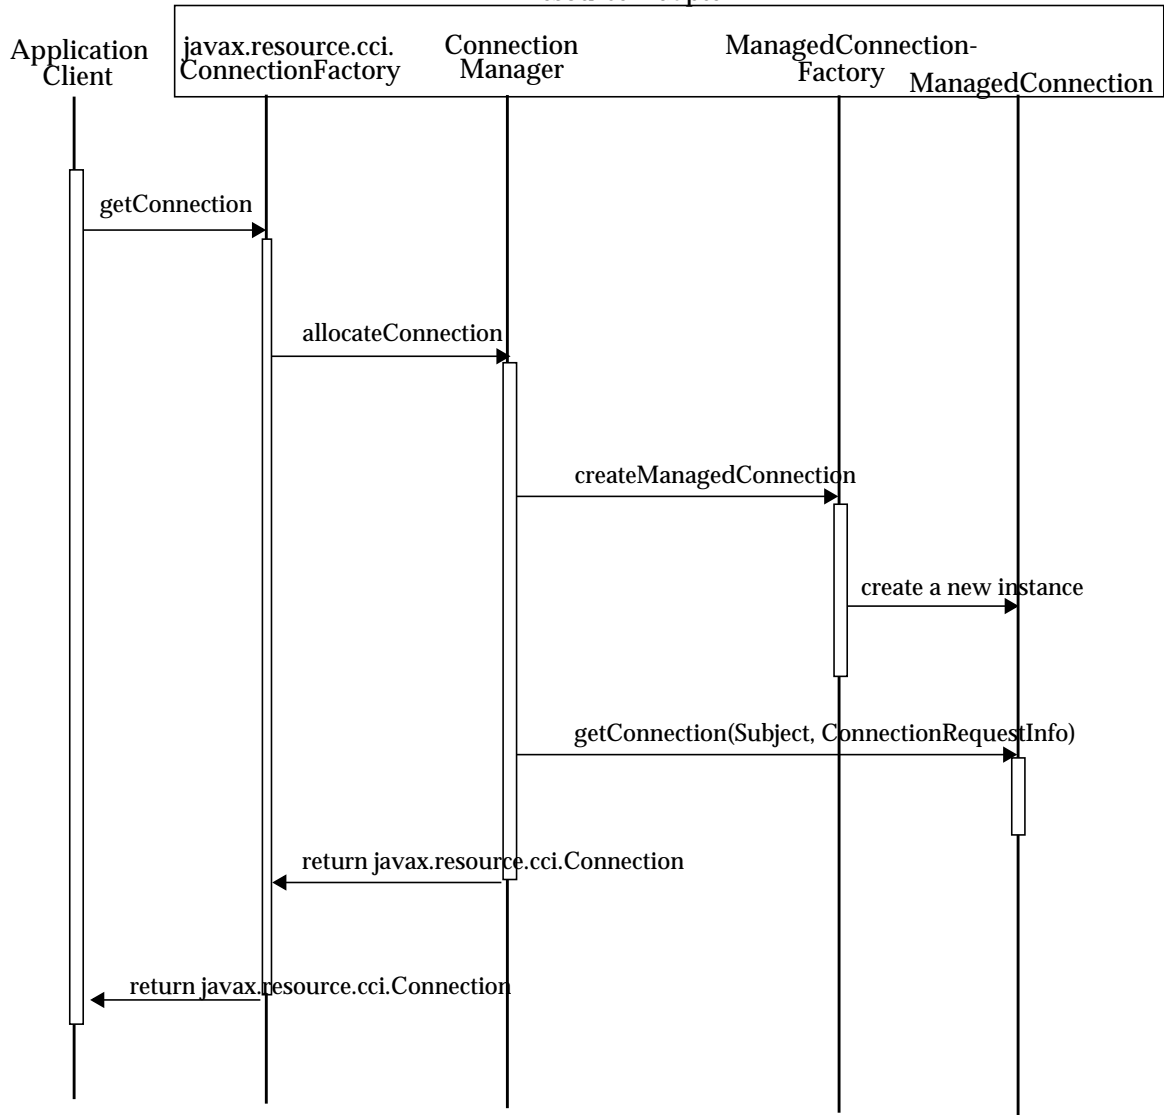

6.9.2 Scenario: Connection Creation in Non-managed Application Scenario

The following object interactions are involved in the scenario shown in FIGURE 6-9 on page 48:

- The application client calls a method on the `javax.resource.cci.ConnectionFactory` instance, returned from the JNDI lookup, to get a connection to the underlying EIS instance.
- The `ConnectionFactory` instance delegates the connection request from the application to the default `ConnectionManager` instance. The resource adapter provides the default `ConnectionManager` implementation.
- The `ConnectionManager` instance creates a new physical connection to the underlying EIS instance by calling the `ManagedConnectionFactory.createManagedConnection` method.
- The `ManagedConnectionFactory` instance handles the `createManagedConnection` method by creating a new physical connection to the underlying EIS, represented by a `ManagedConnection` instance. The `ManagedConnectionFactory` uses the security information, passed as a `Subject` instance, any `ConnectionRequestInfo` instance, and its configured set of properties, such as port number, server name, to create a new `ManagedConnection` instance.
- The `ManagedConnectionFactory` initializes the state of the created `ManagedConnection` instance and returns it to the default `ConnectionManager` instance.
- The `ConnectionManager` instance calls the `ManagedConnection.getConnection` method to get an application-level connection handle. Calling the `getConnection` method does not necessarily create a new physical connection to the EIS instance. Calling `getConnection` produces a temporary handle that is used by an application to access the underlying physical connection. The actual underlying physical connection is represented by a `ManagedConnection` instance.

- The `ConnectionManager` instance returns the connection handle to the `ConnectionFactory` instance, which then returns the connection to the application that initiated the connection request.

FIGURE 6-9 OID: Connection Creation in a Non-managed Application Scenario
Resource Adapter



6.10 Requirements

The following section outlines requirements for the connection management contract.

6.10.1 Resource Adapter

The requirements for a resource adapter are as follows:

- A resource adapter must provide implementations of the following interfaces:
 - `javax.resource.spi.ManagedConnectionFactory`
 - `javax.resource.spi.ManagedConnection`
 - `javax.resource.spi.ManagedConnectionMetaData`
- The `ManagedConnection` implementation provided by a resource adapter must use the following interface and classes to provide support to an application server for connection management and transaction management, as explained later:
 - `javax.resource.spi.ConnectionEvent`
 - `javax.resource.spi.ConnectionEventListener`

To support non-managed environments, a resource adapter is not required to use the above two interfaces to drive its internal object interactions.

- A resource adapter must provide support for basic error logging and tracing by implementing the following methods:
 - `ManagedConnectionFactory.set/getLogWriter`
 - `ManagedConnection.set/getLogWriter`
- A resource adapter must provide a default implementation of the `javax.resource.spi.ConnectionManager` interface. The implementation class comes into play when a resource adapter is used in a non-managed two-tier application scenario. In an application server-managed environment, the resource adapter must not use the default `ConnectionManager` implementation class.

A default implementation of `ConnectionManager` enables the resource adapter to provide services specific to itself. These services can include connection pooling, error logging and tracing, and security management. The default `ConnectionManager` delegates to the `ManagedConnectionFactory` the creation of physical connections to the underlying EIS.

- In a managed environment, with the exception of application client containers, a resource adapter must not asynchronously (that is, using a separate thread other than the application thread) call application objects other than message-driven beans. However, this restriction does not apply to a non-managed scenario, as well as application client containers. A resource adapter deployer may use the

`ResourceAdapter` `JavaBean` to configure the resource adapter during its deployment to set the desired behavior, based on the requirements of the deployment environment.

- A resource adapter is not allowed to support its own internal connection pooling in a managed environment. In this case, the application server is responsible for connection pooling. However, a resource adapter may multiplex connections (one or more `ManagedConnection` instances per physical connection) over a single physical pipe transparent to the application server and components.

In a non-managed two tier application scenario, a resource adapter is allowed to support connection pooling internal to the resource adapter.

6.10.2 Application Server

The requirements for an application server are as follows:

- An application server must use the interfaces defined in the connection management contract to use services provided by a resource adapter. These interfaces are as follows:
 - `javax.resource.spi.ManagedConnectionFactory`
 - `javax.resource.spi.ManagedConnection`
 - `javax.resource.spi.ManagedConnectionMetaData`
- An application server must provide an implementation of the `javax.resource.spi.ConnectionManager` interface. This implementation should not be specific to any particular type of resource adapter, EIS, or connection factory interface.
- An application server must implement the `javax.resource.spi.-ConnectionEventListener` interface and to register `ConnectionEventListener` with a resource adapter to get connection-related event notifications. An application server uses these event notifications to do its pool management, transaction management, and connection cleanup.
- An application server must use the following interfaces (supported by the resource adapter) to provide basic error logging and tracing for its configured set of resource adapters:
 - `ManagedConnectionFactory.set/getLogWriter`
 - `ManagedConnection.set/getLogWriter`
- An application server must use the `javax.resource.spi.ConnectionManager` hook-in mechanism to provide its specific quality of services. The connector architecture does not specify the set of services the application server provides, nor does it specify how the application server implements these services.

Transaction Management

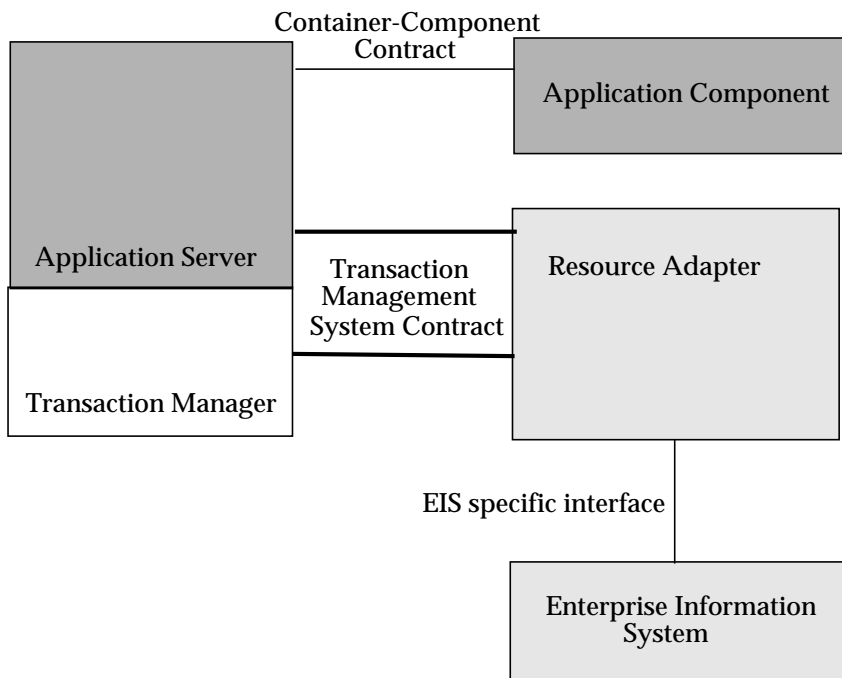
This chapter specifies the transaction management contract between an application server (and supported transaction manager) and an EIS resource manager.

This chapter focuses only on the system-level aspects of transaction management. The J2EE component model specifications describe the application level transaction model. For example, the EJB specification (Section 1., “Enterprise JavaBeans (EJB TM) specification, version 2.1:” on page F-1) specifies the transaction model for EJB components.

7.1 Overview

FIGURE 7-1 shows an application component deployed in a container provided by an application server. The application component performs transactional access to multiple resource managers. The application server uses a transaction manager that takes the responsibility of managing transactions across multiple resource managers.

FIGURE 7-1 Transaction Management Contract



A resource manager can support two types of transactions:

- A transaction that is controlled and coordinated by a transaction manager external to the resource manager. This document refers to such a transaction as JTA or XA transaction.
- A transaction that is managed internal to a resource manager. The coordination of such transactions involves no external transaction managers. This document refers to such transactions as RM local transactions (or local transactions).

A transaction manager coordinates transactions across multiple resource managers. It also provides additional low-level services that enable transactional context to be propagated across systems. The services provided by a transaction manager are not visible directly to the application components.

The connector architecture defines a transaction management contract between an application server and a resource adapter and its underlying resource manager. The transaction management contract has two parts, depending on the type of transaction:

- a JTA `javax.transaction.xa.XAResource` based contract between a transaction manager and a resource manager
- a local transaction management contract

These contracts enable an application server to provide the infrastructure and runtime environment for transaction management. Application components rely on this transaction infrastructure to support their component-level transaction model.


7.2 Transaction Management Scenarios

The following section uses a set of scenarios to present an overview of the transaction management architecture.

7.2.1 Transactions across multiple Resource Managers

In FIGURE 7-2, an application client invokes EJB component X. EJB X accesses transaction programs managed by a TP system and calls EJB Y to access an ERP system.

FIGURE 7-2 Scenario: Transactions Across Multiple Resource Managers



The application server uses a transaction manager to support a transaction management infrastructure that enables an application component to perform transactional access across multiple EIS resource managers. The transaction manager manages transactions across multiple resource managers and supports propagation of the transaction context across distributed systems.

The transaction manager supports a JTA `XAResource`-based transaction management contract with a resource adapter and its underlying resource manager. The ERP system supports JTA transactions by implementing a `XAResource` interface through its resource adapter. The TP system also implements an `XAResource` interface. This interface enables the two resource managers to participate in transactions that are coordinated by an external transaction manager. The transaction manager uses the `XAResource` interface to manage transactions across the two underlying resource managers.

The EJBs X and Y access the ERP and TP system using the respective client access API for the two systems. Behind the scenes, the application server enlists the connections to both systems, obtained from their respective resource adapters, as part of the transaction. When the transaction commits, the transaction manager performs a two-phase commit protocol across the two resource managers, ensuring that all read/write access to resources managed by both the TP system and ERP system is either entirely committed or entirely rolled back.

7.2.2 Local Transaction Management

The transactions are demarcated either by the container (called container-managed demarcation) or by a component (called component-managed demarcation). In component-managed demarcation, an application component can use the JTA `UserTransaction` interface or a transaction demarcation API specific to an EIS (for example, JDBC transaction demarcation using `java.sql.Connection`).

The EJB specification requires an EJB container to support both container-managed and component-managed transaction demarcation models. The JSP and servlet specifications require a web container to support component-managed transaction demarcation.

If multiple resource managers participate in a transaction, the EJB container uses a transaction manager to coordinate the transaction. The contract between the transaction manager and resource manager is defined using the `XAResource` interface.

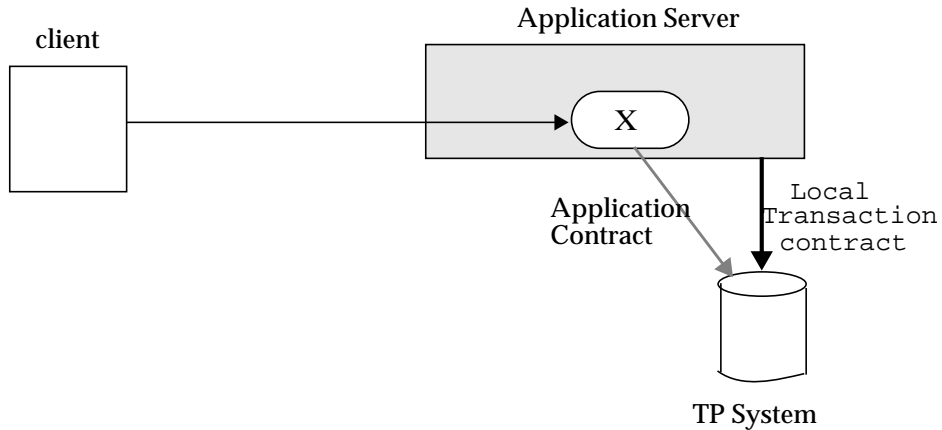
If a single resource manager instance participates in a transaction (either component-managed or container-managed), the container has two choices:

- It uses the transaction manager to manage this transaction. The transaction manager uses one-phase commit-optimization, described in Section 7.6.1 “Scenarios Supported”, to coordinate the transaction for this single resource manager instance.
- The container lets the resource manager coordinate this transaction internally without involving an external transaction manager.

If an application accesses a single resource manager using a XA transaction, it has more performance overhead compared to using a local transaction. The overhead is due to the involvement of an external transaction manager in the coordination of the XA transaction.

To avoid the overhead of using a XA transaction in a single resource manager scenario, the application server may optimize this scenario by using a local transaction instead of a XA transaction. This scenario is shown in FIGURE 7-3.

FIGURE 7-3 Scenario: Local Transaction on a Single Resource Manager

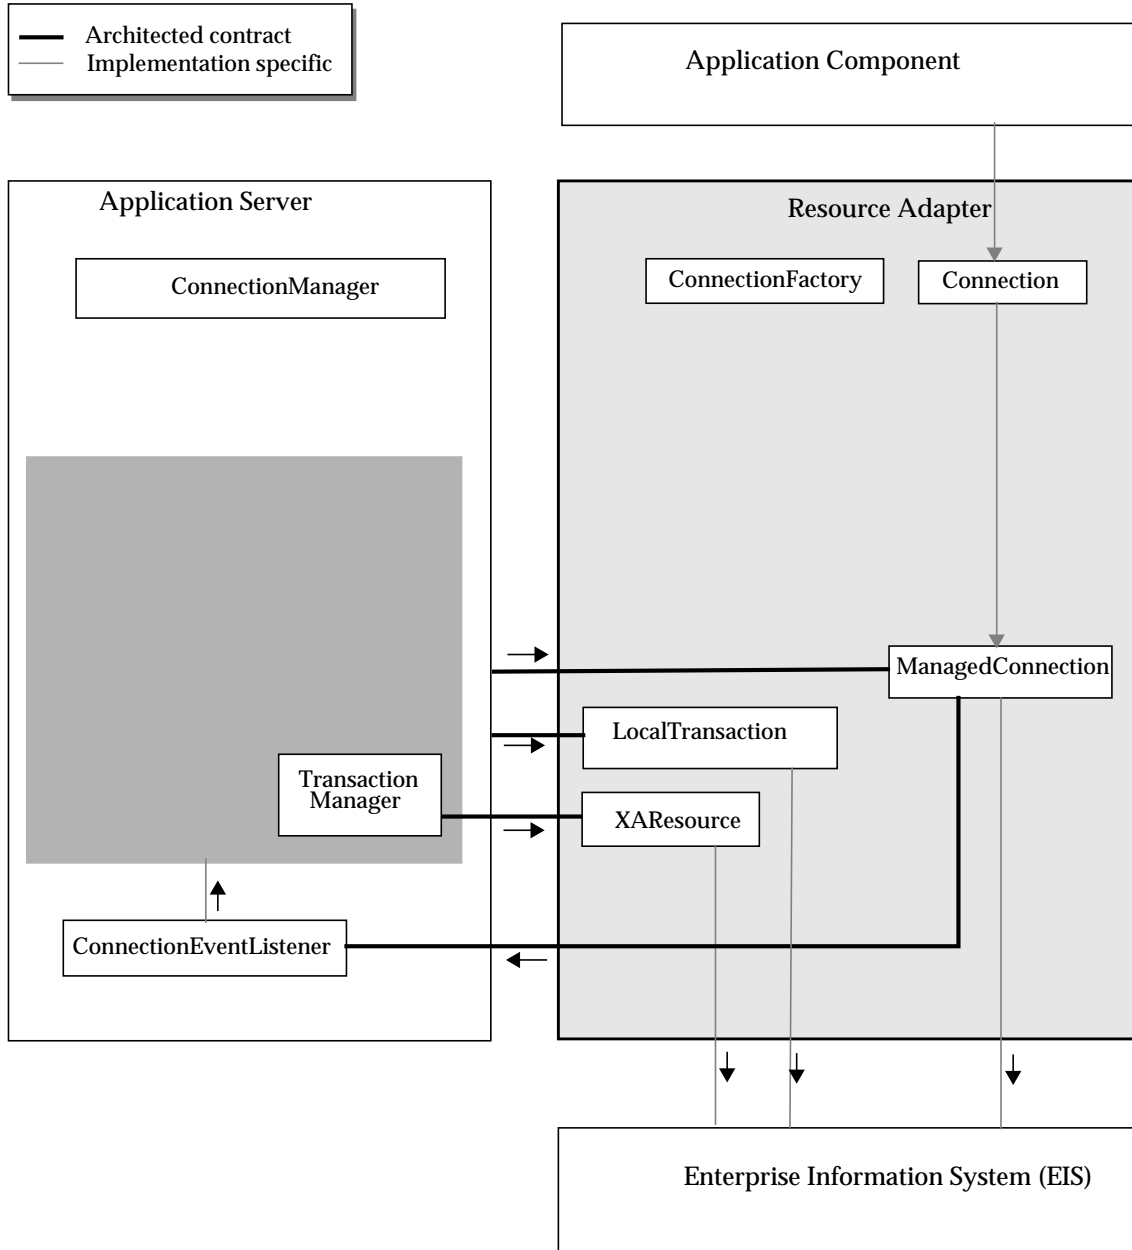


7.3 Transaction Management Contract

This section specifies the transaction management contract. The transaction management contract builds on the connection management contract specified in Chapter 6, “Connection Management.”

FIGURE 7-4 shows the interfaces and flows in the transaction management contract. It does not show the interfaces, classes, and flows that are the same in the connection management contract.

FIGURE 7-4 Architecture Diagram: Transaction Management



7.3.1 Interface: ManagedConnection

The `javax.resource.spi.ManagedConnection` instance represents a physical connection to an EIS and acts as a factory for connection handles.

The following code extract shows the methods on the `ManagedConnection` interface that are defined specifically for the transaction management contract:

```
public interface javax.resource.spi.ManagedConnection {
    public XAResource getXAResource() throws ResourceException;
    public LocalTransaction getLocalTransaction()
        throws ResourceException;
    ...
}
```

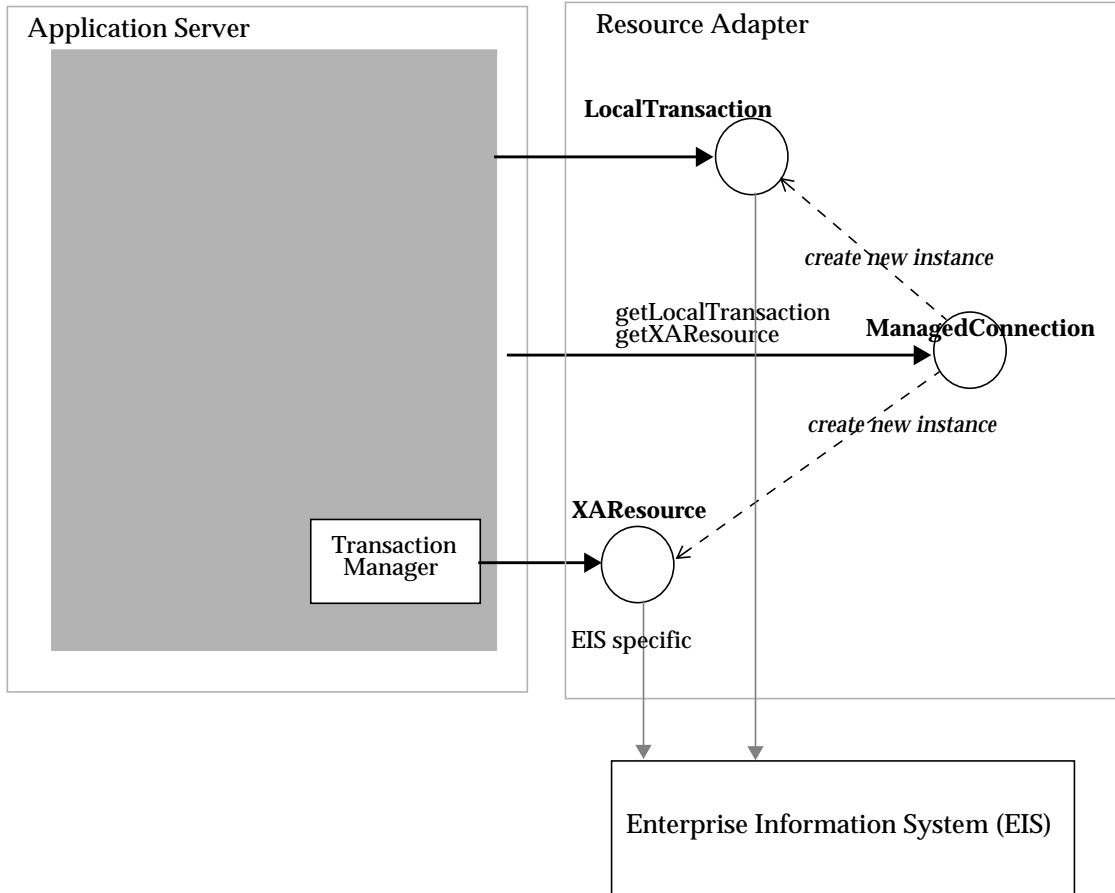
A `ManagedConnection` instance provides access to a pair of interfaces:

`javax.transaction.xa.XAResource` and
`javax.resource.spi.LocalTransaction`.

Depending on the transaction support level of a resource adapter, these methods should raise appropriate exceptions. For example, if the transaction support level for a resource adapter is `NoTransaction`, an invocation of `getXAResource` method should throw a `ResourceException`. Refer to Chapter 19, “Exceptions” for details on the exception hierarchy.

FIGURE 7-5 illustrates this concept:

FIGURE 7-5 ManagedConnection Interface for Transaction Management



The transaction manager uses the `XAResource` interface to associate and dissociate a transaction with the underlying EIS resource manager instance and to perform a two-phase commit protocol. The transaction manager does not directly use the `ManagedConnection` interface. The next section describes the `XAResource` interface in more detail.

The application server uses the `LocalTransaction` interface to manage local transactions.

7.3.2 Interface: XAResource

The `javax.transaction.xa.XAResource` interface is a Java mapping of the industry standard XA interface based on X/Open CAE specification (Section 4., “X/Open CAE Specification -- Distributed Transaction Processing: the XA specification, X/Open document” on page F-1).

The following code extract shows the interface specification for the `XAResource` interface. For more details and API documentation, refer to the JTA (Section 2., “Java Transaction API (JTA) specification, version 1.0.1B” on page F-1) and XA (Section 4., “X/Open CAE Specification -- Distributed Transaction Processing: the XA specification, X/Open document” on page F-1) specifications:

```
public interface javax.transaction.xa.XAResource {
    public void commit(Xid xid, boolean onePhase) throws
XAException;
    public void end(Xid xid, int flags) throws XAException;
    public void forget(Xid xid) throws XAException;
    public int prepare(Xid xid) throws XAException;
    public Xid[] recover(int flag) throws XAException;
    public void rollback(Xid xid) throws XAException;
    public void start(Xid xid, int flags) throws XAException;
}
```

7.3.2.1 Implementation

A resource adapter for an EIS resource manager implements the `XAResource` interface. This interface enables the resource manager to participate in transactions that are controlled and coordinated by an external transaction manager. The transaction manager uses the `XAResource` interface to communicate transaction association, completion, and recovery to the resource manager.

A resource adapter typically implements the `XAResource` interface using a low-level library available for the underlying EIS resource manager. This low-level library either supports a native implementation of the XA interface or provides a proprietary vendor-specific interface for transaction management.

A resource adapter is responsible for maintaining a 1-1 relationship between the `ManagedConnection` and `XAResource` instances. Each time a `ManagedConnection.getXAResource` method is called, the same `XAResource` instance has to be returned.

A transaction manager can use any `XAResource` instance (if it refers to the proper resource manager instance) to initiate transaction completion. The `XAResource` instance used during the transaction completion process need not be the one initially enlisted with the transaction manager for this transaction.

7.3.3 Interface: LocalTransaction

The following code extract shows the `javax.resource.spi.LocalTransaction` interface:

```
public interface javax.resource.spi.LocalTransaction {
    public void begin() throws ResourceException;
    public void commit() throws ResourceException;
    public void rollback() throws ResourceException;
}
```

A resource adapter implements the `LocalTransaction` interface to provide support for local transactions that are performed on the underlying resource manager. An application server uses the `LocalTransaction` interface to manage local transactions for a resource manager.

Section 7.7.1 “Interface: Local Transaction” has more details on the local transaction management contract.

7.4 Relationship to JTA and JTS

The Java Transaction API (JTA) (Section 2., “Java Transaction API (JTA) specification, version 1.0.1B” on page F-1) is a specification of interfaces between a transaction manager and the other parties involved in a distributed transaction processing system: application programs, resource managers, and an application server.

The Java Transaction Service (JTS) API is a Java binding of the Common Object Request Broker Architecture (CORBA) Object Transaction Service (OTS) 1.1 specification. JTS provides transaction interoperability using the standard Internet Inter-ORB Protocol (IIOP) for transaction propagation between servers. The JTS API is intended for vendors who implement transaction processing infrastructure for enterprise middleware. For example, an application server vendor can use a JTS implementation as the underlying transaction manager.

7.4.1 JTA Interfaces

The application server uses the `javax.transaction.TransactionManager` and `javax.transaction.Transaction` interfaces, specified in the JTA specification, for its contract with the transaction manager.

The application server uses the `javax.transaction.TransactionManager` interface to control the transaction boundaries on behalf of the application components that are being managed by the application server. For example, an EJB container manages the transaction states for transactional EJB components. The EJB container uses the `TransactionManager` interface to demarcate transaction boundaries based on the calling thread's transaction context.

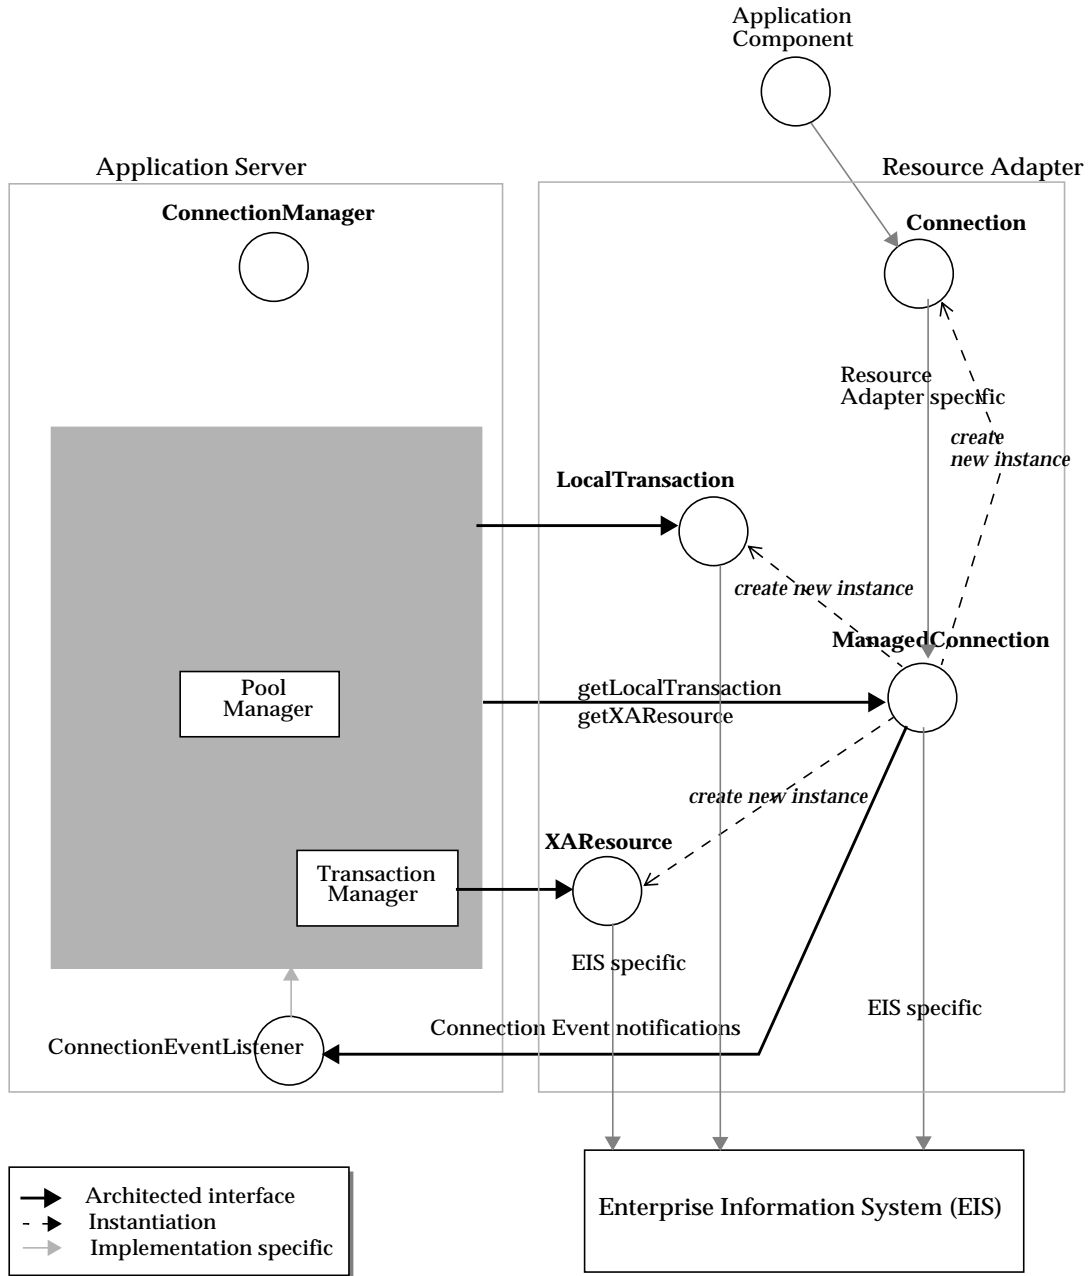
The application server also uses the `javax.transaction.Transaction` interface to enlist and delist transactional connections with the transaction manager. This enables the transaction manager to coordinate transactional work performed by all enlisted resource managers within a transaction.

7.5 Object Diagram

FIGURE 7-6 shows the object instances and their interactions related to transaction management.

Since the transaction management contract builds upon the connection management contract, the following diagram does not show object interactions that have already been discussed in Chapter 6, “Connection Management.”

FIGURE 7-6 Object Diagram: Transaction Management



7.6 XAResource-based Transaction Contract

The following section specifies detailed requirements for a resource manager and a transaction manager for the `XAResource`-based transaction management contract. In this section, the following abbreviations are used: RM (Resource Manager), TM (Transaction Manager), 1PC (one phase commit protocol), and 2PC (two phase commit protocol).

7.6.1 Scenarios Supported

The following table specifies various transaction management scenarios and mentions whether these scenarios are within the scope of the connector architecture.

TABLE 7-1 Transaction Management Scenarios

Description	Supported / NotSupported
TM does two-phase commit (2PC) on RMs that support two phase commit (as defined in RM's requirements for XAResource implementation in the subsection below) Examples of RM: Oracle and DB2 installations that support 2PC in their XAResource implementations.	Supported based on TM's requirement to be JTA/JTS and X/Open compliant, and RM's support for 2PC in the XAResource interface.
TM does one-phase commit (1PC) optimization on the only RM involved in a transaction. RM supports 2PC in its XAResource implementation (as defined in RM's requirements for the XAResource implementation in the subsection below). Example of RM: DB2 installation that supports 2PC in its XAResource implementation.	Supported based on TM's requirement to be JTA/JTS and X/Open compliant, and RM's support for the XAResource interface. Note: This scenario will also work if TM does 2PC on RM.

TABLE 7-1 Transaction Management Scenarios

Description	Supported / NotSupported
TM does one-phase commit optimization on the only RM involved in a transaction. RM does not support 2PC but supports 1PC in its <code>XAResource</code> implementation. Example of RM: ERP system or mainframe TP system that does not support 2PC, but implements 1PC in its <code>XAResource</code> implementation as defined in the RM's requirements for 1PC.	Supported by requiring that TM must support 1PC optimization. A successful transaction coordination of 1PC only RM comes as a result of required 1PC optimization for a TM. The rationale behind this requirement is that this scenario will be an important scenario to support for the connector architecture.
TM does last-resource commit optimization across multiple RMs involved in a transaction—RMs that support 2PC (for example: Oracle and DB2) and a single RM that supports only 1PC (for example: an ERP system).	Out of the scope of the connector architecture specification
More than one RM that support only 1PC involved in a transaction with none or multiple 2PC enabled RMs	Out of the scope of the connector architecture specification

7.6.2 Resource Adapter Requirements

The connector architecture does not require that all resource adapters must support JTA `XAResource` based transaction contract.

If a resource adapter decides to support a `XAResource` based contract, then the connector architecture places certain requirements on a resource adapter and its underlying resource manager (RM).

The following requirements refer to a resource adapter and its resource manager together as a resource manager (RM). The division of responsibility between a resource adapter and its underlying resource manager for supporting the transaction contract is implementation-specific and is out of the scope of the connector architecture.

These requirements assume that a transaction manager (TM) supports JTA/XA and JTS requirements.

The following set of requirements are based on the JTA and XA specifications and should be read in conjunction with these specifications. These detailed requirements are included in this document to clearly specify the requirements from the connector architecture perspective.

7.6.2.1 General

- If a RM supports an `XAResource` contract, then it must support the one-phase commit protocol by implementing `XAResource.commit` when the boolean flag `onePhase` is set to `True`. The RM is not required to implement the two-phase commit protocol support in its `XAResource` implementation.
- However, if a RM supports the two-phase commit protocol, then the RM must use the `XAResource` interface for supporting the two-phase commit protocol.
- An RM is allowed to combine the implementation of 2PC protocol with 1PC optimization by implementing `XAResource.commit(onePhase=True)` in addition to the implementation requirements for 2PC.

7.6.2.2 One-phase Commit

- An RM should allow `XAResource.commit(onePhase=True)` even if it has not received `XAResource.prepare` for the transaction branch.
- If the RM fails to commit a transaction during a 1PC commit, then the RM should throw one of the `XA_RB*` exceptions. In the exception case, an RM should roll back the transaction branch's work and release all held RM resources.
- The RM is responsible for deciding the outcome of a transaction branch on a `XAResource.commit` method. The RM can discard knowledge of the transaction branch once it returns from the `commit` call.
- The RM is not required to maintain knowledge of transaction branches to support failure recovery for the TM.
- If an `XAResource.prepare` method is called on an RM that supports only one-phase commit, then the RM should throw an `XAException` with `XAER_PROTO` or `XA_RB*` flag.
- The RM should return an empty list of XIDs for `XAResource.recover`, because the RM is not required to maintain stable knowledge about transaction branches.

7.6.2.3 Two-phase Commit

- If the RM supports 2PC, then its implementation of 2PC must be compliant with the 2PC protocol definition with presumed rollback as specified in the OSI TP (Transaction Protocol defined by ISO (ISO92)) specification.

- The RM must implement the `XAResource.prepare` method and must be able to report whether it can guarantee its ability to commit the transaction branch. If the RM reports that it can, the RM must hold and record in a stable way all the resources necessary to commit the branch. It must hold all these resources until the TM directs it to commit or rollback the branch.
- An RM that reports a heuristic completion to the TM must not discard its knowledge of the transaction branch. The RM should discard its knowledge of the branch only when the TM calls `XAResource.forget`. The RM must notify the TM of all heuristic decisions.
- On the TM's `XAResource.commit` and `XAResource.rollback` calls, the RM is allowed to report through an `XAException` that it has heuristically completed the transaction branch. This feature is optional.

A TM supporting the OSI TP specification uses the one-phase commit optimization by default to manage an RM that is the only resource involved in the transaction. The mechanism to identify to the TM a particular RM that only supports 1PC is beyond the scope of this specification.

7.6.2.4 Transaction Association and Calling Protocol

- The RM `XAResource` implementation must support `XAResource.start` and `XAResource.end` for association and disassociation of a transaction, as represented by, unique `XID`, with recoverable units of work being done on the RM.
- The RM must ensure that the TM invokes `XAResource` calls in the legal sequence, and must return `XAER_PROTO` or another suitable error if the caller TM violates the state tables, as defined in Chapter 6 of the XA specification (Section 2., “Java Transaction API (JTA) specification, version 1.0.1B” on page F-1).

7.6.2.5 Unilateral Roll-back

- The RM need not wait for global transaction completion to report an error. The RM can return a rollback-only flag as a result of any `XAResource.start` or `XAResource.end` call. This can happen anytime except after a successful `prepare`.
- The RM is allowed to unilaterally rollback and forget a transaction branch any time before it prepares it.

7.6.2.6 Read-Only Optimization

- Support for read-only optimization is optional for RM implementation. An RM can respond to the TM's request to prepare a transaction by asserting that the RM was not asked to update shared resources in this transaction branch. This response concludes the RM's involvement in the transaction, and the RM can release all resources and discard its knowledge of the transaction.

7.6.2.7 XID Support

- The RM must accept XIDs from TMs. The RM is responsible for using the XID to maintain an association between a transaction branch and recoverable units of work done by the application programs.
- The RM must not alter in any way the bits associated in the data portion of an XID. For example, if an RM remotely communicates an XID, it must ensure that the data bits of the XID are not altered by the communication process.

7.6.2.8 Support for Failure Recovery

- A full JTA compliant `XAResource` implementation that supports 2PC must maintain the status of all transaction branches in which it is involved. After responding affirmatively to the TM `prepare` call, an RM should not erase its knowledge of the branch or of the work done in support of the branch until it successfully receives a TM's invocation to commit or rollback the branch.
- If an RM that supports 2PC heuristically completes a branch, it should not forget a branch until the TM explicitly tells it to by calling `XAResource.forget`.
- On the TM's `XAResource.recover` call, an RM that supports 2PC must return a list of all transaction branches that it has prepared or has heuristically completed.
- When an RM recovers from its own failure, it must recover prepared and heuristically completed branches. It should discard its knowledge of all other branches.

7.6.3 Transaction Manager Requirements

The following section specifies requirements of a TM. This section assumes that the TM is compliant with JTA/JTS and X/Open ([Section 4., "X/Open CAE Specification -- Distributed Transaction Processing: the XA specification, X/Open document" on page F-1]) specifications.

7.6.3.1 Interfaces

- The TM must use the `XAResource` interface supported by an RM for transaction coordination and recovery. The TM must be written to handle consistently any information or status that an RM can legally return. The TM must assume that it can support RMs that have different capabilities as allowed by the RM requirements specification section, for instance RMs that make heuristic decisions and RMs that use the read-only optimization. [Requirement derived from Section 7.3, XA specification]

7.6.3.2 XID requirements

- The TM must generate XIDs conforming to the structure defined in section 4.2 on page 19 of the XA specification ([Section 2., “Java Transaction API (JTA) specification, version 1.0.1B” on page F-1]). The generated XIDs must be globally unique and must adequately describe a transaction branch.

7.6.3.3 One-phase Commit Optimization

- The TM must support one-phase commit protocol optimization. The TM uses the 1PC optimization when the TM knows there is only one RM registered in a transaction that is making changes to shared resources. In this optimization, the TM makes its phase 2 commit request to that RM without having made a phase 1 prepare request.
- The TM is not required to record in a stable manner such transactions, and in some failure cases, the TM may not know the outcome of the transaction completion.

7.6.3.4 Implementation Options

- The support of `last-resource optimization` is an implementation-specific option for a TM. A detailed specification of TM and RM requirements for this optimization is outside the scope of the connector architecture.

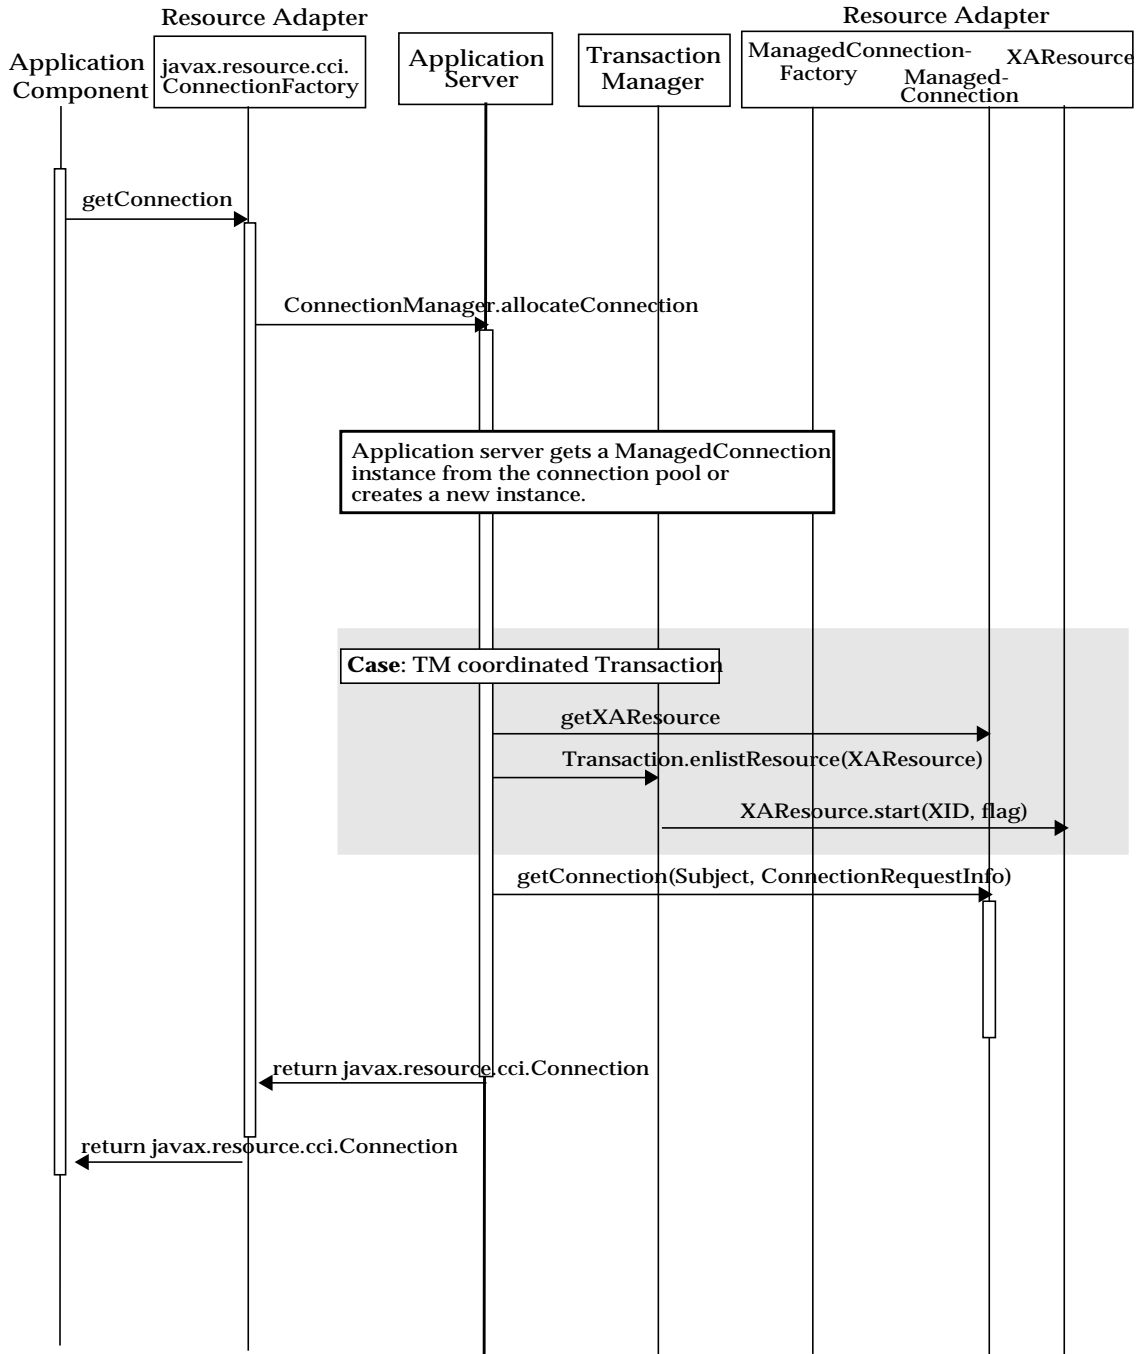
7.6.4 Scenario: Transactional Setup for a ManagedConnection

The following object interactions are involved in the scenario shown in FIGURE 7-7 on page 25.

- The runtime scenario begins with a client method invocation on an EJB instance. This invocation has a transaction context, represented by a unique transaction `xid`, associated with it if the invocation came from a client that was already participating in the transaction. Alternatively, the EJB container starts a transaction before dispatching the client request to the EJB method.
- The EJB instance calls the `getConnection` method on the `ConnectionFactory` instance. The resource adapter delegates the connection request to the application server using the connection management contract. FIGURE 6-6 on page 39 explains this step.
- The application server gains control and handles the connection allocation request.
- To handle the connection allocation request, the application server gets a `Managed-Connection` instance either from the connection pool or creates a new `Managed-Connection` instance. FIGURE 6-6 on page 39 describes this step.
- The application server registers itself as a `ConnectionEventListener` with the `ManagedConnection` instance. This enables the application server to receive notifications for various events on this connection instance. The application server uses these event notifications to manage connection pooling and transactions.
- Based on the current transaction context associated with the connection-requesting thread and the EJB instance, the application server decides whether or not the transaction manager will participate in the coordination of the currently active transaction.
- If the application server decides that the transaction manager will manage the current transaction, it conducts the following transactional setup on the `ManagedConnection` instance:
 - The application server invokes the `ManagedConnection.getXAResource` method to get the `XAResource` instance associated with the `ManagedConnection` instance.
 - The application server enlists the `XAResource` instance with the transaction manager for the current transaction context. The application server uses the `Transaction.enlistResource` method (specified in the JTA specification) to enlist the `XAResource` instance with the transaction manager. This enlistment informs the transaction manager about the resource manager instance participating in the transaction.
 - The transaction manager invokes `XAResource.start` to associate the current transaction with the underlying resource manager instance. This enables the transaction manager to inform the participating resource manager that all units of work performed by the application on the underlying `ManagedConnection` instance should now be associated with this transaction.
- The application server calls the `ManagedConnection.getConnection` method to get a new application-level connection handle. The underlying physical connection is represented by a `ManagedConnection` instance.

- The application server returns the connection handle to the resource adapter. The resource adapter then passes the connection handle to the application component that had initiated the connection request.

FIGURE 7-7 OID: Transactional Setup For Newly Created ManagedConnection Instances



7.6.5 Scenario: Connection Close and JTA Transactional Cleanup

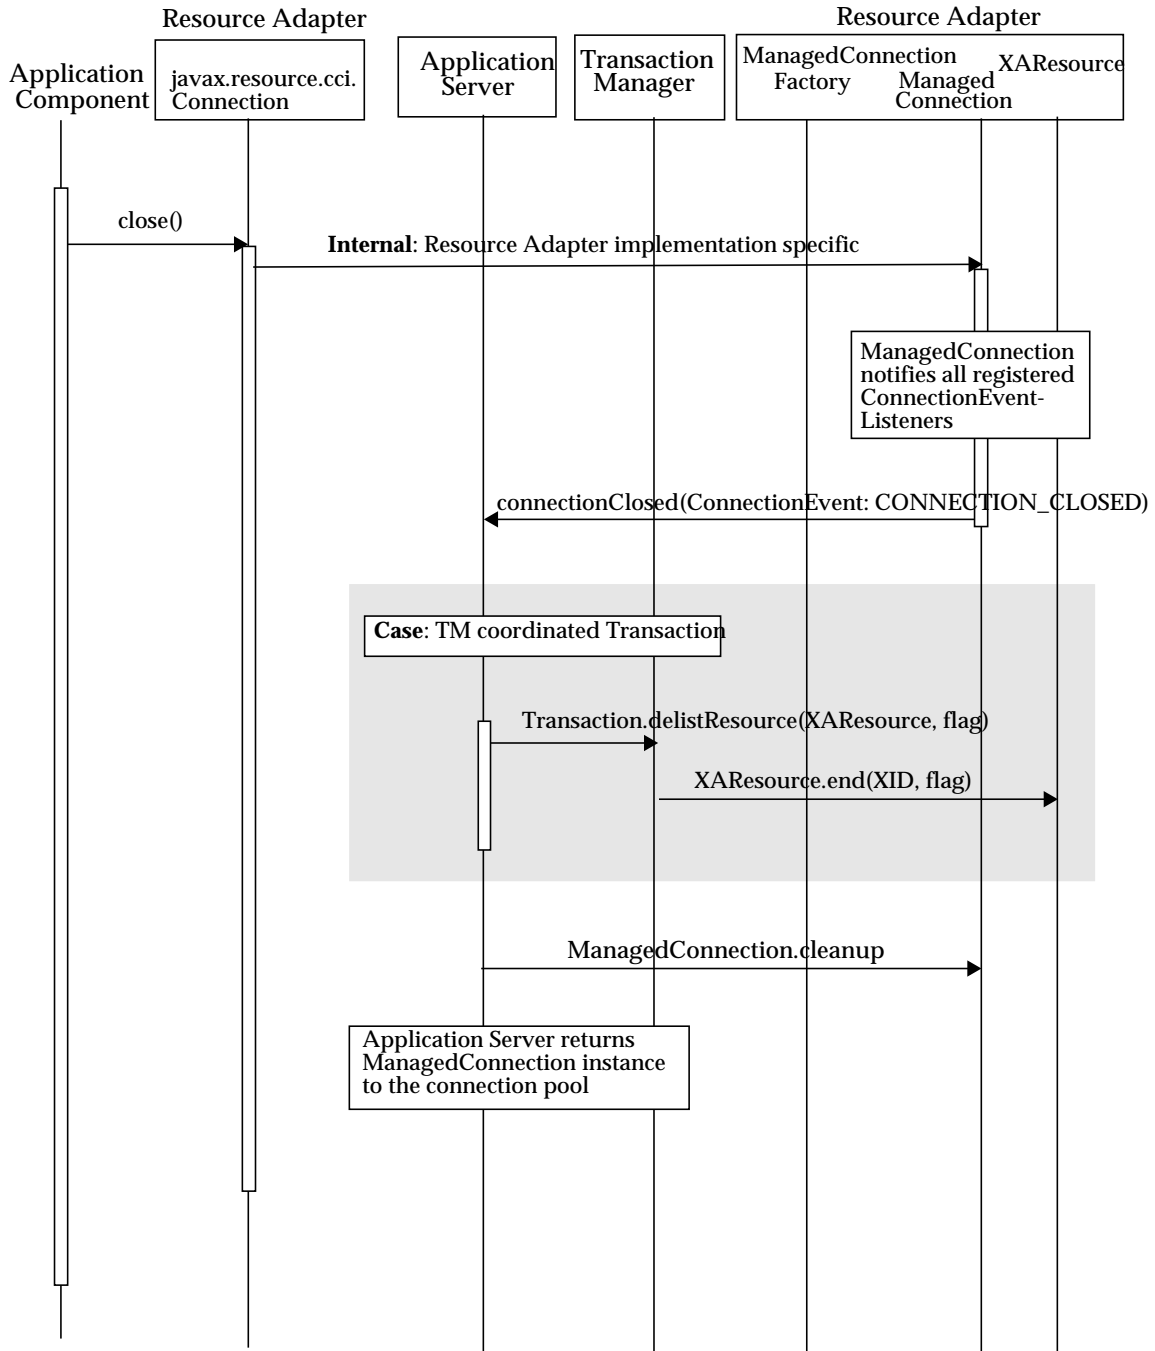
For each `ManagedConnection` instance in the pool, the application server registers a `ConnectionEventListener` instance to receive specific events on the connection. The connection event callback mechanism enables the application server to manage connection pooling and transactions.

FIGURE 6-7 on page 43 describes the following steps when an application component closes a connection:

- The application component releases a `Connection` instance by calling the `close` method. The `Connection` instance delegates the connection close request to its associated `ManagedConnection` instance. A `ManagedConnection` must not alter any state on the physical connection while handling a delegated connection close request.
- The `ManagedConnection` instance notifies all its registered listeners of the application's connection close request using the `ConnectionEventListener.connectionClosed` method. It passes a `ConnectionEvent` instance with the event type set to `CONNECTION_CLOSED`.
- On receiving the connection close notification, the application server performs transactional cleanup for the `ManagedConnection` instance. If the `ManagedConnection` instance was participating in a transaction manager-enlisted JTA transaction, the application server takes the following steps:
 - The application server dissociates the `XAResource` instance, corresponding to the `ManagedConnection` instance, from the transaction manager using the method `Transaction.delistResource`.
 - The transaction manager calls `XAResource.end(Xid, flag)` to inform the resource manager that any further operations on the `ManagedConnection` instance are no longer associated with the transaction, represented by the `Xid` passed in `XAResource.end` call. This method invocation dissociates the transaction from the resource manager instance.
- After the JTA transaction completes, the application server initiates a cleanup of the physical connection instance by calling `ManagedConnection.cleanup` method. After calling the method `cleanup` on the `ManagedConnection` instance, the application server returns the `ManagedConnection` instance to the connection pool.

- The application server can now use the `ManagedConnection` instance to handle future connection allocation requests from either the same or another component instance.

FIGURE 7-8 OID: Connection Close and Transactional Cleanup



7.6.6 OID: Transaction Completion

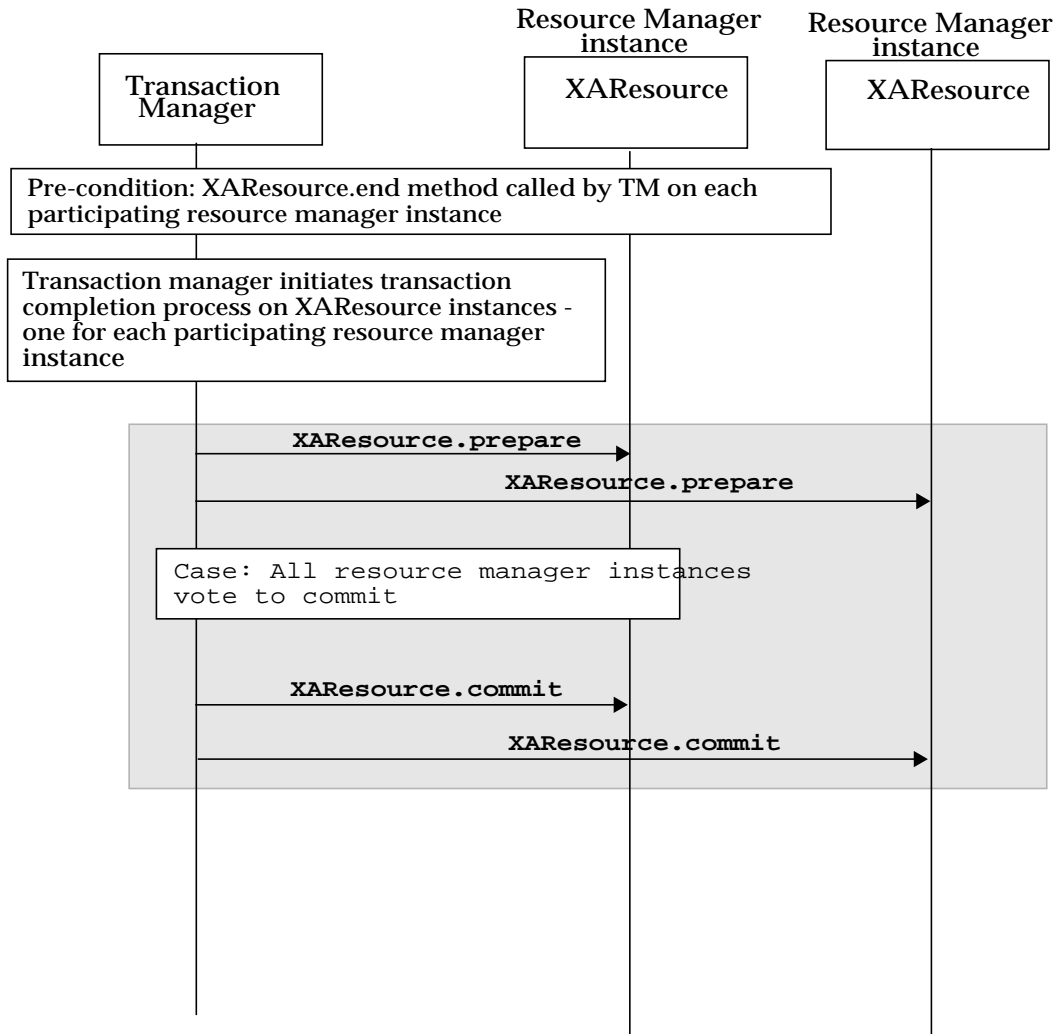
The scenario in FIGURE 7-9 illustrates the steps taken by the transaction manager to commit a transaction across multiple resource manager instances. These steps are executed after the transaction manager calls the `XAResource.end` method for each enlisted resource manager instance.

The following steps happen in this scenario:

- The transaction manager calls `XAResource.prepare` to begin the first phase of the transaction completion protocol. The transaction manager can call any `XAResource` instance is associated with the proper underlying resource manager instance, and is not restricted to the `XAResource` instance initially involved with the transaction. The application server can assume that all `XAResource` instances produced by a `ManagedConnectionFactory` instance refer to the same underlying resource manager instance.

- Assuming that all resource manager instances involved in the transaction agree to commit, the transaction manager calls `XAResource.commit` to commit the transaction. Otherwise, the transaction manager calls `XAResource.rollback`.

FIGURE 7-9 OID: Transaction Completion



7.7 Local Transaction Management Contract

The main motivation for defining a local transaction contract between an application server and a resource manager is to enable an application server to manage resource manager local transactions, hereafter called local transactions.

The local transaction management contract has two parts:

- The application server uses the `javax.resource.spi.LocalTransaction` interface to manage local transactions transparently to an application component. The scenarios in Section 7.10 “Transaction Scenarios” Note – illustrate this part of the local transaction management contract.
- The other part of the contract relates to notifications for local transaction-related events. If the resource adapter supports a local transaction demarcation API, for example, `javax.resource.cci.LocalTransaction` for the Common Client Interface, the resource adapter needs to notify the application server of the events (transaction begin, commit, and rollback) related to the local transaction. An application server uses this part of the contract, as explained in Section 7.8 “Scenarios: Local Transaction Management”.

7.7.1 Interface: Local Transaction

The `javax.resource.spi.LocalTransaction` interface defines the contract between an application server and resource adapter for local transaction management. This interface is defined in Section 7.3.3 “Interface: LocalTransaction”.

7.7.2 Interface: ConnectionEventListener

An application server implements the `javax.resource.spi.ConnectionEventListener` interface. It registers this listener instance with the `ManagedConnection` instance by using `ManagedConnection.addConnectionEventListener` method.

The following code extract specifies the `ConnectionEventListener` interface related to the local transaction management contract:

```
public interface javax.resource.spi.ConnectionEventListener {
    // Local Transaction Management related events
    public void localTransactionStarted(ConnectionEvent event);
    public void localTransactionCommitted(ConnectionEvent event);
    public void localTransactionRolledback(ConnectionEvent
event);

    ...
}
```

The `ManagedConnection` instance notifies its registered listeners for transaction related events by calling the methods `localTransactionStarted`, `localTransactionCommitted`, and `localTransactionRolledback`.

The `ConnectionEvent` class defines the following types of event notifications related to the local transaction management contract:

- `LOCAL_TRANSACTION_STARTED`—a local transaction was started using the `ManagedConnection` instance.
- `LOCAL_TRANSACTION_COMMITTED`—a local transaction was committed using the `ManagedConnection` instance.
- `LOCAL_TRANSACTION_ROLLEDBACK`—a local transaction was rolled back using the `ManagedConnection` instance.

7.7.2.1 Requirements

The connector specification requires an application server to implement the `ConnectionEventListener` interface and handle local transaction related events. This enables the application server to achieve local transaction cleanup and transaction serial interleaving, as illustrated in Section 7.8 “Scenarios: Local Transaction Management”. The connector specification provides the necessary mechanisms for transaction management. Whether these mechanisms are used in an application server depends on the application server’s implementation of the transaction requirements of the J2EE component specifications.

The resource adapter must send local transaction events through the `ConnectionEventListener` interface when an application component starts a local transaction using the application level transaction demarcation interface. An exception to this requirement is when the transaction demarcation API supports the concept of an implicit begin of a local transaction. The JDBC API is an example where there is no explicit local transaction begin method.

However, resource adapters that allow implicit begin of a local transaction, for instance, JDBC drivers, are strongly encouraged to provide support for local transaction events. This may be required in a future release of the specification.

The resource adapter must not send local transaction events for local transactions managed by the container.

7.8 Scenarios: Local Transaction Management

This section illustrates how an application server uses the event notifications from the resource adapter to manage local transactions and to restrict illegal transaction demarcations by an application component.

In these scenarios, an application component starts a local transaction using an application-level transaction demarcation interface, for example, `javax.resource.cci.LocalTransaction` as defined in the CCI, supported by the resource adapter. The resource adapter, in its implementation of the transaction demarcation interface, sends event notifications related to the local transaction, namely, local transaction begin, commit, and rollback. The application server is notified of these local transaction-related events through the `ConnectionEventListener` mechanism.

7.8.1 Local Transaction Cleanup

A stateless session bean with bean-managed transaction demarcation starts a local transaction in a method invocation. It returns from the business method without completing the local transaction.

The application server implements the `ConnectionEventListener` interface. The resource adapter notifies the application server with a `LOCAL_TRANSACTION_STARTED` event when the local transaction is started by the session bean instance.

When the session bean instance returns from the method invocation without completing the local transaction, the application server detects this as an incomplete local transaction because it has not received any matching `LOCAL_TRANSACTION_COMMITTED` or `LOCAL_TRANSACTION_ROLLEDBACK` events from the resource adapter.

On detecting an incomplete local transaction, the application server aborts the transaction, terminates the stateless session bean instance, and throws an exception to the client.

7.8.2 Component Termination

The application server terminates a component instance, for example, because of some system exception in a method invocation.

On termination of a component instance, the application server cleans up all `ManagedConnection` instances being used by this component instance. The cleanup of a connection involves resetting all local transaction and client-specific state. This state is maintained internal to the `ManagedConnection` instance.

The application server initiates a cleanup of a `ManagedConnection` instance by calling `ManagedConnection.cleanup`. After cleanup, the application server returns this connection to the pool to serve future allocation requests.

7.8.3 Transaction Interleaving

The application server uses the connection event listener mechanism, specified through the interfaces `ConnectionEventListener` and `ConnectionEvent`, to flag illegal cases of transaction demarcation. The application server implements the `ConnectionEventListener` interface to support this scenario.

The following subsection illustrates a scenario for component-managed transaction demarcation.

7.8.3.1 Scenario

An EJB component with bean managed transaction demarcation starts a local transaction using the application-level transaction demarcation interface, for example, `javax.resource.cci.LocalTransaction` as defined in the CCI, supported by the resource adapter. It then calls the `UserTransaction.begin` method to start a JTA transaction before it has completed the local transaction.

In this scenario, the EJB component has started but not completed the local transaction. When the application component attempts to start a JTA transaction by invoking the `UserTransaction.begin` method, the application server detects it as a transaction demarcation error and throws an exception from the `UserTransaction.begin` method.

When the application component starts the local transaction, the resource adapter notifies the application server of the `LOCAL_TRANSACTION_STARTED` connection event. When the component invokes the `UserTransaction.begin` method, the application server detects an error condition, because it has not received the matching `LOCAL_TRANSACTION_COMMITTED` or `LOCAL_TRANSACTION_ROLLEDBACK` event from the resource adapter for the currently active local transaction.

7.9 Connection Sharing

Sharing connections typically results in efficient use of resources and better performance. An application can indicate the ability to share its various resource references, or connections, in its deployment descriptor. A connection can be marked either as `shareable` or `unshareable`. The default is `shareable`.

When multiple `shareable` connections `x` and `y` acquired by an application are used within a global transaction scope (for instance, container-managed or bean-managed), the application server must provide a single shared connection behavior under the following conditions:

- `x` and `y` are collocated in a single Java Virtual Machine process address space.
- `x` and `y` are using a single transactional resource manager.
- `x` and `y` have identical properties.
- `x` and `y` are marked as `shareable`.
- `x` and `y` are used within a container-managed or bean-managed transaction scope.

The ability to share is unspecified for connections marked `shareable` that are used outside a global transaction scope. Sharing is not supported for connections obtained from a non-transactional resource adapter, that is, transaction support level is `NoTransaction`.

The intent of the connection sharing requirement is to avoid resource manager lock contentions and read isolation problems, and thus ensure portable behavior for transactional applications. The application server may implement the connection sharing semantics either using a single shared connection or through other mechanisms¹.

1. An application server may not share connections that are marked `shareable`, but still ensure the portability of the transaction application across application servers. For instance, the application server may choose to end XA association of connections when control flows from one application component to another, in order to avoid concurrent access problems on XA Resource Managers, and thus provide single shared connection behavior.

If a connection is marked as `shareable`, it must be transparent to the application whether a single shared connection is used or not. The application must not make assumptions about a single shared connection being used, and hence must use the connection in a shareable manner.

However, a J2EE application component that intends to use a connection in an unshareable way must leave a deployment hint to that effect, which will prevent the connection from being shared by the container. Examples of unshareable usage of a connection include changing the security attributes, isolation levels, character settings, and localization configuration.

Containers must not attempt to share connections that are marked `unshareable`.

J2EE application components may use the optional deployment descriptor element `res-sharing-scope` to indicate whether a connection to a resource manager is shareable or unshareable. Containers must assume connections to be shareable if no deployment hint is provided. Refer to EJB specification (Section 1., “Enterprise JavaBeans (EJBTM) specification, version 2.1:” on page F-1) and the servlet specification (Section 10., “Java Servlet specification, version 2.4:” on page F-1) for a description of the deployment descriptor element.

J2EE application components may cache connection objects and reuse them across multiple transactions. Containers that provide connection sharing should transparently switch such cached connection objects, at dispatch time, to point to an appropriate shared connection with the correct transaction scope. Refer to Section 7.11 “Connection Association” for more details on connection association.

Refer to Section 7.10 “Transaction Scenarios” for a special case of connection sharing as applied to resource adapters that support local transactions.

7.9.1 Sharing Violation Detection

A resource adapter may detect sharing violations. Any operation on a shareable connection which violates shareability is a sharing violation, for example, mutable operations like changing connection attributes, security settings, isolation levels, etc.

When such a mutable operation is performed on a `ManagedConnection`, it may throw a `SharingViolationException` when both the following conditions are true:

- The number of connection handle objects associated with the `ManagedConnection` is more than one.
- The `ManagedConnection` is associated with a transaction, either local or XA.

Further, a resource adapter may reject creation of a connection handle, by throwing a `SharingViolationException`, if the connection is already in an unshareable condition. Any mutable operation performed on a connection makes it unshareable.

7.9.1.1 Scenario 1

Application component A gets a shareable connection to a resource and invokes component B which also gets a shareable connection to the same resource. Both A and B are involved in a common transaction scope, either local or XA. The application server shares the connections acquired by both A and B. From this point onwards, any attempt to change a mutable property, such as isolation level, by either component, results in a `SharingViolationException` being thrown by the resource adapter to the offending component.

7.9.1.2 Scenario 2

Application component A gets a shareable connection to a resource. A is involved in a transaction, either local or XA. A then modifies one of the mutable properties of the resource, such as isolation level. This makes the connection unshareable. The resource adapter does not throw an exception since only one connection handle is present.

Later, A invokes B under the same transaction scope. B also attempts to acquire a shareable connection to the same resource. The application server chooses to share the connection that is already in use by A. At this point, the resource adapter throws a `SharingViolationException` to B since sharing had been attempted on an unshareable connection. The resource adapter does this by remembering that the connection had been made unshareable earlier.

The resource adapter might throw a `SharingViolationException` to B, even if A had closed its connection handle before it invoked B, since the connection acquired by A had become unshareable.

7.10 Transaction Scenarios

This section specifies requirements for various transaction scenarios.

7.10.1 Requirements

The J2EE platform specification (Section 8., “Java 2 Platform Enterprise Edition (J2EETM), Platform specification, version 1. 4:” on page F-1) identifies the following as transactional resources:

- JDBC connections
- JMS sessions

■ Resource adapter connections at the `XATransaction` level

The J2EE platform specification requires that J2EE product providers must transparently support transactions that span multiple components and transactional resources. These requirements must be met regardless of whether a J2EE product is implemented as a single process, multiple processes on the same node, or multiple processes on multiple nodes.

In addition, J2EE product providers must support transactional applications that are comprised of servlets or JSP pages accessing multiple enterprise beans within a single transaction. Each component may also acquire one or more connections to access transactional resources. J2EE product providers must support scenarios where multiple components in an application access transactional resources as part of a single transaction.

The J2EE platform specification requires J2EE platform products to support resource adapters at the `XATransaction` level as a transactional resource. It must be possible to access such resource adapters from multiple application components within a single transaction.

The connector architecture has an additional requirement that is applicable to resource adapters that support local transactions. Note that both `LocalTransaction` and `XATransaction` resource adapters support local transactions.

Application server must use local transactions in a scenario where the following conditions hold:

- Multiple components are involved in a global transaction scope.
- Components use a single resource adapter that is local transaction capable.
- Components get connections to the same EIS instance.
- Components have not specified the `res-sharing-scope` flag as `unshareable`.
This condition accounts for potential sharing of connections in terms of security context, client-specific connection parameters, and EIS specific configuration.

Note that this requirement does not apply to a local transaction that is started by a component using an application level transaction demarcation API that is specific to a resource adapter.

Application server may use connection sharing mechanisms to implement this local transaction requirement.

Application servers must support transaction scenarios where access to a non-transactional resource is combined with access to one or more transactional resources within a single transaction. For example, in a container-managed transaction, an EJB accesses JDBC and JMS resources and also accesses a non-transactional EIS using its resource adapter. If there is a failure during the above scenario, transactional resource managers operating under the transaction should rollback, but the recovery of the non-transactional resource is unspecified in this specification.

The application server is not required to support any additional transaction scenarios beyond the above set of scenarios. A J2EE application should not depend on an application server's support for any optional transaction scenarios. The application should also not depend on whether or not the container detects that a specific optional transaction scenario is illegal. Any errors in optional transaction scenarios are considered application programming errors.

7.10.2 Illustrative Scenarios

The following are examples of optional transaction scenarios. The following section also describes, in a non-prescriptive manner, issues in support for these scenarios by an application server:

- Within a transaction, an EJB component acquires connections to two different resource managers X and Y using their respective non-XA local transaction capable resource adapters.

The container cannot manage a local transaction across two different resource managers. Since resource adapters and underlying resource managers are not XA capable, the container cannot use XA in this case. However, a J2EE application should not depend on the container to detect this illegal scenario.

- Within a transaction, EJB component A acquires a connection to a resource manager X using a non-XA local transaction capable resource adapter. Next, EJB component B under the same transaction context acquires a connection to a different resource manager Y using a non-XA local transaction capable resource adapter

The container cannot manage a local transaction across two different resource managers. Since resource adapters are not XA capable, the container cannot use XA in this case. However, a J2EE application should not depend on the container to detect this illegal scenario.

- Within a transaction, EJB component A acquires a connection to a resource manager X using a non-XA local transaction capable resource adapter. Next, the same EJB (or EJB B) under the same transaction context acquires a connection to a different resource manager Y using an XA capable resource adapter

This scenario may be supported if the transaction manager supports last resource commit optimization. Since this optimization feature is optional and not specified in the connector architecture, a J2EE application should not depend on support for this scenario.

- Within a transaction, EJB A acquires a connection to a resource manager X using an XA capable resource adapter. Next, the same EJB component (or another EJB component B) under the same transaction context acquires a connection to a different resource manager Y using a non-XA local transaction capable resource adapter

This scenario may be supported if the transaction manager supports last resource commit optimization. Since this optimization feature is optional and not specified in the connector architecture, a J2EE application should not depend on support for this scenario.

7.10.3 Scenario: Local Transaction

The following scenario illustrates the use of the connection sharing mechanism to implement requirement for a local transaction to span components.


In this scenario, two EJB components get connections to the same EIS resource manager within a single transaction. Both EJB components use the same local transaction capable resource adapter.

A local transaction is associated with a single physical connection. Both EJB components in this scenario share the same physical connection under the local transaction scope. The container has the responsibility of managing connection sharing as illustrated in the following scenario.

To share a physical connection in the local transaction scope, the container assumes the connection to be shareable unless it has been marked `unshareable` in the `res-sharing-scope`. The container uses connection sharing in a manner that is transparent to application components.

In FIGURE 7-10, the stateful session beans A and B have container-managed transaction demarcation with the transaction attribute set to `Required`. Both A and B access a single EIS resource manager as part of their business logic.

FIGURE 7-10 Scenario to illustrate Local Transaction Management



The following steps happen in this scenario:

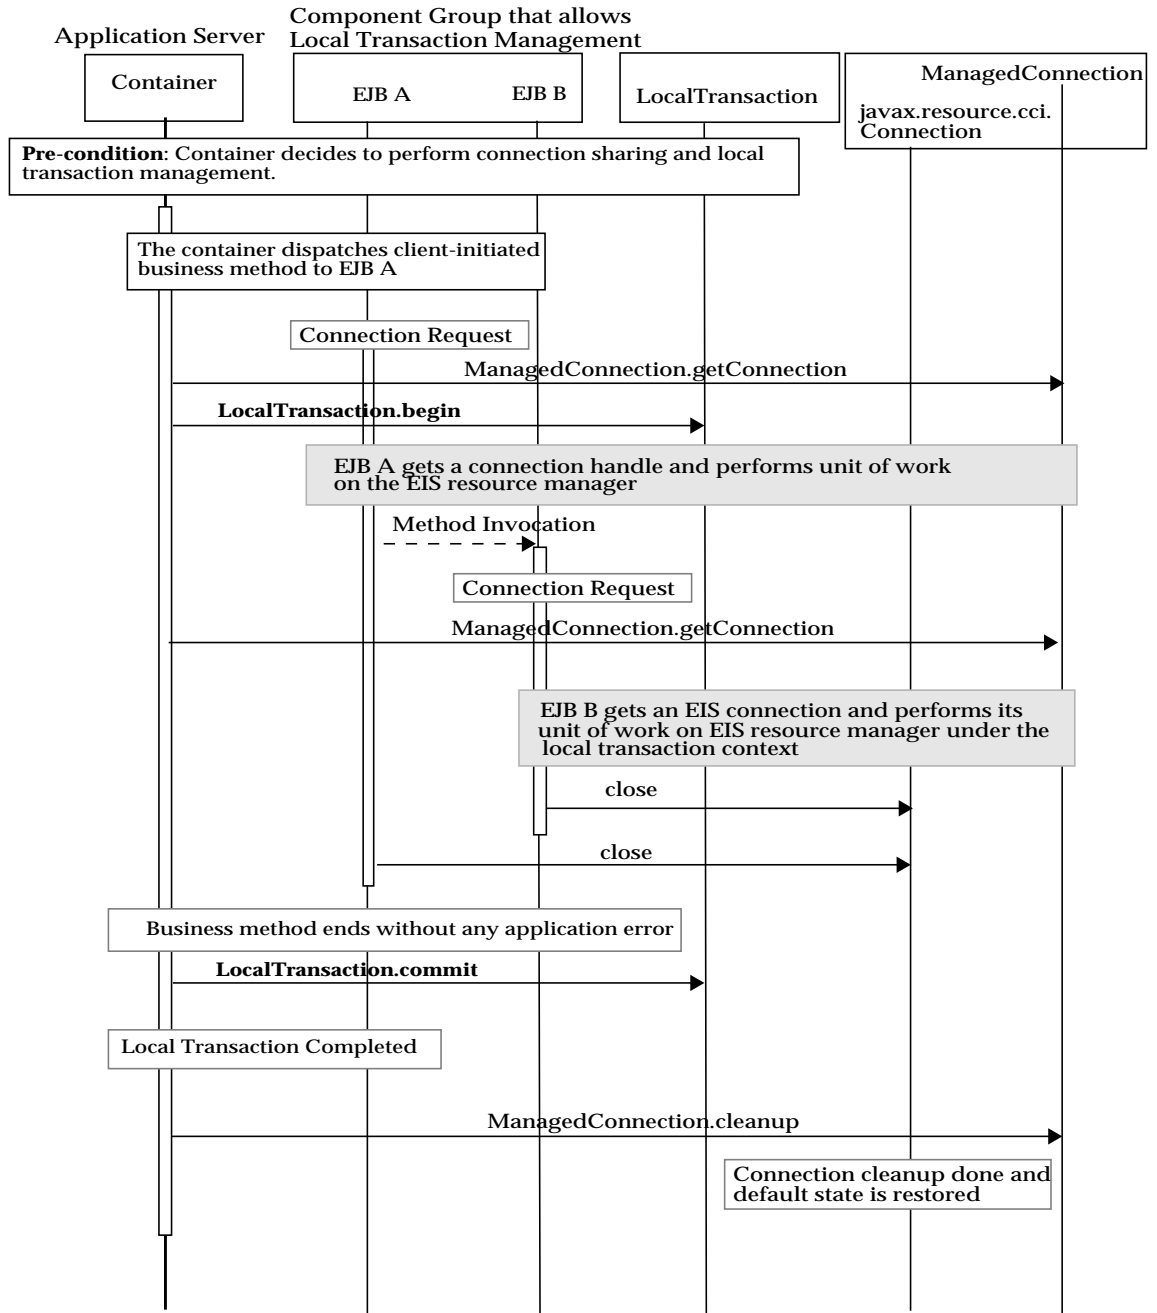
- The client invokes a method on session bean A with no transaction context. In its method implementation, the EJB A acquires a connection to the EIS instance.
- When acquiring the connection, the container starts a local transaction by invoking the `begin` method of the `javax.resource.spi.LocalTransaction` instance. The local transaction is tied to the `ManagedConnection` instance that is associated with the connection handle acquired by the component in the previous step.
- After the local transaction starts, any recoverable unit of work performed by A on the EIS resource manager using the acquired connection is automatically included under the local transaction context.
- Session bean A now invokes a method on the session bean B instance. In this scenario, A does not close the connection handle before invoking the method on B.

Note – A container should ensure that the connection sharing mechanism is equally applicable if A were to close the connection handle before calling the B instance.

- In the invoked method, B makes a request to acquire a connection to the same EIS resource manager.
- The container returns a connection handle using the same `ManagedConnection` instance that was used for handling the connection request from A.
- The container retains the association of the `ManagedConnection` instance with the local transaction context across the method invocation from A to B. This means that any unit of work that B will perform on the EIS resource manager using its acquired connection handle will be automatically included as part of the

current local transaction. The connection state, for example, any open cursors, can also be retained across method invocations when the physical connection is shared.

FIGURE 7-11 OID: Connection Sharing across Component instances



- Before the method invocation on B completes, B calls the `close` method on the connection handle. The container should not initiate any cleanup of the physical connection at this time since there is still an uncompleted local transaction associated with the shared physical connection. In this scenario, the cleanup of a physical connection refers to the dissociation of the local transaction context from the `ManagedConnection` instance.
- When A regains control, A can use the same connection handle, provided A had not called the `close` method on the connection handle, to access EIS resources. All recoverable units of work on the EIS resource manager will be included in the existing local transaction context.

Note – If A closes the connection handle before calling B, and then reacquires the connection handle when regaining control, the container should ensure that the local transaction context stays associated with the shared connection.

- A eventually calls the `close` method on its connection handle. The container gets a connection close event notification based on the scenario described in Section 6.8.3, “Scenario: Connection Event Notifications and Connection Close” on page 6-40.
- Since there is an incomplete local transaction associated with the underlying physical connection, the container does not initiate a cleanup of the `ManagedConnection` on receiving the connection close event notification. The container must still go through the completion process for the local transaction.
- When the business method invocation on A completes successfully without any application error, the container starts the completion protocol for the local transaction. The container calls the `LocalTransaction.commit` method to commit the transaction.
- After the local transaction completes, the container initiates a cleanup of the physical connection instance by calling the `ManagedConnection.cleanup` method.

Note – The container should initiate cleanup of the `ManagedConnection` instance in the case where A does not call the `close` method on the connection handle before returning. The container identifies the need for cleaning up the `ManagedConnection` instance based on the scope of connection sharing.

- On the `cleanup` method invocation, the `ManagedConnection` instance does a cleanup of its local transaction related state and resets itself to a default state.
- The container returns the physical connection to the pool for handling subsequent connection requests.

7.11 Connection Association

According to the connection management contract, a connection handle is created from a `ManagedConnection` instance using the `ManagedConnection.getConnection` method. A connection handle maintains an association with the underlying `ManagedConnection` instance.


7.11.1 Scenario

In the scenario shown in FIGURE 7-12, session bean A acts as a client of entity bean C and makes invocations on methods of entity bean C. Another session bean B also acts as a client of entity bean C. The C is an entity bean that may be shared across multiple clients.

A, B and C get connections to the same EIS. These EJB components have marked `res-sharing-scope` for these connections to be shareable.

A and C define a connection sharing scope. Both A and C share the same physical connection across a transaction that spans methods on A and C. Similarly, B and C define another connection sharing scope. B and C also share the same physical connection across a transaction that spans two components.

FIGURE 7-12 Connection Sharing Scenario




In this scenario, entity bean C obtains an application-level connection handle using the method `getConnection` on the `ConnectionFactory` during its creation. Entity bean C holds the connection handle during its lifetime.

A gets a connection handle and invokes a method on C. At a different time, B gets a connection handle and invokes a method on C.

In both cases, depending on the connection sharing scope, defined in terms of the shared physical `ManagedConnection` instance, in which C is called, the container supports a mechanism to associate the connection handle held by C as part of its state with the current `ManagedConnection` instance.

FIGURE 7-13 State Diagram of Application-Level Connection Handle



7.11.2 Connection Association

The interface `ManagedConnection` defines method `associateConnection` as follows:

```
public interface javax.resource.spi.ManagedConnection {
    public void associateConnection(Object connection)
        throws ResourceException;
    ...
}
```

The container uses the `associateConnection` method to change the association of an application-level connection handle with a `ManagedConnection` instance. The container finds the right `ManagedConnection` instance, depending on the connection sharing scope, and calls the `associateConnection` method. To achieve

this, the container needs to keep track of connection handles acquired by component instances and `ManagedConnection` instances using an implementation-specific mechanism.

The `associateConnection` method implementation for a `ManagedConnection` should dissociate the connection handle passed as a parameter from its currently associated `ManagedConnection` and associate the new connection handle with itself.

Note that the switching of connection associations must happen only for connection handles and `ManagedConnection` instances that correspond to the same `ManagedConnectionFactory` instance. The container should enforce this restriction in an implementation-specific manner. If a container cannot enforce the restriction, the container should not use the connection association mechanism.

7.11.3 Requirements

The container must provide a mechanism to change the association of a connection handle to different `ManagedConnection` instances depending on the connection sharing and transaction scope. This mechanism is used in scenarios where components hold on to connection handles across different local transaction and connection sharing scopes.

The container may use the connection association mechanism in the `XAResource`-based transaction management contract.

The resource adapter must implement the `associateConnection` method to support connection sharing. The container makes a decision on whether or not to use the `associateConnection` method implemented by a resource adapter. The support for this method is required independent of the transaction support level of the resource adapter. Note that the container makes the decision to invoke the `associateConnection` method.

7.12 Local Transaction Optimization

If all the work done as a part of a transaction uses a single resource manager, the application server can use a local transaction in place of an externally coordinated JTA transaction. The use of a local transaction avoids the overhead of initiating a global transaction, and involving the TM for transaction coordination, and leads to more optimized performance.

Since a typical application accesses a single resource manager, the local transaction optimization is a useful performance enhancement for transaction management.

The application server manages local transaction optimization transparent to the J2EE application. Whenever a container-managed or bean-managed transaction is started, the container may attempt local transaction optimization.

When the transaction begins, a container cannot determine beforehand whether or not the unit of work done as part of this transaction will use a single resource manager. The container uses an implementation-specific mechanism to achieve local transaction optimization. For example, the container can choose to start a local transaction when the first resource manager is accessed and lazily start a JTA transaction only when more than one resource managers are accessed in an existing transaction. The mechanism through which the application server and its transaction manager coordinates the initial local transaction and lazily started JTA transactions is outside the scope of the connector specification. Refer to the J2EE platform specification (Section 8., “Java 2 Platform Enterprise Edition (J2EETM), Platform specification, version 1. 4:” on page F-1) for more details on the local transaction optimization.

7.12.1 Requirements

The container is not required to support the local transaction optimization.

7.13 Requirements

The following section outlines the requirements for the transaction management contract.

7.13.1 Resource Adapter

A resource adapter can be classified based on the level of transaction support, as follows:

- **NoTransaction**—The resource adapter supports neither resource manager local nor JTA transactions. It implements neither the `XAResource` nor `LocalTransaction` interfaces.
- **LocalTransaction**—The resource adapter supports resource manager local transactions by implementing the `LocalTransaction` interface. The local transaction management contract is specified in Section 7.7 “Local Transaction Management Contract”.

- `XATransaction`—The resource adapter supports both resource manager local and JTA transactions by implementing the `LocalTransaction` and `XAResource` interfaces. The requirements for supporting the `XAResource`-based contract are specified in Section 7.6 “XAResource-based Transaction Contract”.

Note – Other levels of support (includes any transaction optimizations supported by an underlying resource manager) are outside the scope of the connector architecture.

The above levels reflect the major steps of transaction support that a resource adapter needs to make to allow external transaction coordination. Depending on its transactional capabilities and the requirements of its underlying EIS, a resource adapter can choose to support any one of the above transaction support levels.

7.13.1.1 Auto Commit

When a connection is in an auto-commit mode, an operation on the connection automatically commits after it has been executed. The auto-commit mode must be off if multiple interactions have to be grouped in a single transaction, either local or XA, and committed or rolled back as a unit.

A resource adapter must manage the auto-commit mode as follows:

- A transactional resource adapter, either at `XATransaction` or `LocalTransaction` level, must set the auto-commit mode to false within a transaction, either local or XA, on a connection participating in the transaction. This requirement holds for both container-managed and bean-managed transaction demarcation.
- A transactional resource adapter must set the auto-commit mode to true, on connections that are used outside a transaction.

7.13.2 Application Server

An application server must support resource adapters with all three levels of transaction support—`NoTransaction`, `LocalTransaction`, and `XATransaction`.

The following are the requirements for an application server for the transaction management contract:

- The application server must support a transaction manager that manages transactions using the JTA `XAResource`-based contract. The requirements for a transaction manager to support an `XAResource`-based contract are specified in Section 7.6.3 “Transaction Manager Requirements”.

- The application server must use the `LocalTransaction` interface-based contract to manage local transactions for a resource manager.
- The application server must use the deployment descriptor mechanism to ascertain the transactional capabilities of a resource adapter. Refer to Section 17.3, “Deployment” on page 17-6 for details on the deployment descriptor specification.
- The application server must implement the `ConnectionEventListener` interface to get transaction-related event notifications.

7.14 Connection Optimizations

The following describes two optional connection optimizations:

- Lazy connection association optimization
- Lazy transaction enlistment optimization

7.14.1 Lazy Connection Association Optimization

Application components may acquire connections through a `ConnectionFactory` object (resource-ref) obtained from the JNDI namespace. The connection(s) thus obtained may be closed by the application before method completion, or may be cached by the application for later use.

When a connection is cached by the application component, the cached connection handle is considered active and remains associated with a `ManagedConnection` instance from the application server’s connection pool. If the cached connection handle is used infrequently, then the associated `ManagedConnection` instance remains in hibernation during periods of non-use. This is because the application server does not know when the hibernating `ManagedConnection` instance will be used again by the application.

Such hibernating `ManagedConnection` instances result in suboptimal usage of system resources. Avoiding hibernation of `ManagedConnection` instances leads to more optimal resource utilization and better performance.

The following describes a mechanism that allows an application server to avoid hibernating `ManagedConnection` instances (by dissociating the `ManagedConnection` from its connection handles and using the freed `ManagedConnection` instance for other applications). This mechanism also provides a way to notify the application server when a dissociated connection handle is used by the application, so that it can be associated with an appropriate `ManagedConnection` instance.

FIGURE 7-14 on page 52 describes the processing of a `getConnection` method call initiated by an application component (that is, when the application component first acquires a connection). At a later point in time, the connection may be dissociated by the application server by calling the `dissociateConnections` method on the appropriate `ManagedConnection` instance. This dissociates the `ManagedConnection` instance from all its connection handle objects.

When such a dissociated connection is used by the application (upon method re-entry), it needs to be re-associated with an appropriate `ManagedConnection` instance. FIGURE 7-15 on page 53 describes connection re-association processing. The connection re-association processing depends on the connection notifying the application server upon re-use (lazy re-association trigger). The connection object invokes the `associateConnection` method on the `ConnectionManager` instance in order to lazily re-associate itself with an appropriate `ManagedConnection` instance.

Thus, a dissociate-able connection handle can exist in one of three states: `Active`, `Inactive` or `Closed`. FIGURE 7-16 on page 54 describes the state transitions of a dissociatable connection handle. Note that the state `Inactive` applies only to dissociatable connection handles.

The application server may dissociate connections that are shareable. It must not dissociate connections that are marked unshareable, however, since application-specific state may be retained by a `ManagedConnection` instance.

FIGURE 7-14 Connection Acquisition Processing

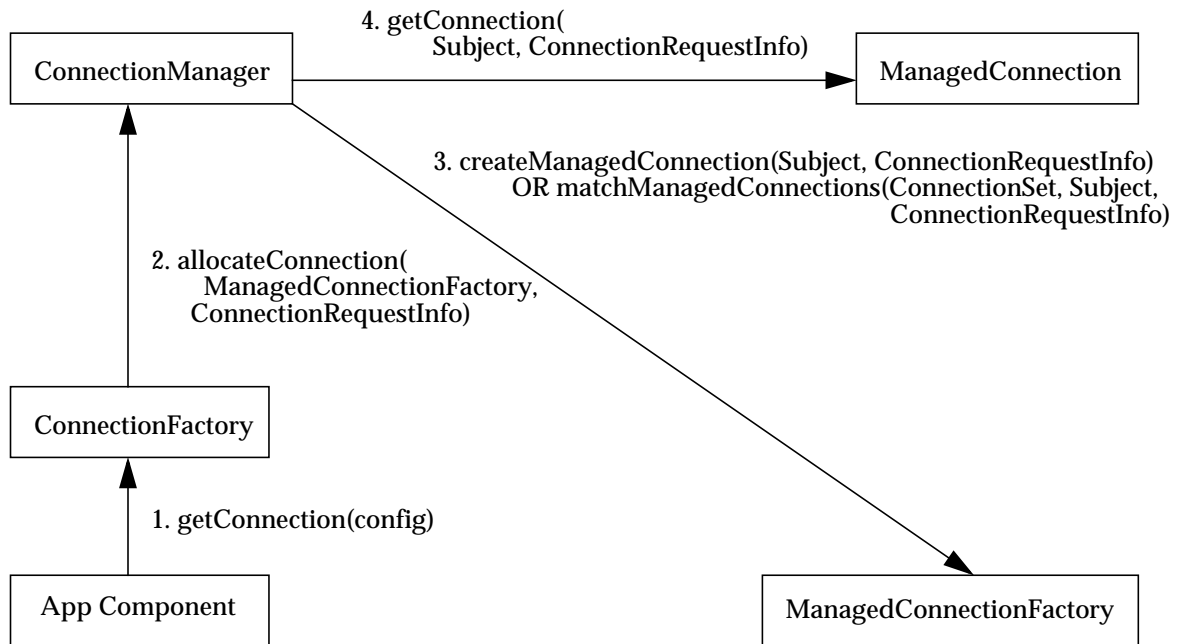


FIGURE 7-15 Connection Re-association Processing

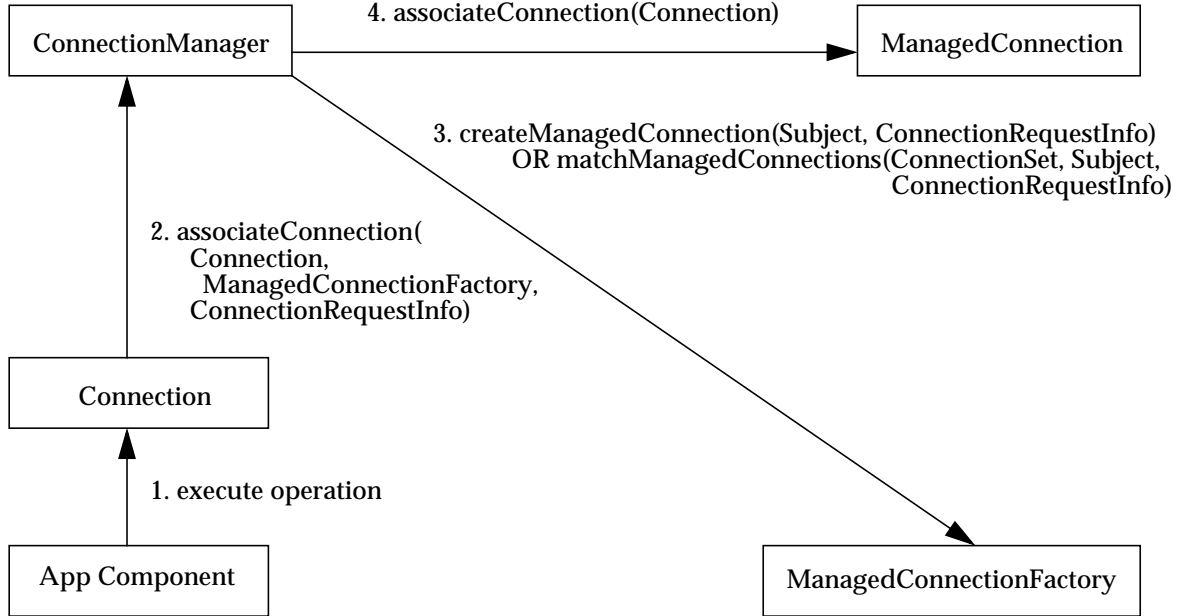
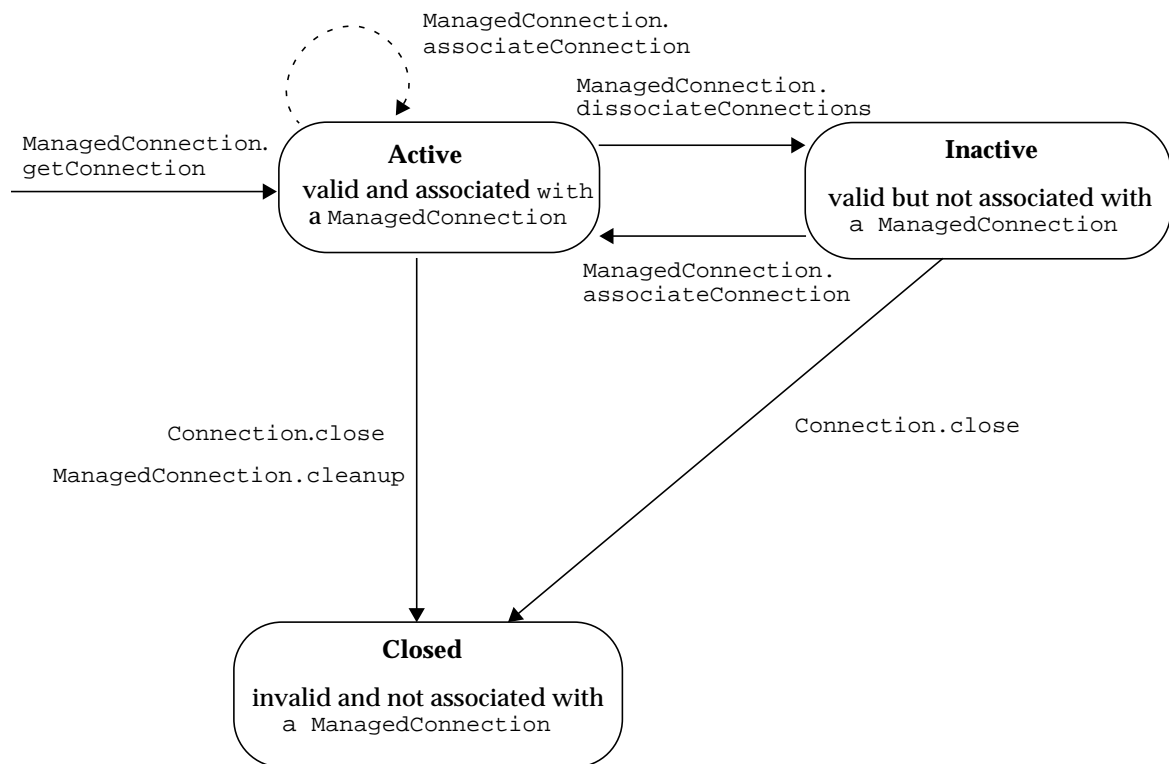


FIGURE 7-16 State Diagram of a Dissociatable Application-level Connection Handle



7.14.1.1 API Additions

```
package javax.resource.spi;

import javax.resource.ResourceException;

interface LazyAssociatableConnectionManager { // application
server
    void associateConnection(
        Object connection, ManagedConnectionFactory mcf,
        ConnectionRequestInfo info)
        throws ResourceException;
}

interface DissociatableManagedConnection { // resource adapter
    void dissociateConnections() throws ResourceException;
}
```

Neither the application server nor the resource adapter must support this optimization.

A resource adapter that does not support this optimization must provide a `ManagedConnection` implementation that does not implement the `DissociatableManagedConnection` interface. This allows an application server to detect that the resource adapter does not support this optimization.

An application server that does not support this optimization must provide a `ConnectionManager` implementation that does not implement the `LazyAssociatableConnectionManager` interface. This allows a resource adapter to detect that the application server does not support this optimization. In reality, a resource adapter will not call this method (in order to re-associate a connection) since an application server that does not support this optimization would never dissociate a connection.

There are no changes to the resource adapter deployment descriptor since the application server can programmatically detect whether a resource adapter supports this optimization or not.

7.14.2 Lazy Transaction Enlistment Optimization

Transactions may be started by an application server before a method call on an application component or it may be started by an application component during a method call. It is also possible that an application server may use a transaction imported from a different server during a method call.

Irrespective of how a transaction is started, an application server enlists all connections (cached or newly acquired by an application component) with the transaction, so that the work done using those connections will be part of the transaction. This enlistment happens before the method call in the case of cached connections and during the method call when connections are newly acquired within the transaction.

But not all the connections that are cached or newly acquired by an application component may be used within a transaction. Since the application server does not know whether these connections would be used within the transaction, it statically (eagerly) enlists all such connections with the transaction. Thus, connections that are not used in a transaction are unnecessarily enlisted, which leads to sub-optimal performance.

The following describes a dynamic mechanism that allows the application server to enlist only those connections that are used within a transaction. A `ManagedConnection` that supports this optimization must invoke the `lazyEnlist` method on the `ConnectionManager` every time it is used outside of a local or XA transaction. The application server uses this method call to lazily enlist the connection in the transaction (if there is one). The application server may delist the `ManagedConnection` instances from the transaction at a later point.

This optimization can be used only on connections that are lazily enlist-able.

7.14.2.1 API Additions

```
package javax.resource.spi;

import javax.resource.ResourceException;
import javax.transaction.xa.Xid;

interface LazyEnlistableConnectionManager { // application server
    void lazyEnlist(ManagedConnection) throws ResourceException;
}

interface LazyEnlistableManagedConnection { // resource adapter
}
```

Neither the application server nor the resource adapter must support this optimization.

A resource adapter that does not support this optimization must provide a `ManagedConnection` implementation which does not implement the `LazyEnlistableManagedConnection` interface. This allows an application server to detect that the resource adapter does not support this optimization.

An application server that does not support this optimization must provide a `ConnectionManager` implementation that does not implement the `LazyEnlistableConnectionManager` interface. This allows a resource adapter to detect that the application server does not support this optimization.

There are no changes to the resource adapter deployment descriptor since the application server can programmatically detect whether a resource adapter supports this optimization or not.

Security Architecture

The following chapter specifies the security architecture for the integration of EISs with the J2EE platform. It adds EIS integration-specific security details to the security requirements specified in other J2EE specifications.

8.1 Overview

It is critical that an enterprise be able to depend on the information in its EIS for its business activities. Any loss or inaccuracy of information or any unauthorized access to the EIS can be extremely costly to an enterprise. There are several mechanisms that can be used to protect an EIS against such security threats, including:

- Identification and authentication of principals, human users to verify they are who they claim to be.
- Authorization and access control to determine whether a principal is allowed to access an application server and/or an EIS.
- Secure communication between an application server and an EIS. Communication over insecure links can be protected using a protocol, for example, Kerberos, that provides authentication, integrity, and confidentiality services. Communication can also be protected by using a secure link protocol, for example, SSL.

8.2 Goals

The security architecture is designed to meet the following goals:

- Extend the end-to-end security model for J2EE applications to include integration with EISs based on the connector architecture.
- Support authentication and authorization of users who are accessing EISs.

- Keep the security architecture technology neutral and enable the specified security contract to be supported by various security technologies.
- Enable the security architecture to support a range of EISs with different levels of security support and existing security environments.
- Support security configuration of a resource adapter in an operational environment.
- Keep the security model for connector architecture-based EIS integration transparent to an application component provider. This includes providing support for single sign-on across multiple EISs.

The security model for EIS integration is not designed to:

- Mandate a specific technology and describe how it can be used to implement the security architecture for connector architecture-based EIS integration.
- Specify and mandate a specific security policy. The security architecture enables an application server and EIS to support the implementation and administration of security policies based on their respective requirements.

8.3 Terminology

The following terms are used in this chapter:

- **Principal:** A principal is an entity that can be authenticated by an authentication mechanism deployed in an enterprise. A principal is identified using a `principal name` and authenticated using `authentication data`. The content and format of the `principal name` and the `authentication data` depend upon the authentication mechanism.
- **Security Attributes:** A principal has a set of security attributes associated with it. These security attributes are related to the authentication and authorization mechanisms. Some examples are security permissions, and credentials for a principal.
- **Credential:** A credential contains or references security information that can authenticate a principal to additional services. A principal acquires a credential upon authentication or from another principal that allows its credential to be used. The latter is termed `principal delegation`.
- **End user:** An end user is an entity, human or service, that acts as a source of a request to an application. An end user is represented as a security principal within a `Subject` as specified in the JAAS framework (“Java Authentication and Authorization Service, version 1.0:” on page F-1).
- **Initiating Principal:** The security principal representing the end-user that interacts directly with the application. An end-user can authenticate using either a web client or an application client.

- **Caller Principal:** A principal that is associated with an application component instance during a method invocation. For example, an EJB instance can call the `getCallerPrincipal` method to get the principal associated with the current security context.
- **Resource Principal:** A security principal under whose security context a connection to an EIS instance is established.
- **Security domain:** A scope within which certain common security mechanisms and policies are established. This specification does not specify the scope of a security domain. An enterprise can contain more than one security domain. Thus an application server and an EIS may either be in the same or different security domains. Appendix D, “Security Scenarios” provides illustrative examples of how security domains can be setup and managed.

In a managed environment, application components are deployed in web or EJB containers. When a method gets invoked on a component, the principal associated with the component instance is termed a caller principal.

The relationship between an initiating principal and a caller principal depends on the principal delegation option for inter-container and inter-component calls. This form of principal delegation is out of the scope of the connector architecture.

The relationship of a resource principal and its security attributes, for example, credentials and access privileges, to an initiating or caller principal depends on how the resource principal has been setup by the system administrator or deployer.

Refer to Section 9.1.1, “Interfaces and Classes” on page 9-1 for details on interfaces and classes that are used to represent a resource principal and its credentials.

8.4 Application Security Model

The following section is a brief summary of the security model from the perspective of an application component provider. Refer to the relevant specifications for more detail.

The application component requests a connection to be established under the security context of a resource principal. The security context includes security attributes—access privileges, authorization level—for a resource principal. Once a connection is successfully established, all application-level invocations to the EIS instance using the connection happen under the security context of the resource principal.

The application component provider has the following two choices related to EIS sign-on:

- Allow the deployer to set up the resource principal and EIS sign-on information. For example, the deployer sets the user name and password for establishing a connection to an EIS instance.
- Perform sign-on to an EIS from the component code by providing explicit security information for a resource principal.

The application component provider uses a deployment descriptor element, for example, `res-auth` for EJB components, to indicate the requirements for one of the above two approaches. If the `res-auth` element is set to `Application`, the component code performs a programmatic sign-on to the EIS. If the `res-auth` element is `Container`, the application server takes on the responsibility of setting up and managing EIS sign-on.

8.4.1 Scenario: Container-Managed Sign-on

The application component provider sets the `res-auth` deployment descriptor element to be `Container` letting the application server take the responsibility of managing EIS sign-on.

The Deployer sets up the principal mapping such that the user account for connecting to the EIS instance is always `eStoreUser`. The Deployer also configures the authentication data, for example, the password, needed to authenticate the `eStoreUser` to the EIS.

The component code invokes the `getConnection` method on the `ConnectionFactory` instance with no security-related parameters. The component relies on the application server to manage sign-on to the EIS instance based on the security information configured by the Deployer.

```
// Method in an application component
Context initctx = new InitialContext();

// perform JNDI lookup to obtain connection factory
javax.resource.cci.ConnectionFactory cxf =
    (javax.resource.cci.ConnectionFactory)initctx.lookup(
        "java:comp/env/eis/MyEIS");

// Invoke factory to obtain a connection. The security
// information is not passed in the getConnection method
javax.resource.cci.Connection cx = cxf.getConnection();
...
```

8.4.2 Scenario: Component-Managed Sign-on

The application component provider sets the `res-auth` element to be `Application`.

The component code performs a programmatic sign-on to the EIS. The application component passes explicit security information, for example, the username and password, to the `getConnection` method of the `ConnectionFactory` instance.

```
// Method in an application component
Context initctx = new InitialContext();

// perform JNDI lookup to obtain connection factory
javax.resource.cci.ConnectionFactory cxf =
    (javax.resource.cci.ConnectionFactory)initctx.lookup(
        "java:comp/env/eis/MyEIS");

// Invoke factory to obtain a connection
com.myeis.ConnectionSpecImpl properties = .. // get a new
ConnectionSpec
properties.setUserName("...");
properties.setPassword("...");
javax.resource.cci.Connection cx = cxf.getConnection(properties);
...
```

8.5 EIS Sign-on

Creating a new physical connection requires a sign-on to an EIS instance. Changing the security context on an existing physical connection can also require EIS sign-on. The latter is termed re-authentication.

An EIS sign-on typically involves one or more of the following steps:

- Determine a resource principal under whose security context a physical connection to an EIS will be established.
- Authenticate a resource principal if it is not already authenticated.
- Establish a secure association between the application server and the EIS. This enables additional security mechanisms, for example, data confidentiality and integrity, to be applied to communication between the two entities.
- Set the access control to EIS resources.

8.5.1 Authentication Mechanism

An application server and an EIS collaborate to ensure resource principals are properly authenticated when the principal connects to the underlying EIS. The connector architecture identifies the following as the commonly-supported authentication mechanisms:

- **BasicPassword**: Basic password based authentication mechanism specific to an EIS
- **Kerbv5**: Kerberos version 5-based authentication mechanism

The `authentication-mechanism-type` element is used in the deployment descriptor to specify whether or not a resource adapter supports a specific authentication mechanism. Refer to Section 17.5.4, “Requirements” on page 17-23 for more details on the specification of the deployment descriptor for a resource adapter.

The connector architecture does not require that a specific authentication mechanism be supported by an application server and an EIS. An application server may support any other authentication mechanisms for EIS sign-on. The connector security architecture is independent of security mechanisms.

8.5.2 Resource Principal

When an application component requests a connection from a resource adapter, the connection request is made under the security context of a resource principal. The Deployer can set a resource principal based on the following options:

- **Configured Identity**: In this case, a resource principal has its own configured identity and security attributes independent of the identity of the initiating or caller principal. The identity of the resource principal can be configured either at deployment time or specified dynamically by a component at the connection creation. The scenario described in Section D.1, “EStore Application” on page D-1 illustrates an example where connections to an EIS are always established under the security context of a valid EIS user account. This happens independent of the initiating or caller principal. For example, if a caller principal is A, then the configured resource principals can be B and C on two different EIS instances, where A, B, and C are independent identities.
- **Principal Mapping**: A resource principal is determined by mapping from the identity and/or security attributes of the initiating or caller principal. In this case, a resource principal does not inherit identity or security attributes of a principal that it has been mapped from. The resource principal gets its identity and security attributes based on the mapping. For example, if the caller principal has identity A, then the mapped resource principal is `mapping(A,EIS1)` and `mapping(A,EIS2)` on two different EIS instances.

- **Caller Impersonation:** A resource principal acts on behalf of an initiating or caller principal. Acting on behalf of a caller principal requires that the caller's identity and credentials be delegated to the EIS. The mechanism by which this is accomplished is specific to a security mechanism and an application server implementation. An example of the impersonation is described in Section D.2, "Employee Self Service Application" on page D-4.

In some scenarios, a caller principal can be a delegate of an initiating principal. In this case, a resource principal transitively impersonates an initiating principal.

The support for principal delegation is typically specific to a security mechanism. For example, Kerberos supports a mechanism for the delegation of authentication. Refer to the Kerberos v5 specification for more details. The security technology specific details are out of the scope of the connector architecture.

- **Credentials Mapping:** This mechanism may be used when an application server and EIS support different authentication domains. For example, the initiating principal has been authenticated and has public key certificate-based credentials. The security environment for the EIS is configured with the Kerberos authentication service. The application server is configured to map the public key certificate-based credentials associated with the initiating principal to the Kerberos credentials. In this case, the resource principal is the same as the caller principal with the mapped credentials.

In the case of credential mapping, the mapped resource principal has the same identity as the initiating or caller principal. For example, a principal with identity `A` has initial credentials `cred(A, mech1)` and has credentials `cred(A, mech2)` after mapping. `mech1` and `mech2` represents different mechanism types.

8.5.3 Authorization Model

Authorization checking to ensure that a principal has access to an EIS resource can be applied at one or more of the following:

- At the EIS.
- At the application server.

Authorization checking at the target EIS can be done in an EIS-specific way and is not specified here. For example, an EIS can define its access control policy in terms of its specific security roles and permissions.

Authorization checking can also be done at the application server level. For example, an application server can allow a principal to create a connection to an EIS only if the principal is authorized to do so. J2EE containers such as EJB and servlet containers support both programmatic and declarative security that can be used to define authorization policies. Programmatic and declarative security are defined in

the individual specifications. Refer to the EJB and servlet specifications for more details. An application component developer developing components for EIS access must follow the requirements defined in these specifications.

8.5.4 Secure Association

The communication between an application server and an EIS can be subject to security threats such as data modification and loss of data. Establishing a secure association counters such threats. A secure association is shared security information that allows a component on the application server to communicate securely with an EIS.

Establishing a secure association includes several steps:

- The resource principal is authenticated to the EIS. This may require that the target principal in the EIS domain authenticate itself back to the application server. A target principal can be set up by the system administrator as a security principal associated with a running EIS instance or specific EIS resource.
- Negotiate quality of protection such as confidentiality and integrity.
- A pair of communicating entities—an application server and an EIS instance—establish a shared security context using the credentials of the resource principal. The security context encapsulates shared state information, required so that communication between the application server and the EIS can be protected through integrity and confidentiality mechanisms. Examples of shared state information are cryptographic keys and message sequence numbers.

A secure association between an application server and an EIS is always established by the resource adapter implementation. Note that a resource adapter library runs within the address space of the application server.

A resource adapter can use any security mechanism to establish the secure association. GSS-API (refer to IETF draft on GSS-API v2[5]) is an example of such a mechanism. Note that the connector architecture does not require use of the GSS-API by a resource adapter or application server.

Configuring a mechanism for establishing secure associations is outside the scope of the connector architecture. This includes setting up the desired quality of protection during secure communication.

Once a secure association is successfully established, the connection is associated with the security context of the resource principal. Subsequently, all application-level invocations to the EIS instance using the connection happen under the security context of the resource principal.

8.6 Roles and Responsibilities

This section describes various roles involved in the security architecture. It also describes responsibilities of each role from the security perspective.

The roles and responsibilities of the Application Component Provider and Deployer are specified in detail in the respective J2EE component model specifications.

8.6.1 Application Component Provider

The following features are common across different J2EE component models from the perspective of an Application Component Provider:

- An Application Component Provider invariably avoids the burden of securing its application and focuses on developing the business functionality of its application.
- A security-aware Application Component Provider can use a simple programmatic interface to manage security at an application level. The programmatic interface enables the Application Component Provider to program access control decisions based on the security context—the principal and role—associated with the caller of a method and to manage programmatic sign-on to an EIS.
- An Application Component Provider specifies security requirements for its application declaratively in a deployment descriptor. The security requirements include security roles, method permissions, and an authentication approach for EIS sign-on.
- More qualified roles—Application Server Vendor, Deployer, System Administrator—have the responsibility of satisfying overall security requirements through the deployment mechanism for resource adapters and components, and managing the security environment.

8.6.2 Deployer

The Deployer specifies security policies that ensure secure access to the underlying EISs from application components. The deployer adapts the *intended* security view of an application for EIS access, specified through a deployment descriptor, to the *actual* security mechanisms and policies used by the application server and EISs in the target operational environment. The Deployer uses tools to accomplish the above task.

The output of the Deployer's work is a security policy descriptor specific to the operational environment. The format of the security policy descriptor is specific to an application server.

The Deployer performs the following deployment tasks for each connection factory reference declared in the deployment descriptor of an application component:

- Provides a connection factory specific security configuration that is needed for opening and managing connections to an EIS instance.
- Binds the connection factory reference in the deployment descriptor of an application component to the JNDI registered reference for the connection factory. Refer to Section 17.5, "JNDI Configuration and Lookup" on page 17-13 for the JNDI configuration of a connection factory during deployment of a resource adapter. The deployer can use the JNDI `LinkRef` mechanism to create a symbolic link to the actual JNDI name of the connection factory.
- If the value of the `res-auth` deployment descriptor element is `Container`, the Deployer is responsible for configuring the security information for EIS sign-on. For example, the Deployer sets up the principal mapping for EIS sign-on.

8.6.3 Application Server

The application server provides a security environment with specific security policies and mechanisms that support the security requirements of the deployed application components and resource adapters, thereby ensuring a secure access to the connected EISs.

The typical responsibilities of an application server are as follows:

- Provide tools to set up security information for a resource principal and EIS sign-on when `res-auth` element is set to `Container`. This includes support for principal delegation and mapping for configuring a resource principal.
- Provide tools to support management and administration of its security domain. For example, security domain administration can include setting up and maintaining both underlying authentication services and trusts between domains, plus managing principals, including identities, keys, and attributes. Such administration is typically security technology specific and is outside the scope of the connector architecture.
- Support a single sign-on mechanism that spans the application server and multiple EISs. The security mechanisms and policies through which single sign-on is achieved are outside the scope of the connector architecture.

The Appendix E, "JAAS Based Security Architecture" specifies how JAAS can be used by an application server to support the requirements of the connector security architecture.

8.6.4 EIS Vendor

EIS provides a security infrastructure and environment that supports the security requirements of the client applications. An EIS can have its own security domain with a specific set of security policies and mechanisms, or it can be set up as part of an enterprise-wide security domain.

8.6.5 Resource Adapter Provider

The resource adapter provider provides a resource adapter that supports the security requirements of the underlying EIS.

The resource adapter implements the security contract specified as part of the connector architecture. Chapter 9, “Security Contract” specifies the security contract and related requirements for a resource adapter.

The resource adapter specifies its security capabilities and requirements through its deployment descriptor. Section 17.5.4, “Requirements” on page 17-23 specifies a standard deployment descriptor for a resource adapter.

8.6.6 System Administrator

The system administrator typically works in close association with administrators of multiple EISs that have been deployed in an operational environment. The system administration tasks can also be performed by the Deployer.

The following tasks are illustrative examples of the responsibilities of the system administrator:

- Set up an operational environment based on the technology and requirements of the authentication service, and if an enterprise directory is supported.
- Configure the user account information for both the application server and the EIS in the enterprise directory. The user account information from the enterprise directory can then be used for authentication of users requesting connectivity to the EIS.
- Establish a password synchronization mechanism between the application server and the EIS. This ensures that the user’s security information is identical on both the application server and the EIS. When an EIS requires authentication, the application server passes the user’s password to the EIS.

Security Contract

This chapter specifies the security contract between the application server and the EIS. It also specifies the responsibilities of the Resource Adapter Provider and the Application Server Vendor for supporting the security contract.

This chapter references the following chapters and documents:

- The security model specified in the J2EE platform specification (Section 8., “Java 2 Platform Enterprise Edition (J2EETM), Platform specification, version 1. 4:” on page F-1).
- Security architecture specified in Chapter 8, “Security Architecture.”
- Security scenarios based on the connector architecture (Refer to Appendix D, “Security Scenarios”).

9.1 Security Contract

The security contract between the application server and the resource adapter extends the connection management contract (described in Chapter 6, “Connection Management”) by adding security-specific details.

This security contract supports EIS sign-on by:

- Passing the connection request from the resource adapter to the application server, enabling the application server to hook-in security services.
- Propagation of the security context, that is, JAAS `Subject` with principal and credentials, from the application server to the resource adapter.

9.1.1 Interfaces and Classes

The security contract includes the following classes and interfaces:

9.1.2 Subject

The following text has been taken from the JAAS specification. For detailed information, refer to the JAAS specification (Section 7., “Java Authentication and Authorization Service, version 1.0:” on page F-1).

A `Subject` represents a grouping of related information for a single entity, such as a person. Such information includes the `Subject`’s identities and its security-related attributes, for example, passwords and cryptographic keys. A `Subject` can have multiple identities. Each identity is represented as a `Principal` within the `Subject`. A `Principal` simply binds a name to a `Subject`.

A `Subject` can also own security-related attributes, which are referred to as `Credentials`. Sensitive credentials that require special protection, such as private cryptographic keys, are stored within a private credential set.

The `Credentials` intended to be shared, such as public key certificates or Kerberos server tickets, are stored within a public credential set. Different permissions are required to access and modify different credential sets.

The `getPrincipals` method retrieves all the principals associated with a `Subject`. The methods `getPublicCredentials` and `getPrivateCredentials` respectively retrieve all the public or private credentials belonging to a `Subject`. The methods defined in the `Set` class modify the returned set of principals and credentials.

9.1.3 Resource Principal

The interface `java.security.Principal` represents a resource principal. The following code extract shows the `Principal` interface:

```
public interface java.security.Principal {
    public boolean equals(Object another);
    public String getName();
    public String toString();
    public int hashCode();
}
```

The method `getName` returns the name of a resource principal.

An application server should use the `Principal` interface, or any derived interface, to pass a resource principal as part of a `Subject` to a resource adapter.

9.1.4 GenericCredential

Note – This interface, introduced in Version 1.0 of this specification, has been deprecated. The preferred way to represent generic credential information is via the `org.ietf.jgss.GSSCredential` interface in J2SE Version 1.4, which provides similar functionality.

The interface `javax.resource.spi.security.GenericCredential` defines a security mechanism-independent interface for accessing the security credential of a resource principal.

The `GenericCredential` interface provides a Java wrapper around an underlying mechanism-specific representation of a security credential. For example, the `GenericCredential` interface can be used to wrap Kerberos credentials.

The connector architecture does not define any standard format and requirements for security mechanism specific credentials. For example, a security credential wrapped by a `GenericCredential` interface can have a native representation specific to an operating system.

Note – A contract for the representation of mechanism-specific credentials must be established between an application server and a resource adapter and is outside the scope of the connector architecture. This includes requirements for the exchange of mechanism-specific credentials between a JAAS module and GSS provider. Refer to Appendix E, “JAAS Based Security Architecture” for details on JAAS-based security architecture.

The `GenericCredential` interface enables a resource adapter to extract information about a security credential. The resource adapter can then manage an EIS sign-on for a resource principal by either:

- Using the credentials in an EIS specific manner if the underlying EIS supports the security mechanism type represented by the `GenericCredential` instance, or,
- Using GSS-AP I (Section 5., “RFC: Generic Security Service API (GSS-API) specification, version 2:” on page F-1) if the resource adapter and underlying EIS instance support GSS-API.

9.1.4.1 Interface

The following code extract shows the `GenericCredential` interface:

```
public interface javax.resource.spi.security.GenericCredential {
    public String getName();
    public String getMechType();
    public byte[] getCredentialData()
        throws javax.resource.spi.SecurityException;

    public boolean equals(Object another);
    public int hashCode();
}
```

The `GenericCredential` interface supports a set of getter methods to obtain information about a security credential.

The method `getName` returns the name of the resource principal associated with a `GenericCredential` instance.

The method `getMechType` returns the mechanism type for the `GenericCredential` instance. The mechanism type definition for `GenericCredential` must be consistent with the Object Identifier (OID) based representation specified in the GSS specification (Section 5., “RFC: Generic Security Service API (GSS-API) specification, version 2:” on page F-1). In the `GenericCredential` interface, the mechanism type is returned as a stringified representation of the OID specification.

The `GenericCredential` interface can be used to get security data for a specific security mechanism. An example is authentication data required for establishing a secure association with an EIS instance on behalf of the associated resource principal. The `getCredentialData` method returns the credential representation as an array of bytes. Note that the connector architecture does not define a standard format for the returned credential data.

9.1.4.2 Implementation

If an application server supports the deployment of a resource adapter which supports `GenericCredential` as part of the security contract, the application server must provide an implementation of the `GenericCredential` interface. Refer to the deployment descriptor specification in Section 17.5.4, “Requirements” on page 17-23 for details on how a resource adapter specifies its support for `GenericCredential`.

9.1.5 GSSCredential

This interface `org.ietf.jgss.GSSCredential` is in J2SE Version 1.4. This provides a mechanism to represent generic credential information. The functionality provided by this interface is similar to the deprecated `GenericCredential` interface.

9.1.5.1 Implementation

If an application server supports the deployment of a resource adapter which supports `GSSCredential` as part of the security contract, the application server must provide an implementation of the `GSSCredential` interface. Refer to the deployment descriptor specification in Section 17.5.4, “Requirements” on page 17-23 for details on how a resource adapter specifies its support for `GSSCredential`.

9.1.6 PasswordCredential

The class `javax.resource.spi.security.PasswordCredential` acts as a holder of username and password information. This class enables an application server to pass the username and password to the resource adapter through the security contract.

The method `getUserName` gets the name of the resource principal. The interface `java.security.Principal` represents a resource principal.

The `PasswordCredential` class must implement the `equals` and `hashCode` methods.

```
public final class javax.resource.spi.security.PasswordCredential
    implements java.io.Serializable {
    public PasswordCredential(String userName, char[] password) {
    ... }
    public String getUserName() { ... }
    public char[] getPassword() { ... }

    public ManagedConnectionFactory getManagedConnectionFactory()
        { ... }
    public void setManagedConnectionFactory(
        ManagedConnectionFactory mcf) { ... }

    public boolean equals(Object other) { ... }
    public int hashCode() { ... }
}
```

The `getManagedConnectionFactory` method returns the `ManagedConnectionFactory` instance for which the user name and password has been set by the application server. Refer to Section 9.1.8 “`ManagedConnectionFactory`” to see how a resource adapter uses this method.

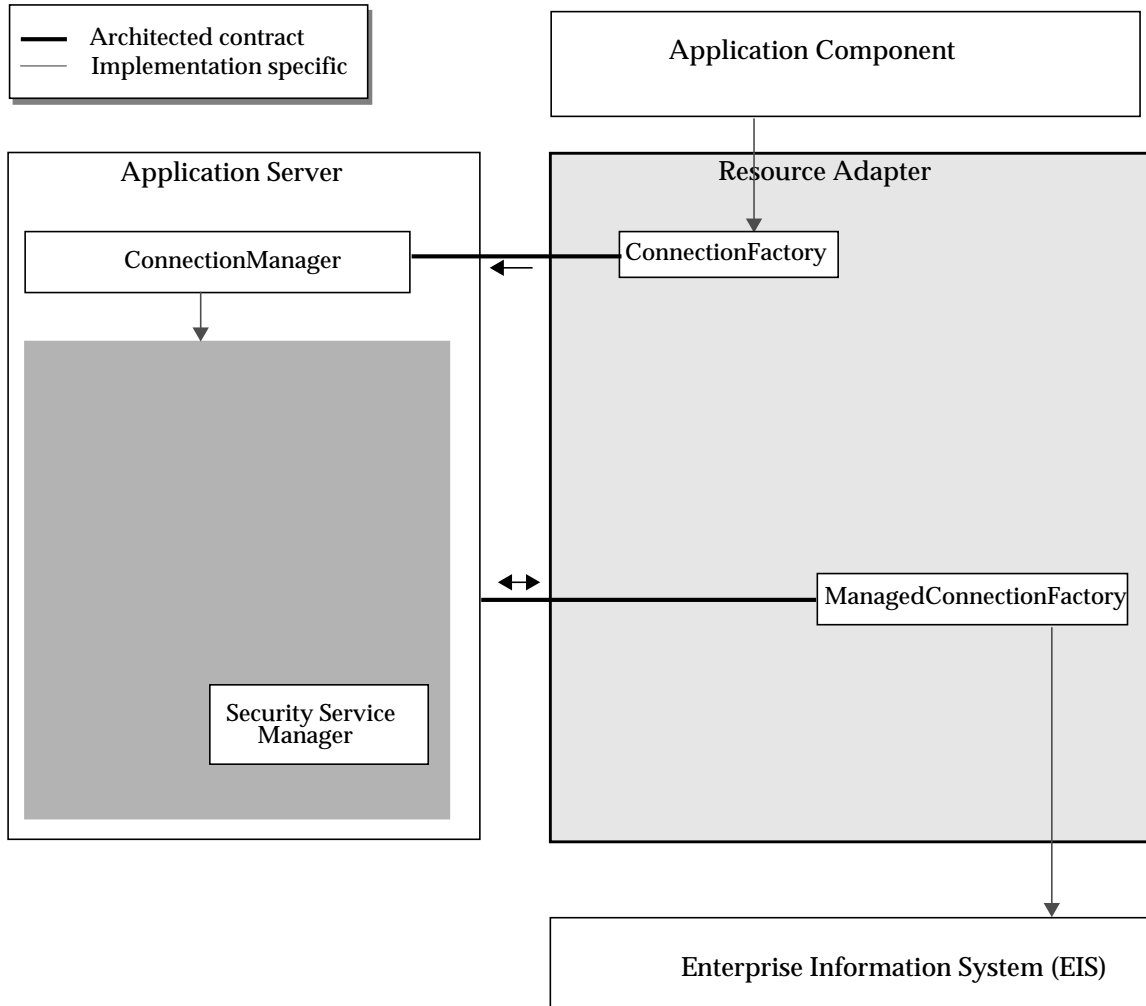
9.1.7 ConnectionManager

The method `allocateConnection` is called by the resource adapter's connection factory instance. This method lets the resource adapter pass a connection request to the application server, so the application server can hook-in security and other services.

```
public interface javax.resource.spi.ConnectionManager
    extends java.io.Serializable {

    public Object allocateConnection(
        ManagedConnectionFactory mcf,
        ConnectionRequestInfo cxRequestInfo)
        throws ResourceException;
}
```

FIGURE 9-1 Security Contract



Depending on whether the application server or application component is configured to be responsible for managing EIS sign-on (refer to Section 8.6.1, “Application Component Provider” on page 8-9), the resource adapter calls the `ConnectionFactoryManager.allocateConnection` method in one of the following ways:

- **Container-managed Sign-on:** The application component passes no security information in the `getConnection` method and the application server is configured to manage EIS sign-on.

The application server provides the required security information for the resource principal through its configured security policies and mechanisms, for example, principal mapping. The application server requests the authentication of the

resource principal to the EIS either itself or passes authentication responsibility to the resource adapter. This aspect is explained later in the specification of the `ManagedConnectionFactory` interface.

- **Component-managed Sign-on:** In this case, the application component provides explicit security information in the `getConnection` method. The resource adapter invokes the `allocateConnection` method by passing security information in the `ConnectionRequestInfo` parameter. Since the security information in the `ConnectionRequestInfo` is opaque to the application server, the application server should rely on the resource adapter to manage EIS sign-on, as explained in the `ManagedConnectionFactory` interface specification under option C.

9.1.8 ManagedConnectionFactory

The following code extract shows the methods of the `ManagedConnectionFactory` interface that are relevant to the security contract:

```
public interface javax.resource.spi.ManagedConnectionFactory
    extends java.io.Serializable {

    public ManagedConnection createManagedConnection(
        javax.security.auth.Subject subject,
        ConnectionRequestInfo cxRequestInfo)
        throws ResourceException;

    ...

}
```

During the JNDI lookup, the `ManagedConnectionFactory` instance is configured by the application server with a set of configuration properties. These properties include default security information and EIS instance-specific information, such as hostname and port number, required for initiating a sign-on to the underlying EIS during the creation of a new physical connection.

The default security configuration on a `ManagedConnectionFactory` can be overridden by security information provided either by a component, in component managed sign-on, or by the container, in container-managed sign-on.

The `createManagedConnection` method is used by the application server when it requests the resource adapter to create a new physical connection to the underlying EIS.

9.1.8.1 Contract for the Application Server

The application server may provide specific security services, such as principal mapping and delegation, and single sign-on, before using the security contract with the resource adapter. For example, the application server can map the caller principal to a resource principal before calling the `createManagedConnection` method to create a new connection under the security context of the resource principal.

In container-managed sign-on, the application server is responsible for creating a `Subject` instance using its implementation-specific security mechanisms and configuration. This should happen before the application server calls the `createManagedConnection` method of the `ManagedConnectionFactory`. The resource adapter is driven by the application server and acts as consumer of security information in the created `Subject`.

If the application server maintains a cache of the security credentials, such as Kerberos ticket granting ticket (TGT), the application server should reuse the credentials as part of the newly created `Subject` instance. For example, the application server uses the `Subject.getPrivateCredentials().add(credential)` method to add a credential to the private credential set.

FIGURE 9-2 Security Contract: Subject Interface and its Containment Hierarchy

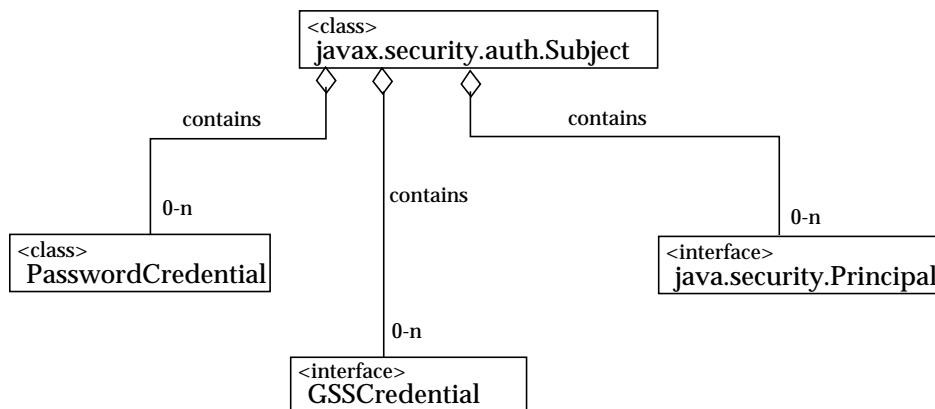


FIGURE 9-2 on page 10 shows the relationship between the `Subject`, `Principal`, `PasswordCredential` and `GSSCredential` interfaces. Note that in the following options A and B defined for `createManagedConnection` method invocation, the `Subject` instance contains a single resource principal, represented as `java.security.Principal`, and multiple credentials.

The application server has the following options for invoking the `createManagedConnection` method:

- **Option A:** The application server invokes the `createManagedConnection` method by passing in a non-null `Subject` instance that carries a *single* resource principal and its corresponding password-based credentials, represented by the class `PasswordCredential` that provides the user name and password. The `PasswordCredential` should be set in the `Subject` instance as part of the private credential set. Note that the passed `Subject` can contain multiple `PasswordCredential` instances.

The resource adapter extracts the username and password from this `Subject` instance by looking for the `PasswordCredential` instance in the `Subject`, and uses this security information to sign-on to the EIS instance during connection creation.

- **Option B:** The application server invokes the `createManagedConnection` method by passing in a non-null `Subject` instance that carries a *single* resource principal and its security credentials. In this option, credentials are represented through the `GSSCredential` interface. A typical example is a `Subject` instance with Kerberos credentials.

For example, an application server may use this option for `createManagedConnection` method invocation when the resource principal is impersonating the caller or initiating principal, and has valid credentials acquired through impersonation. An application server may also use this option for principal mapping scenarios with credentials of a resource principal represented through the `GSSCredential` interface.

Note that sensitive credentials requiring special protection, such as private cryptographic keys, are stored within a private credential set, while credentials intended to be shared, such as public key certificates or Kerberos server tickets, are stored within a public credential set. The two methods `getPrivateCredentials` and `getPublicCredentials` should be used accordingly.

In the case of Kerberos mechanism type, the application server must pass the principal's ticket granting ticket (TGT) to a resource adapter in a private credential set.

The resource adapter uses the resource principal and its credentials from the `Subject` instance to go through the EIS sign-on process before creating a new connection to the EIS.

- **Option C:** The application server invokes the `createManagedConnection` method by passing a null `Subject` instance. The application server must use this option for the component-managed sign-on case. In this option, security information is carried in the `ConnectionRequestInfo` instance. The application server does not provide any security information that can be used by the resource adapter for managing EIS sign-on.

During the deployment of a resource adapter, the application server must be configured to use one of the above specified invocation options. Refer to Chapter 17, “Packaging Requirements” for more details.

9.1.8.2 Contract for Resource Adapter

A resource adapter can do EIS sign-on and connection creation in an implementation-specific way, or it can use the GSS-API. The latter option is specified in Appendix E, “JAAS Based Security Architecture.” A resource adapter has the following options, corresponding to the options for an application server, for handling the invocation of the `createManagedConnection` method:

- **Option A:** The resource adapter explicitly checks whether the passed `Subject` instance carries a `PasswordCredential` instance using the `Subject.getPrivateCredentials` method.

Note that the security contract assumes that a resource adapter has the necessary security permissions to extract a private credential set from a `Subject` instance. The specific mechanism through which such permission is set up is outside the scope of the connector architecture.

If the `Subject` instance contains a `PasswordCredential` instance, the resource adapter extracts the username and password from the `PasswordCredential`. It uses the security information to authenticate the resource principal, corresponding to the username, to the EIS during the creation of a connection. In this case, the resource adapter uses an authentication mechanism that is EIS specific.

Since a `Subject` instance can carry multiple `PasswordCredential` instances, a `ManagedConnectionFactory` should only use a `PasswordCredential` instance that has been specifically passed to it through the security contract. The `getManagedConnectionFactory` method enables a `ManagedConnectionFactory` instance to determine whether or not a `PasswordCredential` instance is to be used for sign-on to the target EIS instance. The `ManagedConnectionFactory` implementation uses the `equals` method to compare itself with the passed instance.

- **Option B:** The resource adapter explicitly checks whether the passed `Subject` instance carries a `GSSCredential` instance using the `getPrivateCredentials` and `getPublicCredentials` methods defined in the `Subject` interface.

In the case of Kerberos mechanism type, the resource adapter must extract Kerberos credentials using the `getPrivateCredentials` method in the `Subject` interface.

The resource adapter uses the resource principal and its credentials, represented by the `GSSCredential` interface, in the `Subject` instance to go through the EIS sign-on process. For example, this option is used for Kerberos-based credentials that have been acquired by the resource principal through impersonation.

A resource adapter uses the getter methods defined in the `GSSCredential` interface to extract information about the credential and its principal. If a resource adapter is using the GSS mechanism, the resource adapter uses a reference to the `GSSCredential` instance in an opaque manner and is not required to understand any mechanism-specific credential representation. However, a resource adapter may need to interpret credential representation if the resource adapter initiates authentication in an implementation-specific manner.

- **Option C:** If the application server invokes the `ManagedConnectionFactory.createManagedConnection` method with a null `Subject` instance, a resource adapter has the following options:
 - The resource adapter should extract security information passed through the `ConnectionRequestInfo` instance. The resource adapter should authenticate the resource principal by combining the configured security information on the `ManagedConnectionFactory` instance with the security information passed through the `ConnectionRequestInfo` instance. The default behavior for the resource adapter is to allow the security information in the `ConnectionRequestInfo` parameter to override the configured security information in the `ManagedConnectionFactory` instance.
 - If the resource adapter does not find any security configuration in the `ConnectionRequestInfo` instance, the resource adapter uses the default security configuration in the `ManagedConnectionFactory` instance.
 - If the EIS does not require authentication, the resource adapter does not need any security information from the `ConnectionRequestInfo` instance, and hence may ignore such security information. This may happen due to a disconnect between the application and the resource adapter.

In the case of option A and option B, a resource adapter should throw a `javax.resource.spi.SecurityException`, if the credential information contained in the `Subject` instance is insufficient to perform authentication. A non-null `Subject` instance with no credentials is not equivalent to a null `Subject` instance, since they indicate different sign-on modes, and hence the resource adapter may handle them differently. A non-null `Subject` instance with no credentials may be interpreted by the resource adapter as follows:

- If the EIS requires authentication, the resource adapter should throw a `javax.resource.spi.SecurityException`. That is, an empty or insufficient credential information is an error.
- If the EIS does not require authentication, the resource adapter does not need any security information from the non-null `Subject` instance, and hence may ignore the `Subject` instance. This may happen due to a disconnect between the application and the resource adapter.

9.1.9 ManagedConnection

A resource adapter can re-authenticate a physical connection (that is, one that already exists in the connection pool under a different security context) to the underlying EIS. A resource adapter performs re-authentication when an application server calls the `getConnection` method with a security context, passed as a `Subject` instance, different from the context previously associated with the physical connection.

If a resource adapter supports re-authentication, the `matchManagedConnections` method in `ManagedConnectionFactory` may return a matched `ManagedConnection` instance with the assumption that the `ManagedConnection.getConnection` method will later switch the security context through re-authentication. Note that the `matchManagedConnections` method should consider a `ManagedConnection` instance as immutable. There is no authentication involved in the `matchManagedConnections` method.

Support for re-authentication depends on whether an underlying EIS supports the re-authentication mechanism for existing physical connections. If a resource adapter does not support re-authentication, the `getConnection` method should throw a `javax.resource.spi.SecurityException` if the passed `Subject` in the `getConnection` method is different from the security context associated with the `ManagedConnection` instance.

```
public interface javax.resource.spi.ManagedConnection {
    public Object getConnection(
        javax.security.auth.Subject subject,
        ConnectionRequestInfo cxRequestInfo)
        throws ResourceException;
    ...
}
```

The `getConnection` method returns a new connection handle. If re-authentication is successful, the resource adapter has changed the security context of the underlying `ManagedConnection` instance to that associated with the passed `Subject` instance.

A resource adapter has the following options for handling `ManagedConnection.getConnection` invocation if it supports re-authentication:

- **Option A:** The resource adapter extracts the `PasswordCredential` instance from the `Subject` and performs an EIS-specific authentication. This option is similar to option A defined in the specification of the method `createManagedConnection` on the interface `ManagedConnectionFactory` (refer to Section 9.1.8 “`ManagedConnectionFactory`”).

- **Option B:** The resource adapter extracts `GSSCredential` instance from the `Subject` and manages authentication either through the GSS mechanism or an implementation-specific mechanism. This option is similar to option B defined in the specification of the method `createManagedConnection` on the interface `ManagedConnectionFactory` (refer to Section 9.1.8 “`ManagedConnectionFactory`”).
- **Option C:** In this case, the `Subject` parameter is `null`. The resource adapter extracts security information from the `ConnectionRequestInfo` (if there is any) and performs authentication in an implementation-specific manner. This option is similar to option C defined in the specification of the method `createManagedConnection` on the interface `ManagedConnectionFactory` (refer to Section 9.1.8 “`ManagedConnectionFactory`”).

9.2 Requirements

The following are the requirements defined by the security contract:

9.2.1 Resource Adapter

The following are the requirements defined for a resource adapter:

- The resource adapter must support the security contract by implementing the method `ManagedConnectionFactory.createManagedConnection`.
- The resource adapter is not required to support re-authentication as part of its `ManagedConnection.getConnection` method implementation.
- If the security information provided by the component or the container is not adequate to authenticate the caller, or if the security information is erroneous, the resource adapter must throw a `SecurityException` to indicate the error condition.
- The resource adapter must specify its support for the security contract as part of its deployment descriptor. The relevant deployment descriptor elements are: `authentication-mechanism`, `authentication-mechanism-type`, `reauthentication-support` and `credential-interface` (refer to Section 17.5.4, “Requirements” on page 17-23 for details).

9.2.2 Application Server

The following are the requirements defined for an application server:

- The application server must use the method `ManagedConnectionFactory.createManagedConnection` to pass the security context to the resource adapter during EIS sign-on.
- The application server must be capable of using options A and C as specified in Section 9.1.8 “`ManagedConnectionFactory`” for the security contract.
- The application server provides an implementation of the `GSSCredential` interface if the following conditions are both true:
 - The application server supports authentication mechanisms, specified as `authentication-mechanism-type` in the deployment descriptor, other than `BasicPassword` mechanism. For example, the application server should implement the `GSSCredential` interface to support the `kerbv5` authentication mechanism type.
 - The application server supports the deployment of resource adapters that are capable of handling `GSSCredential`, and thereby option B as specified in Section 9.1.8 “`ManagedConnectionFactory`”, as part of the security contract.
- The application server must implement the method `allocateConnection` in its `ConnectionManager` implementation.
- The application server must configure its use of the security contract based on the security requirements specified by the resource adapter in its deployment descriptor. For example, if a resource adapter specifies that it supports only `BasicPassword` authentication, the application server should use the security contract to pass a `PasswordCredential` instance to the resource adapter.

Work Management

This chapter specifies a contract between an application server and a resource adapter that allows a resource adapter to do work, such as monitor network endpoints and call application components, by submitting `Work` instances to an application server for execution. The application server dispatches threads to execute submitted `Work` instances. This allows a resource adapter to avoid creating or managing threads directly, provides a mechanism for a resource adapter to perform work, allows an application server to efficiently pool threads, and have more control over its runtime environment. The resource adapter can control the security context and transaction context with which `Work` instances are executed.

10.1 Overview

Some resource adapters merely function as a passive library that executes in the context of an application thread. They do not need to create threads explicitly to do their work. But more sophisticated resource adapters may need threads to function properly. Such resource adapters may use threads to listen to network endpoints, process incoming data, communicate with a network peer, do its internal work, or dispatch calls to application components.

Even though a resource adapter may create Java threads directly and use them to do its work, an application server may prevent it from creating threads for efficiency, security, and manageability reasons. In such situations, a resource adapter needs a mechanism to obtain threads from an application server to do its work.

The work management contract provides such a mechanism which allows a resource adapter to submit `Work` instances to an application server for execution. The application server dispatches threads to execute submitted `Work` instances. This allows a resource adapter to avoid creating or managing threads directly, provides a mechanism for the resource adapter to do its work, and allows an application server more control over its runtime environment.

There are several advantages in allowing an application server to manage threads instead of a resource adapter:

- An application server is optimally designed to manage system resources such as threads. It may pool threads and reuse them efficiently across different resource adapters deployed in its runtime environment.
- A resource adapter may create non-daemon threads that interfere with the orderly shutdown of an application server. It is desirable for an application server to own all the threads to exercise more control over its runtime environment.
- Since an application server knows the overall state of its runtime environment, it may make better decisions on granting threads to a resource adapter, and this leads to better manageability of its runtime environment.
- An application server may need to enforce control over the runtime behavior of its system components, including resource adapters. For example, an application server may choose to intercept operations on a thread object, perform checks, and enforce correct behavior.
- An application server may disallow resource adapters from creating their own threads based on its security policy setting, enforced by a security manager.

10.2 Goals

- Provide a flexible work execution model to handle the thread needs of a resource adapter.
- Provide a mechanism for an application server to pool and reuse threads.
- Exercise more control over thread behavior in a managed environment.

10.3 Work Management Model

A resource adapter obtains a `WorkManager` instance from the `BootstrapContext` instance provided by the application server during its deployment. The resource adapter may create `Work` instances to do its work and submit them to the `WorkManager` along with an optional execution context for execution.

The application server has a pool of free threads waiting for a `Work` instance to be submitted. When a `Work` instance is submitted, one of the free threads picks up the `Work` instance, sets up an appropriate execution context and calls the `run` method on the `Work` instance. The application server is free to choose an appropriate thread to execute the `Work` instance. There is no restriction on the number of `Work` instances submitted by a resource adapter or when `Work` instances may be submitted. When the `run` method on the `Work` instance completes, the application server reuses the thread.

The application server may decide to reclaim active threads based on load conditions. It calls the `release` method on specific `Work` instances from a separate thread. This serves only as a hint to the resource adapter to release the active thread executing the `Work` instance. The resource adapter should periodically monitor such hints and do the necessary internal cleanup to avoid any inconsistencies. It is expected that a resource adapter uses thread resources carefully and releases them when not in use.

The application server is free to implement its own thread pooling strategy. However, the application server must use threads of the same thread priority level to process `Work` instances submitted by a specific resource adapter. This ensures that multiple threads processing `Work` instances from the same resource adapter have equal claim over CPU resources. This assumption helps the resource adapter build its own internal priority-based task queue without having to worry about thread priority levels.

10.3.1 Requirements

- The application server must use threads of the same thread priority level to process `Work` instances submitted by a specific resource adapter.

FIGURE 10-1 Work Management Contract (object diagram)

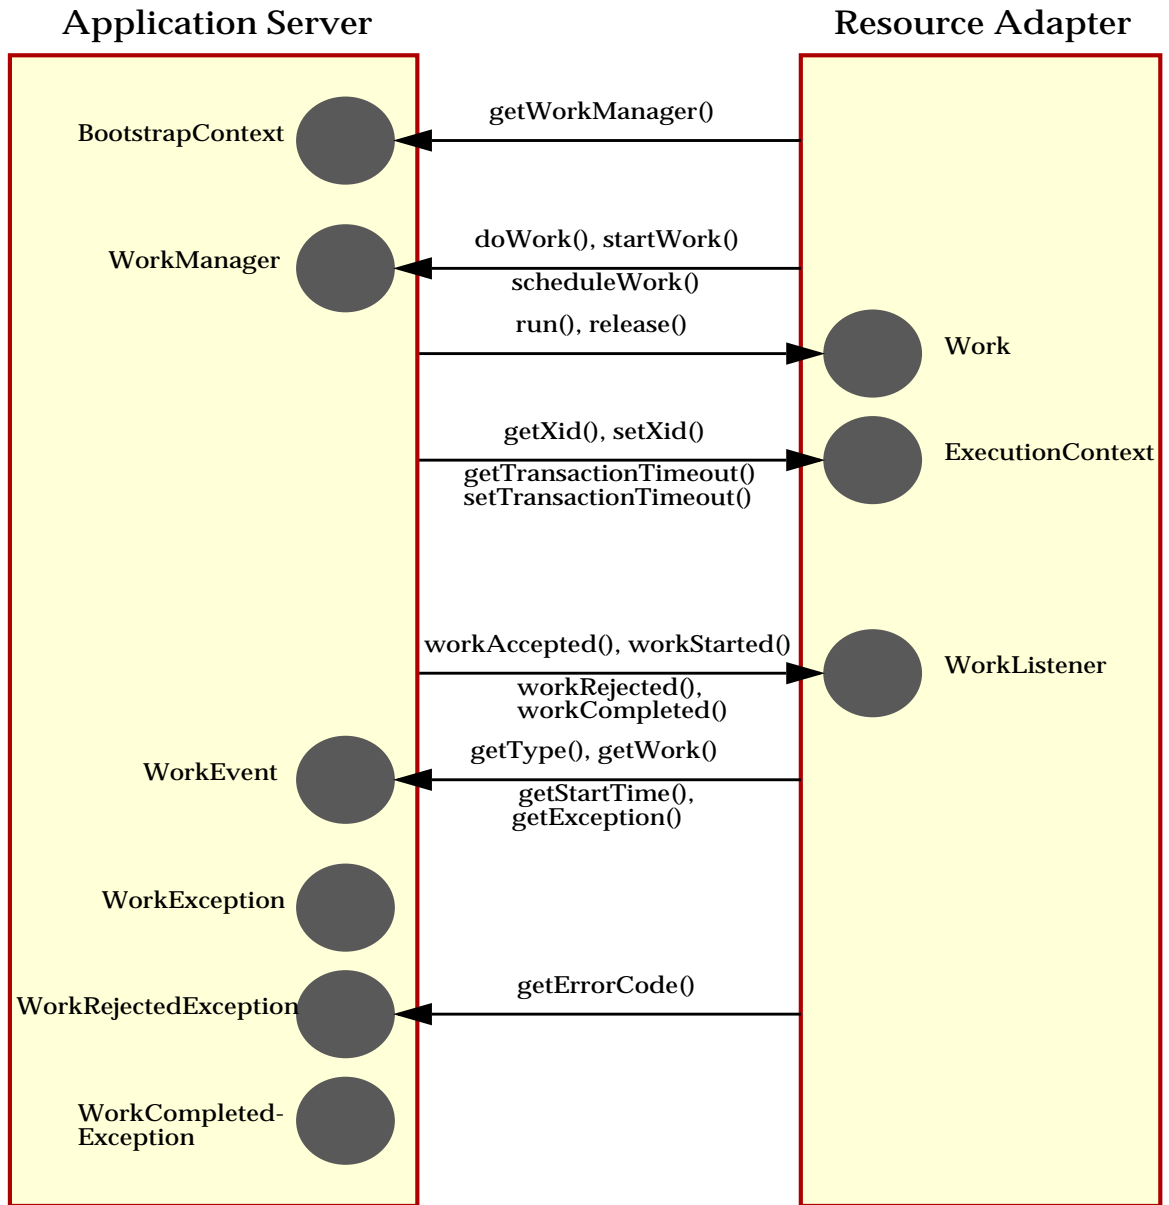
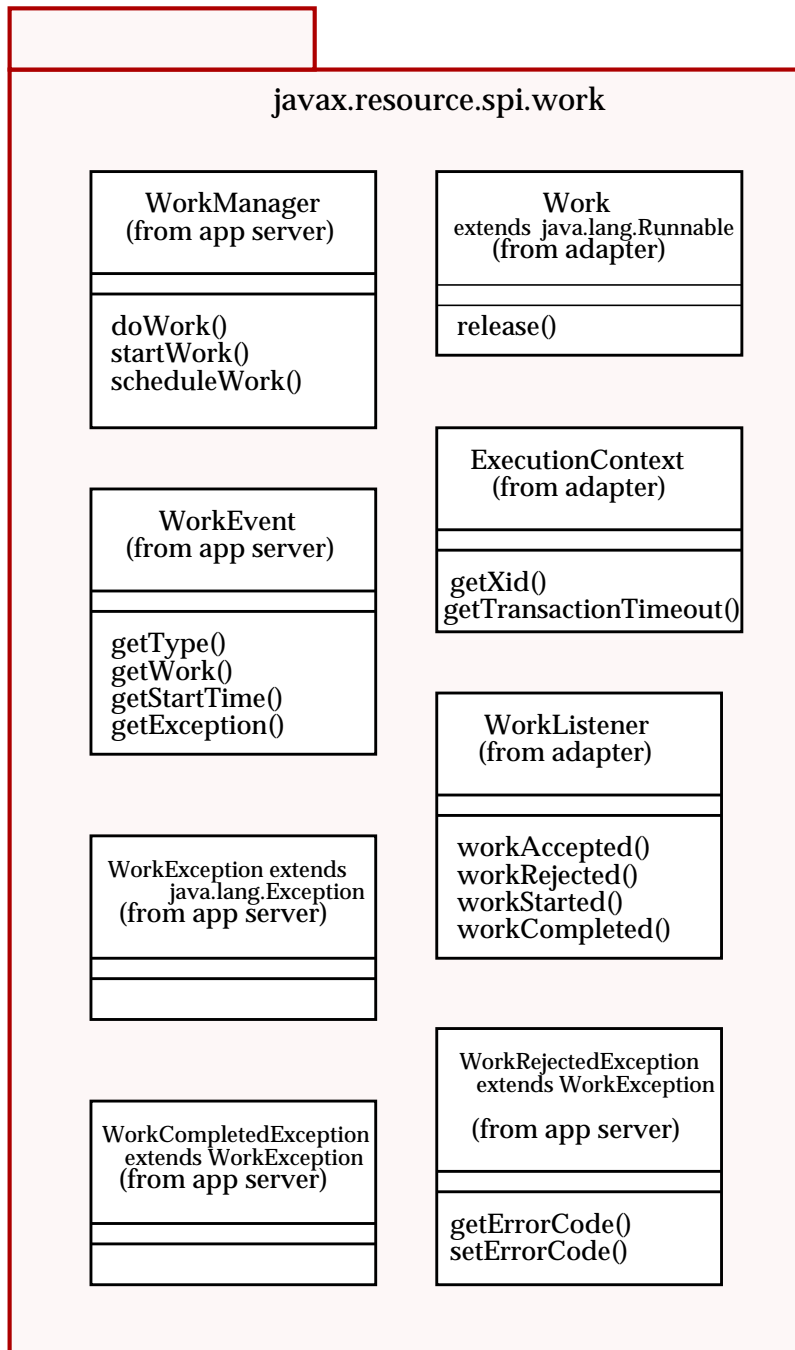


FIGURE 10-2 Work Management Contract (interfaces)



CODE EXAMPLE 10-1 javax.resource.spi.work

```
package javax.resource.spi.work;

import java.lang.Object;
import java.lang.Runnable;
import java.lang.Exception;
import java.lang.Throwable;
import java.util.EventObject;
import java.util.EventListener;

import javax.transaction.xa.Xid;
import javax.resource.ResourceException;
import javax.resource.NotSupportedException;

public interface Work extends Runnable {
    void release();
}

public interface WorkManager {

    long IMMEDIATE = 0L; // immediate action
    long INDEFINITE = Long.MAX_VALUE; // no time constraint
    long UNKNOWN = -1; // indicates an unknown value.

    void doWork(Work work) // startTimeout = INDEFINITE
        throws WorkException;
    void doWork(Work work, long startTimeout, ExecutionContext ctx,
        WorkListener lsnr) throws WorkException;
    long startWork(Work work) // startTimeout = INDEFINITE
        throws WorkException;
    long startWork(Work work, long startTimeout,
        ExecutionContext ctx, WorkListener lsnr)
        throws WorkException;
    void scheduleWork(Work work) // startTimeout = INDEFINITE
        throws WorkException;
    void scheduleWork(Work work, long startTimeout,
        ExecutionContext ctx, WorkListener lsnr)
        throws WorkException;
}

public interface WorkListener extends EventListener {
    void workAccepted(WorkEvent e);
    void workRejected(WorkEvent e);
    void workStarted(WorkEvent e);
    void workCompleted(WorkEvent e);
}

public class WorkAdapter implements WorkListener {
    public void workAccepted(WorkEvent e) {}
}
```

CODE EXAMPLE 10-1 javax.resource.spi.work

```
    public void workRejected(WorkEvent e) {}
    public void workStarted(WorkEvent e) {}
    public void workCompleted(WorkEvent e) {}
}

public class WorkEvent extends EventObject {

    public static final int WORK_ACCEPTED = 1;
    public static final int WORK_REJECTED = 2;
    public static final int WORK_STARTED = 3;
    public static final int WORK_COMPLETED = 4;

    public WorkEvent(Object source, int type, Work work,
        WorkException exc) { ... }

    public WorkEvent(Object source, int type, Work work,
        WorkException exc, long startDuration) { ... }

    public int getType() { ... }
    public Work getWork() { ... }
    public long getStartDuration() { ... }
    public WorkException getException() { ... }
}

public class ExecutionContext {

    public void setXid(xid) { ... }
    public Xid getXid() { ... }
    public long getTransactionTimeout() { ... }
    public void setTransactionTimeout(long seconds)
        throws NotSupportedException { ... }
}

public class WorkException extends ResourceException {

    // Indicates an internal error condition.
    public static final String INTERNAL = "-1";

    // Undefined error code.
    public static final String UNDEFINED = "0";

    // Indicates start timeout expiration.
    public static final String START_TIMED_OUT = "1";

    // Indicates that concurrent work within a transaction is
    // disallowed.
    public static final String TX_CONCURRENT_WORK_DISALLOWED = "2";
```

CODE EXAMPLE 10-1 javax.resource.spi.work

```
// Indicates a failure in recreating the specified transaction.
public static final String TX_RECREATE_FAILED = "3";

public WorkException() { ... }
public WorkException(String message) { ... }
public WorkException(Throwable cause) { ... }
public WorkException(String message, Throwable cause) { ... }
public String getMessage() { ... }
}

public class WorkRejectedException extends WorkException {

    public WorkRejectedException() { ... }
    public WorkRejectedException(String message) { ... }
    public WorkRejectedException(Throwable cause) { ... }
    public WorkRejectedException(String message, Throwable cause)
        { ... }
}

public class WorkCompletedException extends WorkException {

    public WorkCompletedException() { ... }
    public WorkCompletedException(String message) { ... }
    public WorkCompletedException(Throwable cause) { ... }
    public WorkCompletedException(String message, Throwable cause)
        { ... }
}
```

10.3.2 Work Interface

The `Work` interface models a `Work` instance which is executed by a `WorkManager` upon submission. This is implemented by a resource adapter.

```
public interface Work extends Runnable {
    void release();
}
```

- **run method:** The `WorkManager` dispatches a thread that calls the `run` method to begin execution of a `Work` instance. The execution completes when the `run` method returns, with or without an exception. The `Work` instance can treat the calling thread as any Java thread. However, the application server may interpose `java.lang.Thread` methods and perform checks. The `WorkManager` must catch any exception thrown during `Work` processing, which includes execution context setup, and wrap it with a `WorkCompletedException` set to an appropriate error code, which indicates the nature of the error condition.
- **release method:** The `WorkManager` may call the `release` method to request the active `Work` instance to complete execution as soon as possible. This would be called on a separate thread than the one currently executing the `Work` instance. Since this method call causes the `Work` instance to be simultaneously acted upon by multiple threads, the `Work` instance implementation must be thread-safe, and this method must be re-entrant.

The application server thread that calls the `run` method in the `Work` implementation must execute with an unspecified context if no execution context has been specified, or must execute with the specified execution context. It must have at least the same level of security permissions as that of the resource adapter instance. Further, it does not have access to a JNDI context.

Note – The JNDI context of an accessing application is available to a resource adapter via the thread that uses its connection object. Refer to the note in Section 6.4.1, “Managed Application Scenario” on page 6-6. The thread that accesses the connection object could be an application thread, or, could be a `Work` object accessing an application component. In the latter case, the worker thread gains access to the application’s JNDI context during the method call on the component.

Both the `run` and `release` methods in the `Work` implementation may contain synchronization blocks but they must not be declared as synchronized methods.

10.3.3 WorkManager Interface

The `WorkManager` interface provides a mechanism to submit `Work` instances for execution. This is implemented by an application server. A `WorkManager` instance can be obtained by calling the `getWorkManager` method of the `BootstrapContext` instance. The `BootstrapContext` instance is provided by the application server when a resource adapter instance is bootstrapped. The `WorkManager` instance is not required to be unique.

This `WorkManager` facility frees the resource adapter from having to create Java threads directly to do its work. Further, this allows efficient pooling of thread resources by the application server and more control over thread usage.

```
public interface WorkManager {

    long IMMEDIATE = 0L; // immediate action
    long INDEFINITE = Long.MAX_VALUE; // no time constraint
    long UNKNOWN = -1; // unknown start delay duration

    void doWork(Work work) // startTimeout = INDEFINITE
        throws WorkException;
    void doWork(Work work, long startTimeout, ExecutionContext,
        WorkListener) throws WorkException;
    long startWork(Work work) // startTimeout = INDEFINITE
        throws WorkException;
    long startWork(Work work, long startTimeout, ExecutionContext,
        WorkListener) throws WorkException;
    void scheduleWork(Work work) // startTimeout = INDEFINITE
        throws WorkException;
    void scheduleWork(Work work, long startTimeout,
        ExecutionContext, WorkListener) throws WorkException;
}
```

- **doWork method:** This call blocks until the `Work` instance completes execution. The application server may execute a `Work` instance submitted via the `doWork` method using the same calling thread. This method is useful to do work synchronously. For nested `Work` submissions, this provides a first in, first out (FIFO) execution start ordering and last in, first out (LIFO) execution completion ordering guarantee.
- **startWork method:** This call blocks until the `Work` instance starts execution but not until its completion. This returns the time elapsed in milliseconds from `Work` acceptance until the start of execution. Note, this does not offer real-time guarantees. A value of -1 (`WorkManager.UNKNOWN`) must be returned, if the actual start delay duration is unknown. This method is equivalent to the `java.lang.Thread.start` method. For nested `Work` submissions, this provides a FIFO execution start ordering guarantee, but no execution completion ordering guarantee.
- **scheduleWork method:** This call does not block and returns immediately once a `Work` instance has been accepted for processing. This is useful for doing work asynchronously. This does not provide any execution start or execution completion ordering guarantee for nested `Work` submissions.

The optional `startTimeout` parameter specifies a time duration in milliseconds within which the execution of the `Work` instance must start. Otherwise, the `Work` instance is rejected with a `WorkRejectedException` set to an appropriate error code (`WorkException.START_TIMED_OUT`). Note, this does not offer real-time guarantees.

The optional `ExecutionContext` parameter provides an execution context with which the `Work` instance must be executed. The execution context is represented by an `ExecutionContext` instance containing context information. The resource adapter is responsible for populating the `ExecutionContext` instance with an appropriate execution context. The default implementation provides a null context, that is, an `ExecutionContext` instance with null values. A `Work` instance with null context executes with an unspecified context.

The optional `WorkListener` parameter provides a callback event listener object which is notified when the various `Work` processing events (work accepted, work rejected, work started, work completed) occur. Refer to Section 10.3.4 “`WorkListener` Interface and `WorkEvent` Class”.

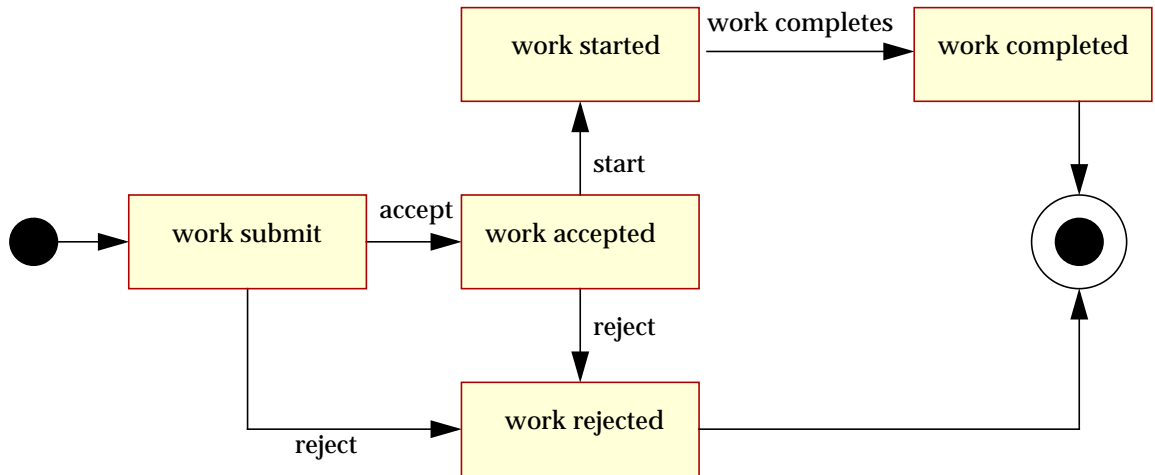
The various stages in `Work` processing are:

10.3.3.1 Work Submit

A `Work` instance is being submitted for execution. The `Work` instance may either be accepted or rejected with a `WorkRejectedException` set to an error code. A submitted `Work` instance, irrespective of the mode of submission: `doWork` method,

`startWork` method or `scheduleWork` method, does not automatically inherit the submitter's execution context. It executes with an unspecified execution context if none is specified, or it executes with the specified context.

FIGURE 10-3 Work Processing Stages and their Outcomes



10.3.3.2 Work Accepted

The submitted `Work` instance has been accepted for further processing. The accepted `Work` instance may either start execution or may be rejected again with a `WorkRejectedException` set to an appropriate error code.

There is no guarantee on when the execution starts unless a start timeout duration is specified. When a start timeout is specified, the `Work` execution must be started within the specified duration, failing which a `WorkRejectedException` set to an error code `WorkException.TIMED_OUT` is thrown. This is not a real-time guarantee. The start delay duration is measured from the moment a `Work` instance is accepted for processing.

10.3.3.3 Work Rejected

The `Work` instance has been rejected. The `Work` instance may be rejected during `Work` submittal or after the `Work` instance has been accepted, but before `Work` instance starts execution. The rejection may be due to internal factors or start timeout expiration. A `WorkRejectedException` with an appropriate error code which indicates the nature of the error condition, is thrown in both cases.

Since the `scheduleWork` method returns after a `Work` instance has been accepted and does not block until a `Work` instance starts, a callback event listener may be used to receive the `WorkRejectedException`. See Section 10.3.4 “`WorkListener` Interface and `WorkEvent` Class” for details.

10.3.3.4 Work Started

The execution of the `Work` instance has started. This means a thread has been allocated for `Work` execution. But this does not guarantee that the allocated thread has been scheduled to run on a CPU resource. Once execution is started, the allocated thread sets up an appropriate execution context and calls the `run` method on the `Work` instance. Note, any exception thrown during execution context setup or while executing the `run` method on the `Work` instance leads to processing completion.

10.3.3.5 Work Completed

The execution of the `Work` instance has been completed. The execution may complete with or without an exception. The `WorkManager` must catch any exception thrown during `Work` processing, which includes execution context setup, and wrap it with a `WorkCompletedException` set to an appropriate error code which indicates the nature of the error condition.

Since the `scheduleWork` method and `startWork` method do not block until execution completion, a callback event listener may be used to receive the `WorkCompletedException`. See Section 10.3.4 “`WorkListener` Interface and `WorkEvent` Class” for details).

10.3.3.6 Requirements

- The application server must implement the `WorkManager` interface.
- The application server must allow nested `Work` submissions.
- Both the `run` and `release` methods must be declared as non-synchronized methods.
- When the application server is unable to recreate an execution context if it is specified for the submitted `Work` instance, it must throw a `WorkCompletedException` set to an appropriate error code.
- The `WorkManager` must catch any exception thrown during `Work` processing, which includes execution context setup and wrap it with a `WorkCompletedException` set to an appropriate error code.

- The application server must execute a submitted `Work` instance with an unspecified context if no execution context has been specified, or must execute it with the specified execution context. That is, a submitted `Work` instance must never inherit the submitter's execution context when no execution context is specified.
- If the application server is unable to start `Work` execution when a start timeout is specified for the submitted `Work` instance, it must reject the `Work` instance with a `WorkRejectedException` set to `WorkException.START_TIMED_OUT`.

- The application server must use a value of -1 (`WorkManager.UNKNOWN`) to indicate an unknown `Work` start delay duration.

FIGURE 10-4 Blocking Durations of Various Work Submissions

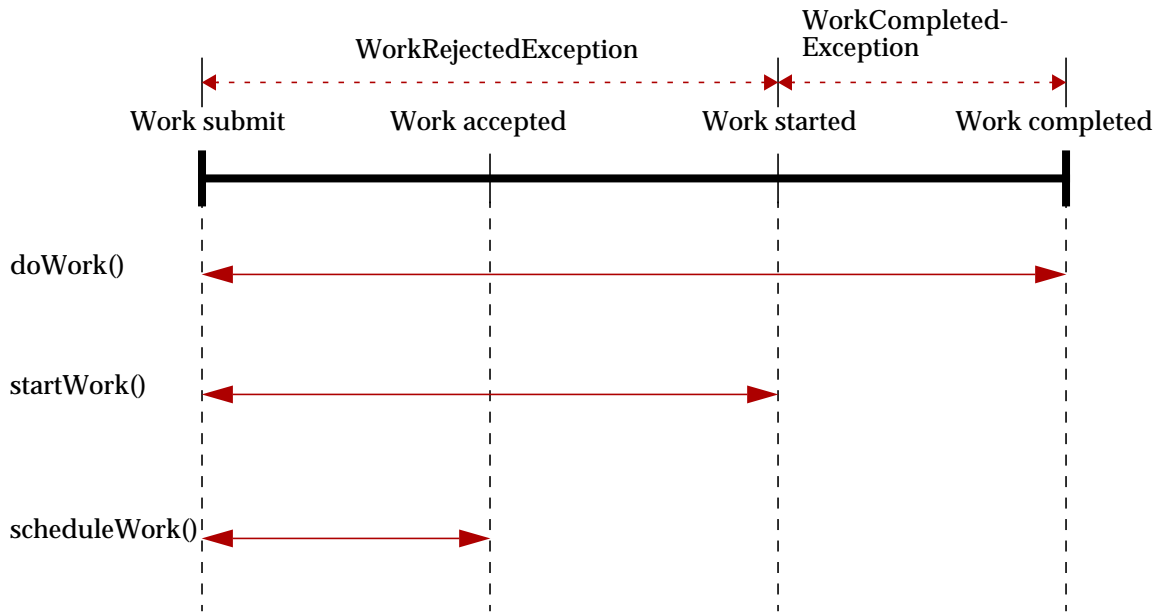
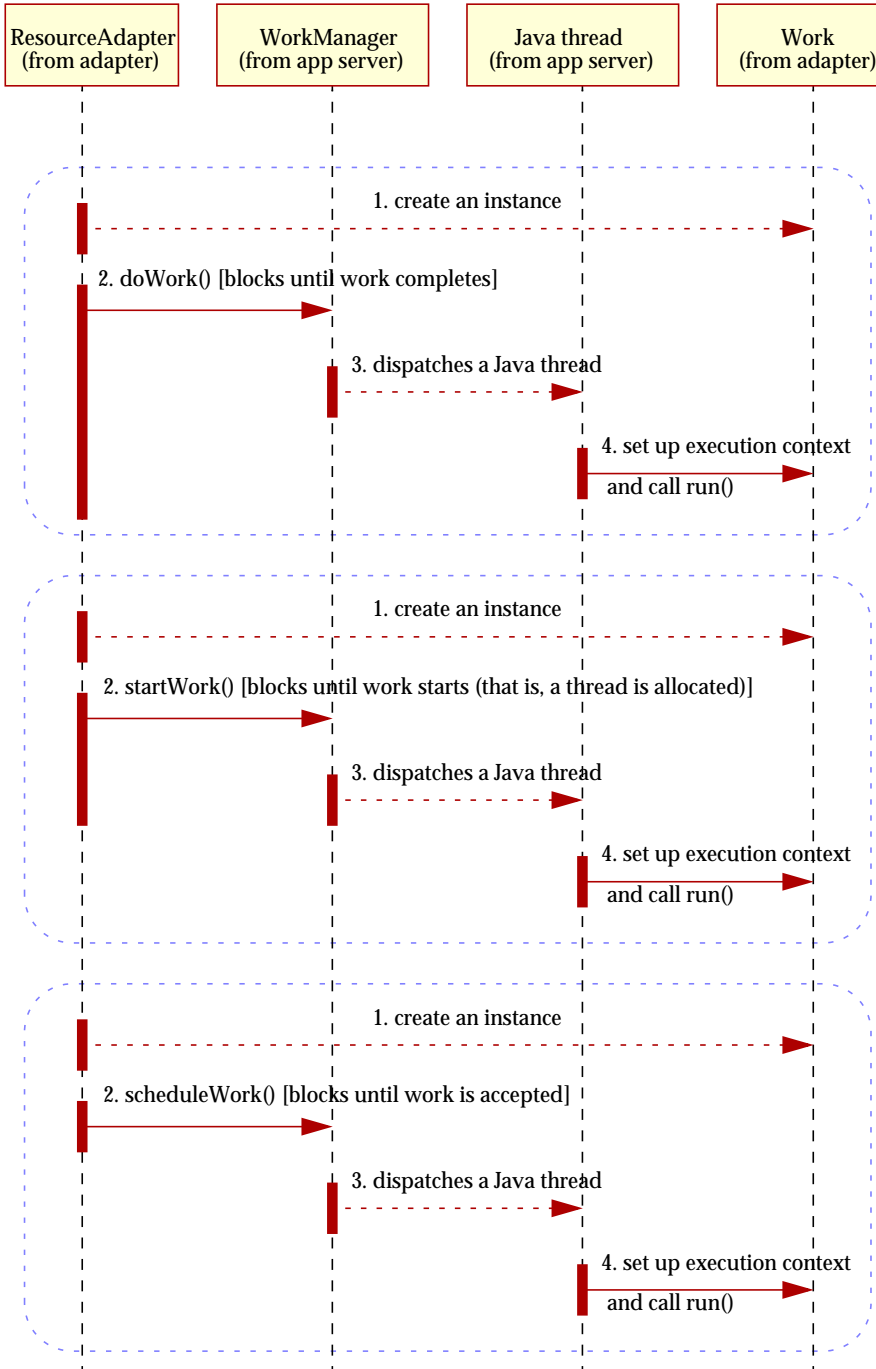


FIGURE 10-5 Work Submission - Blocking Behavior (sequence diagram)



10.3.4 WorkListener Interface and WorkEvent Class

The `WorkListener` interface is optionally implemented by the resource adapter. The `WorkEvent` and `WorkAdapter` classes are defined by the Connector 1.5 specification. The `WorkListener` instance is supplied to the `WorkManager` during `Work` submittal and provides an event listener callback mechanism in order to be notified when the various `Work` processing events, such as work accepted, work rejected, work started, and work completed, occur. When a `WorkListener` is provided by the resource adapter, the application server must send event notifications to the `WorkListener`. These notifications may occur from any thread with an unspecified context.

```
public interface WorkListener extends EventListener {  
    void workAccepted(workEvent);  
    void workRejected(WorkEvent);  
    void workStarted(WorkEvent);  
    void workCompleted(WorkEvent);  
}
```

The `WorkEvent` class and `WorkAdapter` abstract class:

```
public class WorkEvent extends EventObject {

    public static final int WORK_ACCEPTED = 1;
    public static final int WORK_REJECTED = 2;
    public static final int WORK_STARTED = 3;
    public static final int WORK_COMPLETED = 4;

    public WorkEvent(Object source, int type, Work work,
        WorkException exc) { ... }

    public WorkEvent(Object source, int type, Work work,
        WorkException exc, long startDuration) { ... }

    public int getType() { ... }
    public Work getWork() { ... }
    public long getStartDuration() { ... }
    public WorkException getException() { ... }
}

public abstract class WorkAdapter implements WorkListener {
    public void workAccepted(WorkEvent e) {}
    public void workRejected(WorkEvent e) {}
    public void workStarted(WorkEvent e) {}
    public void workCompleted(WorkEvent e) {}
}
```

The `WorkEvent` instance provides the following information:

- The event type.
- The source object, that is, the `Work` instance, on which the event initially occurred.
- A handle to the associated `Work` instance.
- An optional start delay duration in millisecond.
- Any exceptions that were thrown during `Work` processing. Possible exceptions are `WorkRejectedException`, and `WorkCompletedException`.

The type of the event determines the specific contents of a `WorkEvent`.

The `WorkAdapter` class is provided as a convenience for easily creating `WorkListener` instances by extending this class and overriding only those methods of interest. This is a standard event listener pattern in Java.

10.3.4.1 Requirements

- The `WorkListener` instance must not make any thread assumptions and must be thread-safe. That is, a notification can occur from any arbitrary thread with an unspecified context.
- The application server must send `Work` events to the `WorkListener` instance, if any, provided by the resource adapter.
- The `WorkListener` implementation must not make any assumptions on the ordering of notifications.
- The application server must use a value of -1 (`WorkManager.UNKNOWN`) to indicate an unknown `Work` start delay duration.

10.3.5 ExecutionContext Class

```
public class ExecutionContext {  
  
    public void setXid(xid) { ... }  
    public Xid getXid() { ... }  
    public long getTransactionTimeout() { ... }  
    public void setTransactionTimeout(long seconds)  
        throws NotSupportedException { ... }  
  
}
```

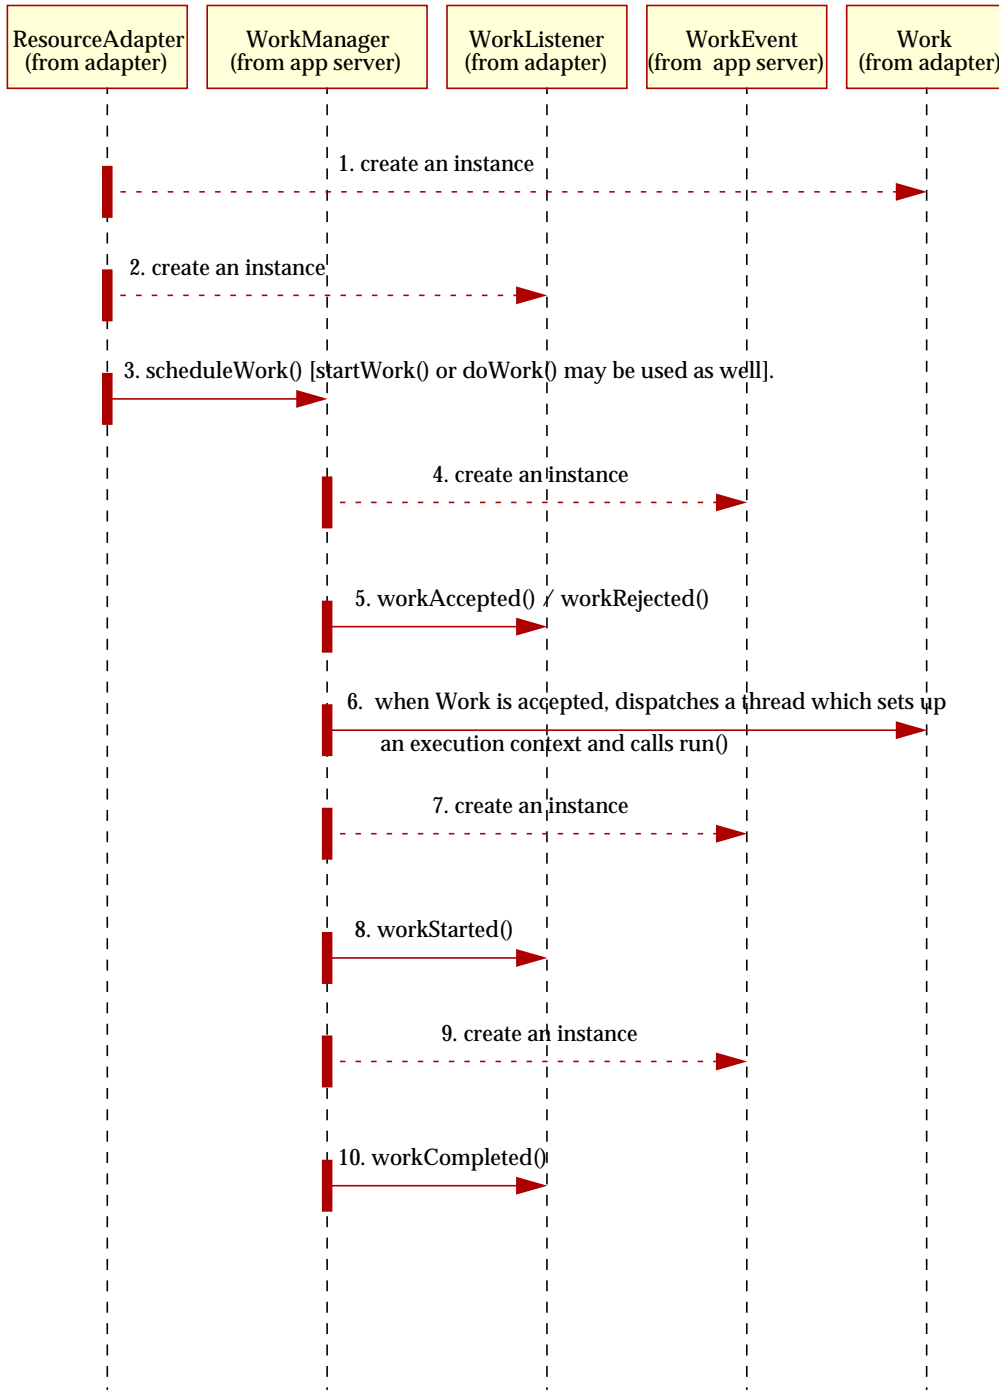
The `ExecutionContext` class allows a resource adapter to specify an execution context, such as a transaction context, with which the `Work` instance must be executed. The resource adapter is responsible for populating the `ExecutionContext` instance with an appropriate execution context. The default implementation provides a null context.

It is better for `ExecutionContext` to be a class rather than an interface because:

- There is no need for a resource adapter to implement this class. It only needs to implement the context information, like transaction context.

- The resource adapter code does not have to change when the `ExecutionContext` class evolves. For example, more context types could be added to the `ExecutionContext` class in the future without forcing resource adapter implementations to change.

FIGURE 10-6 Work Submission - Callback Mechanism (sequence diagram)



10.3.6 Resource Adapter Thread Usage Recommendations

- Resource adapters are strongly recommended to use the work management contract to do work and interact with the application server only from within a `Work` instance, instead of using Java threads directly. This allows the resource adapter to be maximally portable across multiple deployment environments with different security settings.
- Resource adapters are allowed to create Java threads directly as permitted by the server security settings.
- If a resource adapter chooses to use Java threads directly, it is recommended they use the threads as daemon threads, as it does not interfere with an orderly shutdown of the server.

10.3.7 Periodic Execution of Work Instances

A resource adapter may need to periodically execute `Work` instances. It may use the `java.util.Timer` facility available in the Java platform or may use the `BootstrapContext` instance provided by the application server to obtain a `Timer` instance.

A resource adapter may not be able to directly create a `Timer` instance, if it does not have adequate runtime permissions to create threads. This is because the `Timer` instance starts a background thread. In such a case, the resource adapter can instead use the `BootstrapContext` instance to obtain a `Timer` instance from the application server.

```
package javax.resource.spi;

import java.util.Timer;
import javax.resource.spi.UnavailableException;

public interface BootstrapContext {
    ... // other methods
    Timer createTimer() // returns a new or an unshared instance
                       throws UnavailableException;
}
```

When the `createTimer` method of the `BootstrapContext` instance is invoked, the application server provides a new `Timer` instance or an unshared instance (that is, no one else has a reference) with an empty task queue. The application server must throw an `UnavailableException` if a `Timer` instance is unavailable; the resource

adapter may retry later. The application server must throw an `java.lang.UnsupportedOperationException`, if it does not support the Timer service.

Sample code to illustrate periodic Work executions using a Timer instance:

```
package com.xyz.adapter;

import java.util.*;
import javax.resource.spi.*;
import javax.resource.spi.work.WorkManager;

// ResourceAdapter JavaBean
public class MyResourceAdapterImpl implements ResourceAdapter {
    BootstrapContext bootstrapCtx = null;
    public void start(BootstrapContext ctx) {
        bootstrapCtx = ctx;
        ... // other operations
    }
    ... // other methods
}

{ // sample resource adapter code snippet to show Timer usage
    MyResourceAdapterImpl myRA = ... // get ResourceAdapter
    JavaBean
        Timer timer = myRA.bootstrapCtx.createTimer(); // get a Timer
        instance
        WorkManager workManager = myRA.bootstrapCtx.getWorkManager();

        timer.schedule(
            new TimerTask () {
                public void run() {
                    try {
                        workManager.scheduleWork(new MyWork());
                    } catch (WorkException we) { we.printStackTrace(); }
                }
            }, 0, 1000); // one second interval
}
```

10.3.8 Illustration: Using a Work Instance to Listen on Multiple Network Endpoints

J2SE Version 1.4 provides the `java.nio` package that includes a multiplexed, non-blocking I/O facility. Using the `java.nio` package it is possible for a single thread, such as a `Work` instance, to listen on multiple network endpoints or ports. Prior to the `java.nio` facility each network endpoint needed a separate thread to listen to incoming data.

10.3.9 Work Management in a Non-Managed Environment

Although the work management contract is primarily intended for a managed environment, it may still be used in a non-managed environment provided the application that bootstraps a resource adapter instance is capable of functioning as a `WorkManager`.

A resource adapter is free to create Java threads as permitted by the security policy settings of the non-managed environment.

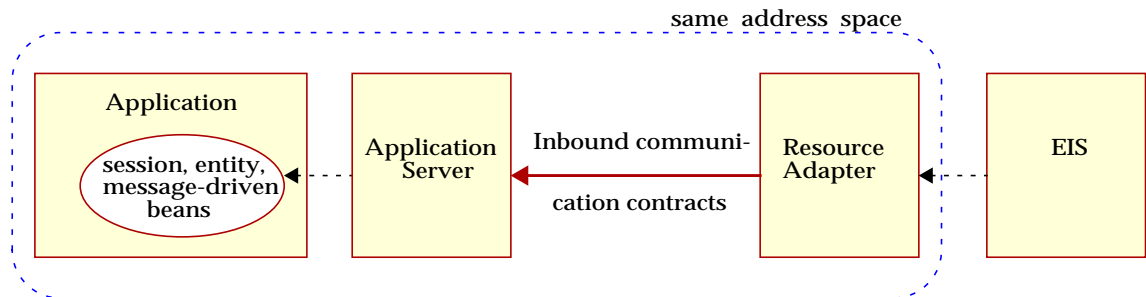
Inbound Communicaton

This chapter provides a high level description of the inbound communication model; that is, the communication from an EIS to an application residing in an application server's EJB container via a resource adapter. This also introduces concepts used in subsequent chapters related to inbound communication: Message Inflow (Chapter 12, "Message Inflow"), EJB Invocation (Chapter 13, "EJB Invocation"), and Transaction Inflow (Chapter 14, "Transaction Inflow").

11.1 Overview

In the inbound communication model, the EIS initiates all communication to an application. In this case, the application may be composed of EJBs (session, entity and message-driven beans) and resides in an EJB container.

FIGURE 11-1 Inbound Communication Model



In order to enable inbound communication, a mechanism to invoke EJBs (session, entity and message-driven beans) from a resource adapter is necessary. Further, a mechanism is needed to propagate transaction information from an EIS to an application residing in an EJB container.

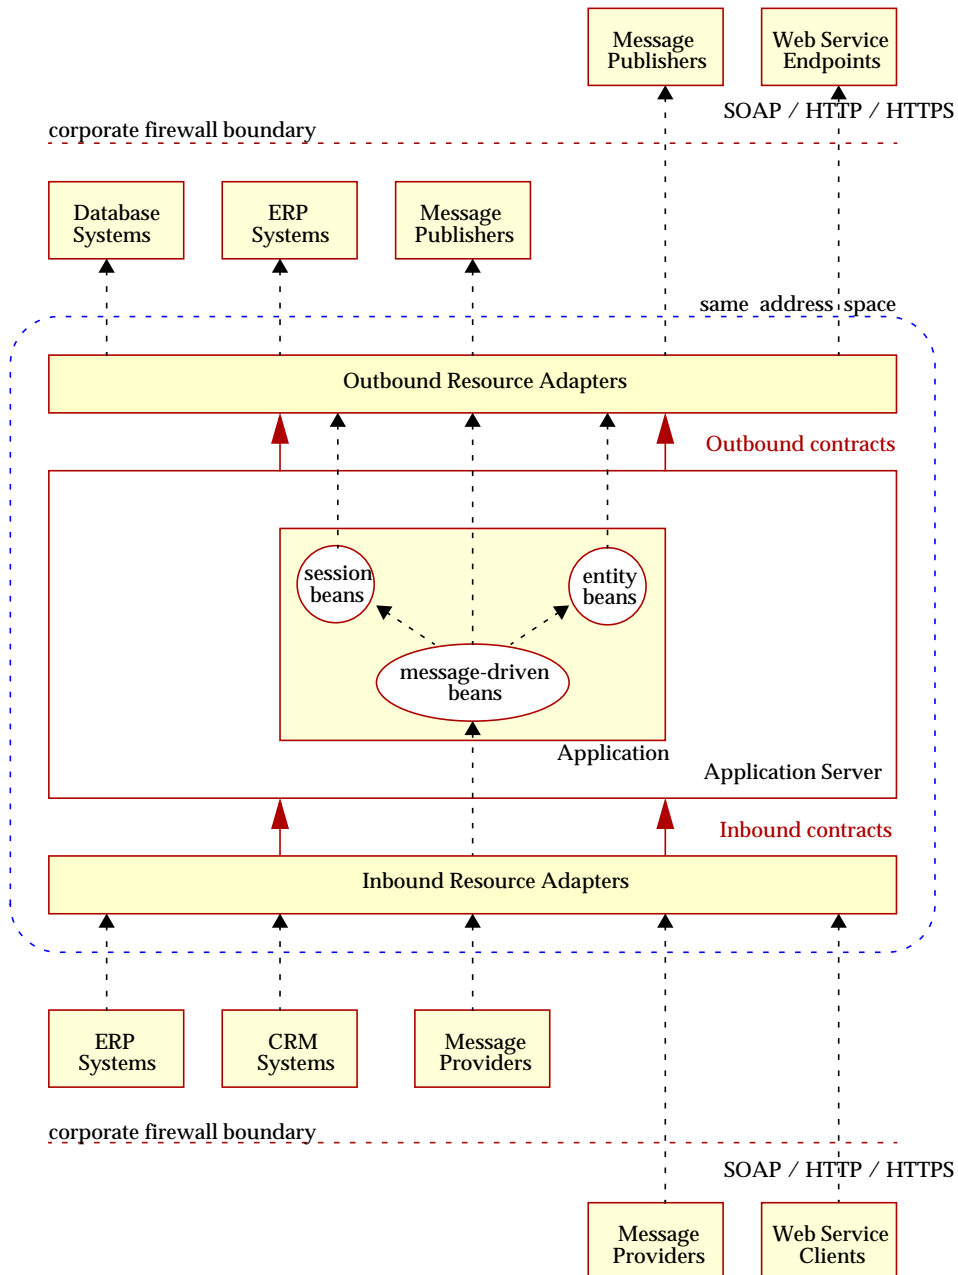
Chapter 12, “Message Inflow” describes a mechanism to invoke message-driven beans from a resource adapter. Chapter 14, “Transaction Inflow” provides a mechanism to import transaction information from an EIS into an EJB container.

11.2 An Illustrative Use Case

Wombat Inc. is a finance company which has a variety of software systems as part of its enterprise infrastructure. The software systems include databases, enterprise resource planning (ERP) and customer relationship management (CRM) systems, messaging middleware, mainframe systems, as well as several J2EE application

servers which host business logic written as EJBs (session, entity and message-driven beans). Further, there are web service interactions that occur as part of the overall corporate workflow.

FIGURE 11-2 Inbound Communication Model (an illustrative use case)



In order to integrate the various disparate software systems, and to allow them to communicate with each other, Wombat Inc. did the following:

- Used the application servers to hold the integration as well as business logic, developed as EJBs.
- Purchased resource adapters and deployed them on the application servers, in order to provide bi-directional connectivity between the applications residing on the application servers and the various software systems.

Thus, using the resource adapter as a connectivity enabler, Wombat Inc. was able to integrate the disparate software systems in its enterprise infrastructure.

Message Inflow

This chapter specifies a standard, generic contract between an application server and a resource adapter that allows a resource adapter to asynchronously deliver messages to message endpoints residing in the application server independent of the specific messaging style, messaging semantics and messaging infrastructure used to deliver messages. This contract also serves as the standard message provider pluggability contract that allows a wide range of message providers to be plugged into any J2EE compatible application server via a resource adapter.

12.1 Overview

Asynchronous message delivery or event notification is a widely used application communication paradigm. Some of the characteristics of the asynchronous message-based communication paradigm are:

- The message producer may not be directly aware of message consumers. There may be one or more consumers interested in the message.
- Message delivery is solicited; that is, a message consumer has to express interest in receiving messages.
- The messaging infrastructure is type-agnostic; that is, it treats messages as a Binary Large Object (BLOB). It stores and routes messages reliably, to interested message consumers, depending on Quality of Service (QoS) capabilities.
- The interaction is inherently loosely coupled. The message producer and the consumer do not share any execution context.
- The message producer generally is not interested in the outcome of message processing by consumers. However, it is possible that the provider may care to know if the message has been consumed or not.
- The message delivery always involves a message routing infrastructure, which offers varying QoS capabilities for storing (persistence) and routing messages reliably.


The J2EE application programming model offers a rich set of components: EJBs (session, entity and message-driven beans), JSPs, and servlets for applications to use. The message-driven bean is an asynchronous message consumer, or message endpoint.

J2EE applications may use two different patterns to interact with a message provider:

- It may directly use specific messaging APIs, such as Java Messaging Service (JMS), to send and synchronously receive messages. This is achieved using the standard connector contracts for connection management. See Chapter 6, “Connection Management.” Any message provider may provide a connector resource adapter that supplies connection objects for use by applications to send and synchronously receive messages using the specific messaging API.
- It may use message-driven beans to asynchronously receive messages via a message provider. The EJB specification (Section 1., “Enterprise JavaBeans (EJB™) specification, version 2.1:” on page F-1) describes the message-driven bean component contract in detail.

While the above patterns allow a J2EE application to send and receive messages, they do not provide a standard system-level contract to plug-in message providers to an application server and to deliver messages to message endpoints, or message-driven beans, residing in the application server. Without a standard pluggability contract, an application server would have to use special contracts to interact with various message providers, and a message provider has to do the same to interact with different application servers, which is a *mxn* problem.

FIGURE 12-1 Message Inflow Contract



Thus, there is a need for a standard, generic contract between an application server and a message provider which allows a message provider to deliver messages to message endpoints (message-driven beans) residing in the application server independent of the specific messaging style, messaging semantics, and messaging infrastructure used to deliver messages. Such a contract also serves as the standard

message provider pluggability contract which allows a wide range of message providers to be plugged into any J2EE compatible application server via a resource adapter.

12.2 Goals

- Provide a standard, generic mechanism to plug in a wide range of message providers, including JMS, into a J2EE compatible application server via a resource adapter and dispatch messages to message endpoints. This will allow J2EE components to act as consumers of messages with no required changes to the client programming models. Further, the J2EE components will be able to access messages with no awareness that a resource adapter is delivering the message.
- This generic contract must be capable of supporting various messaging delivery guarantees provided by different messaging styles, as well as allow concurrent delivery of messages.

12.3 Message Inflow Model

FIGURE 12-2 Message Inflow Contract (object diagram)

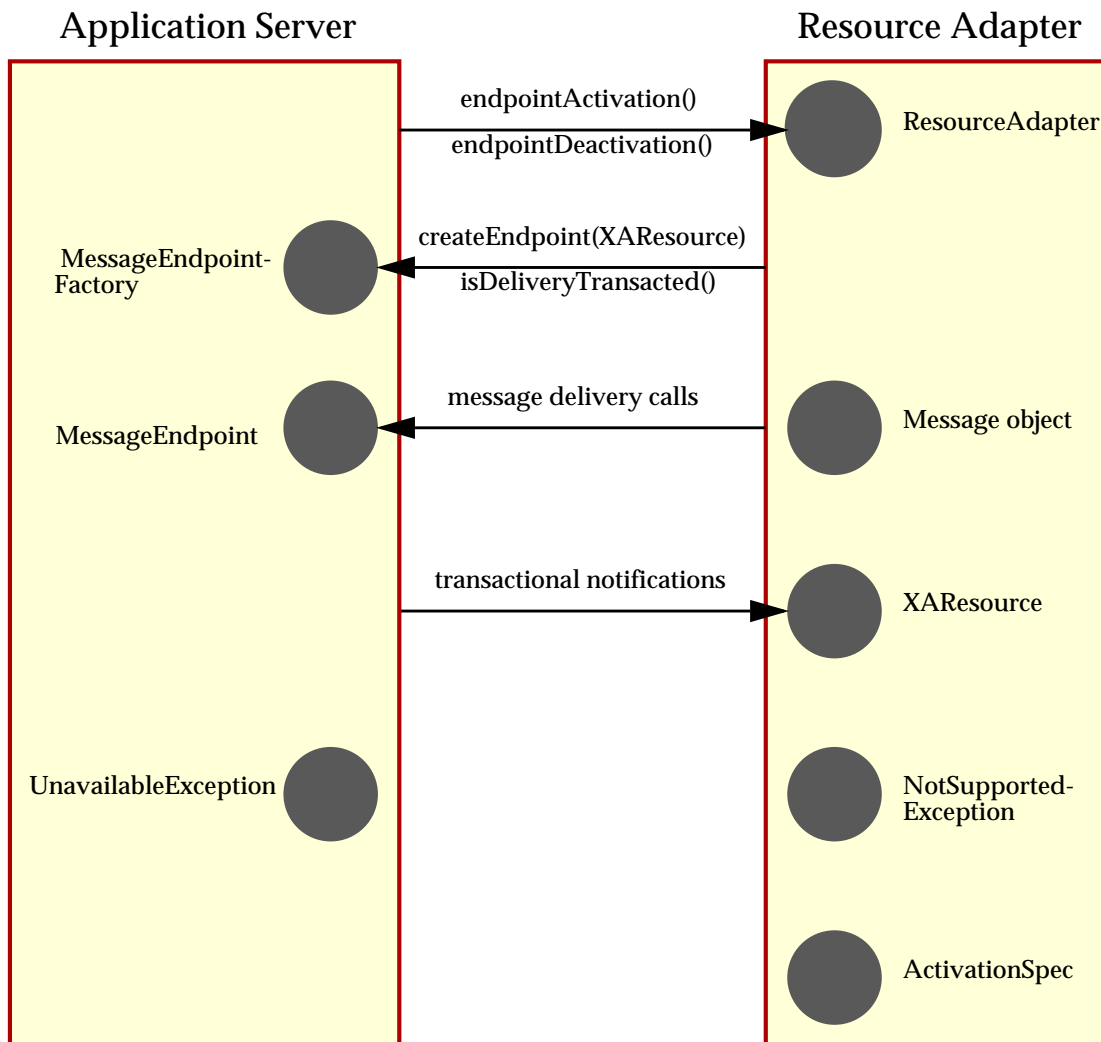
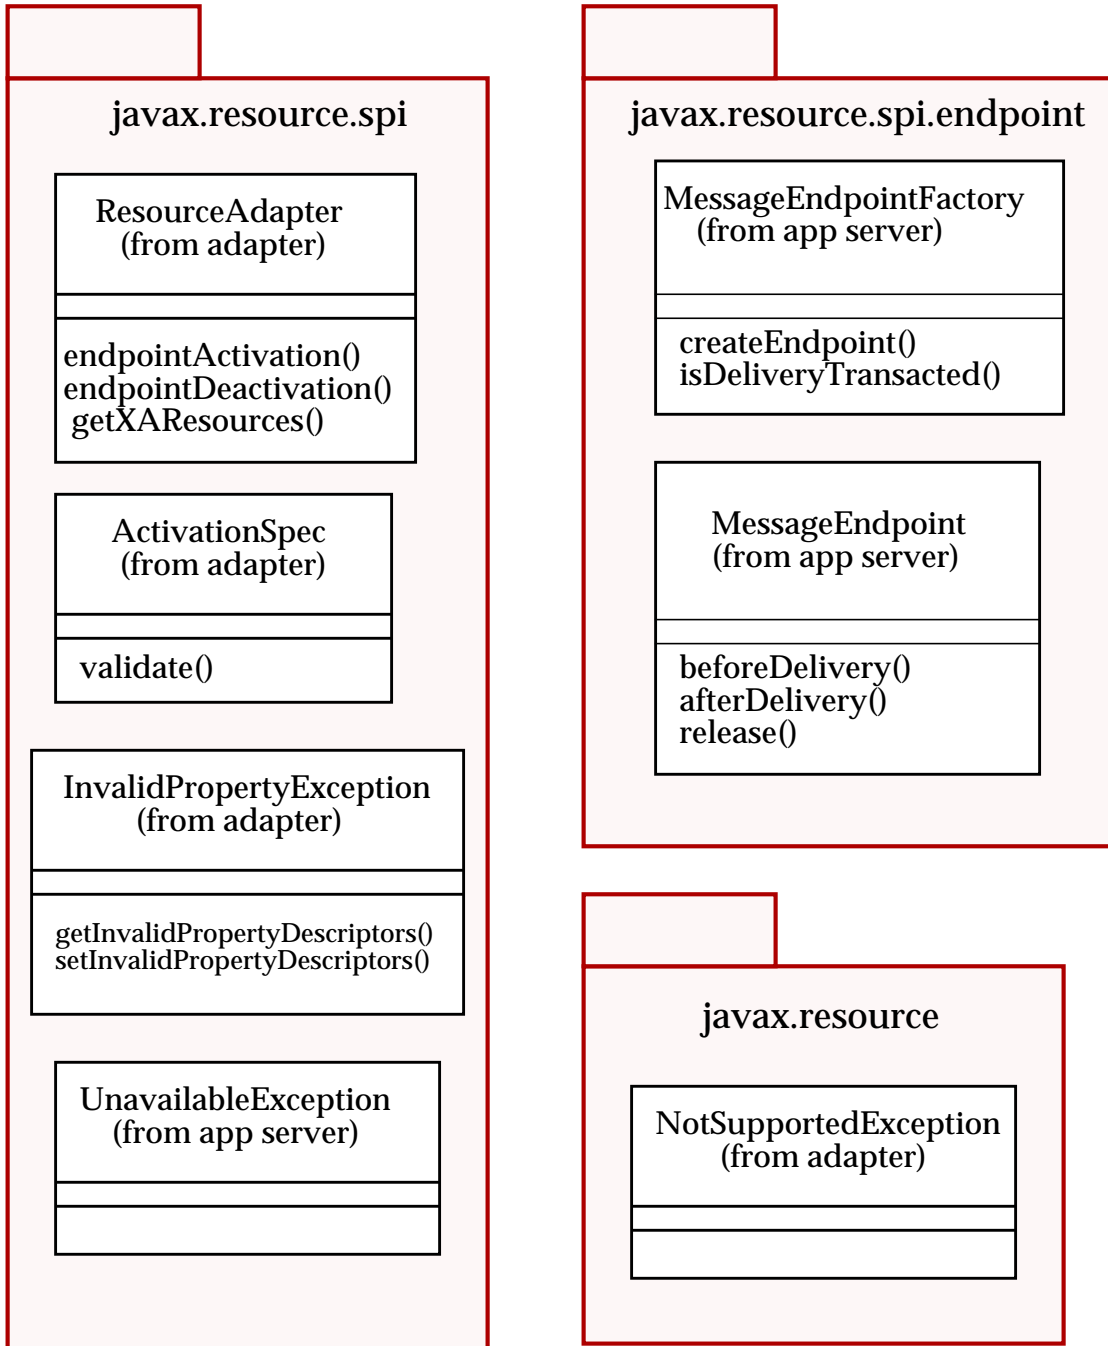


FIGURE 12-3 Message Inflow Contract (interfaces)



CODE EXAMPLE 12-1 javax.resource.spi

```
package javax.resource.spi;

import java.beans.PropertyDescriptor;
import javax.resource.NotSupportedException;
import javax.resource.spi.endpoint.MessageEndpointFactory;

public interface ResourceAdapter {
    ... // other methods
    void endpointActivation(MessageEndpointFactory,
        ActivationSpec) throws ResourceException;
    void endpointDeactivation(MessageEndpointFactory,
        ActivationSpec);
    XAResource[] getXAResources(ActivationSpec[] specs)
        throws ResourceException;
}

public interface ActivationSpec { // JavaBean
    void validate() throws InvalidPropertyException;
}

public class InvalidPropertyException extends ResourceException {
    public InvalidPropertyException() { ... }
    public InvalidPropertyException(String message) { ... }
    public InvalidPropertyException(String message,
        String errorCode) { ... }
    public void setInvalidPropertyDescriptors(
        PropertyDescriptor[] invalidProperties) { ... }
    public PropertyDescriptor[] getInvalidPropertyDescriptors() {
        ... }
}

public class UnavailableException extends ResourceException {
    public UnavailableException() { ... }
    public UnavailableException(String message) { ... }
    public UnavailableException(Throwable cause) { ... }
    public UnavailableException(String message, Throwable cause) {
        ... }
}
```

CODE EXAMPLE 12-2 javax.resource.spi.endpoint

```
package javax.resource.spi.endpoint;

import java.lang.Exception;
import java.lang.Throwable;
import java.lang.NoSuchMethodException;
import javax.transaction.xa.XAResource;
```

CODE EXAMPLE 12-2 javax.resource.spi.endpoint

```
import javax.resource.ResourceException;
import javax.resource.spi.UnavailableException;

public interface MessageEndpointFactory {
    MessageEndpoint createEndpoint(XAResource)
        throws UnavailableException;
    boolean isDeliveryTransacted(java.lang.reflect.Method)
        throws NoSuchMethodException;
}

public interface MessageEndpoint {
    void beforeDelivery(java.lang.reflect.Method)
        throws NoSuchMethodException, ResourceException;
    void afterDelivery() throws ResourceException;
    void release();
}
```

The `ResourceAdapter` interface supports methods used for endpoint activations and deactivations. The `endpointActivation` method is called by the application server when a message endpoint is activated and the `endpointDeactivation` method is called by the application server when a message endpoint is deactivated. The resource adapter is supplied a `MessageEndpointFactory` instance and a configured `ActivationSpec` instance during endpoint activations and deactivations. The resource adapter may reject an activation by throwing a `NotSupportedException`, if the activation information is incorrect.

The resource adapter uses the `MessageEndpointFactory` instance to obtain message endpoint instances for delivering messages either serially or concurrently. The `MessageEndpointFactory` may be used for obtaining any number of message endpoint instances. The `createEndPoint` method call may throw an `UnavailableException` for several reasons:

- The application server has not completed endpoint activation.
- The application server may decide to limit concurrent message deliveries.
- The application server is about to shutdown.
- The application server may have encountered an internal error condition.

In some cases where the offending condition is temporary, the application server may decide to block the `createEndPoint` method call instead of throwing an `UnavailableException`.

The `MessageEndpointFactory` may also be used to find out whether message deliveries to a target method on a message listener interface that is implemented by a message endpoint will be transacted or not via the `isDeliveryTransacted` method. The message delivery preferences must not change during the lifetime of a message endpoint.

A resource adapter capable of message delivery to message endpoints must provide an `ActivationSpec` JavaBean class for each supported endpoint message listener type. The `ActivationSpec` JavaBean has a set of configurable properties specific to the messaging style and the message provider. An instance of the `ActivationSpec` JavaBean is configured by a message endpoint, or application, deployer to setup the necessary configuration information for the endpoint activation, and passed on to the resource adapter via the application server during endpoint deployment.

The resource adapter is expected to know the endpoint message listener type, either by using the `ActivationSpec` JavaBean contents or based on the `ActivationSpec` JavaBean class, and deliver messages to the endpoint. The resource adapter may optionally pass an `XAResource` instance while creating a message endpoint in order to receive transactional notifications when a message delivery is transacted.

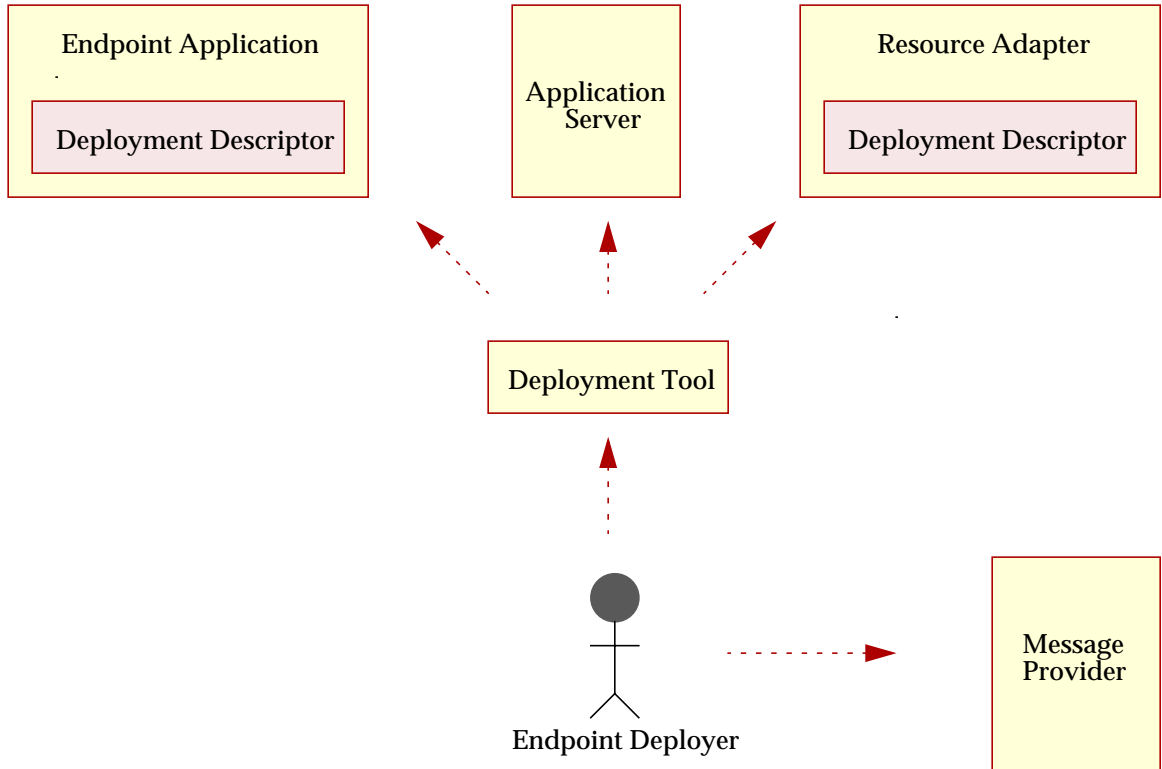
The following steps in sequential order represent the various stages in the message endpoint lifecycle, during which message inflow contracts are used:

- Endpoint¹ deployment.
- Message delivery (transacted and non-transacted).
- Endpoint undeployment.

1. Endpoint in this chapter refers to a message endpoint (for example, a message-driven bean application).

12.4 Endpoint Deployment

FIGURE 12-4 Endpoint (Message-Driven Bean) Deployment (Actors)



There are several actors involved in the deployment of a message endpoint:

- A message endpoint that is to be deployed on an application server.
- A resource adapter capable of message delivery. The resource adapter is typically provided by a message provider or a third party, and is used to plug an external message provider into an application server. The resource adapter may be standalone that may be shared by different applications or it may be packaged along with an endpoint application.
- An application server that provides the runtime environment for the application.
- A deployer of the application, a human, who understands the application's needs, and is also aware of the details of the runtime environment in which the application will be deployed.

- A message provider, or messaging infrastructure, that is the source for messages. A message provider may provide special tools that can be used by the deployer to setup the message provider for message delivery.

The roles and responsibilities of the various actors are as follows:

12.4.1 Message Endpoint

The message endpoint is a message-driven bean application which is to be deployed on the application server. It asynchronously consumes messages from a message provider. It is also possible for the application to send and synchronously receive messages by directly using messaging-style specific APIs.

The message-driven bean developer provides activation configuration information in the message-driven bean deployment descriptor. This includes messaging style specific configuration details, and possibly message provider-specific details as well, which is used by the message-driven bean deployer to setup the activation.

The EJB specification (Section 1., “Enterprise JavaBeans (EJB™) specification, version 2.1:” on page F-1) has more details on the message-driven bean deployment descriptor element `activation-config` used to hold the activation configuration information. For example, the deployment descriptor of a message-driven bean which consumes from a JMS resource adapter may contain:

CODE EXAMPLE 12-3 Message-Driven Bean Deployment Descriptor

```
<!-- message-driven bean deployment descriptor -->
...
<activation-config>
  <activation-config-property>
    <activation-config-property-name>
      destinationType
    </activation-config-property-name>
    <activation-config-property-value>
      javax.jms.Topic
    </activation-config-property-value>
  </activation-config-property>
  <activation-config-property>
    <activation-config-property-name>
      SubscriptionDurability
    </activation-config-property-name>
    <activation-config-property-value>
      Durable
    </activation-config-property-value>
  </activation-config-property>
  <activation-config-property>
    <activation-config-property-name>
```

CODE EXAMPLE 12-3 Message-Driven Bean Deployment Descriptor

```
        MessageSelector
        </activation-config-property-name>
        <activation-config-property-value>
            JMSType = 'car' AND color = 'blue'
        </activation-config-property-value>
    </activation-config-property>
    ...
</activation-config>
...
```

The EJB specification does not specify messaging style-specific descriptor elements contained within the `activation-config` element. It is the responsibility of each individual messaging specification or product to specify the standard descriptor elements specific to the messaging style for representing the activation configuration information.

12.4.2 Resource Adapter

The resource adapter is a system component located in the application server's address space (that is, it has already been deployed) that provides connectivity to message providers and is capable of delivering messages to message endpoints residing in the application server. The resource adapter is typically provided by a message provider or a third party, and is used to plug an external message provider into an application server. The resource adapter may be standalone, shared by different applications, or may be packaged along with an endpoint application.

The resource adapter provides the following information via the resource adapter deployment descriptor that is used by the endpoint deployer to setup endpoint activation:

12.4.2.1 List of Supported Message Listener Types

The resource adapter provides a list of endpoint message listener types it supports. Each type is represented as a fully qualified name of the Java type of the message listener interface.

12.4.2.2 ActivationSpec JavaBean

The resource adapter provides the Java class name of an `ActivationSpec` JavaBean, one for each supported message listener type, containing a set of configurable properties that is used to specify endpoint activation configuration information during endpoint deployment. Refer to Section 16.3, "JavaBean

Requirements” on page 16-2. An `ActivationSpec` JavaBean instance is created during endpoint deployment, and the instance is configured by the endpoint deployer.

During configuration, an `ActivationSpec` JavaBean instance may check the validity of the configuration settings provided by the endpoint deployer. The `ActivationSpec` has a `validate` method which may be used during endpoint deployment to validate the overall activation configuration information provided by the endpoint deployer. This helps to catch activation configuration errors earlier on without having to wait until endpoint activation time for configuration validation. The implementation of this self-validation check behavior is optional.

Note, the `ActivationSpec` JavaBean instance must not make any assumptions about the availability of a live resource adapter instance.

The resource adapter may also provide in its deployment descriptor, using the `required-config-property` element, an optional list of configuration property names required for each activation specification. This information may be used during deployment to ensure that the required configuration properties are specified. An endpoint activation should fail if the required property values are not specified.

In the case of JMS message providers, the `destination` property value (refer to Appendix B, “Activation Configuration for Message Inflow to JMS Endpoints”) may also be an object that implements the `javax.jms.Destination` interface. In such a case, the resource adapter should provide an administered object (refer to Section 12.4.2.3, “Administered Objects” on page 12-12) that implements the `javax.jms.Destination` interface. The specific type of the JMS destination is specified by the `destinationType` property value. The JMS `ActivationSpec` JavaBean properties should be standardized by the JMS community.

12.4.2.3 Administered Objects

The resource adapter may provide the Java class name and the interface type of an optional set of JavaBean classes representing various administered objects. Refer to Section 16.3, “JavaBean Requirements” on page 16-2. Administered objects are specific to a messaging style or message provider.

For example, some messaging styles may need applications to use special administered objects for sending and synchronously receiving messages via connection objects using messaging-style specific APIs. It is also possible that administered objects may be used to perform transformations on an asynchronously received message in a message provider-specific way.

Note, administered objects are not used for setting up asynchronous message deliveries to message endpoints. The `ActivationSpec` JavaBean is used to hold all the necessary activation information needed for asynchronous message delivery setup.

12.4.2.4 Configuring Administered Objects

- Create an administered object JavaBean instance. This will initialize the instance with the defaults specified via the JavaBean mechanism.
- Apply the administered object class configuration properties specified in the resource adapter deployment descriptor, on the administered object instance. This may override some of the default values specified via the JavaBean mechanism.
- The deployer may further override the values of the administered object before deployment.

12.4.3 Endpoint Deployer

The endpoint deployer is a human who has the responsibility to deploy the message endpoint, or application, on an application server. The deployer is expected to know the needs of the application and be aware of the details of the runtime environment in which the application will be deployed.

The deployer selects a suitable resource adapter that matches the needs of the application depending on endpoint message listener type, QoS capabilities, and so on. The deployer configures an `ActivationSpec` JavaBean instance based on the information provided by the application developer or assembler, which is contained in the endpoint deployment descriptor. The deployer may also use additional message provider-specific information to configure the `ActivationSpec` JavaBean instance.

The deployer also configures a set of administered objects, if necessary. The resource adapter provides the JavaBean classes for such administered objects. The deployer may also interact with a message provider to do the necessary setup for message delivery.

Then the deployer deploys the application on the application server. As part of the deployment procedure, the deployer provides all the configured JavaBean instances to the application server, and also specifies the chosen resource adapter instance to be used for message delivery.

12.4.4 Application Server

The application server provides the runtime environment for the message endpoint. It activates message endpoints when they are deployed. All such deployed endpoints are automatically reactivated when an application server restarts after a normal shutdown or system crash. When an application is undeployed, the application server deactivates the endpoint.

When an endpoint is activated, the application server calls the chosen resource adapter via the `endpointActivation` method and passes on a `MessageEndpointFactory` instance and the `ActivationSpec` `JavaBean`, which was configured by the endpoint deployer. The application server does not interpret the contents of the `ActivationSpec` `JavaBean` and treats it as an opaque entity. The resource adapter may reject an endpoint activation by throwing a `NotSupportedException` during the `endpointActivation` method call. This is due to incorrect activation information.

The resource adapter uses the `MessageEndpointFactory` to create message endpoint instances to deliver messages either serially or concurrently. There is no limit to the number of such endpoints that may be created to deliver messages. However, in practice the application server may decide to limit concurrency by rejecting attempts to create new endpoints by throwing an `UnavailableException`. The application server may also attempt to block a message delivery method call in order to limit concurrency and perform flow control.

Note, a resource adapter may attempt to deliver messages during the `endpointActivation` method call. It is up to the application server to decide whether to allow message delivery before activation is completed. If the application server chooses to prevent message delivery during endpoint activation, it may block the `createEndpoint` method call until the activation is completed or throw an `UnavailableException`.

The resource adapter may pass an `XAResource` instance while creating a message endpoint in order to receive transactional notifications when a message delivery is transacted. The application server must notify the resource adapter via the `XAResource` instance if a message delivery is transacted.

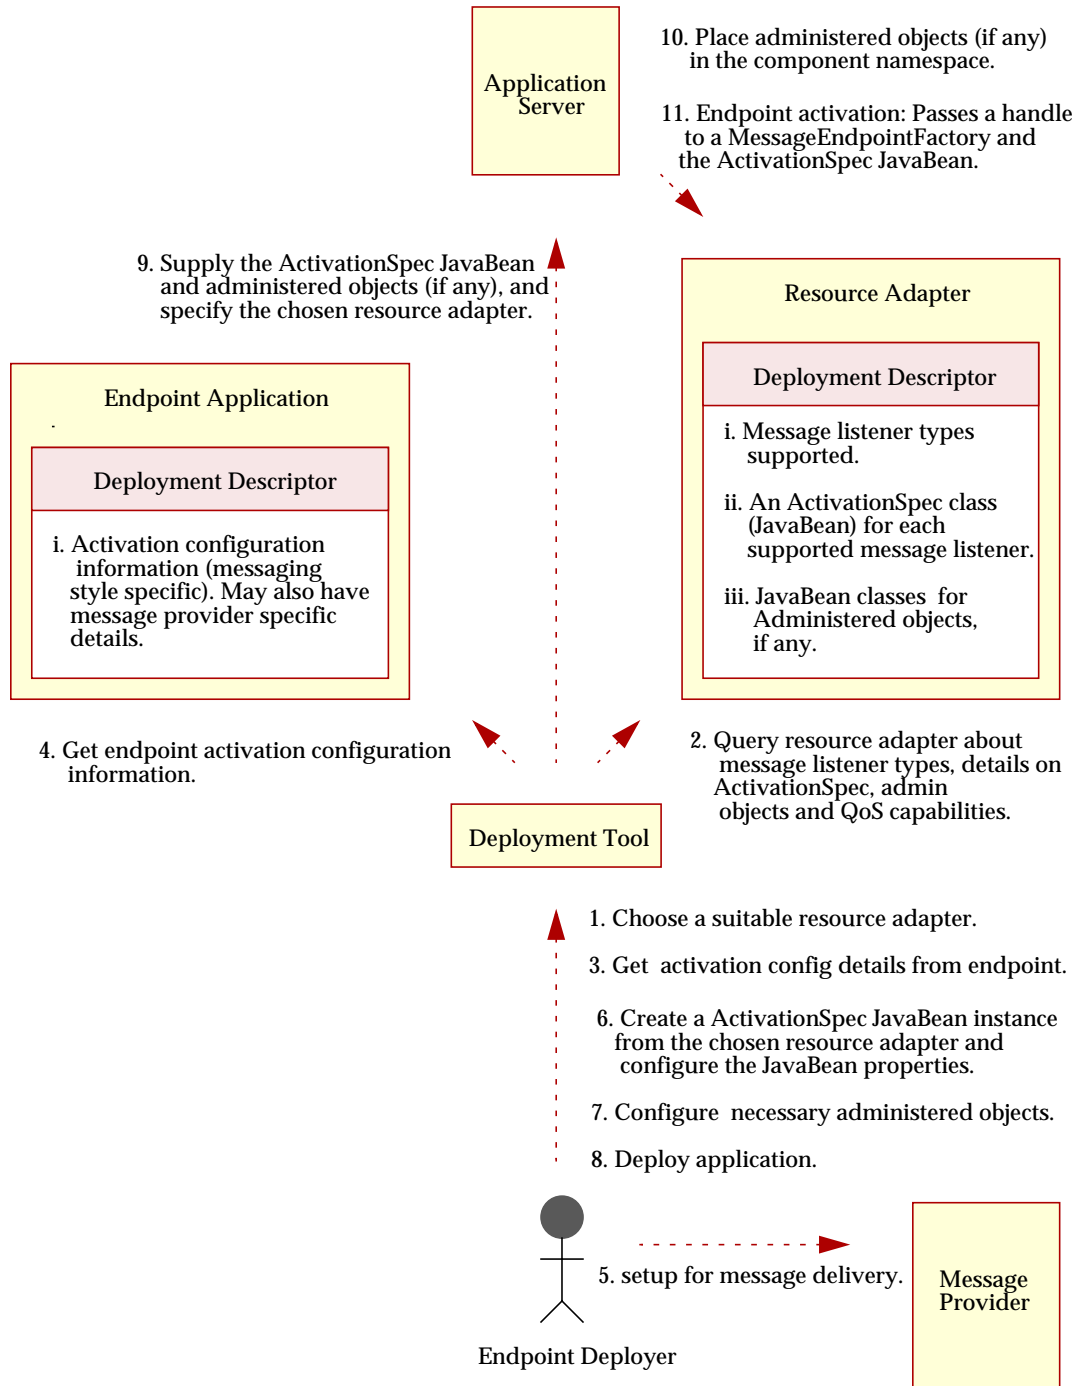
During endpoint deployment, the application server places the configured administered objects, if any, supplied by the endpoint deployer in the component namespace `java:comp/env`. The endpoint deployer specifies a location in the component namespace where each administered object should reside. The configured administered objects residing in the component namespace are used by the endpoint application in a messaging style-specific manner.

When an endpoint is deactivated, the application server notifies the resource adapter via the `endpointDeactivation` method call. The application server must pass the same `MessageEndpointFactory` instance and the `ActivationSpec` `JavaBean` instance that was used during endpoint activation.

12.4.5 Message Provider

A message provider, or messaging infrastructure, is typically an external system that is the source for messages. Message providers may vary in their QoS capabilities. A message provider may provide special tools that can be directly used by the endpoint deployer to setup the message provider for message delivery.

FIGURE 12-5 Endpoint (message-driven bean) Deployment Steps



12.4.6 Endpoint Deployment Steps

The sequence of steps involved in endpoint deployment involving the various actors is as follows:

- The endpoint deployer obtains a list of resource adapters capable of delivering messages to the message endpoint, and chooses a suitable one. The decision is based on the message listener type supported by the resource adapter and its QoS capabilities. However, it is possible that the message endpoint application may already contain a suitable resource adapter. In such a case, the resource adapter is deployed along with the endpoint application and is used to deliver messages to the specific endpoint application.
- The deployer obtains the activation configuration information provided by the endpoint developer available in the endpoint deployment descriptor.
- The deployer may need to setup the message provider for message delivery to the endpoint. This may be done using a message provider specific tool.
- The deployer obtains an `ActivationSpec` `JavaBean` from the selected resource adapter and configures it. The configuration information is messaging style-specific and may include message provider specific details.
- The deployer configures the `JavaBean` instances of administered objects, if any are necessary.
- The deployer provides the configured `JavaBean` instances to the application server, and also specifies the resource adapter chosen for message delivery. Note, the contract between a deployment tool and an application server is out of scope for this specification.
- The application server places the administered objects, if any, in the `java:comp/env` component namespace for use by the message endpoint.
- The application server activates the message endpoint by calling the chosen resource adapter via the `endpointActivation` method and passes a `MessageEndpointFactory` instance and the configured `ActivationSpec` `JavaBean` instance provided by the deployer. The resource adapter may reject the endpoint activation by throwing a `NotSupportedException`, which is due to incorrect activation information.

12.4.7 Requirements

- A resource adapter that is capable of delivering messages to message endpoints must provide a list of endpoint message listener types it supports, and also must provide an `ActivationSpec` `JavaBean` class for each message listener type it supports. This information must be part of the resource adapter deployment descriptor.
- `ActivationSpec` and all administered objects must be `JavaBeans`.
- A resource adapter must allow an application server to make concurrent `endpointActivation` method or `endpointDeactivation` method calls.

- The endpoint application's `activation-config` properties, specified in the endpoint deployment descriptor, should be a subset of the `ActivationSpec` JavaBean's properties. There must be a one-to-one correspondence between the `activation-config` property names and the `ActivationSpec` JavaBean's property names. This allows automatic merging of the `activation-config` properties with an `ActivationSpec` JavaBean instance during endpoint deployment. Any specified `activation-config` property which does not have a matching property in the `ActivationSpec` JavaBean should be treated as an error.
- When an application server notifies a resource adapter during endpoint deactivation, it must pass the same `MessageEndpointFactory` instance and the `ActivationSpec` JavaBean instance that was used during endpoint activation.
- Any exception thrown by the `endpointDeactivation` method call must be ignored. After this method call the endpoint is deemed inactive.
- All deployed endpoints must be automatically reactivated by the application server when it restarts after a normal shutdown or system crash.
- Before a resource adapter is undeployed, the application server must deactivate all active endpoints consuming messages from that specific resource adapter.

12.4.8 Structure of a Message Listener Interface

A message listener interface implemented by a message endpoint, a message-driven bean, is allowed to have multiple methods. Each method of a message listener interface is allowed to have multiple arguments, a return value, and throw checked application exceptions or unchecked system exceptions.

Checked exceptions are thrown only by a message listener implementation. The message-driven bean container must propagate to the resource adapter any checked exception that occurs during message dispatch.

Unchecked exceptions (`java.lang.RuntimeException` and `java.lang.Error`) may be thrown by either the message listener implementation or may be thrown by the application server code during message dispatch. The application server must wrap such an unchecked exception within a `javax.ejb.EJBException`, which is a `java.lang.RuntimeException`, and throw the `javax.ejb.EJBException` to the resource adapter.

The EJB specification describes in detail the structural requirements of a message listener interface implemented by a message-driven bean.

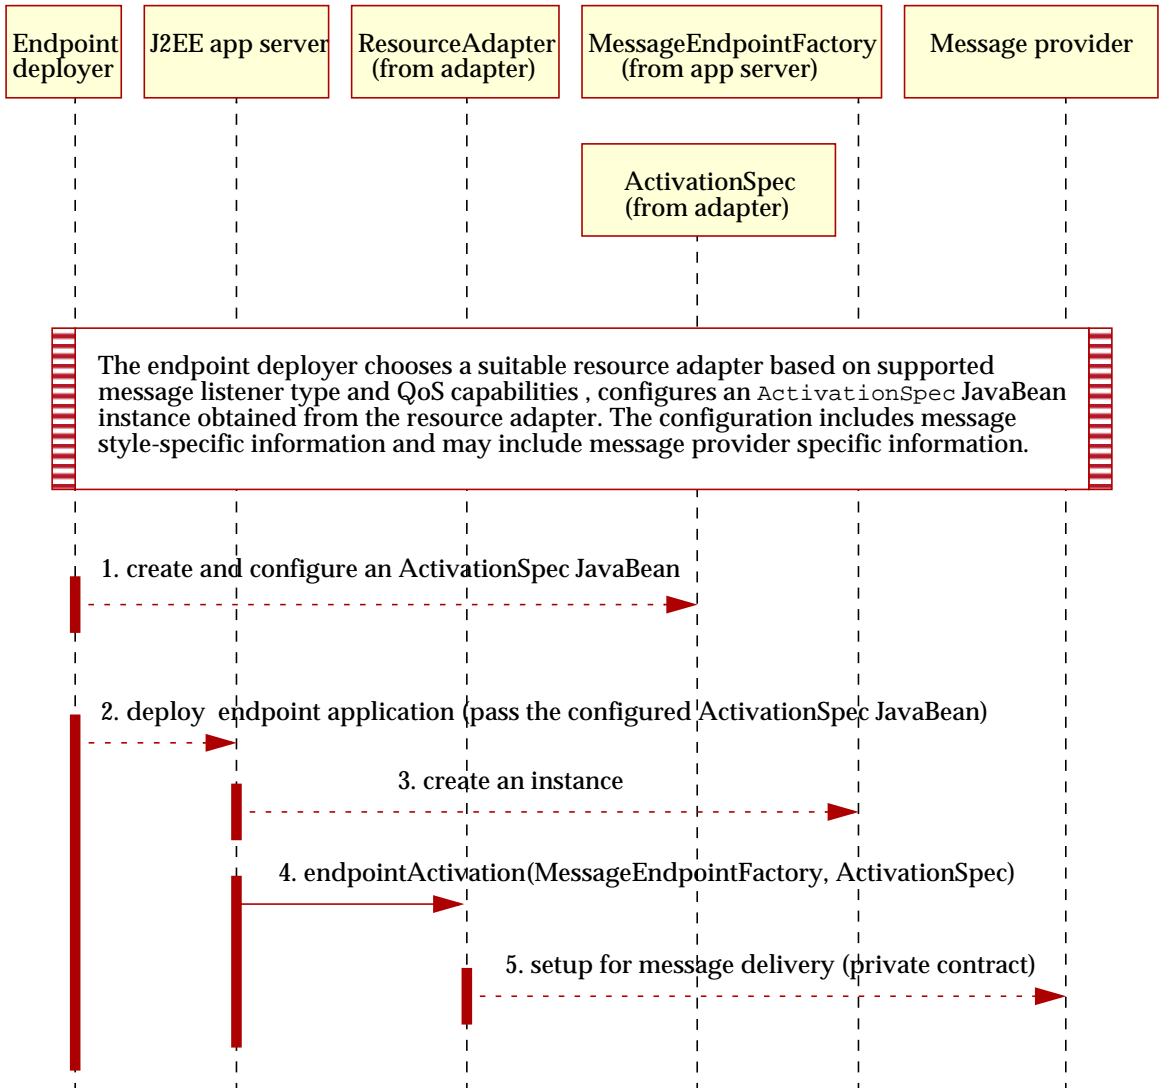
12.4.9 Multiple Endpoint Activations With Similar Activation Configuration

Since multiple endpoints, that is, different applications, with similar activation configuration may be deployed in a single application server, the application server may call the `endpointActivation` method on a resource adapter instance multiple times with similar activation configuration. The resource adapter must treat multiple endpoint activations with similar activation configuration separately. When messages start arriving, the resource adapter must, for each active endpoint, deliver a separate copy of incoming messages, even if there are multiple endpoints with similar activation configuration.

12.4.9.1 Requirements

- The application server must supply a unique `MessageEndpointFactory` instance for each activation.
- Refer to Section 16.4, “Equality Constraints” on page 16-3 for equality constraints on `MessageEndpointFactory` and `ActivationSpec` implementations.
- The resource adapter must treat multiple endpoints with similar activation configuration separately and guarantee message delivery semantics.
- The resource adapter must treat each `ActivationSpec` `JavaBean` uniquely irrespective of its contents. That is, the resource adapter must not treat two separate `ActivationSpec` `JavaBeans` as equals.

FIGURE 12-6 Endpoint Deployment (sequence diagram)



12.5 Message Delivery

Once endpoints are activated, they are ready to receive messages. When messages arrive, the resource adapter uses the `MessageEndpointFactory` to create an endpoint instance. The resource adapter narrows the endpoint instance to the actual message listener type (it knows the endpoint type from the `ActivationSpec`), and delivers the message to the endpoint instance. The EJB specification (Section 1., “Enterprise JavaBeans (EJB™) specification, version 2.1:” on page F-1) prescribes the rules for the message listener interface structure. The same endpoint instance may be used again to deliver subsequent messages serially, but it must not use the same endpoint instance concurrently.

Note that the endpoint instance supplied by the `createEndPoint` method call is a proxy which implements the endpoint message listener type and the `MessageEndpoint` interface and it is not the actual endpoint. This is necessary because the application server may need to intercept the message delivery in order to inject transactions, depending on the actual endpoint preferences, and to perform other checks.

The proxy endpoint instance is implemented by the application server and is used to intercept the message delivery, performs checks, inject transactions, and so on, and to route the message to the actual message endpoint instance. The resource adapter does not have to make this distinction, and should treat the proxy endpoint instance as the actual endpoint instance.

The resource adapter may use a proxy endpoint instance to serially deliver messages. The resource adapter must not use a proxy endpoint instance concurrently from different threads. The proxy endpoint may throw a `java.lang.IllegalStateException` when invoked concurrently. However, a proxy endpoint instance may be used by different threads in a serial fashion.

The resource adapter may call the `release` method on the proxy endpoint instance to indicate that it no longer needs the proxy instance. This hint may be used by the application server for proxy endpoint pooling decisions. This method call frees the state of the proxy instance. The released proxy instance may be reused for further proxy endpoint requests from the same resource adapter. A proxy endpoint instance must not be reused across multiple resource adapter instances.

Between the time a proxy endpoint instance is released and before it is granted back to the same resource adapter (via a `createEndpoint` method call), the proxy endpoint instance is considered to be in a free and available state. Any attempted use of a free proxy must result in a `java.lang.IllegalStateException` thrown by the application server.

The application server may start a transaction before delivering the message to the actual endpoint depending on the endpoint preferences. The resource adapter may optionally pass an `XAResource` instance via the `createEndPoint` method in order to receive transaction notifications for those transactions started by an application server before message delivery.

12.5.1 Sample Resource Adapter Code To Illustrate Message Delivery

CODE EXAMPLE 12-4 Message Delivery In a Resource Adapter

```
// Reader Thread(s)
{
    // 1. Strip off msg header and parse message description
    // 2. Choose a set of endpoints which match message description
    // 3. Place message in appropriate buffer / queue
    // 4. Notify worker threads
}

// Worker Thread(s)
{
    // 1. Wake up on notification (message arrival)
    // 2. Pick up the message and locate the MessageEndpointFactory
    // associated with the subscription
    Message msg = ...;
    MessageEndpointFactory endpointFactory = ...;
    MyXAResource xaResource = ...; // for transacted delivery

    // 4. Obtain a message endpoint and narrow it to the
    // correct type.

    // ActivationSpec has endpoint message listener type
    // information.
    Object obj = endpointFactory.createEndpoint(xaResource);
    javax.jms.MessageListener endpoint =
        (javax.jms.MessageListener) obj;

    // 5. Link the XAResource with the endpoint. This allows the
    // XAResource object to know which endpoint/message delivery
    // is transacted when it receives transaction notifications.
    // This may be unnecessary depending on the implementation.
    xaResource.setEndpoint(endpoint);

    // Note: It may be possible to make the XAResource object
    // thread-safe/reentrant and reuse the same object for receiving
    // transaction notifications for different endpoints.
    // The XAResource object may use thread-local storage to
    // remember state, and thus avoid creating multiple
    // XAResource objects.

    // 6. Deliver the message.
    endpoint.onMessage(msg);

    // 7. Wait for notification of incoming messages
```


CODE EXAMPLE 12-4 Message Delivery In a Resource Adapter

```
    // and repeat the above steps on message arrival.
}

package com.wombat.ra;

import javax.transaction.xa.*;

class MyXAResource implements javax.transaction.xa.XAResource {

    public void start(Xid xid) throws XAException {
        // Associate the message delivery with the transaction id.
        // That is, create the tuple (msg id, transaction id) in
        // memory.
    }

    public int prepare(Xid xid) throws XAException {
        // Forward the tuple (message id, transaction id) to the
        // message provider. The provider must persist this
        // information, which is used during crash recovery by the
        // application server. During crash recovery,
        // the application calls the message provider, via the
        // recover method on an XAResource object, queries for
        // a list of in-doubt transactions and completes them.
        // Upon successful completion, return "ready_to_commit"
        // vote, else return "rollback_only" vote.
    }

    public void commit(Xid xid, boolean onePhase)
        throws XAException {
        // forward the transaction id to the message provider. This
        // serves as the acknowledgement that a message was
        // delivered.
    }

    public void rollback(Xid xid) throws XAException {
        // forward the transaction id to the message provider. This
        // indicates to the provider that the message was not
        // delivered.
    }

    ...; // other methods
}
```

12.5.1.1 Requirements

- The application server's proxy endpoint instance must implement the endpoint message listener type and the `MessageEndpoint` interface.
- The application server must pass by reference all the method parameter objects passed by the resource adapter during a message delivery method call on a proxy endpoint. The application server must not copy or clone the passed method parameter objects during message delivery to the actual endpoint instance.
- If the application server starts a new transaction, depending on endpoint preferences, before delivering a message to an endpoint instance, it must send all transaction notifications to the `XAResource` instance optionally supplied by the resource adapter as part of the `createEndPoint` method call.
- A resource adapter must not attempt to deliver messages concurrently to a single endpoint instance. The application server must reject concurrent usage of an endpoint instance.

12.5.2 Message Redelivery Upon Crash Recovery

An application server may crash during message delivery. In the case of message deliveries transacted by the application server, the application server must notify the commit decision to the message provider during crash recovery.

During crash recovery:

- The application server must first restart resource adapter instances by calling the `start` method on each persisted `ResourceAdapter` `JavaBean`, each corresponding to a resource adapter instance that was active prior to the crash.
- The application server must call the `getXAResources` method on each `ResourceAdapter` `JavaBean`, and pass in an array of `ActivationSpec` `JavaBeans`, each of which corresponds to a deployed endpoint application that was active prior to the system crash. This method need not be called if there were no endpoint applications that were active prior to the crash.
- Upon being called by the application server during crash recovery via the `getXAResources` method, the resource adapter must return an array of `XAResource` objects, each of which represents a unique resource manager.

The resource adapter may return `null` if it does not implement the `XAResource` interface. Otherwise, it must return an array of `XAResource` objects, each of which represents a unique resource manager that was used by the endpoint applications. The resource adapter may throw a `ResourceException` if it encounters an error condition.

Since each returned `XAResource` object represents a unique resource manager, the number of returned `XAResource` objects must be less than or equal to the number of `ActivationSpec` instances specified.

- Since it is possible that multiple resource adapters may use the same resource manager, there may be more than one `XAResource` object in the collection representing a resource manager. The application server may still need to narrow the collection of `XAResource` objects to a unique set of resource managers by using the `isSameRM` method on the `XAResource` object.
- The application server must use the `XAResource` objects to query each resource manager for a list of in-doubt in an already prepared state awaiting a commit decision transactions. Then, it must complete each pending transaction by sending the commit decision to the participating resource managers. Note, it is possible that a resource manager may not have pending in-doubt transactions.

The crash recovery procedure ensures that the message provider gets notified about the outcome of all message deliveries that were in an in-doubt transaction state at the time of the crash. Upon such notification, the message provider, depending on the delivery outcome, may decide to redeliver the undelivered messages to the various endpoints when they are reactivated.

12.5.3 Durable Message Delivery Setup

Once message endpoints are activated, they are ready to receive messages from a message provider. Message delivery setup may either be durable or non-durable.

In the case of non-durable message deliveries, messages are lost during application server downtime. When the application server becomes functional again, it automatically reactivates all message endpoints that were previously deployed, and message delivery starts again. But the messages that were produced during the downtime are lost. This is because messages are not persisted by the message provider and redelivered when the message endpoints are reactivated.

In the case of durable activations, messages are not lost during application server downtime. When the application server becomes functional again, it automatically reactivates all message endpoints that were previously deployed, and message delivery starts again. The messages that were produced during the downtime are persisted by the message provider and redelivered when the message endpoints are reactivated. It is the responsibility of the message provider to persist undelivered messages and redeliver them when the endpoint is available; that is, when the endpoint is reactivated by the app server.

Durability of message delivery may be an attribute of the activation setup, and thus it must be captured as part of the endpoint activation information. No additional contracts are required to support durable activations. Activation durability can be specified by a endpoint deployer via the `ActivationSpec` JavaBean. Note, some message providers may not support durable message deliveries and hence it is a QoS capability offered by the message provider.

12.5.4 Concurrent Delivery of Messages

During message endpoint activation, the application server supplies a `MessageEndpointFactory` to the resource adapter. The `MessageEndpointFactory` is used to get an endpoint instance via the `createEndPoint` method call. Each call results in a new or an unused endpoint instance that may be used to deliver messages concurrently; that is, for each active message endpoint, there may be multiple endpoint instances consuming messages concurrently.

Thus, for each message endpoint, depending on traffic, the resource adapter may choose to deliver messages serially using a single endpoint instance or concurrently using multiple endpoint instances.

There is no limit to the number of such endpoint instances that may be created, although the application server may limit the concurrency by either throwing an `UnavailableException` or by blocking the `createEndPoint` method call.

The application server may also attempt to block a message delivery method call in order to limit concurrency and perform flow control.

12.5.4.1 Requirements

- The application server must return a new or an unused endpoint instance for every `createEndPoint` method call on a `MessageEndpointFactory`.

12.5.5 Delivery Semantics and Acknowledgement

When the resource adapter delivers a message to an endpoint instance, which is really a proxy endpoint instance, the application server intercepts the message delivery to perform checks, inject transactions, and so on, and routes the message to the actual message endpoint instance.

The application server may start a transaction before delivering the message to the actual endpoint depending on the endpoint preferences. In the case of a transacted delivery, the resource adapter may use the transaction notifications received via the `XAResource` object to send back an acknowledgement to its message provider.

In the case of non-transacted delivery, that is, the application server does not start a transaction, the resource adapter has to rely on the successful completion of the message delivery call in order to send back an acknowledgement to its provider.

12.5.6 Transacted Delivery (Using Container-Managed Transaction)

Depending on the endpoint preferences, the application server brackets the message delivery to an endpoint instance with a Java Transaction API (JTA) transaction.

- This ensures that all the work done by the endpoint instance is enlisted as part of the transaction.
- This provides atomic message delivery/message consumption; that is, if the transaction were to be aborted by the application server due to an exceptional condition, all the work done by the endpoint instance is aborted, and the delivery is undone. If this does not occur, the transaction is committed, all the work done by the endpoint instance is committed and the delivery is completed.

The application server notifies the resource adapter while beginning and completing transactions by using the `XAResource` instance optionally supplied via the `createEndPoint` method call.

- This allows the adapter to know the outcome of a transacted delivery, and also influence the outcome of the transaction via the `prepare` method call on the `XAResource` instance.
- This allows the adapter to send back an acknowledgement to its message provider contingent on successful delivery; that is, when notified via the `commit` method call on the `XAResource` instance.
- This also allows the adapter to be notified of the correct delivery outcome upon failure recovery processing; that is, if the system crashes when the transaction is in-doubt, that is, when the transaction has already been prepared, the application server upon recovery correctly completes the transaction and notifies the adapter of the outcome of the transaction. Thus, the adapter can send back an acknowledgement to its message provider after failure recovery, if the message had been successfully delivered.

A resource adapter may optionally provide an `XAResource` instance via the `createEndPoint` method call in order to receive transactional notifications for those transactions started by an application server before message delivery. The resource adapter may find out whether message deliveries to a target method on a message endpoint will be transacted or not via the `isDeliveryTransacted` method in the `MessageEndPointFactory` instance, and decide whether to provide an `XAResource` instance via the `createEndPoint` method. Note, this does not require the resource adapter to support the transaction inflow contract (see Chapter 14, “Transaction Inflow”).

There are two delivery options available to the resource adapter for transacted deliveries:

- Option A, traditional XA style: The resource adapter optionally provides a `XAResource` instance via the `createEndPoint` method in order to receive XA transaction notifications for transacted message deliveries. In this case, the

application server fully controls the transaction boundaries and the resource adapter is merely a participant (the `XAResource` Resource Manager (RM)). See FIGURE 12-7 on page 35.

- Option B, `beforeDelivery/afterDelivery`: The resource adapter still optionally provides an `XAResource` instance via the `createEndPoint` method in order to receive XA transaction notifications for transacted message deliveries. But the resource adapter controls the transaction boundaries via the `beforeDelivery/afterDelivery` calls, in spite of being only a participant, an `XAResource` RM.

During the `beforeDelivery` call from the resource adapter, depending on the transactional preferences of the intended target method (specified via the `java.lang.reflect.Method` `method` parameter), the application server starts a transaction and enlists the `XAResource` instance in the transaction. The processing (by the application server) of the actual message delivery method call on the endpoint must be independent of the class loader associated with the descriptive method object (parameter).

During the `afterDelivery` call from the resource adapter, the application server completes the transaction and sends transaction completion notifications to the `XAResource` instance. The actual message delivery happens in between the `beforeDelivery` and `afterDelivery` calls.

In this case, the resource adapter controls when the transaction is started and completed by the application server, even though the application server decides on the outcome of the transaction. This allows resource adapters more flexibility in handling message deliveries. For example, the resource adapter may choose to dequeue a message from within the container-managed transaction so that the message dequeue is automatically undone if the container-managed transaction aborts.

There must not be more than one message delivery in-between a single `beforeDelivery` and `afterDelivery` method call pair. The application server must reject `beforeDelivery` or `afterDelivery` calls that are out of sequence by throwing an `IllegalStateException`.

The `beforeDelivery` and `afterDelivery` method calls are considered part of a single message delivery call. For each message delivery, the `beforeDelivery`, `afterDelivery` methods and the actual message delivery method, must be called from a single thread of control.

Further, the application server must set the thread context class loader to the endpoint application class loader during the `beforeDelivery` call and must reset it during a corresponding `afterDelivery` call. This allows a resource adapter to use the application class loader to load application specific classes while deserializing, or reconstructing, a message object. Note, setting of the thread context class loader during the `beforeDelivery` call is independent of whether an `XAResource` instance is provided by the resource adapter.

For each message delivery to an endpoint instance, the application server must match an `afterDelivery` call with a corresponding `beforeDelivery` call; that is, for each message delivery to an endpoint instance, `beforeDelivery` and `afterDelivery` calls are treated as a pair. See FIGURE 12-8 on page 36.

Thus, in the case of transacted deliveries:

- If a resource adapter does not provide an `XAResource` instance, it does not get XA transaction notifications.
- If a resource adapter provides an `XAResource` instance, it gets XA transaction notifications.
- If a resource adapter calls `beforeDelivery` and `afterDelivery` methods in addition to providing an `XAResource` instance, it not only receives XA transaction notifications but also gains control over when the transaction is started and completed. The `beforeDelivery` and `afterDelivery` calls have no effect when the resource adapter does not provide an `XAResource` instance or when the delivery is not transacted.

These various delivery options provide more choices to the resource adapter and allow a wide range of resource adapter and messaging provider implementations to be plugged-in. The application server must support both delivery options, option A and option B.

The `release` method call on a proxy endpoint instance releases the state of the proxy instance and makes it available for reuse. If the `release` method is called while a message delivery is in-progress, the application server must throw a `java.lang.IllegalStateException`, since concurrent calls on a proxy endpoint instance is disallowed. In the case of option B, if the `release` method is called in-between `beforeDelivery` and `afterDelivery` method calls, any transaction started during the corresponding `beforeDelivery` method call must be aborted by the application server.

12.5.7 Non-transacted Delivery

1. The application server does not bracket the message delivery to an endpoint instance within a JTA transaction.
2. The resource adapter relies on the successful return of the message delivery call on the endpoint instance for delivery confirmation and may send out an acknowledgement to its message provider if appropriate.
3. Any exception thrown by an endpoint instance during message delivery is taken as a failed delivery. The application server must propagate any exception thrown during message delivery to the resource adapter.

4. The application server does not notify the resource adapter about the delivery outcome upon crash recovery. Note, system failures may happen before the application server calls the actual endpoint instance, or while the actual endpoint instance is doing work, or after the endpoint has completed its work but before the message delivery on the endpoint returns.
- The application server does not have delivery status information available during failure recovery, nor does it know what state the actual endpoint instance was in when the failure happened. Consequently, it is hard to model exactly once delivery semantics for non-transacted dispatches.

12.5.8 Transacted Delivery Using an Imported Transaction

It is possible that a resource adapter may attempt message delivery to an endpoint instance with a transaction initiated by a message source, or message provider; that is, the message source initiates a transaction, and pushes a message to the resource adapter from within the transaction. The resource adapter in turn imports the transaction and attempts message delivery on an endpoint instance from within the source managed transaction.

The resource adapter must use the transaction inflow contract (see Chapter 14, “Transaction Inflow”) to import transactions initiated by a message source.

It must be possible to serially deliver one or more messages to one or more endpoint instances belonging to one or more endpoint applications within a single transaction, and be able to commit or abort the transaction as a single unit.

That is, it must be possible for a resource adapter to serially submit one or more `Work` objects (associated with a single transaction) that deliver messages to one or more endpoint instances belonging to one or more endpoint applications. If the enclosing transaction successfully commits, the messages are deemed to have been successfully delivered. If the enclosing transaction aborts, the messages that were delivered as part of the transaction are canceled.

12.5.9 Requirements

An application server must implement the following behavior for transacted and non-transacted message delivery to an endpoint instance. Before invoking the actual endpoint instance the application server must do the checks prescribed in Table 12-1 shown below, depending on the endpoint transaction preferences and the presence of a source managed transaction:

TABLE 12-1 Application Server Behavior For Transacted and Non-transacted Message Delivery

	Source managed transaction	No source managed transaction
endpoint instance requires transacted message delivery	Use the source managed transaction. Ignore the <code>XAResource</code> instance supplied by the resource adapter, if any.	Start a new transaction. Notify the <code>XAResource</code> instance supplied by the resource adapter, if any.
endpoint does not need transacted message delivery	Suspend the source managed transaction. Ignore the <code>XAResource</code> instance supplied by the resource adapter, if any.	No action. Ignore the <code>XAResource</code> instance supplied by the resource adapter, if any.

The application server must propagate any exception thrown during a message delivery to the resource adapter irrespective of whether the delivery is transacted or not.

For transacted deliveries, the application server must support both delivery options, option A and option B.

12.6 Endpoint Undeployment

- When a message endpoint is undeployed, the application server notifies the resource adapter via the `endpointDeactivation` method. The application server must pass the same `MessageEndpointFactory` instance and the `ActivationSpec` `JavaBean` instance that was used during the endpoint activation.

- The resource adapter removes the endpoint information from its internal state and in turn may notify the message provider.

FIGURE 12-7 Transacted Message Delivery: Option A (sequence diagram)

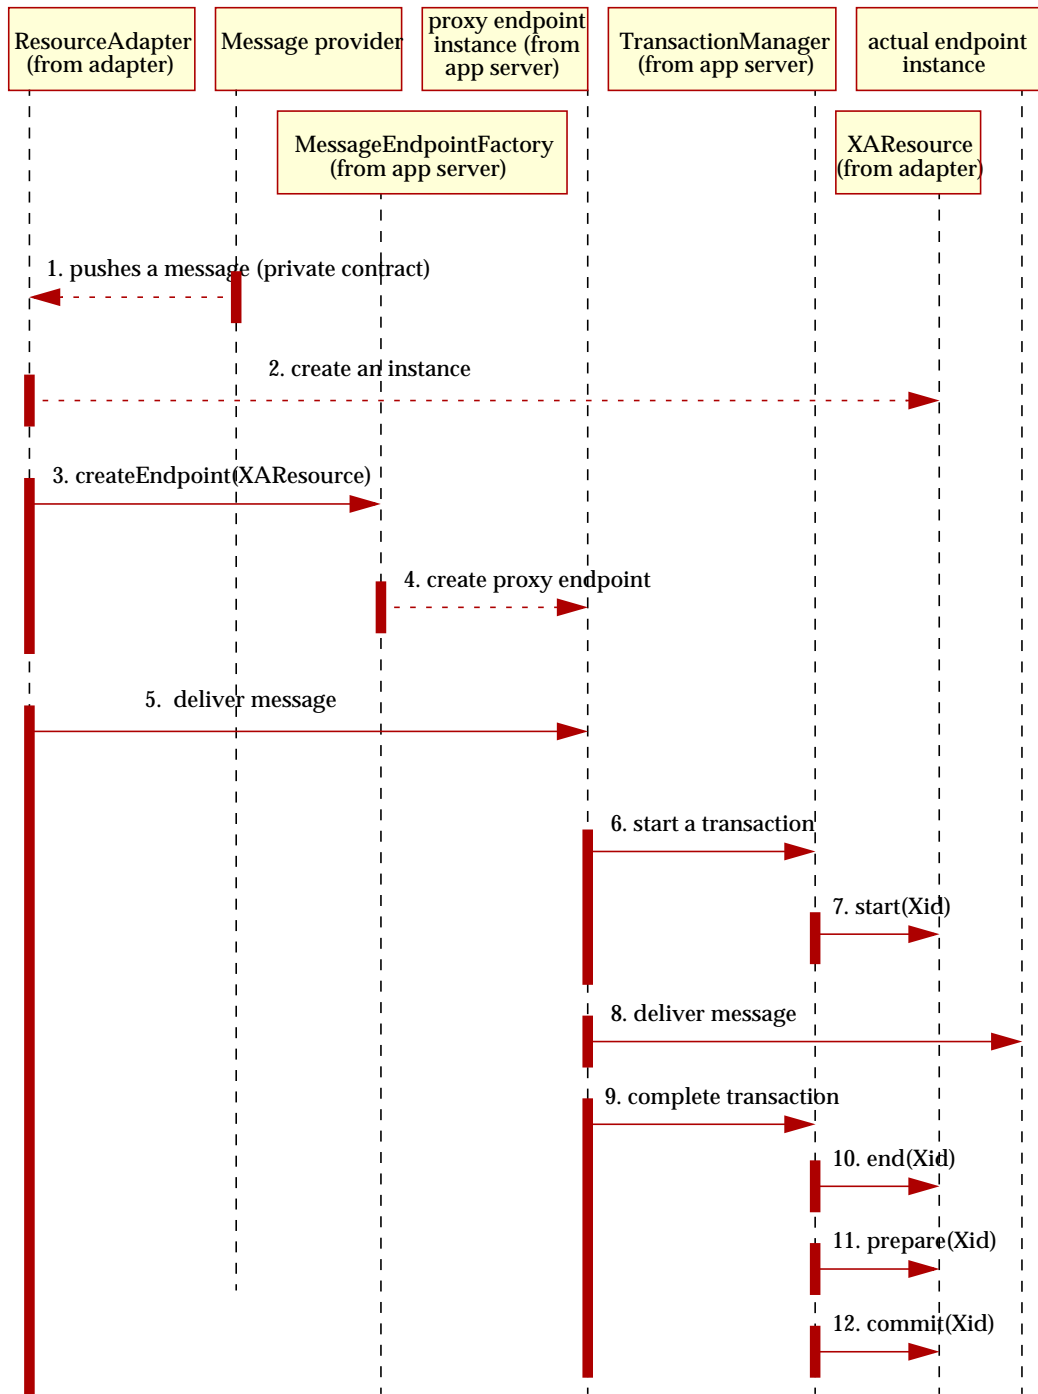


FIGURE 12-8 Transacted Message Delivery: Option B (sequence diagram)

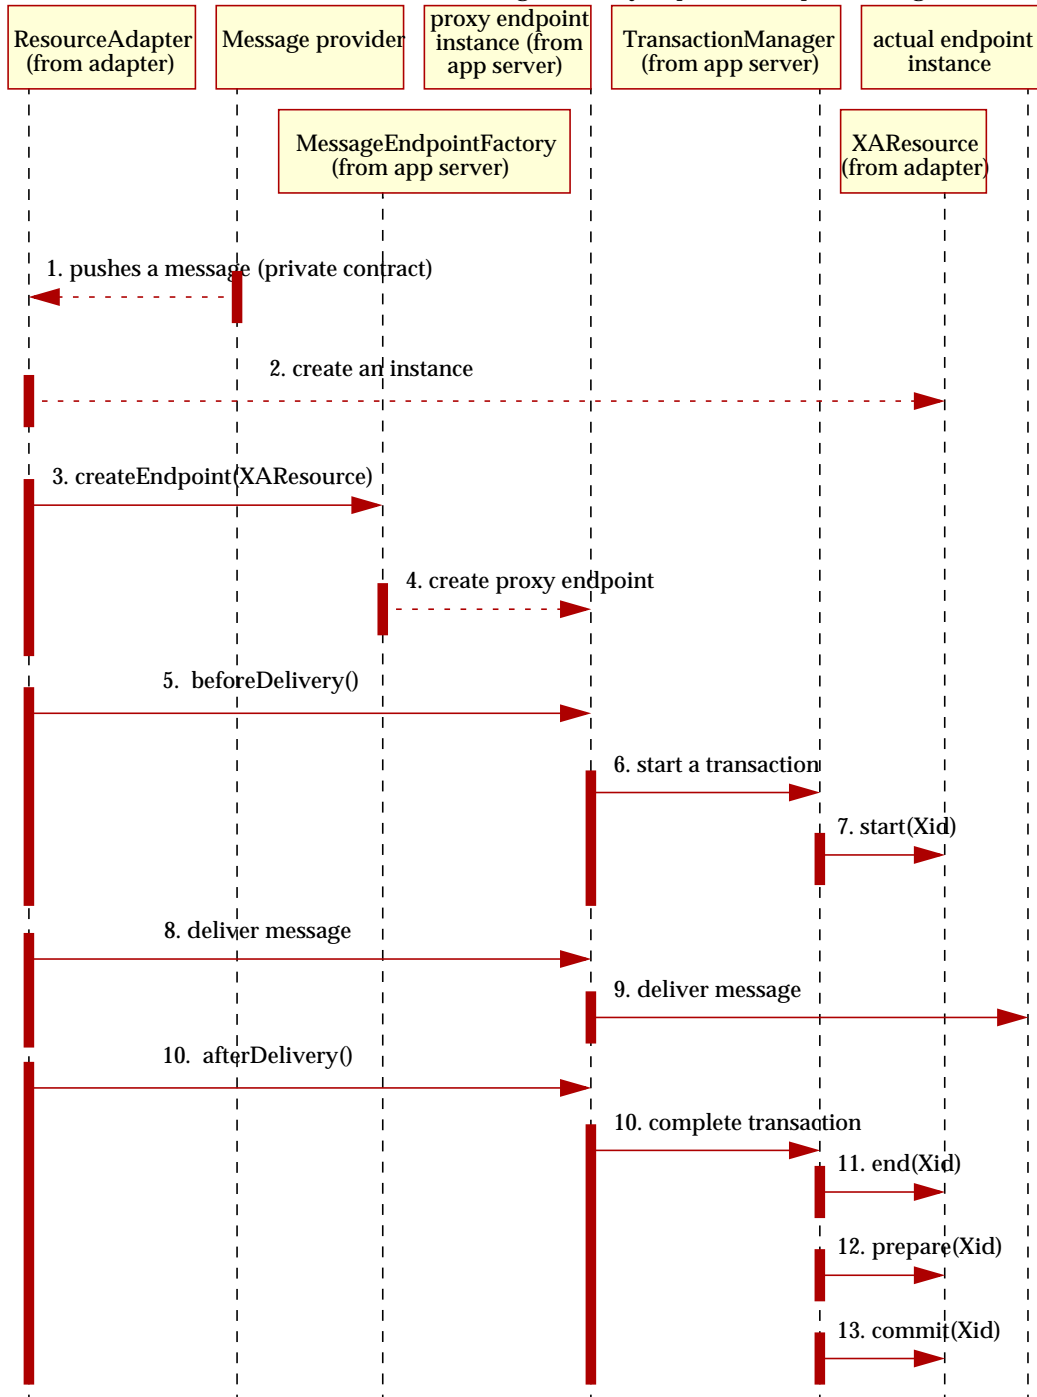


FIGURE 12-9 Non-transacted Message Delivery (sequence diagram)

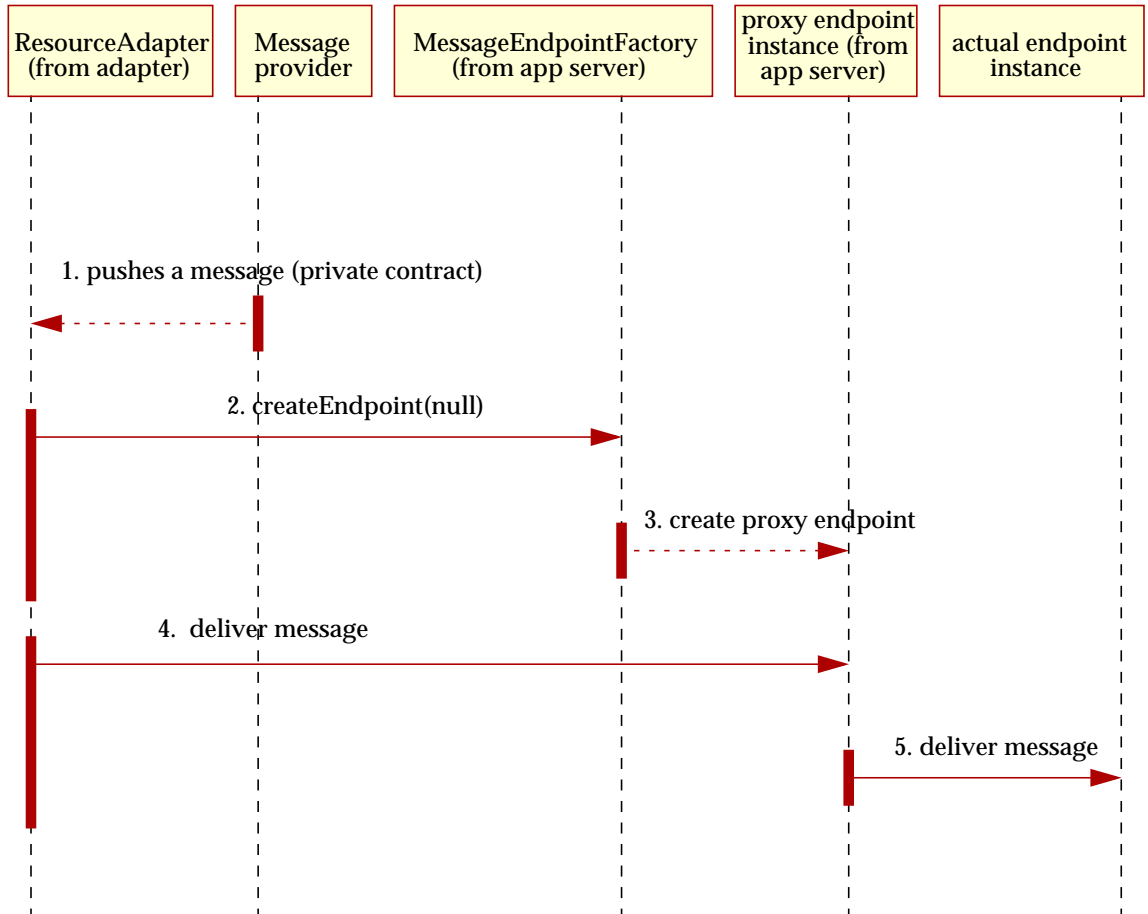
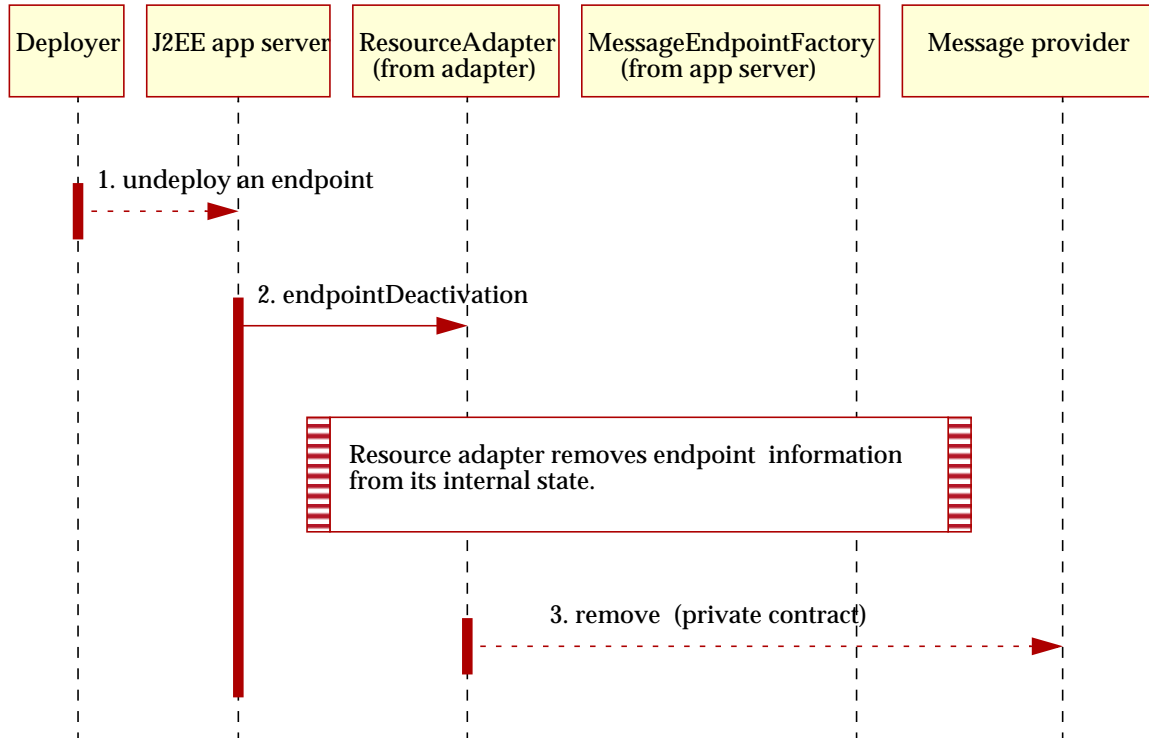


FIGURE 12-10 Endpoint Undeployment (sequence diagram)



12.7 Java Message Service (JMS) Use Case

For illustration purposes, a JMS use case involving a JMS resource adapter is discussed. The intent is to show how

- The JMS resource adapter uses the generic message inflow contract and asynchronously delivers messages to message-driven beans via the `onMessage` method on the `javax.jms.MessageListener` interface.
- The JMS resource adapter is used by an EJB application to send and synchronously receive messages via a `javax.jms.Connection` object.

This illustrates how a JMS provider is plugged into a J2EE application server using the standard connector contracts.

12.7.0.1 A Sample JMS Resource Adapter Deployment Descriptor

CODE EXAMPLE 12-5 Sample JMS Resource Adapter Deployment Descriptor

```
<?xml version="1.0" encoding="UTF-8"?>

<connector xmlns="http://java.sun.com/xml/ns/j2ee"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" version="1.5">
  <display-name>Wombat-JMSAdapter</display-name>
  <vendor-name>Wombat Software Inc.</vendor-name>
  <eis-type>JMS Provider</eis-type>
  <resourceadapter-version>1.0</resourceadapter-version>

  <resourceadapter>
    <resourceadapter-class>
      com.wombat.connector.jms.JMSAdapterImpl
    </resourceadapter-class>

    <!-- ResourceAdapter default configuration properties -->
    <config-property>
      <config-property-name>ServerName</config-property-name>
      <config-property-type>java.lang.String</config-property-type>
      <config-property-value>WombatServer</config-property-value>
    </config-property>
    <config-property>
      <config-property-name>PortNumber</config-property-name>
      <config-property-type>java.lang.String</config-property-type>
      <config-property-value>1050</config-property-value>
    </config-property>
    <config-property>
      <config-property-name>OperationalMode</config-property-name>
      <config-property-type>java.lang.String</config-property-type>
      <config-property-value>Managed</config-property-value>
    </config-property>
    <config-property>
      <config-property-name>ContainerType</config-property-name>
      <config-property-type>java.lang.String</config-property-type>
      <config-property-value>EJB-WEB</config-property-value>
    </config-property>

    <outbound-resourceadapter>

      <connection-definition>
        <managedconnectionfactory-class>
          com.wombat.connector.jms.QueueManagedConnectionFactoryImpl
        </managedconnectionfactory-class>

        <!-- ManagedConnectionFactory default configuration properties -->
        <config-property>
```

CODE EXAMPLE 12-5 Sample JMS Resource Adapter Deployment Descriptor

```
<config-property-name>ServerName</config-property-name>
<config-property-type>java.lang.String</config-property-type>
<config-property-value>
    WombatQueueServer
</config-property-value>

<config-property>
    <config-property-name>PortNumber</config-property-name>
    <config-property-type>java.lang.String</config-property-type>
    <config-property-value>1051</config-property-value>
</config-property>

<connectionfactory-interface>
    javax.jms.QueueConnectionFactory
</connectionfactory-interface>
<connectionfactory-impl-class>
    com.wombat.connector.jms.QueueConnectionFactoryImpl
</connectionfactory-impl-class>

<connection-interface>
    java.jms.QueueConnection
</connection-interface>
<connection-impl-class>
    com.wombat.connector.jms.QueueConnectionImpl
</connection-impl-class>
</connection-definition>

<connection-definition>
    <managedconnectionfactory-class>
        com.wombat.connector.jms.TopicManagedConnectionFactoryImpl
    </managedconnectionfactory-class>

    <!-- ManagedConnectionFactory default configuration properties -->
    <config-property>
        <config-property-name>ServerName</config-property-name>
        <config-property-type>java.lang.String</config-property-type>
        <config-property-value>
            WombatTopicServer
        </config-property-value>
    </config-property>
    <config-property>
        <config-property-name>PortNumber</config-property-name>
        <config-property-type>java.lang.String</config-property-type>
        <config-property-value>1052</config-property-value>
    </config-property>

    <connectionfactory-interface>
```


CODE EXAMPLE 12-5 Sample JMS Resource Adapter Deployment Descriptor

```
        javax.jms.TopicConnectionFactory
    </connectionfactory-interface>
    <connectionfactory-impl-class>
        com.wombat.connector.jms.TopicConnectionFactoryImpl
    </connectionfactory-impl-class>

    <connection-interface>
        java.jms.TopicConnection
    </connection-interface>
    <connection-impl-class>
        com.wombat.connector.jms.TopicConnectionImpl
    </connection-impl-class>
</connection-definition>

<connection-definition>
    <managedconnectionfactory-class>
        com.wombat.connector.jms.ManagedConnectionFactoryImpl
    </managedconnectionfactory-class>

    <!--
    This ManagedConnectionFactory JavaBean inherits the ResourceAdapter
    JavaBean configuration properties, and does not override any
    of the global defaults.
    -->

    <connectionfactory-interface>
        javax.jms.ConnectionFactory
    </connectionfactory-interface>
    <connectionfactory-impl-class>
        com.wombat.connector.jms.ConnectionFactoryImpl
    </connectionfactory-impl-class>

    <connection-interface>
        java.jms.Connection
    </connection-interface>
    <connection-impl-class>
        com.wombat.connector.jms.ConnectionImpl
    </connection-impl-class>
</connection-definition>

    <transaction-support>XATransaction</transaction-support>
    <reauthentication-support>false</reauthentication-support>

</outbound-resourceadapter>

<inbound-resourceadapter>
    <messageadapter>
```

CODE EXAMPLE 12-5 Sample JMS Resource Adapter Deployment Descriptor

```
<messagelistener>
  <messagelistener-type>
    javax.jms.MessageListener
  </messagelistener-type>

  <activationspec>
    <activationspec-class>
      com.wombat.connector.jms.ActivationSpecImpl
    </activationspec-class>

    <!--
      The endpoint deployer configures the ActivationSpec JavaBean
      and may override some of the global defaults inherited from
      the ResourceAdapter JavaBean. For example, the ServerName
      and the PortNumber properties.
    -->

    <!-- required config property names for ActivationSpec -->
    <required-config-property>
      <config-property-name>Destination</config-property-name>
    </required-config-property>

  </activationspec>

</messagelistener>
</messageadapter>
</inbound-resourceadapter>

<adminobject>
  <adminobject-interface>javax.jms.Queue</adminobject-interface>
  <adminobject-class>
    com.wombat.connector.jms.QueueImpl
  </adminobject-class>
</adminobject>
<adminobject>
  <adminobject-interface>javax.jms.Topic</adminobject-interface>
  <adminobject-class>
    com.wombat.connector.jms.TopicImpl
  </adminobject-class>
</adminobject>
<adminobject>
  <adminobject-interface>
    javax.jms.Destination
  </adminobject-interface>
  <adminobject-class>
    com.wombat.connector.jms.DestinationImpl
  </adminobject-class>
```

CODE EXAMPLE 12-5 Sample JMS Resource Adapter Deployment Descriptor

```
        </adminobject>

    </resourceadapter>
</connector>
```

12.7.0.2 A Sample JMS ActivationSpec Implementation

```
package com.wombat.connector.jms;

import java.io.Serializable;
import javax.resource.spi.ActivationSpec;
import javax.resource.spi.InvalidPropertyException;

public class ActivationSpecImpl implements ActivationSpec,
    Serializable {

    public setAcknowledgeMode(String mode) { ... }
    public String getAcknowledgeMode() { ... }

    public setSubscriptionDurability(String durability) { ... }
    public String getSubscriptionDurability() { ... }

    public setMessageSelector(String selector) { ... }
    public String getMessageSelector() { ... }

    public setDestinationType(String destType) { ... }
    public String getDestinationType() { ... }

    public setDestination(String dest) { ... }
    public String getDestination() { ... }

    public setSubscriptionName(String name) { ... }
    public String getSubscriptionName() { ... }

    public setClientId(String id) { ... }
    public String getClientId() { ... }

    public void validate() throws InvalidPropertyException { ... }
}
```

12.7.0.3 A Sample EJB 2.0 Message-driven Bean Deployment

Descriptor

```
<?xml version="1.0" encoding="US-ASCII"?>

<!DOCTYPE ejb-jar PUBLIC "-//Sun Microsystems, Inc.//DTD
Enterprise
JavaBeans 2.0//EN" 'http://java.sun.com/dtd/ejb-jar_2_0.dtd'>

<ejb-jar>
  <display-name>MDB1</display-name>
  <enterprise-beans>
    <message-driven>
      <display-name>MDB_DURABLE_CMT</display-name>
      <ejb-name>MDB_DURABLE_CMT</ejb-name>
      <ejb-class>msgbean.MsgBean</ejb-class>
      <transaction-type>Container</transaction-type>
      <message-driven-destination>
        <destination-type>javax.jms.Topic</destination-type>
        <subscription-durability>Durable</subscription-durability>
      </message-driven-destination>
      <message-selector>JMSType = 'car' AND color = 'blue'
      </message-selector>
    </message-driven>
  </enterprise-beans>
</ejb-jar>
```

12.7.0.4 A Sample EJB 2.1 Message-driven Bean Deployment Descriptor

CODE EXAMPLE 12-6 Message-Driven Bean Deployment Descriptor

```
<?xml version="1.0" encoding="UTF-8"?>

<ejb-jar xmlns="http://java.sun.com/xml/ns/j2ee"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  version="2.1">
  <display-name>MDB1</display-name>
  <enterprise-beans>
    <message-driven>
      <display-name>MDB_DURABLE_CMT</display-name>
      <ejb-name>MDB_DURABLE_CMT</ejb-name>
      <ejb-class>msgbean.MsgBean</ejb-class>
      <messaging-type>javax.jms.MessageListener
```

CODE EXAMPLE 12-6 Message-Driven Bean Deployment Descriptor

```
        </messaging-type>
    <transaction-type>Container</transaction-type>
    <activation-config>
        <activation-config-property>
            <activation-config-property-name>
                destinationType
            </activation-config-property-name>
            <activation-config-property-value>
                javax.jms.Topic
            </activation-config-property-value>
        </activation-config-property>
        <activation-config-property>
            <activation-config-property-name>
                SubscriptionDurability
            </activation-config-property-name>
            <activation-config-property-value>
                Durable
            </activation-config-property-value>
        </activation-config-property>
        <activation-config-property>
            <activation-config-property-name>
                MessageSelector
            </activation-config-property-name>
            <activation-config-property-value>
                JMSType = 'car' AND color = 'blue'
            </activation-config-property-value>
        </activation-config-property>
    </activation-config>
</message-driven>
</enterprise-beans>
</ejb-jar>
```

12.7.0.5 A Sample EJB 2.1 EJB Deployment Descriptor

CODE EXAMPLE 12-7 Sample EJB 2.1 Deployment Descriptor

```
<?xml version="1.0" encoding="UTF-8"?>
<ejb-jar xmlns="http://java.sun.com/xml/ns/j2ee"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  version="2.1">
  <display-name>Ejb1</display-name>
  <enterprise-beans>
    <session>
      <display-name>JMSBean</display-name>
      <ejb-name>JMSBean</ejb-name>
      <home>jms.JMSHome</home>
      <remote>jms.JMS</remote>
      <ejb-class>jms.JMSEJB</ejb-class>
      <session-type>Stateful</session-type>
      <transaction-type>Container</transaction-type>
      <resource-ref>
        <description>description</description>
        <res-ref-name>jms/MyQueueConnectionFactory</res-ref-name>
        <res-type>javax.jms.QueueConnectionFactory</res-type>
        <res-auth>Application</res-auth>
        <res-sharing-scope>Shareable</res-sharing-scope>
      </resource-ref>
      <resource-env-ref>
        <resource-env-ref-name>jms/QueueName</resource-env-ref-name>
        <resource-env-ref-type>javax.jms.Queue</resource-env-ref-type>
      </resource-env-ref>
    </session>
  </enterprise-beans>
</ejb-jar>
```

12.7.1 Message-driven Bean Asynchronously Receiving Messages

12.7.1.1 Message-Driven Bean Deployment

- A message-driven bean application developer or assembler supplies a deployment descriptor which specifies a destination type, message selector, and subscription durability information needed to setup subscription to a certain destination, Queue or Topic. Note, this information is a hint which is used by the message-driven bean deployer to setup the subscription.

- The message-driven bean deployer selects an appropriate JMS resource adapter based on the quality of service and creates an `ActivationSpec` `JavaBean` instance and configures the required property "Destination" as well as other properties related to the JMS messaging style and the specific resource adapter.
- The endpoint deployer may need to interact with the JMS provider to setup an appropriate "Destination" and other steps necessary to complete message-driven bean deployment.
- The deployer deploys the message-driven bean application. During deployment, the deployer provides the configured `ActivationSpec` `JavaBean` to the application server, along with information about the chosen JMS resource adapter.
- The application server calls the `endpointActivation` method on the JMS resource adapter and passes the configured `ActivationSpec` `JavaBean` instance and a `MessageEndpointFactory` instance. During the `endpointActivation` method call the JMS adapter interacts with its provider to setup message delivery to the message-driven bean. This completes the endpoint activation, and the message-driven bean is ready to receive messages.

12.7.1.2 Message Delivery

- When messages start arriving, the JMS adapter uses the `MessageEndpointFactory` instance to get an endpoint instance and delivers messages to the endpoint via the `javax.jms.MessageListener.onMessage` method.
- The application server interposes the message delivery and injects transactions based on the message-driven bean preferences, container-managed transaction or bean-managed transaction, before delivering the message to a message-driven bean instance.
- When a delivery is transacted, the application server notifies the JMS resource adapter using the `XAResource` object. The JMS resource adapter may use the notifications to send acknowledgements to its message provider.
- The JMS resource adapter, depending on the traffic, may attempt concurrent delivery of messages by using multiple endpoint instances obtained via `MessageEndpointFactory`. The application server appropriately handles concurrent message deliveries and dispatches messages to separate message-driven bean instances.

12.7.1.3 Message-Driven Bean Undeployment

- When the message-driven bean is undeployed, the application server calls the `endpointDeactivation` method on the JMS resource adapter to deactivate the message endpoint.
- The JMS adapter in turn notifies its message provider.

12.7.2 EJB Using JMS API to Send and Synchronously Receive Messages Via a JMS Resource Adapter

- The JMS resource adapter provides `javax.jms.Connection` objects which expose the JMS API to the EJB application. The EJB directly uses the JMS API to send and synchronously receive messages. The `javax.jms.Connection` objects are obtained from a `ConnectionFactory` supplied by the JMS resource adapter.
- Based on the EJB deployment descriptor information (`resource-ref`'s and `resource-env-ref`'s), the EJB deployer configures the appropriate `ConnectionFactory` objects (`resource-ref`'s) in the component name space (`java:comp/env`). The application deployer also configures the necessary `Queue` or `Topic` administered objects (`resource-env-ref`'s) in the component name space. The JMS resource adapter provides the implementation of the various `ConnectionFactory` and administered objects.
- At runtime, the component does a JNDI lookup of a `ConnectionFactory` object from its component name space (`java:comp/env`), and uses it to create a `javax.jms.Connection` object which is used for sending and synchronously receiving messages. Similarly, the component uses the JNDI lookup mechanism to obtain the configured `Queue` or `Topic` administered objects.

12.7.2.1 Using JMS API to Send Messages

CODE EXAMPLE 12-8 Sending Messages Using the JMS API

```
// get JNDI handle
Context jndiContext = new InitialContext();

// get connection factory
ConnectionFactory connectionFactory = (ConnectionFactory)
    jndiContext.lookup("QueueConnectionFactory");

// get connection from factory
Connection connection = connectionFactory.getConnection();

// get session from connection
Session session = connection.createSession(true,
    AUTO_ACKNOWLEDGE);

// get destination from JNDI
Queue stockQueue = (Queue) jndiContext.lookup("StockQueue");

// create a message producer
MessageProducer sender = session.createProducer(stockQueue);

// create a message
TextMessage message = session.createTextMessage();
message.setText(msgData);

// send the message
sender.send(message);
```

12.7.2.2 J2EE Component Using JMS API to Synchronously Receive Messages

CODE EXAMPLE 12-9 Synchronously Receiving Messages in a J2EE Component

```
// get JNDI handle
Context jndiContext = new InitialContext();

// get connection factory
ConnectionFactory connectionFactory = (ConnectionFactory)
    jndiContext.lookup("QueueConnectionFactory");

// get connection from factory
Connection connection = connectionFactory.getConnection();

// get session from connection
```

CODE EXAMPLE 12-9 Synchronously Receiving Messages in a J2EE Component

```
Session session = connection.createSession(true,
    AUTO_ACKNOWLEDGE);

// get destination from JNDI
Queue stockQueue = (Queue) jndiContext.lookup("StockQueue");

// create a message consumer
MessageConsumer receiver = session.createConsumer(stockQueue);

// enable connection to receive messages
connection.start();

// synchronously receive the message
TextMessage message = (TextMessage) receiver.receive(message);
```

12.8 A Non-JMS Use Case

This illustration is intended to show that it is possible to plug a wide range of message providers into a J2EE application server via the standard connector contracts, such that it is possible for an application to either asynchronously receive messages via the message inflow contract or to use a connection object to send and synchronously receive messages.

12.8.1 Resource Adapter Deployment Descriptor

This is an example deployment descriptor for a resource adapter that can provide both inbound and outbound communication with a particular EIS.

On the inbound side, it can deliver messages to a message-driven bean that implements a `com.kangaroo.MessageListener`. Note, the deployment descriptor has a `messagelistener`-type element with the value `com.kangaroo.MessageListener`. The `activation-spec-class` is of type `com.kangaroo.MyEISActivationSpecImpl`. This `ActivationSpec` `JavaBean` has a single required property `PortNumber`, that is required to establish a connection to the remote EIS. When the EIS data is received, the resource adapter will convert it to a `com.kangaroo.Message` and deliver it to a message-driven bean instance.

The resource adapter also provides a `ManagedConnectionFactory` implementation for outbound communication to the EIS. This also takes a single configuration parameter called `PortNumber`.

CODE EXAMPLE 12-10 Deployment Descriptor for a Resource Adapter

```
<?xml version="1.0" encoding="UTF-8"?>

<connector xmlns="http://java.sun.com/xml/ns/j2ee"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  version="1.5">

  <display-name>KangarooAdapter</display-name>
  <vendor-name>Kangaroo Software Inc.</vendor-name>
  <eis-type>JMS Provider</eis-type>
  <resourceadapter-version>1.0</resourceadapter-version>

  <resourceadapter>
    <resourceadapter-class>
      com.kangaroo.MyEISAdapterImpl
    </resourceadapter-class>

    <!-- ResourceAdapter default configuration properties -->
    <config-property>
      <description>URL for EIS instance</description>
      <config-property-name>EIS_URL</config-property-name>
      <config-property-type>java.lang.String</config-property-type>
      <config-property-value>TBD</config-property-value>
    </config-property>

    <outbound-resourceadapter>
      <connection-definition>
        <managedconnectionfactory-class>
          com.Kangaroo.MyManagedConnectionFactoryImpl
        </managedconnectionfactory-class>

        <!-- ManagedConnectionFactory default configuration properties -->
        <config-property>
          <config-property-name>PortNumber</config-property-name>
          <config-property-type>java.lang.String</config-property-type>
          <config-property-value>1051</config-property-value>
        </config-property>

        <connectionfactory-interface>
          javax.resource.cci.ConnectionFactory
        </connectionfactory-interface>
        <connectionfactory-impl-class>
          com.Kangaroo.MyConnectionFactoryImpl
        </connectionfactory-impl-class>
      </connection-definition>
    </outbound-resourceadapter>
  </resourceadapter>
</connector>
```

CODE EXAMPLE 12-10 Deployment Descriptor for a Resource Adapter

```
        <connection-interface>
            javax.resource.cci.Connection
        </connection-interface>
        <connection-impl-class>
            com.Kangaroo.MyConnectionImpl
        </connection-impl-class>
    </connection-definition>

    <transaction-support>NoTransaction</transaction-support>
    <reauthentication-support>false</reauthentication-support>
</outbound-resourceadapter>

<inbound-resourceadapter>
    <messageadapter>
        <messagelistener>
            <messagelistener-type>
                com.kangaroo.MessageListener
            </messagelistener-type>
            <activationspec>
                <activationspec-class>
                    com.Kangaroo.MyEISActivationSpecImpl
                </activationspec-class>
                <required-config-property>
                    <config-property-name>PortNumber</config-property-name>
                </required-config-property>
            </activationspec>
        </messagelistener>
    </messageadapter>
</inbound-resourceadapter>

</resourceadapter>
</connector>
```

12.8.2 Resource Adapter Deployment

Before use, the resource adapter needs to be deployed on the application server. During resource adapter deployment, the deployer configures a `ResourceAdapter` JavaBean instance and deploys the resource adapter.

12.8.3 Message-Driven Bean Asynchronously Receiving Notifications From an EIS

12.8.3.1 The Message-Driven Bean Deployment Descriptor

CODE EXAMPLE 12-11 Deployment Descriptor for a Message-Driven Bean

```
<?xml version="1.0" encoding="US-ASCII"?>

<!DOCTYPE ejb-jar PUBLIC "-//Sun Microsystems, Inc.//DTD
Enterprise JavaBeans 2.1//EN" ?http://java.sun.com/dtd/ejb-jar_2_1.dtd?>

<ejb-jar>
  <display-name>Ejb1</display-name>
  <enterprise-beans>
    <message-driven>
      <display-name>EIS Receiver Bean</display-name>
      <ejb-name>EISReceiver</ejb-name>
      <ejb-class>myapp.EISReceiverBean</ejb-class>
      <messaging-type>com.kangaroo.MessageListener</messaging-type>
      <transaction-type>Container</transaction-type>
      <activation-config>
        <activation-config-property>
          <activation-config-property-name>
            functionName
          </activation-config-property-name>
          <activation-config-property-value>
            CustomerChangeNotification
          </activation-config-property-value>
        </activation-config-property>
        <activation-config-property>
          <activation-config-property-name>
            CustomerName
          </activation-config-property-name>
          <activation-config-property-value>
            Wombat Inc.
          </activation-config-property-value>
        </activation-config-property>
      </activation-config>
    </message-driven>
  </enterprise-beans>
</ejb-jar>
```

When the message-driven bean is deployed, the bean deployer chooses an appropriate resource adapter based on the type of the message listener it supports. In this case, the deployer chooses the resource adapter with the `ResourceAdapter`

`class com.vendor.MyEISAdapterImpl` since it supports the `com.kangaroo.MessageListener` type.

Then the deployer creates an instance of `com.vendor.MyEISActivationSpecImpl` and populates it with values. The `ActivationSpec` `JavaBean` instance will also contain values of properties that are set in the `activation-config` section of the bean's deployment descriptor. In the example above, the properties are `FunctionName`, with the value `CustomerChangeNotification`, and `CustomerName`, with the value `Wombat, Inc.`, which the deployer may choose to override.

Finally, the deployer provides the configured `ActivationSpec` `JavaBean` instance to the application server.

12.8.4 Message-Driven Bean and Resource Adapter Activation

When the application server is started, it will activate the resource adapter by calling its `start` method. The application server will create an instance of the message-driven bean with class name `myapp.EISReceiverBean`. Then the application server will call the `endpointActivation` method on the resource adapter instance and pass in the configured `ActivationSpec` instance associated with the deployed message-driven bean, and a `MessageEndpointFactory` instance. The resource adapter will use the information in the `ActivationSpec` to establish a subscription to the requested data from the EIS.

12.8.5 Message Delivery

When a notification arrives from the EIS, the resource adapter has the responsibility of converting its data to a `com.kangaroo.Message` if it is not already in this format. The resource adapter will then use the `MessageEndpointFactory` to deliver the notification to the message-driven bean. Note that, rather than calling the `MessageEndpointFactory` directly, the resource adapter is likely to instantiate a `Work` object, and pass it to the application server via the `WorkManager` interface. When the `doWork` method of the `WorkManager` is called the dispatch will occur. This will allow the resource adapter to continue to process incoming messages without blocking until message-driven bean dispatch has completed.

EJB Invocation

This chapter describes how to invoke session and entity beans from a resource adapter.

13.1 Overview

A resource adapter may need to call session or entity beans for several reasons:

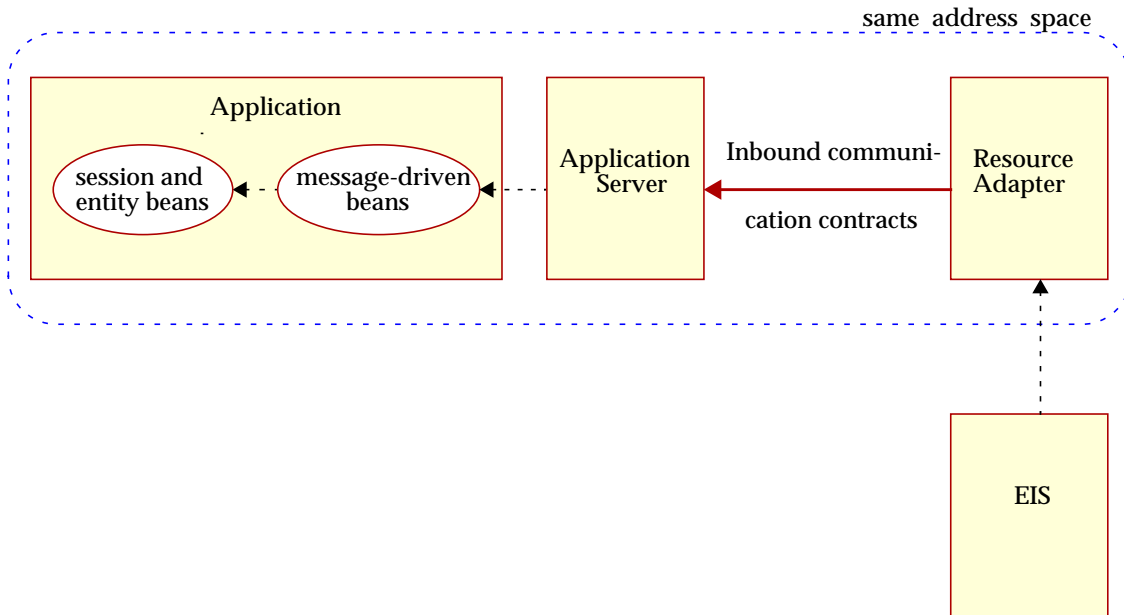
- To dispatch calls from an EIS to a bean in order to execute business logic.
- To use EJB container-managed persistence (CMP) mechanism for persistence.

In order to dispatch calls to a session or entity bean, the resource adapter needs to know the target bean type, the method name, and the method parameters. Upon receiving a request from the EIS via a remote protocol, the resource adapter's dispatch logic needs to:

- Choose an appropriate bean and a target method name based on the request received from the EIS.
- Unmarshall, that is, deserialize, the request parameters received from the EIS and call the target bean method.

13.2 EJB Invocation Model

FIGURE 13-1 EJB Invocation Model



For session or entity bean invocations, the resource adapter's bean dispatch logic can use the bean client view model via a message-driven bean. The EJB specification (Section 1., "Enterprise JavaBeans (EJBTM) specification, version 2.1:" on page F-1) defines the EJB client view, and describes how the client view is used to access session or entity beans. The EJB client view is available to a message-driven bean.

The resource adapter could structure its code such that its bean dispatch logic is written as a message-driven bean. The message-driven bean chooses an appropriate session or entity bean and a target method, unmarshalls the request parameters and invokes the chosen bean based on the request information received from the EIS.

The resource adapter can use the message inflow contract to call a message-driven bean, and use the message-driven bean to dispatch calls to session and entity beans using the EJB client view model. The EJB specification allows a request-response style message-driven bean call which could be used for synchronous RPC-style calls. The message-driven bean could be packaged either with the resource adapter or separately.

Thus, the message-driven bean could be used as a replaceable unit of the resource adapter which serves the job of a bean dispatcher. The message inflow contract allows the creation of multiple endpoint instances (message-driven beans) at runtime, and hence it is possible to do concurrent bean dispatches.

Further, the transaction inflow mechanism (described in Chapter 14, “Transaction Inflow”) allows the resource adapter to use the transaction information obtained from the EIS for bean invocations. Note, however, the application server may suspend the imported transaction depending on the transaction preference of the target bean method.

13.3 An Illustrative Use Case

Wombat Inc. is a resource adapter vendor. The resource adapter supports inbound communication from an EIS to application components residing in an application server container. The resource adapter uses the message inflow contract to call message-driven beans which serve as a dispatcher for session and entity bean invocations. In this case, Wombat Inc. supplies both the resource adapter and the message-driven beans, even though these could be supplied by different vendors.

The EIS uses multiple concurrent conversations in its interactions with the resource adapter. Each conversation may involve multiple serial requests. The resource adapter has a set of `Work` objects (threads), each of which is used for carrying on a specific conversation. The resource adapter manages the multiple concurrent conversations, and calls a message-driven bean instance whenever a request message arrives as part of a conversation.

The following code sample shows a possible message-driven bean implementation:

CODE EXAMPLE 13-1 A Message-Driven Bean Implementation

```
package com.wombat.ra;

import javax.ejb.MessageDrivenBean;
import javax.naming.InitialContext;

public class WombatMDB
    implements MessageDrivenBean, WombatMessageListener {

    public static int CONV_START = 0;
    public static int CONV_CONTINUE = 1;
    public static int CONV_END = 2;

    private Context jndiContext = null;
    private ConvBeanHome chome = null;
```

CODE EXAMPLE 13-1 A Message-Driven Bean Implementation

```
public void ejbCreate() {
    jndiContext = new InitialContext();
    chome = (ConvBeanHome)
        jndi.lookup("java:comp/env/ConvBeanHome");
}

ConvResponse onMessage(ConvRequest requestMsg) {
    // get conversation id and request type from the request
    // message
    int convId = ...;
    int type = ...;

    if (type == CONV_START) {
        // create entity EJB for holding the specific
        // conversation state
        ConvBean cbean = chome.create(convId);
    } else if (type == CONV_CONTINUE) {
        ConvBean cbean = chome.findByPrimaryKey(convId);

        // unmarshall EJB method parameters
        ...;

        // invoke EJB and return response
        Object resp = cbean.myBusinessMethod(params);
        ConvResponse cresp = Utility.convert(resp);
        return cresponse;
    } else if (type == CONV_END) {
        cbean.remove();
    }

    return null;
}

public void ejbRemove() {
    jndiContext = null;
    chome = null;
}
}
```

The resource adapter uses the message-driven bean as a generic dispatcher for session and entity bean invocations, and relies on the application server to efficiently pool message-driven bean instances. Each message-driven bean call should be just as efficient as a method call on a resource adapter local object.

13.3.0.1 Message-driven Bean Dispatcher Pattern

When a worker thread from a resource adapter accesses a message-driven bean method, the JNDI context of the bean is available to the thread, although only during the method call on the bean.

The resource adapter may take advantage of this, and use the bean as a dispatcher. That is, the resource adapter may park the thread within the bean method inside a while loop, and use it to process resource adapter specific data structures passed into the bean method as method parameters, and also use the JNDI context of the bean to access resources and other components.

In this case, the bean becomes a special Java object that has access to JNDI context, which the resource adapter may use. This usage pattern illustrates a tight coupling between the resource adapter and the message-driven bean, and it is likely that the resource adapter would provide the bean implementation as well.

Transaction Inflow

This chapter specifies a contract between an application server and a resource adapter that allows a resource adapter to propagate an imported transaction to an application server, so that the application server and subsequent participants can do work as part of the imported transaction. This contract also allows a resource adapter to flow-in transaction completion and crash recovery calls initiated by an EIS, and ensures that the atomicity, consistency, isolation and durability (ACID) properties of the imported transaction are preserved.

14.1 Overview

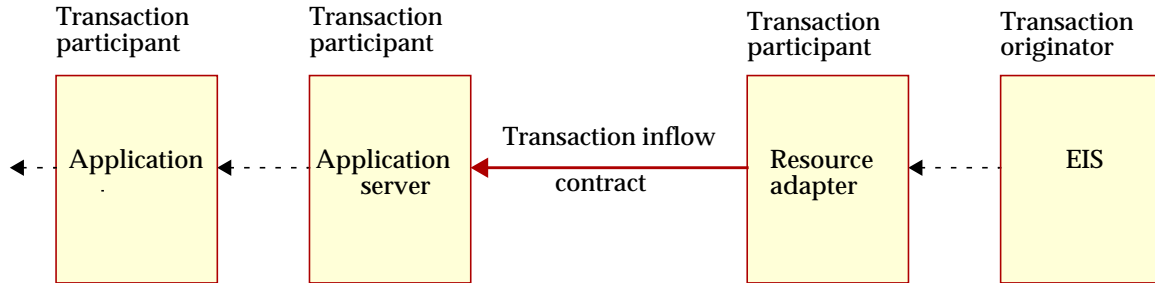
A resource adapter may need to import an incoming external transaction context obtained from a remote protocol message and do work as part of the imported transaction. The work done by the resource adapter as part of the imported transaction may involve interactions with the application server and the application components.

The resource adapter is expected to know the wire protocol and the transaction context format and be able to import an incoming transaction in an EIS-specific way. The resource adapter needs to propagate the imported transaction to the application server and also flow-in transaction completion and crash recovery calls initiated by the EIS. In order for the resource adapter to accomplish this, it needs the following:

- A standard form to represent the transaction context imported by the resource adapter.
- A mechanism to associate the work done by the resource adapter as part of the imported transaction.

- A mechanism to treat the application server like a resource manager in order to make it participate in the 2-phase commit and crash recovery flows initiated by the external transaction originator, the EIS.

FIGURE 14-1 Transaction Inflow Contract



14.2 Goals

- Provide a standard mechanism for a resource adapter to propagate an imported transaction to an application server.
- Provide a standard mechanism for a resource adapter to flow-in transaction completion and crash recovery calls from an EIS.
- Ensure that the ACID properties of the transaction imported by a resource adapter are preserved.

14.3 Use Case Scenario

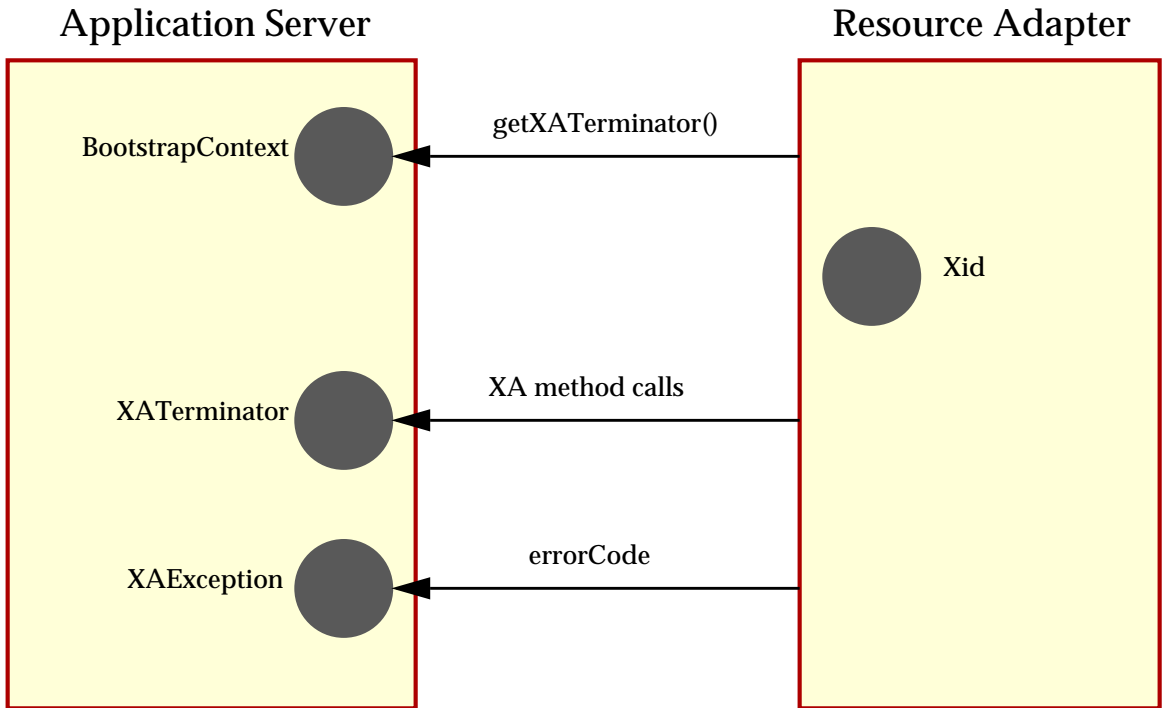
An EIS initiates a transaction and calls application components residing in an application server via a resource adapter. The EIS propagates the transaction context as part of each call to the resource adapter, which is used by the application server to recreate the transaction context before the application components are called. The work done by the application components is automatically enlisted as part of the imported transaction. When the EIS completes the transaction, the transaction completion notifications flow to the application server via the resource adapter, and the transaction is completed.

The transaction inflow contract may be used in various situations. For example:

- A message provider may use the contract to push messages to a resource adapter for delivery to application components.

- The contract may be leveraged to make the application components do work as part of a transaction initiated by a legacy EIS¹.

FIGURE 14-2 Transaction Inflow Contract (object diagram)



1. Note that application components may not always do work as part of an EIS-initiated transaction, for example, when the declarative transaction attribute of an enterprise bean's method is *RequiresNew*, *Never*, *NotSupported*, or if bean-managed transaction demarcation is used.

14.4 Transaction Inflow Model

```
package javax.resource.spi;

import javax.transaction.xa.Xid;
import java.transaction.xa.XAException;
import javax.resource.spi.work.WorkManager;

public interface BootstrapContext {
    WorkManager getWorkManager();
    XATerminator getXATerminator();
    ... // other methods
}

public interface XATerminator {
    public void commit(Xid xid, boolean onePhase) throws
XAException;
    public void forget(Xid xid) throws XAException;
    public int prepare(Xid xid) throws XAException;
    public Xid[] recover(int flag) throws XAException;
    public void rollback(Xid xid) throws XAException;
}
```

The `BootstrapContext` interface allows the resource adapter to obtain a `WorkManager` instance and an `XATerminator` instance. These instances are not required to be unique. The resource adapter uses the `WorkManager` instance to submit `Work` instances for execution, and uses the `XATerminator` instance for transaction completion and crash recovery flows.

14.4.1 Processing of Transactional Calls

The steps involved in propagating an imported transaction from a resource adapter to the application server in order to do transactional work is as follows:

- The EIS makes a transactional call to the resource adapter. The resource adapter is expected to know the EIS-specific transaction context structure and the wire protocol. The resource adapter imports the transaction context that arrived along with the incoming message.
- The resource adapter represents the imported transaction context in a standard form using the `javax.transaction.xa.Xid` instance.

- The resource adapter constructs a `Work` instance, which is expected to do work as part of the transactional message, and also creates an `ExecutionContext` instance containing the constructed `Xid`. It then submits the `Work` instance along with the `ExecutionContext` instance to the application server's `WorkManager` for execution.
- The application server's `WorkManager` accepts the submitted `Work` instance and recreates the execution context for the `Work` instance. That is, the work to be done is enlisted as part of the imported transaction. It then calls the `run` method on the `Work` object.
Note, however, all the work done by the `Work` object may not be part of the transaction. For example, the application server may suspend the imported transaction depending on the transaction preference of the bean method that may be invoked.

The above steps may be repeated any number of times for a particular transaction from any resource adapter. However, the application server must disallow transactional `Work` submissions with a `WorkCompletedException` set to an appropriate error code, irrespective of which resource adapter it comes from, under the following circumstances:

- If a `Work` instance associated with the transaction is already present. That is, concurrent work within an imported transaction is disallowed. The error code to indicate this is `WorkException.TX_CONCURRENT_WORK_DISALLOWED`.
- The application server is unable to recreate the transaction. That is, it fails in its attempt to enlist the `Work` instance with the transaction. The error code to indicate this is `WorkException.TX_RECREATE_FAILED`.

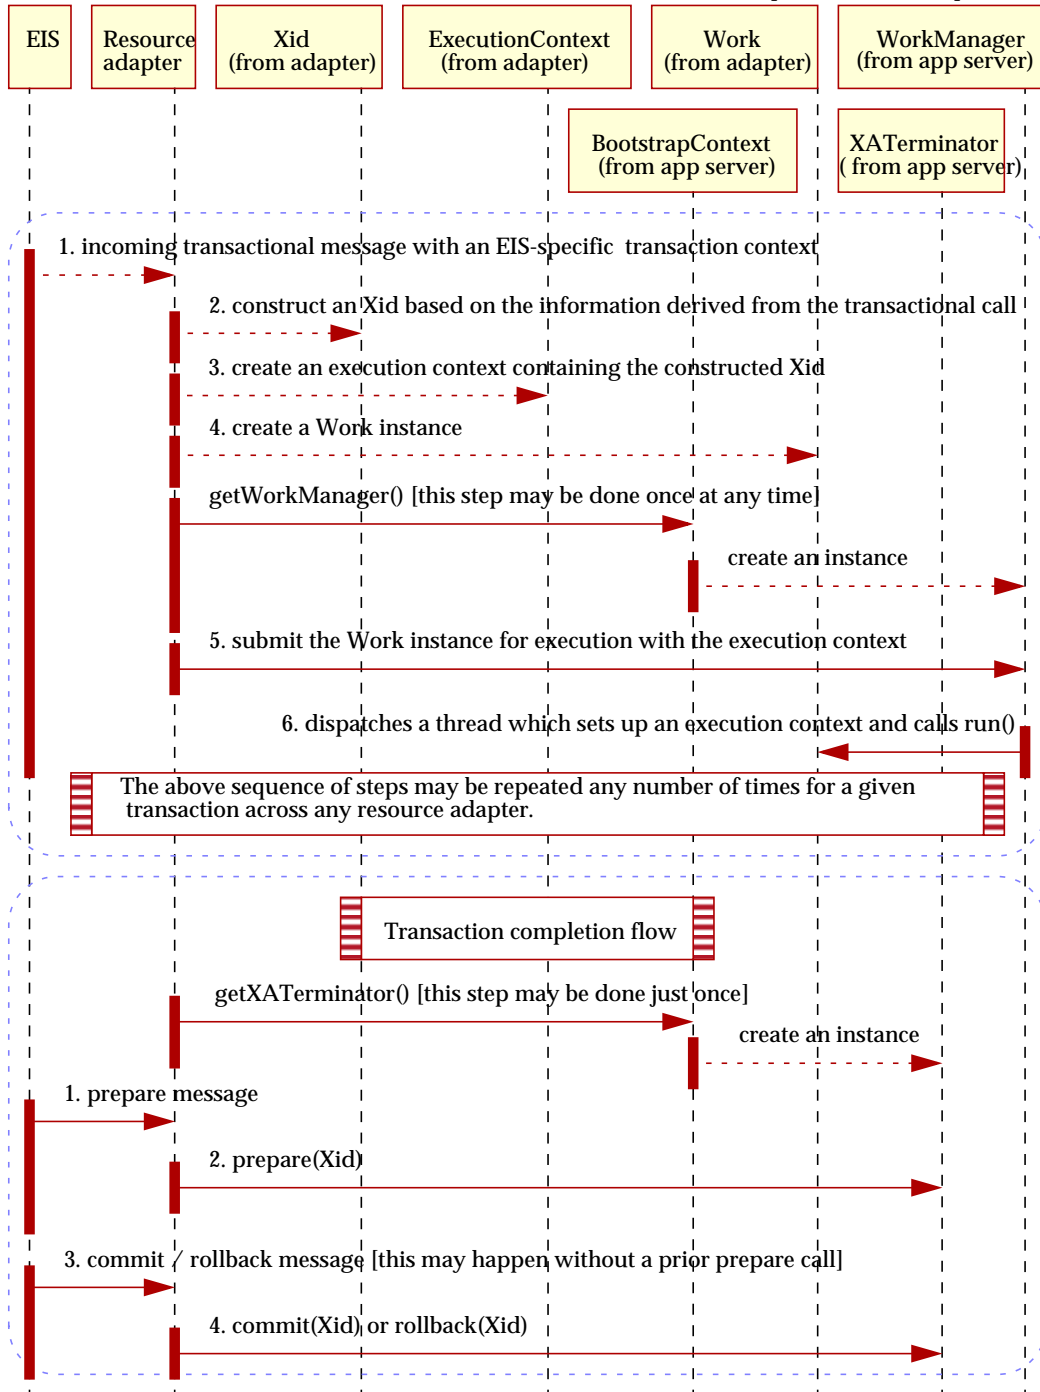
14.4.2 Transaction Completion Processing

The steps involved in completing of the imported transaction initiated by the external EIS are as follows:

- The EIS sends a prepare message for a particular transaction.
- The resource adapter obtains an `XATerminator` instance from the application server via the `getXATerminator` method of the `BootstrapContext` instance. Note, this step may be done at any time, and the obtained `XATerminator` instance may be used for transaction completion flows across multiple imported transactions. The `XATerminator` implementation should be thread-safe and re-entrant.
- The resource adapter calls the `prepare` method of the `XATerminator` instance with an appropriate `Xid` instance, and returns the outcome of the prepare operation to the EIS.

- When the EIS sends a commit message for the transaction, the resource adapter calls the `commit` method of the `XATerminator` instance with an appropriate `Xid` instance. Note, it is possible for the `commit` method to be called without a prior `prepare` method call in the case of one-phase commit.

FIGURE 14-3 Transactional Calls and Transaction Completion Flow (sequence diagram)



14.4.3 Crash Recovery Processing

- If the EIS detects the failure of the application server while the transaction is active (that is, transaction completion has not begun), it does not wait for the application server to recover or do any recovery processing. The transactional work done by the application server site is presumed aborted.
- If the EIS detects the failure of the application server while the transaction is in-doubt (that is, the transaction has already been prepared), the EIS retries completion by attempting to re-establish network communication until it succeeds. When the application server recovers from the crash, it should recover the state of all transactions that were successfully prepared before the crash, and complete them upon receiving a `commit` method or `rollback` method call from the EIS via the resource adapter.
- If the resource adapter detects the failure of the EIS while the transaction is active (that is, transaction completion has not begun), it aborts all active transactions that originated from the EIS. The resource adapter should keep a list of active transactions and abort them upon EIS failure.

- If the resource adapter detects the failure of the EIS while the transaction is in-doubt (that is, the transaction has already been prepared), it waits for the EIS to recover. When the EIS recovers, it re-establishes network communication with the resource adapter, and queries it for a list of in-doubt transactions. It then completes the in-doubt transactions.

FIGURE 14-4 Crash Recovery Flows When Application Server Crashes (sequence diagram)

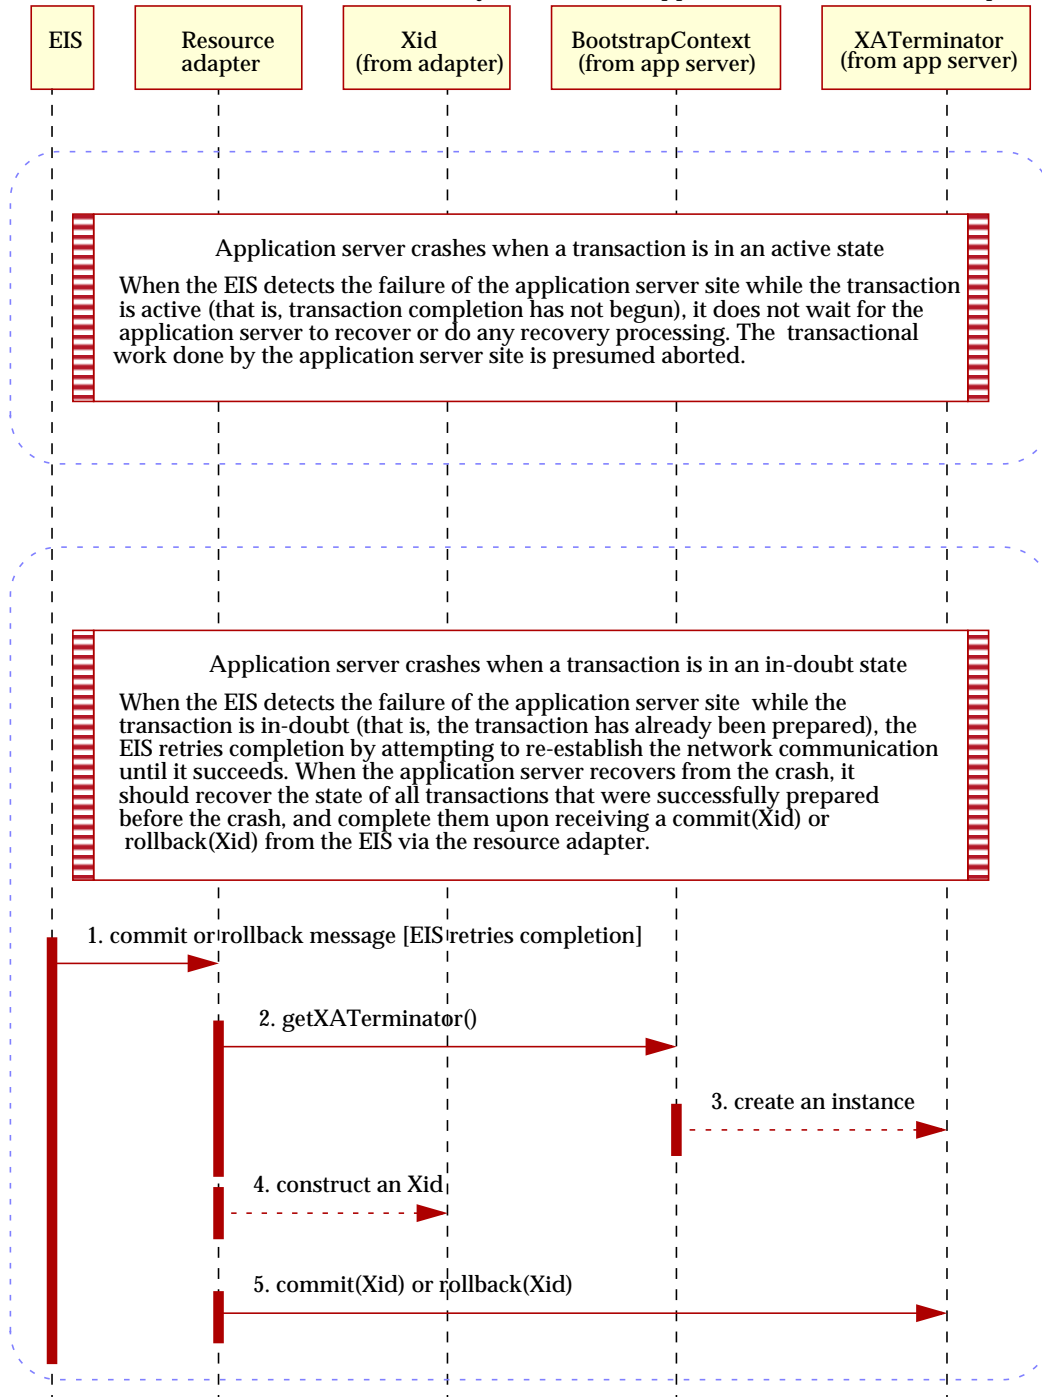
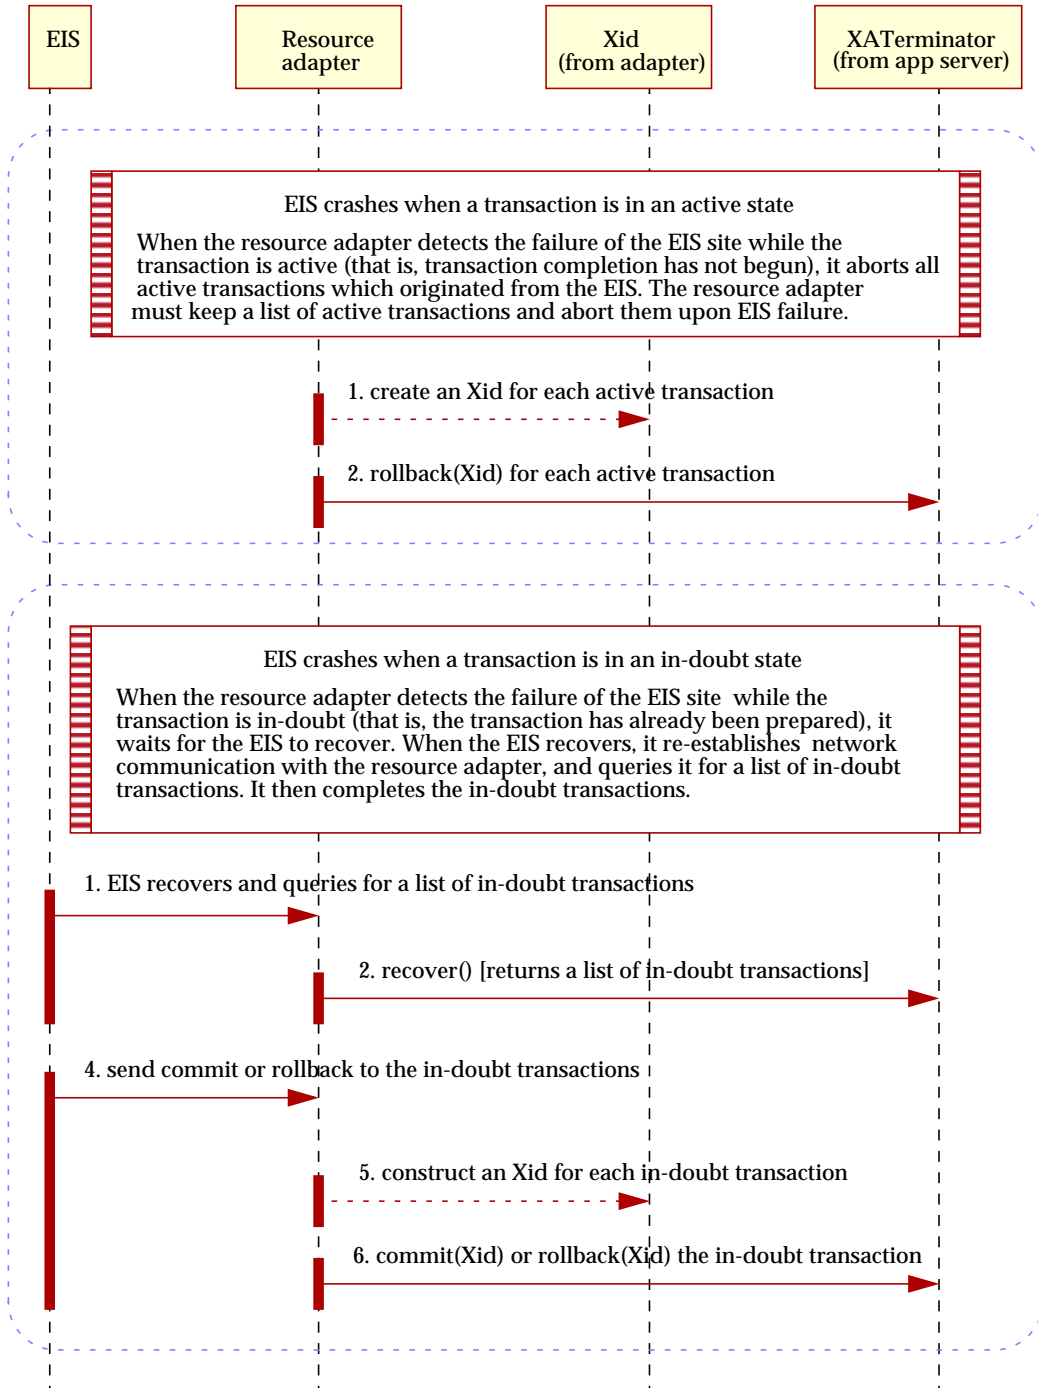


FIGURE 14-5 Crash Recovery Flows When EIS Crashes (sequence diagram)



14.4.4 Requirements

- An application server must implement the transaction inflow contract. That is, it must allow `Work` submissions with a transaction context, an `Xid`, and provide a valid `XATerminator` instance when called via the `getXATerminator` method of the `BootstrapContext` instance.
- A resource adapter may optionally choose to use the transaction inflow contract. But, a resource adapter that uses the transaction inflow contract to import an EIS transaction and do transactional work must implement the prescribed transaction inflow contract.
- The `XATerminator` instance provided by the application server must be thread-safe and re-entrant. The resource adapter may use an `XATerminator` instance across different transactions concurrently.
- When the application server is unable to recreate the transaction context, if any, specified for a `Work` instance, it must throw a `WorkCompletedException` set to the error code `WorkException.TX_RECREATE_FAILED`.
- For a particular imported transaction, at any given time, there must be at most one `Work` instance associated with the transaction. The associated `Work` instance may be in any state, that is, waiting for execution to begin or already executing. However, it must be possible for several `Work` instances to do work on a transaction as long as there is at most one `Work` instance associated with the transaction at any time. It must also be possible for different resource adapters to participate in the same transaction.

The application server must disallow `Work` submissions with a `WorkCompletedException` set to the error code `WorkException.TX_CONCURRENT_WORK_DISALLOWED`, if there is already a `Work` instance associated with the transaction, irrespective of which resource adapter is involved in the `Work` submission. This must be done using the `getGlobalTransactionId` method of the `Xid` object present in the execution context of the submitted `Work` instance. The `Xid`'s branch identifier must be ignored.

The application server must not try to serialize `Work` processing based on transaction information.

- The application server must reject `Work` submissions for a transaction whose completion is in-progress, with a `WorkCompletedException` set to the error code `WorkException.TX_CONCURRENT_WORK_DISALLOWED`.
- The application server must reject transaction completion or crash recovery calls for a specific transaction with a `javax.transaction.xa.XAException`, when a `Work` instance associated with the transaction is present. The application server must not block or serialize transaction completion or crash recovery calls waiting for a `Work` instance associated with the transaction to complete.
- The application server must reject multiple transaction completion or crash recovery calls for the same transaction with a `javax.transaction.xa.XAException`.

- The application server must reject transaction completion or crash recovery calls with a `javax.transaction.xa.XAException` upon any errors.
- The application server should recover the state of all in-doubt transactions upon failure recovery.

14.4.5 Non-Requirements

- The application server is not responsible for ensuring transaction ids of the imported transactions from different EISs are unique. Each EIS is expected to use unique transaction ids.
- It is possible that a rogue resource adapter or EIS may provide non-unique `xids`, or attempt to complete transactions that it does not own. The application server is not required to detect the above cases. It is also not required to detect transactional, transaction completion, or crash recovery calls from a rogue EIS.

14.4.6 Recommendations

- The resource adapter should keep a list of active transactions and abort them upon detecting EIS failure.

14.4.7 Transaction Inflow in a Non-managed Environment

Though the transaction inflow contract is primarily intended for a managed environment, it may be used in a non-managed environment provided the application that bootstraps a resource adapter instance is capable of functioning as a resource manager.

In a non-managed environment, support for the transaction inflow contract is not required. That is, the `getXATerminator` method of the `BootstrapContext` instance may return a null instance.

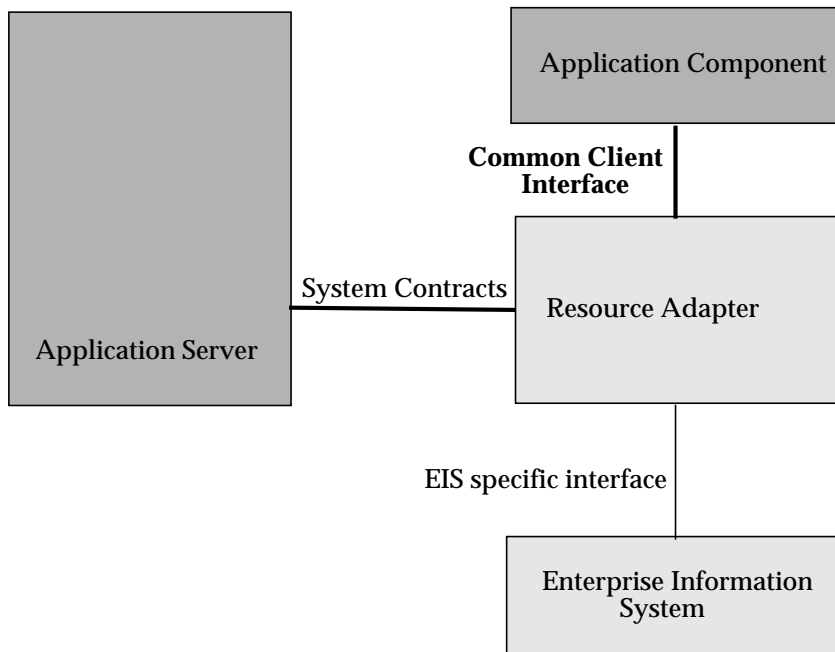
Common Client Interface

The following chapter specifies the Common Client Interface (CCI).

15.1 Overview

The CCI defines a standard client API for application components. The CCI enables application components and Enterprise Application Integration (EAI) frameworks to drive interactions across heterogeneous EISs using a common client API. FIGURE 15-1 shows a high-level view of the CCI and its relationship to other application components.

FIGURE 15-1 Common Client Interface



15.2 Goals

The CCI is designed with the following goals:

- It defines a remote function-call interface that focuses on executing functions on an EIS and retrieving the results. The CCI can form a base level API for EIS access on which higher level functionality can be built.
- It is targeted primarily towards application development tools and EAI frameworks.

- Although it is simple, it has sufficient functionality and an extensible application programming model.
- It provides an API that both leverages and is consistent with various facilities defined by the J2SE and J2EE platforms.
- It is independent of a specific EIS. For example, it does not use data types specific to an EIS. However, the CCI can be capable of being driven by EIS-specific metadata from a repository.

An important goal for the CCI is to complement existing standard JDBC API and not to replace this API. The CCI defines a common client API that is parallel to the JDBC for EISs that are not relational databases.

Since the CCI is targeted primarily towards application development tools and EAI vendors, it is not intended to discourage the use of JDBC APIs by these vendors. For example, an EAI vendor will typically combine JDBC with CCI by using the JDBC API to access relational databases and using CCI to access other EISs.

15.3 Scenarios

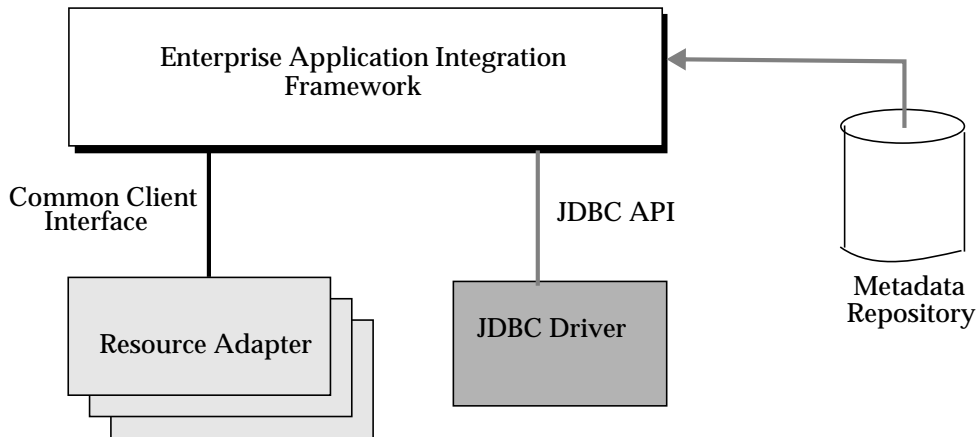
The following scenarios illustrate the use of CCI by enterprise tools and Enterprise Application Integration (EAI) vendors:

15.3.1 Enterprise Application Integration Framework

The EAI vendor uses the Common Client Interface as a standard way to plug-in resource adapters for heterogeneous EISs. The vendor provides an application integration framework on top of the functionality provided by the resource adapters. The framework uses the standard CCI interfaces to drive interactions with the connected EISs.

FIGURE 15-2 also shows the use of JDBC by the EAI framework for connecting to and accessing relational databases.

FIGURE 15-2 Scenario: EAI Framework



15.3.2 Metadata Repository and API

An EAI or application development tool uses a metadata repository to drive CCI-based interactions with heterogeneous EISs. See FIGURE 15-2 and FIGURE 15-3 for illustrative examples. A repository may maintain meta information about functions, with type mapping information and data structures for the invocation parameters, existing on an EIS system.

Note – The specification of a standard repository API and metadata format is outside the scope of the current version of the connector architecture.

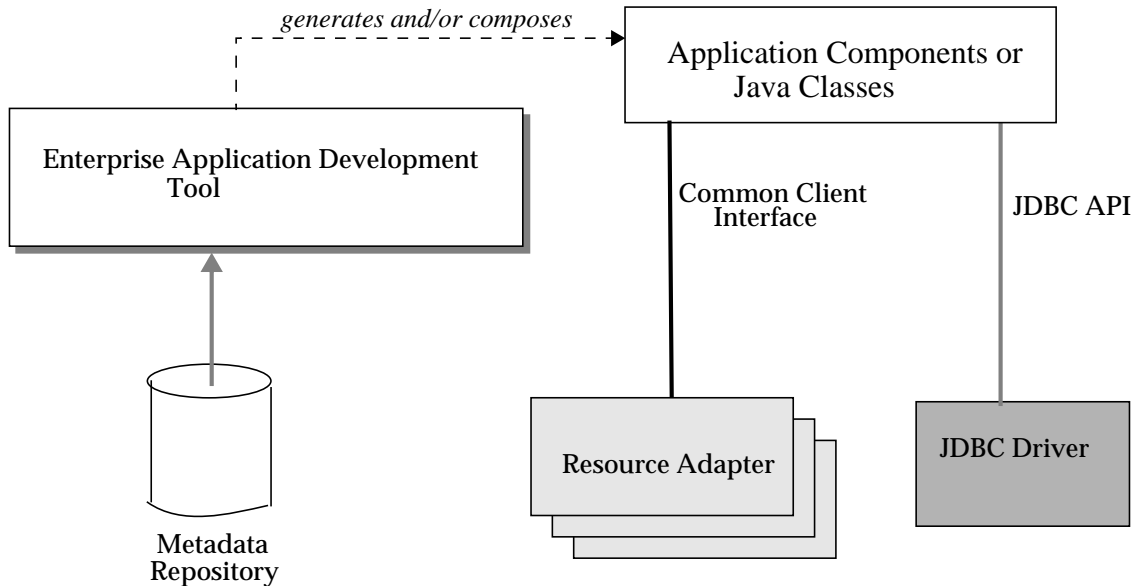
15.3.3 Enterprise Application Development Tool

The CCI functions as a plug-in contract for an application development tool that develops additional functionality around a resource adapter.

The application development tool generates Java classes based on the meta information accessed from a metadata repository. These Java classes encapsulate CCI-based interactions and expose a simple application programming model, typically based on the JavaBeans framework, to the application developers. An application component uses the generated Java classes for EIS access.

An application development tool can also compose or generate an application component that uses the generated Java classes for EIS access. See FIGURE 15-3.

FIGURE 15-3 Scenario: Enterprise Application Development Tool



15.4 Common Client Interface

The CCI is divided in to the following parts:

- Connection-related interfaces that represent a connection factory and an application level connection:
 - `javax.resource.cci.ConnectionFactory`
 - `javax.resource.cci.Connection`
 - `javax.resource.cci.ConnectionSpec`
 - `javax.resource.cci.LocalTransaction`
- Interaction-related interfaces that enable a component to drive an interaction, specified through an `InteractionSpec`, with an EIS instance:
 - `javax.resource.cci.Interaction`
 - `javax.resource.cci.InteractionSpec`
- Service endpoint message listener interface:
 - `javax.resource.cci.MessageListener`

- Data representation-related interfaces that are used to represent data structures involved in an interaction with an EIS instance:
 - `javax.resource.cci.Record`
 - `javax.resource.cci.MappedRecord`
 - `javax.resource.cci.IndexedRecord`
 - `javax.resource.cci.RecordFactory`
 - `javax.resource.cci.Streamable`
 - `javax.resource.cci.ResultSet`
 - `java.sql.ResultSetMetaData`
- Metadata related-interfaces that provide basic meta information about a resource adapter implementation and an EIS connection:
 - `javax.resource.cci.ConnectionMetaData`
 - `javax.resource.cci.ResourceAdapterMetaData`
 - `javax.resource.cci.ResultSetInfo`
- Additional classes:
 - `javax.resource.ResourceException`
 - `javax.resource.cci.ResourceWarning`

See FIGURE 15-4 on page 8 for the class diagram for CCI.

15.4.1 Requirements

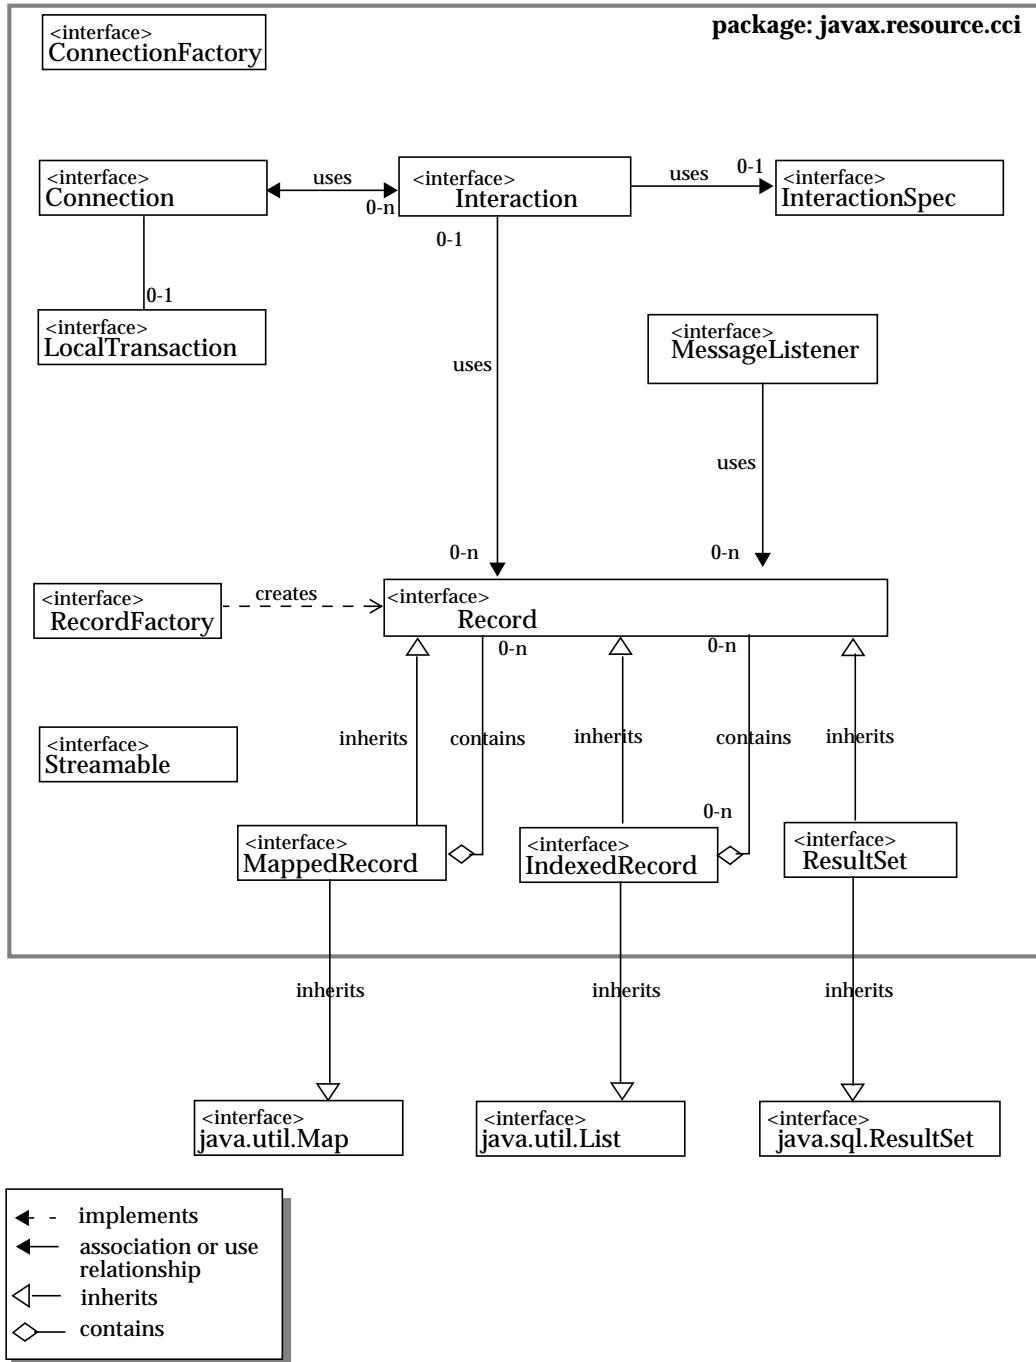
A resource adapter provider provides an implementation of the CCI interfaces as part of its resource adapter implementation. The connector architecture does not mandate that a resource adapter support the CCI interfaces as its client API.

Note – A resource adapter is allowed to support a client API specific to its underlying EIS. An example of an EIS-specific client APIs is JDBC API for relational databases.

The connector architecture also allows a third party vendor to provide an implementation of CCI interfaces above a resource adapter. For example, a base resource adapter supports the system contracts and provides an EIS specific client API. A third party tools vendor may provide the CCI implementation above this base resource adapter.

The connector architecture also allows a resource adapter implementation to support all interfaces except the data representation-related interfaces. In this case, a third party vendor provides both the development-time and runtime aspects of data structures required to drive interactions with an EIS instance. The section on the `Record` interface specification describes this case in more detail.

FIGURE 15-4 Class Diagram: Common Client Interface



15.5 Connection Interfaces

The following section specifies interfaces for the connection factory and application level connection.

15.5.1 ConnectionFactory

The `javax.resource.cci.ConnectionFactory` provides an interface for getting a connection to an EIS instance. A component looks up a `ConnectionFactory` instance from the JNDI namespace and then uses it to get a connection to the EIS instance.

The following code extract shows the `ConnectionFactory` interface:

```
public interface javax.resource.cci.ConnectionFactory
    extends java.io.Serializable,
    javax.resource.Referenceable {

    public RecordFactory getRecordFactory()
        throws ResourceException;

    public Connection getConnection()
        throws ResourceException;
    public Connection getConnection(
        javax.resource.cci.ConnectionSpec properties)
        throws ResourceException;

    public ResourceAdapterMetaData getMetaData()
        throws ResourceException;

}
```

The `getConnection` method gets a connection to an EIS instance. The `getConnection` variant with no parameters is used when a component requires the container to manage EIS sign-on. In this case of the container-managed sign-on, the component does not pass any security information.

A component may also use the `getConnection` variant with a `javax.resource.cci.ConnectionSpec` parameter, if it needs to pass any resource adapter specific security information and connection parameters. In the

component-managed sign-on case, an application component passes security information, such as username and password, through the `ConnectionSpec` instance.

It is important to note that the properties passed through the `getConnection` method should be client-specific, such as username, password, and language, and not be related to the configuration of a target EIS instance, such as port number or server name. The `ManagedConnectionFactory` instance is configured with a complete set of properties required for the creation of a connection to an EIS instance. Configured properties on a `ManagedConnectionFactory` can be overridden by client-specific properties passed by an application component through the `getConnection` method. Refer to Section 17.4.2, “`ManagedConnectionFactory`” on page 17-11 for configuration of a `ManagedConnectionFactory`.

Note that in a managed environment, the `getConnection` method with no parameters is the recommended model for getting a connection. The container manages the EIS sign-on in this case.

The `ConnectionFactory` interface also provides a method to get a `RecordFactory` instance. The `ConnectionFactory` implementation class may throw a `javax.resource.NotSupportedException` from the method `getRecordFactory`.

15.5.1.1 Requirements

An implementation class for `ConnectionFactory` must implement the `java.io.Serializable` interface to support JNDI registration. A `ConnectionFactory` implementation class is also required to implement `javax.resource.Referenceable`. Note that the `javax.resource.Referenceable` interface extends the `javax.naming.Referenceable` interface. Refer to Section 17.5, “JNDI Configuration and Lookup” on page 17-13 for more details on JNDI based requirements for the `ConnectionFactory` implementation.

An implementation class for `ConnectionFactory` must provide a default constructor.

15.5.2 ConnectionSpec

The interface `javax.resource.cci.ConnectionSpec` is used by an application component to pass connection request-specific properties to the `getConnection` method.

The `ConnectionSpec` interface has been introduced to increase the toolability of the CCI. The `ConnectionSpec` interface must be implemented as a `JavaBean`. Refer to Section 16.3, “`JavaBean Requirements`” on page 16-2.

The following code extract shows the `ConnectionSpec` interface.

```
public interface javax.resource.cci.ConnectionSpec {  
}
```

The CCI specification defines a set of standard properties for a `ConnectionSpec`. The properties are defined either on a derived interface or an implementation class of an empty `ConnectionSpec` interface. In addition, a resource adapter may define additional properties specific to its underlying EIS.

The following standard properties are defined by the CCI specification for `ConnectionSpec`:

TABLE 15-1 Standard Properties for `ConnectionSpec`

Property	Description
UserName	The name of the user establishing a connection to an EIS instance.
Password	The password for the user establishing a connection.

An important point to note is about the relationship between `ConnectionSpec` and `ConnectionRequestInfo`. The `ConnectionSpec` is used at the application level and is defined under the scope of CCI while `ConnectionRequestInfo` is defined as part of the system contracts. Separate interfaces have been defined to ensure the separation between CCI interfaces and system contracts. `ConnectionRequestInfo` has no explicit dependency on CCI. Note that a resource adapter may not implement CCI but it must implement system contracts. The specification of a standard repository API and metadata format is outside the scope of the current version of the connector architecture. The mapping between CCI's `ConnectionSpec` and `ConnectionRequestInfo` is achieved in an implementation-specific manner by a resource adapter.

15.5.3 Connection

A `javax.resource.cci.Connection` represents an application level connection handle that is used by a component to access an EIS instance. The actual physical connection associated with a `Connection` instance is represented by a `ManagedConnection`.

A component gets a `Connection` instance by using the `getConnection` method of a `ConnectionFactory` instance. A `Connection` instance may be associated with zero or more `Interaction` instances.

The following code extract shows the `Connection` interface:

```
public interface javax.resource.cci.Connection {
    public Interaction createInteraction() throws
        ResourceException;

    public ConnectionMetaData getMetaData() throws
        ResourceException;
    public ResultSetInfo getResultSetInfo() throws
        ResourceException;

    public LocalTransaction getLocalTransaction()
        throws ResourceException;

    public void close() throws ResourceException;
}
```

The `createInteraction` method creates an `Interaction` instance associated with the `Connection` instance. An `Interaction` enables a component to access EIS data and functions.

The `getMetaData` method returns information about the EIS instance associated with a `Connection` instance. The EIS instance-specific information is represented by the `ConnectionMetaData` interface.

The `getResultSetInfo` method returns information on the result set functionality supported by the connected EIS instance. If the CCI implementation does not support result set functionality, then the method `getResultSetInfo` must throw a `NotSupportedException`.

The `close` method initiates a close of the connection. The OID in FIGURE 6-6 on page 39 describes the resulting behavior of such an application level connection close.

The `getLocalTransaction` method returns a `LocalTransaction` instance that enables a component to demarcate resource manager local transactions. If a resource adapter does not allow a component to demarcate local transactions using the `LocalTransaction` interface, the `getLocalTransaction` method must throw a `NotSupportedException`.

15.5.3.1 Auto Commit

When a `Connection` is in an auto-commit mode, an `Interaction`, associated with the `Connection`, automatically commits after it has been executed. The auto-commit mode must be turned off if multiple interactions have to be grouped in a single transaction and committed or rolled back as a unit.

CCI does not provide explicit `set/getAutoCommit` methods in the `Connection` interface. This simplifies the application programming model for the transaction management.

A resource adapter must manage the auto-commit mode as follows:

- A transactional resource adapter either at the `XATransaction` or `LocalTransaction` level must set the auto-commit mode of `Connection` instances participating in a transaction to off within the transaction. This requirement holds for true both container-managed and bean-managed transaction demarcation.
- A transactional resource adapter must set the auto-commit mode of `Connection` instances to on when used outside a transaction.

These requirements are independent of whether a transaction is managed as a local or XA transaction. A transactional resource adapter should implement this requirement in an implementation-specific manner.

A non-transactional resource adapter at the `NoTransaction` level, is not required to support the auto-commit mode for `Connection`.

15.6 Interaction Interfaces

The following section specifies interfaces that enable a component to drive an interaction with an EIS instance and to demarcate resource manager local transactions.

15.6.1 Interaction

The `javax.resource.cci.Interaction` enables a component to execute EIS functions. An `Interaction` instance supports the following interactions with an EIS instance:

- An `execute` method that takes an input `Record`, output `Record`, and an `InteractionSpec`. This method executes the EIS function represented by the `InteractionSpec` and updates the output `Record`.

- An `execute` method that takes an input `Record` and an `InteractionSpec`. This method implementation executes the EIS function represented by the `InteractionSpec` and produces the output `Record` as a return value.

If an `Interaction` implementation does not support a variant of the `execute` method, the method must throw a `javax.resource.NotSupportedException`.

Refer to Section 15.10.2 “Interaction and Record” for details on how input and output records are created and used in the above variants of the `execute` method.

An `Interaction` instance is created from a `Connection` and must maintain its association with the `Connection` instance. The `close` method releases all resources maintained by the resource adapter for the `Interaction`. The `close` of an `Interaction` instance should not close the associated `Connection` instance.

The following code extract shows the `Interaction` interface:

```
public interface javax.resource.cci.Interaction {

    public Connection getConnection();

    public void close() throws ResourceException;

    public boolean execute(InteractionSpec ispec,
                          Record input,
                          Record output) throws ResourceException;

    public Record execute(InteractionSpec ispec,
                          Record input) throws ResourceException;

    ...

}
```

15.6.2 InteractionSpec

A `javax.resource.cci.InteractionSpec` holds properties for driving an `Interaction` with an EIS instance. An `InteractionSpec` uses an `Interaction` to execute the specified function on an underlying EIS.

The CCI specification defines a set of standard properties for an `InteractionSpec`. The properties are defined either on a derived interface or an implementation class of an empty `InteractionSpec` interface. The following code extract shows the `InteractionSpec` interface.

```
public interface javax.resource.cci.InteractionSpec
    extends java.io.Serializable {

    // Standard Interaction Verbs
    public static final int SYNC_SEND = 0;
    public static final int SYNC_SEND_RECEIVE = 1;
    public static final int SYNC_RECEIVE = 2;

}
```

An `InteractionSpec` implementation is not required to support a standard property if that property does not apply to its underlying EIS. The `InteractionSpec` implementation class must provide getter and setter methods for each of its supported properties. The getter and setter methods convention should be based on the JavaBeans design pattern.

15.6.2.1 Standard Properties

The standard properties are as follows:

- **FunctionName:** A string representing the name of an EIS function. Some examples are the name of a transaction program in a CICS system or the name of a business object or function module in an ERP system. The format of the name is specific to an EIS and is outside the scope of the CCI specification.
- **InteractionVerb:** An integer representing the mode of interaction with an EIS instance as specified by the `InteractionSpec`. The values of the interaction verb may be one of the following:
 - **SYNC_SEND:** The execution of an `Interaction` does only a send to the target EIS instance. The input record is sent to the EIS instance without any synchronous response in terms of an output `Record` or `ResultSet`.
 - **SYNC_SEND_RECEIVE:** The execution of an `Interaction` sends a request to the EIS instance and receives a response synchronously. The input record is sent to the EIS instance with the output received either as `Record` or a `ResultSet`.
 - **SYNC_RECEIVE:** The execution of an `Interaction` results in a synchronous receive of an output `Record`. For instance, a session bean gets a method invocation and it uses this `SYNC_RECEIVE` form of interaction to retrieve messages that have been delivered to a message queue.

The default `InteractionVerb` property is `SYNC_SEND_RECEIVE`.

If the `InteractionVerb` property is not defined for an `InteractionSpec`, the default mode for an interaction is `SYNC_SEND_RECEIVE`.

Other forms of interaction verbs are outside the scope of the CCI specification.

The CCI does not support asynchronous delivery of messages to the component instances. The message inflow contract should be used for asynchronous delivery of messages.

- `ExecutionTimeout`: An integer representing the number of milliseconds an `Interaction` waits for an `EIS` to execute the specified function.

15.6.2.2 ResultSet Properties

The following standard properties give hints to an `Interaction` instance about the `ResultSet` requirements:

- `FetchSize`: An integer representing the number of rows that should be fetched from an `EIS` when more rows are needed for a result set. If the value is zero, the hint is ignored. The default value is zero.
- `FetchDirection`: An integer representing the direction in which the rows in a result set are processed. The valid integer values are defined in the `java.sql.ResultSet` interface. The default value is `ResultSet.FETCH_FORWARD`.
- `MaxFieldSize`: An integer representing the maximum number of bytes allowed for any value in a column of a result set or a value in a `Record`.
- `ResultSetType`: An integer representing the type of the result set produced by an execution of the `InteractionSpec`. The `java.sql.ResultSet` interface defines the result set types.
- `ResultSetConcurrency`: An integer representing the concurrency type of the result set produced by the execution of the `InteractionSpec`. The `java.sql.ResultSet` interface defines the concurrency types for a result set.

Note that if a CCI implementation cannot support the specified requirements for a result set, it should choose an appropriate alternative and raise a `SQLWarning` from the `ResultSet` methods to indicate this condition. Refer to Section 15.11 “`ResultSet`” for more details.

A component can determine the actual scrolling ability and concurrency type of a result set by invoking the `getType` and `getConcurrencyType` methods of the `ResultSet` interface.

15.6.2.3 Additional Properties

An `InteractionSpec` implementation may define additional properties besides the standard properties. Note that the format and type of the additional properties is specific to an `EIS` and is outside the scope of the CCI specification.

15.6.2.4 Implementation

The `InteractionSpec` interface must be implemented as a `JavaBean` to support tools. The properties on the `InteractionSpec` implementation class must be defined through the getter and setter methods design pattern.

The CCI implementation may, though is not required to, provide a `BeanInfo` class for the `InteractionSpec` implementation. This class provides explicit information about the properties supported by the `InteractionSpec`.

An implementation class for the `InteractionSpec` interface must implement the `java.io.Serializable` interface.

The specified properties must be implemented as either bound or constrained properties. Refer to the `JavaBeans` specification (refer to <http://java.sun.com/products/javabeans>) for details on bound and constrained properties.

15.6.2.5 Administered Object

An `InteractionSpec` instance may be, though it is not required to be, registered as an administered object in the JNDI namespace. This enables a component provider to access `InteractionSpec` instances using logical names, called resource environment references. Resource environment references are special entries in the component's environment. The deployer binds a resource environment reference to an `InteractionSpec` administered object in the operational environment.

The EJB specification (Section 1., "Enterprise JavaBeans (EJBTM) specification, version 2.1:" on page F-1) specifies resource environment references in more detail.

15.6.2.6 Illustrative Scenario

The development tool introspects the `InteractionSpec` implementation class and shows a property sheet with all the configurable properties. The developer then configures the properties for an `InteractionSpec` instance.

At runtime, the configured `InteractionSpec` instance is used to specify properties for the execution of an `Interaction`. The runtime environment may lookup an `InteractionSpec` instance using a logical name from the JNDI namespace.

15.6.3 LocalTransaction

The `javax.resource.cci.LocalTransaction` defines a transaction demarcation interface for resource manager local transactions. An application component uses the `LocalTransaction` interface to demarcate local transactions. Refer to Chapter 7, “Transaction Management” for more details on local transactions.

Note that this interface is used for local transaction demarcation at the application level, while the `javax.resource.spi.LocalTransaction` interface is defined as part of the system contracts and is used by a container for local transaction management.

The following code extract shows the `LocalTransaction` interface:

```
public interface javax.resource.cci.LocalTransaction {
    public void begin() throws ResourceException;
    public void commit() throws ResourceException;
    public void rollback() throws ResourceException;
}
```

15.6.3.1 Requirements

A CCI implementation may, though is not required to, implement the `LocalTransaction` interface.

If the `LocalTransaction` interface is supported by a CCI implementation, the `Connection.getLocalTransaction` method must return a `LocalTransaction` instance. A component may then use the returned `LocalTransaction` to demarcate a resource manager local transaction on the underlying EIS instance.

A resource adapter is allowed to implement the `javax.resource.spi.LocalTransaction` interface without implementing the application-level `javax.resource.cci.LocalTransaction` interface. In this case, a container uses the system contract-level `LocalTransaction` interface for managing local transactions. Refer to Section 7.7, “Local Transaction Management Contract” on page 7-31 for more details on local transaction management.

15.7 Basic Metadata Interfaces

The following section specifies the interfaces that provide basic meta information about a resource adapter implementation and an EIS connection.

15.7.1 ConnectionMetaData

The `javax.resource.cci.ConnectionMetaData` interface provides information about an EIS instance connected through a `Connection` instance. A component calls the `Connection.getMetaData` method to get a `ConnectionMetaData` instance.

The following code extract shows the `ConnectionMetaData` interface:

```
public interface javax.resource.cci.ConnectionMetaData {
    public String getEISProductName() throws ResourceException;
    public String getEISProductVersion() throws
        ResourceException;
    public String getUserName() throws ResourceException;
}
```

The `getEISProductName` and `getEISProductVersion` methods return information about the EIS instance.

The `getUserName` method returns the user name for an active connection as known to the underlying EIS instance. The name corresponds the resource principal under whose security context a connection to the EIS instance has been established.

15.7.1.1 Implementation

A CCI implementation must provide an implementation class for the `ConnectionMetaData` interface.

A resource adapter provider or third-party vendor may extend the `ConnectionMetaData` interface to provide additional information. Note that the format and type of the additional information is specific to an EIS and is outside the scope of the CCI specification.

15.7.2 ResourceAdapterMetaData

The `javax.resource.cci.ResourceAdapterMetaData` interface provides information about the capabilities of a resource adapter implementation. Note that this interface does not provide information about an EIS instance that is connected through a resource adapter.

A component uses the `ConnectionFactory.getMetaData` method to get metadata information about a resource adapter. The `getMetaData` method does not require that an active connection to an EIS instance be established.

The following code extract shows the `ResourceAdapterMetaData` interface:

```
public interface javax.resource.cci.ResourceAdapterMetaData {
    public String getAdapterVersion();
    public String getAdapterVendorName();
    public String getAdapterName();
    public String getAdapterShortDescription();

    public String getSpecVersion();

    public String[] getInteractionSpecsSupported();
    public boolean supportsExecuteWithInputAndOutputRecord();
    public boolean supportsExecuteWithInputRecordOnly();

    public boolean supportsLocalTransactionDemarcation();
}
```

The `getSpecVersion` method returns a string representation of the version of the connector architecture specification that is supported by the resource adapter.

The `getInteractionSpecsSupported` method returns an array of fully-qualified names of `InteractionSpec` types supported by the CCI implementation for this resource adapter. Note that the fully-qualified class name is for the implementation class of an `InteractionSpec`. This method may be used by tools vendors to find information on the supported `InteractionSpec` types. The method should return an array of length 0 if the CCI implementation does not define specific `InteractionSpec` types.

The `supportsExecuteWithInputAndOutputRecord` and `supportsExecuteWithInputRecordOnly` methods are used by tools vendors to find information about the `Interaction` implementation. It is important to note that the `Interaction` implementation must support at least one variant of the execute methods.

The `supportsExecuteWithInputAndOutputRecord` method returns true if the implementation class for the `Interaction` interface implements the public `boolean execute(InteractionSpec ispec, Record input, Record output)` method. If not, the method returns false.

The `supportsExecuteWithInputRecordOnly` method returns true if the implementation class for the `Interaction` interface implements the public `Record execute(InteractionSpec ispec, Record input)` method. If not, the method returns false.

The `supportsLocalTransactionDemarcation` method returns true if the resource adapter implements the `LocalTransaction` interface and supports local transaction demarcation on the underlying EIS instance through the `LocalTransaction` interface.

The `ResourceAdapterMetaData` interface may be extended to provide more information specific to a resource adapter implementation.

15.8 Service Endpoint Message Listener Interface

The `MessageListener` interface serves as a request-response message listener type that message endpoints (refer to Chapter 12, “Message Inflow”) may implement. This allows an EIS to communicate with an endpoint using a request-response style.

```
interface javax.resource.cci.MessageListener {
    Record onMessage(Record inputData) throws ResourceException;
}
```

15.9 Exception Interfaces

The following section specifies `ResourceException` class defined by the CCI.

15.9.1 ResourceException

The `javax.resource.ResourceException` class is used as the root of the exception hierarchy for CCI. A `ResourceException` provides the following information:

- A resource adapter-specific string describing the error. This string is a standard Java exception message and is available through the `getMessage` method.
- A resource adapter-specific error code.
- A reference to another exception. A `ResourceException` is often the result of a lower level problem. If appropriate, this lower level exception, a `java.lang.Exception` or its derived exception type, can be linked to a `ResourceException` instance. Note, this has been deprecated in favor of the J2SE release 1.4 exception chaining facility.

A CCI implementation can extend the `ResourceException` interface to throw more specific exceptions. It may also chain instances of `java.lang.Exception` or its subtypes to a `ResourceException`.

15.9.2 ResourceWarning

The `javax.resource.cci.ResourceWarning` class provides information on the warnings related to interactions with EIS. A `ResourceWarning` is silently chained to an `Interaction` instance that has caused the warning to be reported.

The `Interaction.getWarnings` method enables a component to access the first `ResourceWarning` in a chain of warnings. Other `ResourceWarning` instances are chained to the first returned `ResourceWarning` instance.

15.10 Record

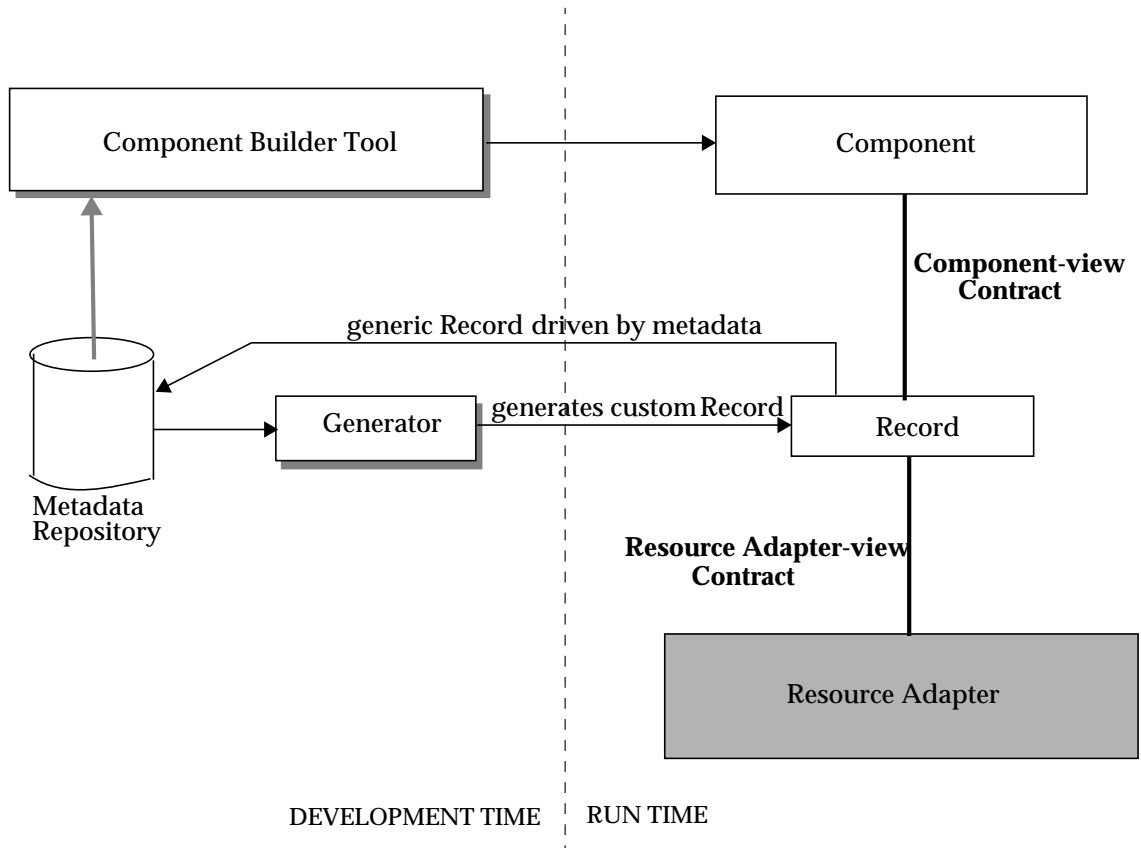
A `Record` is the Java representation of a data structure used as input or output to an EIS function.

A `Record` has both development-time and runtime aspects. See FIGURE 15-5 for an illustration of this. An implementation of a `Record` is either:

- A custom `Record` implementation that gets generated at the development time by a tool. The generation of a custom implementation is based on the meta information accessed by the tool from a metadata repository. The type mapping and data representation is generated as part of the custom `Record` implementation. So the custom `Record` implementation typically does not need to access the metadata repository at runtime.
- A generic `Record` implementation that uses a metadata repository at runtime for meta information. For example, a generic type of `Record` may access the type mapping information from the repository at runtime.

Note: The specification of a standard repository API and metadata format is outside the scope of the current version of the connector architecture.

FIGURE 15-5 Record at Development-time and Runtime



The meta information used in a `Record` representation and type mapping may be available in a metadata repository as:

- Meta information expressed in an EIS-specific format. For example, an ERP system has its own descriptive format for its meta information.
- Formatted in structures based on the programming language that has been used for writing the target EIS function, such as, COBOL structures used by CICS transaction programs.
- A standard representation of data structures as required for EIS functions. The standard representation is typically aggregated in a metadata repository based on the meta information extracted from multiple EISs.

A resource adapter may provide an implementation of all CCI interfaces except the data representation-related interfaces, namely, `Record` and `RecordFactory`. In this case, a third party vendor provides both development-time and runtime support for

the `Record` and `RecordFactory` interfaces. This requires that a `Record` implementation must support both component-view and resource adapter-view contracts, as specified in the following subsections.

15.10.1 Component-view Contract

The component-view contract provides a standard contract for using a `Record` for components and component building tools. A `Record` implementation must support the component-view contract.

The application programming model for a `Record` is as follows:

- A component creates an instance of a generated implementation class for a custom record. The implementation class represents an EIS-specific data structure.
- A component uses the `RecordFactory` interface to create an instance of the generic `Record` implementation class. The implementation class of a generic `Record` is independent of any EIS-specific data structure.

Note – A related CCI issue is the level of support in the CCI data representation interfaces (namely, `Record`, `MappedRecord`, and `IndexedRecord`) for the type mapping facility. The issue has to be addressed based on the following parameters:

- There is no standardized mapping across various type systems. For example, the existing type systems range from Java, CORBA, COM, COBOL and many more. It is difficult to standardize the type specification and mappings across such a diverse set of type systems within the connector architecture scope.
- Building a limited type mapping facility into the CCI data representation interfaces will constrain the use of CCI data representation interfaces across different types of EISs. For example, it may be difficult to support EISs that have complex structured types with a limited type mapping support.
- Building an extensive type mapping facility into the current version of CCI data representation interfaces will limit the future extensibility of these interfaces. This applies specifically to the support for standards that are emerging for XML-based data representation. An important goal for CCI data representation interfaces is to support XML-based facilities. This goal is difficult to achieve in the current scope of the connector architecture.

This specification proposes that the type mapping support for the CCI be kept open for future versions. A future version of this specification may standardize type mappings.

15.10.1.1 Type Mapping

Type mapping for EIS-specific types to Java types is not directly exposed to an application component. For example in the case of a custom `Record` implementation, the getter and setter methods, defined in a `Record` and exposed to an application component, return the correct Java types for the values extracted from the `Record`. The custom `Record` implementation internally handles all the type mapping.

In the case of a generic `Record` implementation, the type mapping is done in the generic `Record` by means of the type mapping information obtained from the metadata repository. Since the component uses generic methods on the `Record` interface, the component code does the required type casting.

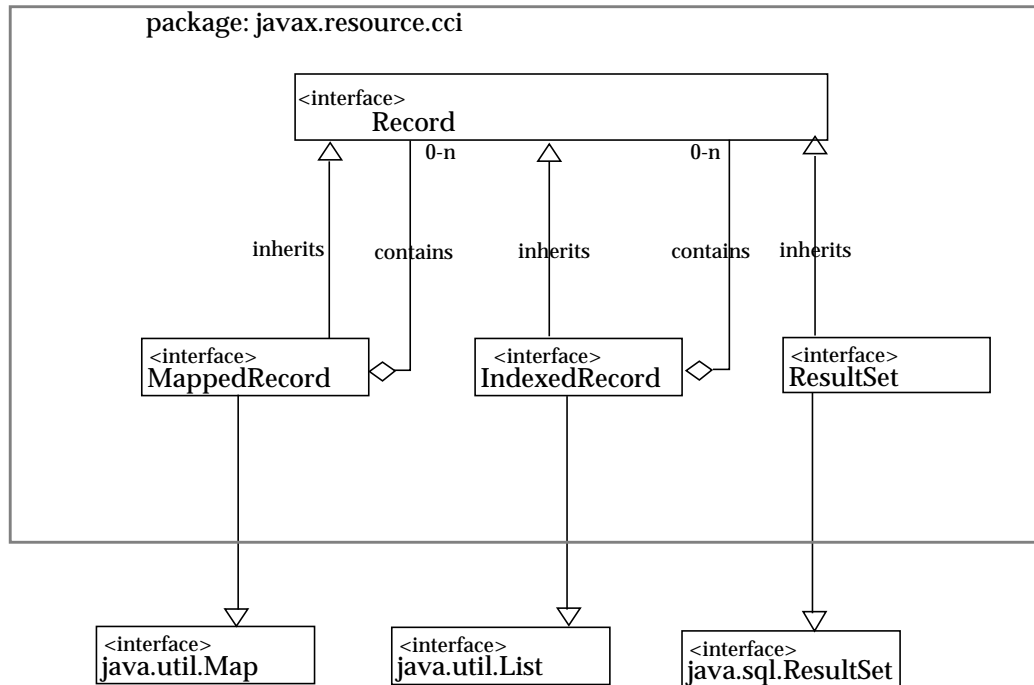
The compatibility of Java types and EIS types should be based on a type mapping that is defined specific to a class of EISs. For example, an ERP system from vendor X specifies a type mapping specific to its own EIS. Another example is type mapping between Java and COBOL types. Note that the JDBC specification specifies a standard mapping of SQL data types to Java types specific to relational databases.

In cases of both custom and generic `Records`, the type mapping information is provided by a metadata repository either at development-time or runtime.

15.10.1.2 Record Interface

The `javax.resource.cci.Record` interface is the base interface for the representation of a record. A `Record` instance is used as an input or output to the `execute` methods defined in an `Interaction`. See FIGURE 15-6.

FIGURE 15-6 Component-view Contract



The `Record` interface may be extended to form one of the following representations:

- `javax.resource.cci.MappedRecord`: A key-value pair based collection represents a record. This interface is based on `java.util.Map`.
- `javax.resource.cci.IndexedRecord`: An ordered and indexed collection represents a record. This interface is based on `java.util.List`.
- `javax.resource.cci.ResultSet`: This interface extends both `java.sql.ResultSet` and `javax.resource.cci.Record`. A result set represents tabular data. Section 15.11 “`ResultSet`” specifies the requirements for the `ResultSet` interface in detail.
- A `JavaBean` based representation of an EIS data structure: An example is a custom record generated to represent a purchase order in an ERP system or an invoice in a mainframe TP system.

Refer to Section 15.12 “Code Samples” for code samples that illustrate the use of `record`.

MappedRecord or IndexedRecord may contain another Record. This means that MappedRecord and IndexedRecord can be used to create a hierarchical structure of any arbitrary depth.

MappedRecord and IndexedRecord can be used to represent either a generic or custom record.

A basic Java type is used as the leaf element of a hierarchical structure represented by a MappedRecord or IndexedRecord.

A generated custom Record may also contain other records to form a hierarchical structure.

The following code extract shows the Record interface:

```
public interface javax.resource.cci.Record
    extends java.lang.Cloneable, java.io.Serializable {

    public String getRecordName();
    public void setRecordName(String name);

    public void setRecordShortDescription(String description);
    public String getRecordShortDescription();

    public boolean equals(Object other);
    public int hashCode();

    public Object clone() throws CloneNotSupportedException;
}
```

The Record interface defines the following set of standard properties:

- **Name of a Record:** Note that the CCI does not define a standard format for naming a Record. The name format is specific to an EIS type.
- **Description of a Record:** This property is used primarily by tools to show a description of a Record instance.

15.10.1.3 MappedRecord and IndexedRecord Interfaces

The `javax.resource.cci.MappedRecord` interface is used for representing a key-value map based collection of record elements. The `MappedRecord` interface extends both the `Record` and `java.util.Map` interface.

```
public interface javax.resource.cci.MappedRecord
    extends Record, java.util.Map,
    java.io.Serializable {
}
```

The `javax.resource.cci.IndexedRecord` interface represents an ordered collection of record elements based on the `java.util.List` interface. This interface allows a component to access record elements by their integer index, position in the list, and search for elements in the list.

```
public interface javax.resource.cci.IndexedRecord
    extends Record, java.util.List,
    java.io.Serializable {
}
```

15.10.1.4 RecordFactory

The `javax.resource.cci.RecordFactory` interface is used for creating `MappedRecord` and `IndexedRecord` instances. Note that the `RecordFactory` is only used for creating generic record instances. A CCI implementation provides an implementation class for the `RecordFactory` interface.

The following code extract shows the `RecordFactory` interface:

```
public interface javax.resource.cci.RecordFactory {
    public MappedRecord createMappedRecord(String recordName)
        throws ResourceException;

    public IndexedRecord createIndexedRecord(String
        recordName) throws ResourceException;
}
```


The methods `createMappedRecord` and `createIndexedRecord` take the name of the record that is to be created by the `RecordFactory`. The name of the record acts as a pointer to the meta information stored in the metadata repository for a specific record type. The format of the name is outside the scope of the CCI specification and specific to a CCI implementation and/or metadata repository.

A `RecordFactory` implementation should be capable of using the name of the desired `Record` and accessing meta information for the creation of the `Record`.

15.10.2 Interaction and Record

Records should be used as follows for the two variants of the `execute` method of the `Interaction` interface:

```
boolean execute(InteractionSpec, Record input, Record output)
```

- A custom record instance is used as an input or output to the `execute` method. A custom record implementation class is generated by an application development tool or EAI framework based on the meta information.
- The `RecordFactory` interface is used to create a generic `MappedRecord` or `IndexedRecord` instance. The generic record is used as input or output to the `execute` method.

```
Record execute(InteractionSpec, Record input)
```

- The input record can be either a custom or generic record.
- The returned record is a generic record instance created by the implementation of the `execute` method. The generic record instance may represent a `ResultSet` or a hierarchical structure as represented through the `MappedRecord` and `IndexedRecord` interfaces.

When the `Interaction.execute` method is called, a generic record instance may use the connection associated with the `Interaction` instance to access the metadata from the underlying EIS. If there is a separate metadata repository, then the generic record gets the metadata from the repository. The generic record implementation may use the above illustrative mechanism to achieve the necessary type mapping.

The generic record implementation encapsulates the above behavior and interacts with `Interaction` implementation in the `execute` method to get the active connection, if so needed. The contract between the generic record and `Interaction` implementation classes is specific to a CCI implementation.

15.10.3 Resource Adapter-view Contract

A resource adapter views the data represented by a `Record` either as:

- A stream of bytes through the `Streamable` interface, or,

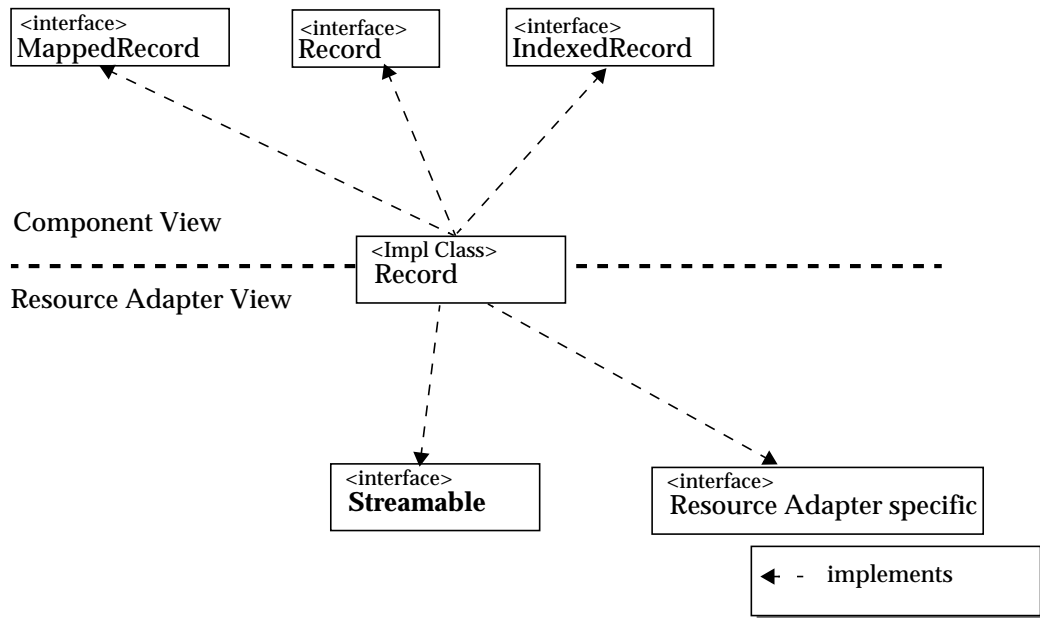
- A format specific to a resource adapter. For example, a resource adapter may extract or set the data for a `Record` using an interface defined specifically for the resource adapter.

A resource adapter-specific interface for viewing the `Record` representation is outside the scope of the CCI specification. A resource adapter must describe the resource adapter-specific interface to the users, typically tools vendors, of the resource adapter-view contract.

15.10.3.1 Streamable Interface

The `javax.resource.cci.Streamable` interface enables a resource adapter to extract data from an input `Record` or set data into an output `Record` as a stream of bytes. See FIGURE 15-7.

FIGURE 15-7 Streamable Interface



The `Streamable` interface provides a resource adapter's view of the data set in a `Record` instance by a component. A component uses `Record` or any derived interfaces to manage records.

A component does not directly use the `Streamable` interface. The interface is used by a resource adapter implementation.

The following code extract shows the `Streamable` interface:

```
public interface javax.resource.cci.Streamable {
    public void read(InputStream istream) throws IOException;
    public void write(OutputStream ostream) throws IOException;
}
```

The `read` extracts method data from an `InputStream` and initializes fields of a `Streamable` object. The `write` method writes fields of a `Streamable` object to an `OutputStream`. The implementations of both the `read` and `write` methods for a `Streamable` object must call the `read` and `write` methods respectively on the super class if there is one.

An implementation class of `Record` may choose to implement the `Streamable` interface or support a resource adapter-specific interface to manage record data.

15.11 ResultSet

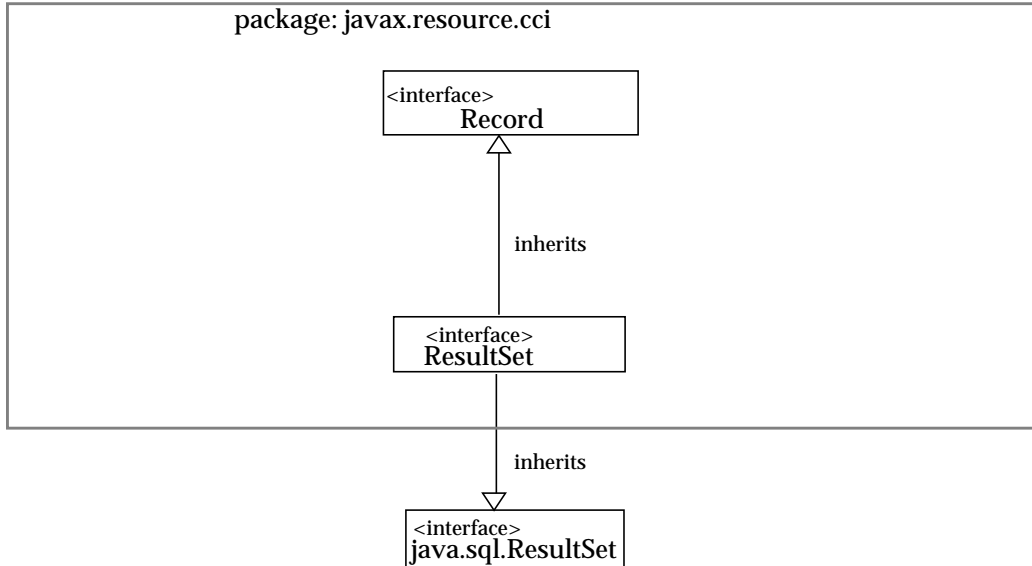
A result set represents tabular data that is retrieved from an EIS instance by the execution of an interaction. The `execute` method on the `Interaction` interface can return a `ResultSet` instance.

The CCI `ResultSet` interface is based on the JDBC `ResultSet` interface. The `ResultSet` extends the `java.sql.ResultSet` and `javax.resource.cci.Record` interfaces.

The following code extract shows the `ResultSet` interface:

```
public interface javax.resource.cci.ResultSet
    extends Record, java.sql.ResultSet {
}
```

FIGURE 15-8 `ResultSet` Interface



The following section specifies the requirements for a CCI `ResultSet` implementation.

Refer to the JDBC (Section 3., “JDBC API specification, version 3.0” on page F-1) specification and Java docs for more details on the `java.sql.ResultSet` interface. The following section specifies only a brief outline of the `ResultSet` interface. It focuses on the differences between the implementation requirements set by the CCI and JDBC. Note that the JDBC semantics for a `ResultSet` hold for the cases that are not explicitly mentioned in the following section.

CCI uses the JDBC `ResultSet` interface because:

- JDBC `ResultSet` is a standard, established, and well-documented interface for accessing and updating tabular data.
- JDBC `ResultSet` interface is defined in the core `java.sql` package. An introduction of an independent CCI-specific `ResultSet` interface (that is, different from the JDBC `ResultSet` interface) may create confusion in terms of differences in the programming model and functionality.

- The use of the JDBC `ResultSet` interface enables a tool or EAI vendor to leverage existing facilities that have been for the JDBC `ResultSet`.

Note – A CCI implementation is not required to support the `javax.resource.cci.ResultSet` interface. If a CCI implementation does not support result set functionality, it should not support interfaces and methods that are associated with the result set functionality. An example is the `java.sql.ResultSetMetaData` interface.

15.11.1 ResultSet Interface

The `ResultSet` interface provides a set of getter methods for retrieving column values from the current row. A column value can be retrieved using either the index number of the column or the name of the column. The columns are numbered starting at one. For maximum portability, result set columns within each row should be read left-to-right, and each column should be read only once.

The `ResultSet` interface also defines a set of `updateXXX` methods for updating the column values of the current row.

15.11.1.1 Type Mapping

A `ResultSet` implementation should attempt to convert the underlying EIS-specific data type to the Java type as specified in the `XXX` part of the `getXXX` method and return a suitable Java value.

A `ResultSet` implementation must establish a type mapping between the EIS specific data types and Java types. The type mapping is specific to an EIS.

The CCI specification does not specify standard type mappings specific to each type of EIS.

15.11.1.2 ResultSet Types

The CCI `ResultSet`, similar to the JDBC `ResultSet`, supports the following types of result set: `forward-only`, `scroll-insensitive`, and `scroll-sensitive`.

A forward-only result set is non-scrollable; its cursor moves only forward, from top to bottom. The view of data in the result set depends on whether the EIS instance materializes results incrementally.

A scroll-insensitive result set is scrollable; its cursor can move forward or backward and can be moved to a particular row or to a row whose position is relative to the current row. This type of result set is not sensitive to any changes made by another transaction or result sets in the same transaction that are made while the result set is open. This type of result set provides a static view of the underlying data with respect to changes made by other result sets. The order and values of rows are set at the time of the creation of a scroll-insensitive result set.

A scroll-sensitive result set is scrollable. It is sensitive to changes that are made while the result set is open. This type of result set provides a more dynamic view of the underlying data.

A component can use the `ownUpdatesAreVisible`, `ownDeletesAreVisible`, and `ownInsertsAreVisible` methods of the `ResultSetInfo` interface to determine whether a result set can see its own changes while the result set is open. For example, a result set's own changes are visible if the updated column values can be retrieved by calling the `getXXX` method after the corresponding `updateXXX` method. Refer to the JDBC (Section 3., "JDBC API specification, version 3.0" on page F-1) specification for more details on this feature.

15.11.1.3 Scrolling

The CCI `ResultSet` supports the same scrolling ability as the JDBC `ResultSet`.

If a resource adapter implements the cursor movement methods, its result sets are scrollable. A scrollable result set created by executing an `Interaction` can move through its contents in both a forward (first-to-last) or backward (last-to-first) direction. A scrollable result set also supports relative and absolute positioning.

The CCI `ResultSet`, similar to the JDBC `ResultSet`, maintains a cursor that indicates the row in the result set that is currently being accessed. The cursor maintained on a `forward-only` result set can only move forward through the contents of the result set. The rows are accessed in a first-to-last order. A scrollable result set can also be moved in a backward direction (last-to-first) and to a particular row.

Note that a CCI `ResultSet` implementation should only provide support for scrollable result sets if the underlying EIS supports such a facility.

15.11.1.4 Concurrency Types

A component can set the concurrency type of a CCI `ResultSet` to be either `read-only` or `updatable`. These types are consistent with the concurrency types defined by the JDBC `ResultSet`.

A result set that uses read-only concurrency does not allow updates of its content, while an updatable result set allows updates to its contents. An updatable result set may hold a write lock on the underlying data item and thus reduce concurrency.

Refer to the JDBC specification (Section 3., “JDBC API specification, version 3.0” on page F-1) for detailed information and examples.

15.11.1.5 Updatability

A result set of concurrency type `CONCUR_UPDATABLE` supports the ability to update, insert, and delete its rows. The CCI support for this type of result set is similar to the `JDBC ResultSet`.

The methods `updateXXX` on the `ResultSet` interface are used to modify the values of an individual column in the current row. These methods do not update the underlying EIS. The `updateRow` method must be called to update data on the underlying EIS. A resource adapter may discard changes made by a component if the component moves the cursor from the current row before calling the method `updateRow`.

Refer to the JDBC specification (Section 3., “JDBC API specification, version 3.0” on page F-1) for more information.

15.11.1.6 Persistence of Java Objects

The `ResultSet` interface provides the `getObject` method to enable a component to retrieve column values as Java objects. The type of the Java object returned from the `getObject` method is compatible with the type mapping supported by a resource adapter-specific to its underlying EIS. The `updateObject` method enables a component to update a column value using a Java object.

15.11.1.7 Support for SQL Types

It is optional for a CCI `ResultSet` to support the SQL type `JAVA_OBJECT` as defined in `java.sql.Types`. The JDBC specification specifies the JDBC support for persistence of Java objects.

The support for the following SQL types as defined in `java.sql.Types` is optional for a CCI `ResultSet` implementation:

- Binary large object (BLOB)
- Character large object (CLOB)
- SQL ARRAY type
- SQL REF type
- SQL DISTINCT type

- SQL STRUCT type

If an implementation of the `CCI ResultSet` interface does not support these types, it must throw a `java.sql.SQLException` indicating that the method is not supported, or `java.lang.UnsupportedOperationException` from the following methods:

- `getBlob`
- `getClob`
- `getArray`
- `getRef`

15.11.1.8 Support for Customized SQL Type Mapping

The CCI is not required to support customized mapping of SQL structured and distinct types to Java classes. The JDBC API defines support for such customization mechanisms.

The `CCI ResultSet` should throw a `java.sql.SQLException` indicating that the method is not supported or `java.lang.UnsupportedOperationException` from the `getObject` method that takes a `java.util.Map` parameter.

15.11.2 ResultSetMetaData

The `java.sql.ResultSetMetaData` interface provides information about the columns in a `ResultSet` instance. A component uses `ResultSet.getMetaData` method to get information about a `ResultSet`.

Refer to the JDBC Javadocs for a detailed specification of the `ResultSetMetaData` interface.

15.11.3 ResultSetInfo

The `javax.resource.cci.ResultSetInfo` interface provides information on the support provided for `ResultSet` functionality by a connected EIS instance. A component calls the `Connection.getResultInfo` method to get the `ResultSetInfo` instance.

A CCI implementation is not required to support the `javax.resource.cci.ResultSetInfo` interface. The implementation of this interface is provided only if the CCI supports the `ResultSet` facility.

The following code extract shows the `ResultSetInfo` interface:

```
public interface javax.resource.cci.ResultSetInfo {
    public boolean updatesAreDetected(int type)
        throws ResourceException;
    public boolean insertsAreDetected(int type)
        throws ResourceException;
    public boolean deletesAreDetected(int type)
        throws ResourceException;

    public boolean supportsResultSetType(int type)
        throws ResourceException;
    public boolean supportsResultSetConcurrency(int type,
        int concurrency)
        throws ResourceException;

    public boolean ownUpdatesAreVisible(int type)
        throws ResourceException;
    public boolean ownInsertsAreVisible(int type)
        throws ResourceException;
    public boolean ownDeletesAreVisible(int type)
        throws ResourceException;

    public boolean othersUpdatesAreVisible(int type)
        throws ResourceException;
    public boolean othersInsertsAreVisible(int type)
        throws ResourceException;
    public boolean othersDeletesAreVisible(int type)
        throws ResourceException;
}
```

The type parameter to the above methods represents the type of the `ResultSet`, defined as `TYPE_XXX` in the `ResultSet` interface.

Note that these methods should throw a `ResourceException` in the following cases:

- A resource adapter and the connected EIS instance cannot provide any meaningful values for these properties.
- The CCI implementation does not support the `ResultSet` functionality. In this case, a `NotSupportedException` should be thrown from invocations on the above methods.

A component uses the `rowUpdated`, `rowInserted`, and `rowDeleted` methods of the `ResultSet` interface to determine whether a row has been affected by a visible update, insert, or delete is the result set is open. The `updatesAreDetected`, `insertsAreDetected` and `deletesAreDetected` methods enable a component to find out whether or not changes to a `ResultSet` are detected.

A component uses the `ownUpdatesAreVisible`, `ownDeletesAreVisible` and `ownInsertsAreVisible` methods to determine whether a `ResultSet` can see its own changes when the result set is open.

A component uses the `supportsResultSetType` method to check the `ResultSet` types supported by a resource adapter and its underlying EIS instance.

The `supportsResultSetTypeConcurrency` method provides information on the `ResultSet` concurrency types supported by a resource adapter and its underlying EIS instance.

15.12 Code Samples

The following code extracts illustrate the application programming model based on the CCI.

An application development tool or EAI framework normally hides all the CCI-based programming details from an application developer. For example, an application development tool generates a set of Java classes that abstract the CCI-based application programming model and offers a simple programming model to an application developer.

15.12.1 Connection

- Get a `Connection` to an EIS instance after a lookup of a `ConnectionFactory` instance from the JNDI namespace. In this case, the component allows the container to manage the EIS sign-on.

```
javax.naming.Context nc = new InitialContext();
javax.resource.cci.ConnectionFactory cf =
    (ConnectionFactory)nc.lookup(
        "java:comp/env/eis/ConnectionFactory");
javax.resource.cci.Connection cx = cf.getConnection();
```

- Create an `Interaction` instance:

```
javax.resource.cci.Interaction ix = cx.createInteraction();
```

15.12.2 InteractionSpec

- Create a new instance of the respective `InteractionSpec` class or look up a pre-configured `InteractionSpec` in the runtime environment using JNDI.

```
com.wombat.cci.InteractionSpecImpl ixSpec = // ...

ixSpec.setFunctionName("<EIS_SPECIFIC_FUNCTION_NAME>");
ixSpec.setInteractionVerb(InteractionSpec.SYNC_SEND_RECEIVE);
...
```

15.12.3 Mapped Record

- Get a `RecordFactory` instance:

```
javax.resource.cci.RecordFactory rf = // ... get a RecordFactory
```

- Create a generic `MappedRecord` using the `RecordFactory` instance. This record instance acts as an input to the execution of an interaction. The name of the `Record` acts as a pointer to the meta information, stored in the metadata repository, for a specific record type.

```
javax.resource.cci.MappedRecord input =
    rf.createMappedRecord("<NAME_OF_RECORD>");
```

- Populate the generic `MappedRecord` instance with input values. The component code adds the values based on the meta information it has accessed from the metadata repository.

```
input.put("<key: element1>", new String("<VALUE>"));
input.put("<key: element2>", ...);
...
```

- Create a generic `IndexedRecord` to hold the output values that are set by the execution of the interaction.

```
javax.resource.cci.IndexedRecord output =
    rf.createIndexedRecord("<NAME_OF_RECORD>");
```

- Execute the Interaction:

```
boolean ret = ix.execute(ixSpec, input, output);
```

- Extract data from the output `IndexedRecord`. Note that the type mapping is done in the generic `IndexedRecord` by means of the type mapping information in the metadata repository. Since the component uses generic methods on the `IndexedRecord`, the component code does the required type casting.

```
java.util.Iterator iterator = output.iterator();
while (iterator.hasNext()) {
    // Get a record element and extract value
}
```

15.12.4 ResultSet

- Set the requirements for the `ResultSet` returned by the execution of an `Interaction`. This step is optional. Default values are used if the requirements are not explicitly set:

```
com.wombat.cci.InteractionSpecImpl ixSpec = .. // get an
                                              // InteractionSpec;

ixSpec.setFetchSize(20);
ixSpec.setResultSetType(ResultSet.TYPE_SCROLL_INSENSITIVE);
```

- Execute an `Interaction` that returns a `ResultSet`:

```
javax.resource.cci.ResultSet rs = (javax.resource.cci.ResultSet)
    ix.execute(ixSpec, input);
```

- Iterate over the `ResultSet`. The example here positions the cursor on the first row and then iterates forward through the contents of the `ResultSet`. The `getXXX` methods are used to retrieve column values:

```
rs.beforeFirst();
while (rs.next()) {
    // get the column values for the current row using getXXX
    // method
}
```

- The following example shows a backward iteration through the `ResultSet`:

```
rs.afterLast();
while (rs.previous()) {
    // get the column values for the current row using getXXX
    // method
}
```

15.12.5 Custom Record

- Extend the `Record` interface to represent an EIS-specific custom `Record`. The `CustomerRecord` interface supports a simple getter-setter design pattern for its field values. A development tool generates the implementation class of the `CustomerRecord`.

```
public interface CustomerRecord extends
    javax.resource.cci.Record,
        javax.resource.cci.Streamable {

    public void setName(String name);
    public void setId(String custId);
    public void setAddress(String address);

    public String getName();
    public String getId();
    public String getAddress();
}
```

- Create an empty `CustomerRecord` instance to hold output from the execution of an `Interaction`.

```
CustomerRecord customer = ... // create an instance
```

- Create a `PurchaseOrderRecord` instance as an input to the `Interaction` and set the properties on this instance. The `PurchaseOrderRecord` is another example of a custom `Record`.

```
PurchaseOrderRecord purchaseOrder = ... // create an instance
purchaseOrder.setProductName("...");
purchaseOrder.setQuantity("...");
...
```

- Execute an Interaction that populates the output CustomerRecord instance.

```
// Execute the Interaction
boolean ret = ix.execute(ixSpec, purchaseOrder, customer);

// Check the CustomerRecord
System.out.println( customer.getName() + ":" +
                    customer.getId() + ":" +
                    customer.getAddress());
```

API Requirements

This chapter specifies the API requirements for the resource adapter and application server implementations.

16.1 Requirements of the Application Server

- The application server must support the deployment of a resource adapter in EJB and Web containers.
- The application server must support all the connector architecture API requirements in EJB and Web containers.
- A single resource adapter instance may be shared by both a Web container and an EJB container.
- The application server must support all versions of the resource adapter DTDs (Document Type Definitions) and the resource adapter XML Schema Definition. This ensures that resource adapters written to previous versions of this specification can be deployed on products supporting the current version of this specification.

16.2 Requirements of the Resource adapter

The following matrix specifies the required (+) and optional (?) API requirements on a resource adapter.

LM - Lifecycle management contract

WM - Work management contract

MI - Message Inflow contract

TI - Transaction Inflow contract

CM - Connection management contract

TM- Transaction management contract

SM - Security management contract

CCI - Common Client Interface

TABLE 16-1 Resource Adapter API Requirements

	LM	WM	MI	TI	CM	TM	SM	CCI
Outbound	?	?			+	+	+	?
Inbound	+	?	+	?				
Bi-directional	+	?	+	?	+	+	+	?

Note, the message inflow contract must be supported by an inbound resource adapter.

16.3 JavaBean Requirements

The various JavaBean implementations provided by a resource adapter must adhere to the following rules:

- A JavaBean implementation must contain a null constructor.
- A JavaBean implementation must provide getter and setter methods, to access and modify the public properties of the JavaBean instance.

Note, for JavaBean serialization, implementing the `java.io.Serializable` interface is not necessary. The XML long-term persistence mechanism introduced in J2SE 1.4 can save the state of a JavaBean in an XML format that is resilient to version changes in the implementation of that JavaBean. Refer to J2SE 1.4 (Section 12., “Java 2 Platform, Standard Edition (J2SETM), API specification, version 1.4:” on page F-1) classes `java.beans.XMLEncoder`, `java.beans.XMLDecoder`, and `java.beans.PersistenceDelegate`.

For details, refer to JavaBeans specification (Section 14., “JavaBeansTM Specification 1.01 Final Release” on page F-2).

16.4 Equality Constraints

This section specifies the equality constraints on object implementations of the various types defined by this specification.

Equality based on Java object identity

The candidate objects are implementations of `MessageEndpointFactory`, `ActivationSpec`, `ManagedConnection` types.

These objects, in general, should not override the default `equals` and `hashCode` methods. However, if these methods are overridden, they must preserve the equality constraints based on Java object identity; that is, no two objects are considered equal.

Equality based on config properties and class information

The candidate objects are implementations of `ResourceAdapter`, `ManagedConnectionFactory`, `ConnectionRequestInfo`, `java.security.Principal`, `org.ietf.jgss.GSSCredential`, `GenericCredential`, `PasswordCredential`, and `Record` types.

These objects must override the default `equals` and `hashCode` methods, and provide an equality behavior based on the configuration properties and class information. That is, any two objects can be equal only if their configuration properties match and they have the same class implementation.

Packaging Requirements


This chapter specifies requirements for packaging and deploying a resource adapter. These requirements support a modular, portable deployment of a resource adapter into a J2EE compliant application server.

17.1 Overview

A resource adapter provider develops a set of Java interfaces and classes as part of its implementation of a resource adapter. These Java classes implement connector architecture-specified contracts and EIS-specific functionality provided by the resource adapter. The development of a resource adapter may also require the use of native libraries specific to the underlying EIS.

The Java interfaces and classes are packaged together (with required native libraries, help files, documentation, and other resources) with a deployment descriptor to create a `Resource Adapter Module`. A deployment descriptor defines the contract between a resource adapter provider and a deployer for the deployment of a resource adapter.

FIGURE 17-1 Packaging and Deployment Lifecycle of a Resource adapter



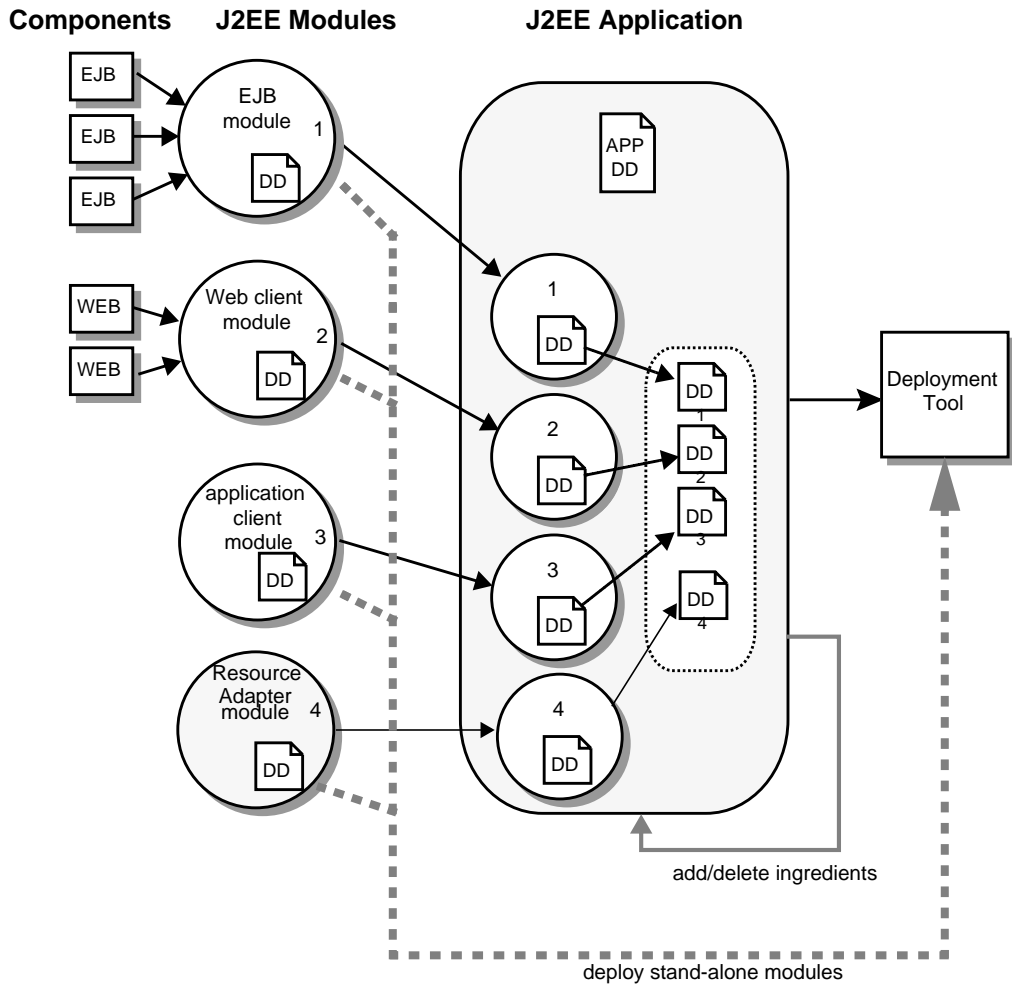
A resource adapter module corresponds to a J2EE module in terms of the J2EE composition hierarchy. Refer to the J2EE Platform specification (Section 8., “Java 2 Platform Enterprise Edition (J2EETM), Platform specification, version 1. 4:” on page F-1) for more details on the deployment of J2EE modules and applications. A J2EE module represents the basic unit of composition of a J2EE application. Examples of J2EE modules include EJB modules, application client modules, and web client modules.

A resource adapter module must be deployed either:

- Directly into an application server as a stand-alone unit or,
- Deployed with a J2EE application that consists of one or more J2EE modules in addition to a resource adapter module. The J2EE specification specifies requirements for the assembly and packaging of J2EE applications.

FIGURE 17-2 shows the composition model of a resource adapter module with other J2EE modules.

FIGURE 17-2 Deployment of a Resource Adapter Module



The stand-alone deployment of a resource adapter module into an application server is typically done to support scenarios in which multiple J2EE applications share a single resource adapter module. However, in certain scenarios, a resource adapter

module is required only by components within a single J2EE application. The deployment option of a resource adapter module bundled with a J2EE application supports the latter scenario.

At deployment time, a resource adapter deployer deploys a resource adapter module to an application server.

17.2 Packaging

The file format for a packaged resource adapter module defines the contract between a resource adapter provider and deployer.

A packaged resource adapter includes the following elements:

- Java classes and interfaces that are required for the implementation of both the connector architecture contracts and the functionality of the resource adapter.
- Utility Java classes for the resource adapter.
- Platform-dependent native libraries required by the resource adapter.
- Help files and documentation.
- Descriptive meta information that ties the above elements together.

17.2.0.1 Resource Adapter Archive

A resource adapter must be packaged using the `Java Archive (JAR)` format in to an `RAR (ResourceAdapter Archive)`. For example, a resource adapter for `EIS A` can be packaged as an archive with a filename `eisA.rar`.

The `RAR` file must contain a deployment descriptor based on the format specified in Section 17.5.4 “Requirements”. The deployment descriptor must be stored with the name `META-INF/ra.xml` in the `RAR` file.

The Java interfaces, implementation, and utility classes required by the resource adapter must be packaged as one or more `JAR` files as part of the resource adapter module. A `JAR` file must use the `.jar` file extension.

The platform-specific libraries required by the resource adapter must be packaged with the resource adapter module.

17.2.0.2 RAR Contents

The following describes the contents of a RAR file, their respective location within the RAR file and whether they are required.

TABLE 17-1 Description of RAR File Contents

Contents of RAR file	Requirements	Relative location within RAR file
Deployment descriptor	Required	META-INF/ra.xml
howto.html, image files, locale files, etc.	Optional	Arbitrary (that is, could be at root level or at a sub-level).
JAR files	Optional	Arbitrary
Platform-specific native libraries	Optional	Arbitrary

17.2.0.3 Sample Directory Structure

The following lists the files in a sample resource adapter module:

```
META-INF/ra.xml
howto.html
images/icon.jpg
ra.jar
cci.jar
win.dll
solaris.so
```

In the above example, `ra.xml` is the deployment descriptor. `ra.jar` and `cci.jar` contain Java interfaces and implementation classes for the resource adapter. `win.dll` and `solaris.so` are examples of native libraries.

Note that a resource adapter module can be structured such that various elements are partitioned using subdirectories.

17.2.0.4 Requirements

- When a standalone resource adapter RAR is deployed, the resource adapter must be made available to all J2EE applications in the application server.
- When a resource adapter RAR packaged within a J2EE application EAR is deployed, the resource adapter must be made available only to the J2EE application with which it is packaged.

17.3 Deployment

A deployment descriptor defines the contract between a resource adapter provider and a deployer. It captures the declarative information that is intended for the deployer to enable deployment of a resource adapter in a target operational environment.

A resource adapter module must be deployed based on the deployment requirements specified by the resource adapter provider in the deployment descriptor. Section 17.5.4 “Requirements” specifies the XML DTD for the deployment descriptor for a resource adapter module.

The J2EE Deployment API Specification (Section 13., “J2EETM Deployment API Specification” on page F-2) describes the general deployment procedure in detail.

17.3.1 Resource Adapter Provider

The resource adapter provider is responsible for specifying the deployment descriptor for a resource adapter.

The resource adapter provider must specify the following information in the deployment descriptor:

- **General information:** The resource adapter provider should specify the following general information about a resource adapter:
 - Name of the resource adapter.
 - Description of the resource adapter.
 - URI of a UI icon for the resource adapter.
 - Name of the vendor who provides the resource adapter.
 - Licensing requirement and description. Note that the management of licensing is outside the scope of the connector architecture.
 - Type of the EIS system supported. For example, the name of a specific database, ERP system, or mainframe TP system without any versioning information.
 - Version of the connector architecture specification, represented as a string, supported by the resource adapter.
 - Version of the resource adapter represented as a string.
- **ResourceAdapter class:** The resource adapter provider must specify the fully qualified name of a Java class that implements the `javax.resource.spi.ResourceAdapter` interface. The implementation of this class must be a `JavaBean`. A `ResourceAdapter` `JavaBean` is configured by the resource adapter deployer during deployment. The application server must instantiate exactly one `ResourceAdapter` `JavaBean` per functional resource

adapter instance. The application server must create at least one functional resource adapter instance per resource adapter deployment. The configuration properties are specific to a resource adapter.

- **ResourceAdapter class configuration properties:** The resource adapter provider may optionally provide a set of configuration properties for the `ResourceAdapter` instance, which may be used by the resource adapter deployer to configure a `ResourceAdapter` JavaBean instance.
- **Outbound resource adapter information**
 - **ManagedConnectionFactory class:** The resource adapter provider must specify the fully qualified name of the Java class that implements the `javax.resource.spi.ManagedConnectionFactory` interface. The implementation must be a JavaBean.

Typically, a `ManagedConnectionFactory` class is used to produce `ConnectionFactory` and `Connection` objects of a particular type. In order to produce objects of different types, a separate `ManagedConnectionFactory` class can be used for each supported type. The deployment descriptor element `connection-definition` can be used to specify different `ManagedConnectionFactory` classes, each pertaining to a particular type.

- **ConnectionFactory interface and implementation class:** The resource adapter provider must specify the fully-qualified name of the Java interface and implementation class for each connection factory supported by the resource adapter.
- **Connection interface and implementation class:** The resource adapter provider must specify the fully-qualified name of the Java interface and implementation class for each connection supported by the resource adapter.
- **Transactional Support:** The resource adapter provider must specify the level of transaction support provided by the resource adapter implementation. The level of transaction support must be any one of the following:
`NoTransaction`, `LocalTransaction`, or `XATransaction`. Note that this support is specified for a resource adapter and not for the underlying EIS instance.

NoTransaction: The resource adapter does not support either the resource manager local or JTA transactions. It does not implement either `XAResource` or `LocalTransaction` interfaces.

LocalTransaction: The resource adapter supports resource manager local transactions by implementing the `LocalTransaction` interface. The local transaction management contract is specified in Section 7.7, “Local Transaction Management Contract” on page 7-31.

XATransaction: The resource adapter supports both resource manager local and JTA transactions by implementing the `LocalTransaction` and `XAResource` interfaces respectively. The requirements for supporting the `XAResource` based contract are specified in Section 7.6, “XAResource-based Transaction Contract” on page 7-16.

- **Configurable properties per ManagedConnectionFactory instance:** The resource adapter provider specifies the name, type, description, and an optional default value for the properties that have to be configured on a per `ManagedConnectionFactory` instance.

Each `ManagedConnectionFactory` instance creates connections to a specific EIS instance based on the properties configured on the `ManagedConnectionFactory` instance. The configurable properties are specified only once in the deployment descriptor, even though a resource adapter can be used to configure multiple `ManagedConnectionFactory` instances that create connections to different instances of the same underlying EIS type.

- **Authentication Mechanism:** The resource adapter provider must specify all authentication mechanisms supported by the resource adapter. This includes the support provided by the resource adapter implementation but not by the underlying EIS instance. The standard values are: `BasicPassword` and `Kerbv5`. A resource adapter may support one or more of these authentication mechanisms.
 - `BasicPassword`: user-password based authentication mechanism that is specific to an EIS.
 - `Kerbv5`: Kerberos version 5 based authentication mechanism.

If no authentication mechanism is specified as part of the deployment descriptor, the resource adapter supports no standard security authentication mechanism as part of the security contract.

- **Reauthentication support:** The resource adapter provider must specify whether a resource adapter supports re-authentication of an existing physical connection.
- **Extended Security Permissions:** The security permissions listed in the deployment descriptor are different from those required by the default permission set. Refer to Chapter 18, “Runtime Environment” for more details on security permissions.
- **Inbound resource adapter information**
 - **Message listener type:** The resource adapter provider must specify one or more message listener types supported by a messaging resource adapter. The message listener type is the fully qualified name of the Java type of a message listener interface.
 - **ActivationSpec class:** The resource adapter provider must specify the fully qualified Java class name of the activation specification class. The implementation of this class must be a `JavaBean`. An `ActivationSpec` specifies an activation specification per message listener type. The `ActivationSpec` is configured by a message endpoint deployer during application deployment.

- **Required ActivationSpec properties:** The resource adapter provider may optionally specify a set of required properties for an `ActivationSpec`. This is useful in validating the `ActivationSpec` during endpoint application deployment.
- **Administered objects:** The resource adapter provider must specify the fully qualified name of the Java type of the interface implemented by an administered object, which must be a `JavaBean`, and its Java class name. Administered objects are specific to a messaging style or message provider. There may be zero to more administered objects specified.

The deployment descriptor specified by the resource adapter provider for its resource adapter must be consistent with the XML DTD specified in Section 17.5.4 “Requirements”.

Note – The connector architecture does not specify standard deployment properties for the configuration of non-Java parts, such as native libraries, of a resource adapter. This applies only to the properties of the non-Java part not exposed through the Java part of the resource adapter. The non-Java part of a resource adapter should be configured using mechanisms specific to a resource adapter.

17.3.2 Deployer

During resource adapter deployment, the deployer is responsible for configuring a resource adapter. The configuration of a resource adapter is based on the properties defined in the deployment descriptor as part of the resource adapter module.

17.3.2.1 Stand-Alone Resource Adapter Module

During deployment, the deployer configures and deploys a resource adapter based on the deployment descriptor information. The deployer may choose to override the information in the deployment descriptor.

17.3.2.2 Resource Adapter Module with J2EE Application

Refer to the J2EE platform specification (Section 8., “Java 2 Platform Enterprise Edition (J2EETM), Platform specification, version 1. 4:” on page F-1) for the requirements specified for the deployment of a J2EE application.

17.3.2.3 Configuration

The deployer must perform the following tasks to configure a resource adapter:

- Configure a `ResourceAdapter` JavaBean instance. The configuration properties are specific to a resource adapter.

In the case of outbound resource adapters,

- Configure one or more property sets (one property set per `ManagedConnectionFactory` instance) for creating connections to various underlying EIS instances. The deployer creates a property set to set valid values for various configurable fields. The configuration of each field is based on the name, type and description of the field specified in the deployment descriptor.

Each property set represents a specific configuration to be set on a `ManagedConnectionFactory` instance for creating connections to a specific EIS instance. Since a resource adapter may be used to create connections to multiple instances of the same EIS, there can be multiple property sets for a single resource adapter, one for each configured `ManagedConnectionFactory` instance.

- Configure application server mechanisms for transaction management based on the level of transaction support specified by the resource adapter.
- Configure security in the target operational environment based on the security requirements specified by the resource adapter in its deployment descriptor.

17.3.2.4 Security Configuration

The security configuration is based on:

- Whether or not the resource adapter supports a specific authentication mechanism and credentials interface. The deployment descriptor includes an element `authentication-mechanism` that specifies a supported authentication mechanism and the corresponding credentials interface.
- Whether or not the application server is configured to support a specific mechanism type. For example, if the application server is not configured for the Kerberos mechanism, it is not capable of passing Kerberos credentials to the resource adapter as part of the security contract.

During the deployment, the deployer may, though is not required to, check whether or not an underlying EIS supports the same capabilities, such as transaction support and authentication mechanisms, as the corresponding resource adapter.

For example, if a resource adapter provides implementation support for Kerberos based authentication but the underlying EIS instance does not support Kerberos, the deployer may decide not to configure Kerberos for authentication to this EIS instance. However if the deployer does not perform such checks during deployment, any invalid configurations should lead to runtime exceptions.

17.4 Interfaces/Classes

This section specifies the Java classes and interfaces related to the configuration of a resource adapter in an operational environment.

17.4.1 ResourceAdapter

The Java class which implements the interface

`javax.resource.spi.ResourceAdapter` must be a `JavaBean`. A `ResourceAdapter JavaBean` represents exactly one functional resource adapter unit or instance. The application server must instantiate exactly one `ResourceAdapter JavaBean` per functional resource adapter instance. The application server must create at least one functional resource adapter instance per resource adapter deployment. A `ResourceAdapter JavaBean` instance is configured by the resource adapter deployer during deployment. The configuration properties are specific to a resource adapter.

The resource adapter provider may optionally provide a set of configuration properties, specified in the resource adapter deployment descriptor, for the `ResourceAdapter` instance, which is used by the resource adapter deployer to configure the `ResourceAdapter JavaBean` instance during deployment. The deployer may override the configuration information in the deployment descriptor while configuring the `ResourceAdapter JavaBean` instance.

17.4.1.1 Requirements

The `ResourceAdapter` implementation must be a `JavaBean`.

17.4.2 ManagedConnectionFactory

The class that implements the `ManagedConnectionFactory` interface supports a set of properties. These properties provide information required by the `ManagedConnectionFactory` for the creation of physical connections to the underlying EIS.

A resource adapter must implement the `ManagedConnectionFactory` interface as a `JavaBean`. As a `JavaBean` implementor, the resource adapter can also provide a `BeanInfo` class that implements the `java.beans.BeanInfo` interface and provides explicit information about the methods and properties supported by the `ManagedConnectionFactory` implementation class.

The implementation of `ManagedConnectionFactory` as a JavaBean improves the ability of tools that are based on the JavaBeans framework to manage the configuration of `ManagedConnectionFactory` instances.

17.4.2.1 Requirements

The `ManagedConnectionFactory` implementation must be a JavaBean. Any specified `ManagedConnectionFactory` property in the deployment descriptor which does not have a matching property in the `ManagedConnectionFactory` JavaBean should be treated as an error.

17.4.3 Properties Conventions

The `ManagedConnectionFactory` implementation class must provide getter and setter methods for each of its supported properties. The supported properties must be consistent with the specification of configurable properties specified in the deployment descriptor.

The getter and setter methods convention must be based on the JavaBeans design pattern. These methods are defined in the implementation class and not in the `ManagedConnectionFactory` interface. This requirement keeps the `ManagedConnectionFactory` interface independent of any resource adapter or EIS-specific properties.

17.4.4 Standard Properties

The connector architecture identifies a standard set of properties common across various types of resource adapters and EISs. A resource adapter is not required to support a standard property if that property does not apply to its configuration.

These standard properties are defined as follows:

TABLE 17-2 Standard Properties of the Connector Architecture

Property	Description
<code>ServerName</code>	Name of the server for the EIS instance.
<code>PortNumber</code>	Port number for establishing a connection to an EIS instance.
<code>UserName</code>	Name of the user establishing a connection to an EIS instance.
<code>Password</code>	Password for the user establishing a connection.
<code>ConnectionURL</code>	URL for the EIS instance to which it connects.

In addition to these standard properties, a `ManagedConnectionFactory` implementation class may support properties specific to a resource adapter and its underlying EIS.

All properties are administered by the deployer and are not visible to an application component provider.

The specified properties are required to be implemented as either bound or constrained properties. Refer to the JavaBeans specification (refer to <http://java.sun.com/products/javabeans>) for details on bound and constrained properties.

In the XML deployment descriptor, any bounds or well-defined values of properties should be described in the `description` element.

17.5 JNDI Configuration and Lookup

This section specifies requirements for the configuration of the JNDI environment for a resource adapter.

In both managed and non-managed application scenarios, an application component or application client must look up a connection factory instance in the component's environment using the JNDI interface. The application component then uses the connection factory instance to get a connection to the underlying EIS. Section 6.4, "Application Programming Model" on page 6-6 specifies the application programming model in more detail.

The following code extract shows the JNDI lookup of a `javax.resource.cci.ConnectionFactory` instance.

```
// Application Component/Client Code
obtain the initial JNDI context
Context initctx = new InitialContext();

// perform JNDI lookup to obtain connection factory
javax.resource.cci.ConnectionFactory cxf =

(javax.resource.cci.ConnectionFactory)initctx.lookup("java:comp/
env/eis/MyEIS");

javax.resource.cci.Connection cx = cxf.getConnection();
```

17.5.1 Responsibilities

In both managed and non-managed environments, registration of a connection factory instance in the JNDI namespace must use either the JNDI `Reference` or `Serializable` mechanism.

The choice between the two JNDI mechanisms depends on:

- Whether or not the JNDI provider being used supports a specific mechanism.
- Whether or not the application server and resource adapter provide the necessary support, specified in the respective requirements.
- Constraints on the size of serialized objects that can be stored in the JNDI namespace. The reference mechanism allows only a reference to the actual object to be stored in the JNDI namespace. This is preferable to the serializable mechanism, which stores the whole serialized object in the namespace.

The following section specifies the responsibilities of the roles involved in the JNDI configuration of a resource adapter.

17.5.1.1 Deployer

The deployer is responsible for configuring connection factory instances in the JNDI environment. The deployer should manage the JNDI namespace such that the same programming model, as shown in Section 17.5 “JNDI Configuration and Lookup”, for the JNDI-based connection factory lookup is supported in both managed and non-managed environments.

17.5.1.2 Resource Adapter

The implementation class for a connection factory interface must implement both the `java.io.Serializable` and `javax.resource.Referenceable` interfaces to support JNDI registration.

The following code extract shows the `javax.resource.Referenceable` interface:

```
public interface javax.resource.Referenceable
    extends javax.naming.Referenceable {
    public void setReference(Reference ref);
}
```

The `ManagedConnectionFactory` implementation class must implement the `java.io.Serializable` interface.

To support the `Reference` mechanism in a non-managed environment, a resource adapter or a helper class must provide an implementation of the `javax.naming.spi.ObjectFactory` interface.

17.5.1.3 Application Server

The implementation class for `javax.resource.spi.ConnectionManager` must implement the `java.io.Serializable` interface.

An application server must provide an implementation class for the `javax.naming.spi.ObjectFactory` interface to support JNDI `Reference` mechanism-based connection factory lookup. The implementation of this interface is application server-specific.

Section 17.5.3 “Scenario: Referenceable” specifies more details on `Reference` mechanism-based JNDI configuration in a managed environment.

17.5.2 Scenario: Serializable

The implementation classes for both the `javax.resource.cci.ConnectionFactory` and `javax.resource.spi.ManagedConnectionFactory` interfaces implement the `java.io.Serializable` interface.

The deployment code retrieves the configuration properties from the XML deployment descriptor for the resource adapter. The deployment code then creates an instance of the `ManagedConnectionFactory` implementation class and configures the properties of the instance.

```
// Deployment Code
// Create an instance of the ManagedConnectionFactory class
com.myeis.ManagedConnectionFactoryImpl mcf =
    new com.myeis.ManagedConnectionFactoryImpl();

// Set the properties of the ManagedConnectionFactory instance
// Note: Properties are defined in the implementation class and
// not in the
// javax.resource.spi.ManagedConnectionFactory interface
mcf.setServerName("...");
mcf.setPortNumber("...");
...
```

Note that in a non-managed environment, an application developer writes the deployment code. In a managed environment, the deployment tool typically hides the deployment code.

The deployment code uses the `ManagedConnectionFactory` instance to create a connection factory instance. The code then registers the connection factory instance in the JNDI namespace.

```
// Deployment Code
// In a managed environment, create a ConnectionManager specific to
// the application server. Note that in a non-managed environment,
// ConnectionManager will be specific to the resource adapter.
com.wombatserver.ConnectionManager cm =
    new com.wombatserver.ConnectionManager(...);

// Create an instance of a connection factory
Object cxf = mcf.createConnectionFactory(cm);

// Get the JNDI context
javax.naming.Context ctx = new javax.naming.InitialContext(env);

// Bind to the JNDI namespace specifying a factory name
ctx.bind("...", cxf);
```

When an application component does a JNDI lookup of a connection factory instance, the returned connection factory instance should get associated with a configured `ManagedConnectionFactory` instance and a `ConnectionManager` instance. The implementation class for connection factory should achieve the association between these instances in an implementation-specific manner.

The following section illustrates JNDI configuration in a managed environment based on the Reference mechanism. This section uses the CCI interfaces `javax.resource.cci.ConnectionFactory` and `javax.resource.cci.Connection` as the connection factory and connection interfaces respectively.

17.5.3 Scenario: Referenceable

The implementation class for the `ConnectionFactory` interface implements `javax.resource.Referenceable` shown in the following code extract. Refer to the JNDI specification for more details on the `Referenceable` interface.

```
public class com.myeis.ConnectionFactoryImpl implements
    javax.resource.Referenceable,
    java.io.Serializable,
    javax.resource.cci.ConnectionFactory {

    // Reference to this ConnectionFactory
    javax.naming.Reference reference;

    // setReference is called by the deployment code
    public void setReference(Reference ref) {
        reference = ref;
    }

    // getReference is called by the JNDI provider during
    // Context.bind
    public Reference getReference() throws NamingException {
        return reference;
    }
    ...
}
```

The `getReference` method on the `ConnectionFactory` implementation class must return a non-null value or throw `javax.naming.NamingException`.

17.5.3.1 ObjectFactory Implementation

An application server provides a class (in an application server-specific implementation) that implements the `javax.naming.spi.ObjectFactory` interface. Refer to the JNDI specification for more details on the `ObjectFactory` interface.

In the `ObjectFactory.getObjectInstance` method, the information carried by the `Reference` parameter (set in the `ConnectionFactoryImpl.setReference` method) is used to lookup the property set to be configured on the target `ManagedConnectionFactory` instance.

The mapping from a `Reference` instance to multiple configured property sets enables an application server to configure multiple `ManagedConnectionFactory` instances with respective property sets. An application server maintains the property set configuration in an implementation-specific way based on the deployment descriptor specification.

The implementation and structure of `Reference` is specific to an application server. The following code extract is an illustrative example. It illustrates an implementation of the `ObjectFactory.getObjectInstance` method:

```
public class com.wombatserver.ApplicationServerJNDIHandler
    implements javax.naming.spi.ObjectFactory {
    // ...
    public Object getObjectInstance(Object obj, Name name,
        Context ctx, Hashtable env)
        throws Exception {

        javax.naming.Reference ref =
            (javax.naming.Reference)obj;

        // Using the information carried by the Reference
        // instance,
        // (<referenceName, logicalName> in this example) lookup
        // a configured property set and then configure a
        // ManagedConnectionFactory instance with specified
        // properties.
        ... // [implementation specific]
        //
        // For example, instantiation of the
        // ManagedConnectionFactory
        // class and invocation of its setter method
        // can be done using the Java Reflection mechanism.

        javax.resource.spi.ManagedConnectionFactory mcf = ...

        // Create a Connection Manager instance specific to the
        // application server
        com.wombatserver.ConnectionManager cxManager = ...

        // Create a connection factory instance.
        // The ConnectionManager instance provided by the
        // application
        // server gets associated with the created
        // connection factory instance
        return mcf.createConnectionFactory(cxManager);
    }
    ...
}
```

17.5.3.2 Deployment

The following deployment code shows the registration of a reference to a connection factory instance in the JNDI namespace:

```
// Deployment Code
javax.naming.Context ctx = new javax.naming.InitialContext(env);

// Create an instance of the connection factory
com.myeis.ConnectionFactoryImpl cf =
    new com.myeis.ConnectionFactoryImpl();

// Create a reference for the ConnectionFactory instance
javax.naming.Reference ref = new javax.naming.Reference(
    ConnectionFactoryImpl.class.getName(),
    new javax.naming.StringRefAddr(
        "<referenceName>", "<logicalName>"),
    ApplicationServerJNDIHandler.class.getName(),
    null);

cf.setReference(ref);

// Bind to the JNDI namespace specifying a name for the connection
// factory
ctx.bind("...", cf);
```

Note that the deployment code should be designed as generic, though the above example does not show it that way. The code should dynamically create an instance of a connection factory, create a `Reference` instance, and then set the reference.

The `Context.bind` method registers a `Reference` to the connection factory instance in the JNDI namespace.

17.5.3.3 Scenario: Connection Factory Lookup

The following steps occur when an application component calls the method `JNDIContext.lookup` to lookup a connection factory instance:

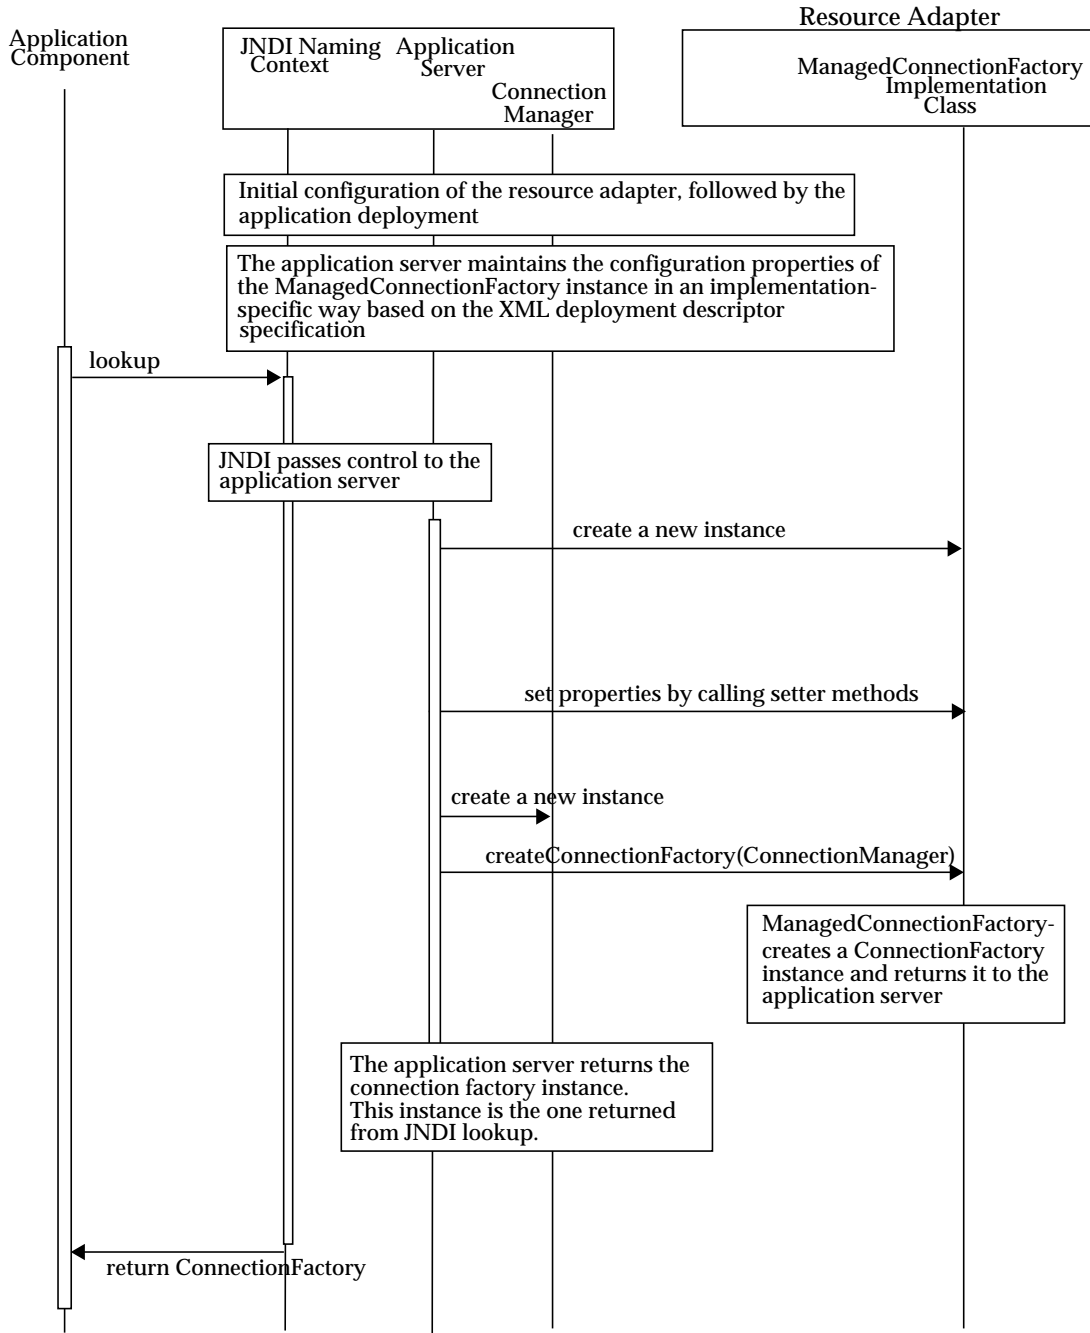
- JNDI passes control to the application server. The `ObjectFactory.getObjectInstance` method implemented by the application server is called.
- The application server creates a new instance of the `ManagedConnectionFactory` implementation class provided by the resource adapter.
- The application server calls setter methods on the `ManagedConnectionFactory` instance to set various configuration properties of this instance. These properties provide information required by the `ManagedConnectionFactory` instance to

create physical connections to the underlying EIS. The application server uses an existing property set configured during the deployment of a resource adapter to set the required properties of the `ManagedConnectionFactory` instance.

- After the newly created `ManagedConnectionFactory` instance has been configured with its properties set, the application server creates a new `ConnectionManager` instance.
- The application server calls the `createConnectionFactory` method of the `ManagedConnectionFactory` instance, passing in the `ConnectionManager` instance from the previous step, to get a `ConnectionFactory` instance.

- The application server returns the connection factory instance to the JNDI provider, so that this instance can be returned as a result of the JNDI lookup. The application component gets the `ConnectionFactory` instance as a result of the JNDI lookup.

FIGURE 17-3 OID: Lookup of a ConnectionFactory Instance from JNDI



17.5.4 Requirements

The default configuration values for the various JavaBean classes specified in the resource adapter deployment descriptor via the `config-property` element override and take precedence over the defaults specified for the same classes by the resource adapter developer through the JavaBean mechanism. Note, a deployer may finally override such default configuration information while configuring the various JavaBean instances.

17.6 Resource Adapter XML Schema Definition

This section specifies the XML Schema Definition (XSD) for the deployment descriptor for a resource adapter. Some of the types used in this XSD are defined in the J2EE platform specification (Section 8., “Java 2 Platform Enterprise Edition (J2EETM), Platform specification, version 1. 4:” on page F-1). The comments in the XSD specify additional requirements for syntax and semantics that cannot be specified by using the XML Schema language. Note, the `description-group` element defined in `j2ee_1_4.xsd` allows multiple descriptions, in order to allow specifying the same description for different locales.

CODE EXAMPLE 17-1 Schema Definition for the Deployment Descriptor for a Resource Adapter

```
<?xml version="1.0" encoding="UTF-8"?>
<xsd:schema xmlns="http://www.w3.org/2001/XMLSchema"
  targetNamespace="http://java.sun.com/xml/ns/j2ee"
  xmlns:j2ee="http://java.sun.com/xml/ns/j2ee"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema"
  elementFormDefault="qualified"
  attributeFormDefault="unqualified"
  version="1.5">
  <xsd:annotation>
    <xsd:documentation>
      @(#)connector_1_5.xsds1.27 06/17/03
    </xsd:documentation>
  </xsd:annotation>

  <xsd:annotation>
    <xsd:documentation>
      <![CDATA[

This is the XML Schema for the Connector 1.5 deployment
descriptor. The deployment descriptor must be named
```

CODE EXAMPLE 17-1 Schema Definition for the Deployment Descriptor for a Resource Adapter

"META-INF/ra.xml" in the connector's rar file. All Connector deployment descriptors must indicate the connector resource adapter schema by using the J2EE namespace:

```
http://java.sun.com/xml/ns/j2ee
```

and by indicating the version of the schema by using the version element as shown below:

```
<connector xmlns="http://java.sun.com/xml/ns/j2ee"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://java.sun.com/xml/ns/j2ee
http://java.sun.com/xml/ns/j2ee/connector_1_5.xsd"
  version="1.5">
  ...
</connector>
```

The instance documents may indicate the published version of the schema using the `xsi:schemaLocation` attribute for J2EE namespace with the following location:

```
http://java.sun.com/xml/ns/j2ee/connector_1_5.xsd
```

```
]]>
</xsd:documentation>
</xsd:annotation>
```

```
<xsd:annotation>
  <xsd:documentation>
```

The following conventions apply to all J2EE deployment descriptor elements unless indicated otherwise.

- In elements that specify a pathname to a file within the same JAR file, relative filenames (i.e., those not starting with "/") are considered relative to the root of the JAR file's namespace. Absolute filenames (i.e., those starting with "/") also specify names in the root of the JAR file's namespace. In general, relative names are preferred. The exception is .war files where absolute names are preferred for consistency with the Servlet API.

```
</xsd:documentation>
</xsd:annotation>

<xsd:include schemaLocation="j2ee_1_4.xsd"/>
```

CODE EXAMPLE 17-1 Schema Definition for the Deployment Descriptor for a Resource Adapter

```
<!-- ***** -->

<xsd:element name="connector" type="j2ee:connectorType">
  <xsd:annotation>
    <xsd:documentation>

      The connector element is the root element of the deployment
      descriptor for the resource adapter. This element includes
      general information - vendor name, resource adapter version,
      icon - about the resource adapter module. It also includes
      information specific to the implementation of the resource
      adapter library as specified through the element
      resourceadapter.

    </xsd:documentation>
  </xsd:annotation>
</xsd:element>

<!-- ***** -->

<xsd:complexType name="activationSpecType">
  <xsd:annotation>
    <xsd:documentation>

      The activationSpecType specifies an activation
      specification. The information includes fully qualified
      Java class name of an activation specification and a set of
      required configuration property names.

    </xsd:documentation>
  </xsd:annotation>

  <xsd:sequence>
    <xsd:element name="activationSpec-class"
      type="j2ee:fully-qualified-classType">
      <xsd:annotation>
        <xsd:documentation>
          <![CDATA[

            The element activationSpec-class specifies the fully
            qualified Java class name of the activation
            specification class. This class must implement the
            javax.resource.spi.ActivationSpec interface. The
            implementation of this class is required to be a
```

CODE EXAMPLE 17-1 Schema Definition for the Deployment Descriptor for a Resource Adapter

```

        JavaBean.

        Example:
        <activation-spec-class>com.wombat.ActivationSpecImpl
        </activation-spec-class>

        ]]>
    </xsd:documentation>
</xsd:annotation>
</xsd:element>
<xsd:element name="required-config-property"
    type="j2ee:required-config-propertyType"
    minOccurs="0"
    maxOccurs="unbounded"/>
</xsd:sequence>
<xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>

<!-- ***** -->

<xsd:complexType name="adminObjectType">
    <xsd:annotation>
        <xsd:documentation>

The adminObjectType specifies information about an
administered object. Administered objects are specific to a
messaging style or message provider. This contains
information on the Java type of the interface implemented by
an administered object, its Java class name and its
configuration properties.

        </xsd:documentation>
    </xsd:annotation>

    <xsd:sequence>
        <xsd:element name="adminobject-interface"
            type="j2ee:fully-qualified-classType">
    </xsd:annotation>
        <xsd:documentation>
            <![CDATA[

                The element adminobject-interface specifies the
                fully qualified name of the Java type of the
                interface implemented by an administered object.

                Example:
                <adminobject-interface>javax.jms.Destination

```

CODE EXAMPLE 17-1 Schema Definition for the Deployment Descriptor for a Resource Adapter

```

    </adminobject-interface>

    ]]>
    </xsd:documentation>
</xsd:annotation>

</xsd:element>
<xsd:element name="adminobject-class"
    type="j2ee:fully-qualified-classType">
<xsd:annotation>
    <xsd:documentation>
        <![CDATA[

            The element adminobject-class specifies the fully
            qualified Java class name of an administered object.

            Example:
            <adminobject-class>com.wombat.DestinationImpl
            </adminobject-class>

        ]]>
    </xsd:documentation>
</xsd:annotation>
</xsd:element>

    <xsd:element name="config-property"
        type="j2ee:config-propertyType"
        minOccurs="0"
        maxOccurs="unbounded"/>
</xsd:sequence>
    <xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>

<!-- ***** -->

<xsd:complexType name="authentication-mechanismType">
    <xsd:annotation>
        <xsd:documentation>

            The authentication-mechanismType specifies an authentication
            mechanism supported by the resource adapter. Note that this
            support is for the resource adapter and not for the
            underlying EIS instance. The optional description specifies
            any resource adapter specific requirement for the support of
            security contract and authentication mechanism.

            Note that BasicPassword mechanism type should support the

```

CODE EXAMPLE 17-1 Schema Definition for the Deployment Descriptor for a Resource Adapter

javax.resource.spi.security.PasswordCredential interface.
 The Kerbv5 mechanism type should support the
 org.ietf.jgss.GSSCredential interface or the deprecated
 javax.resource.spi.security.GenericCredential interface.

```

    </xsd:documentation>
  </xsd:annotation>

```

```

<xsd:sequence>
  <xsd:element name="description"
    type="j2ee:descriptionType"
    minOccurs="0"
    maxOccurs="unbounded"/>
  <xsd:element name="authentication-mechanism-type"
    type="j2ee:xsdStringType">

```

```

<xsd:annotation>
  <xsd:documentation>
    <![CDATA[

```

The element authentication-mechanism-type specifies
 type of an authentication mechanism.

The example values are:

```

<authentication-mechanism-type>BasicPassword
</authentication-mechanism-type>

```

```

<authentication-mechanism-type>Kervb5
</authentication-mechanism-type>

```

Any additional security mechanisms are outside the
 scope of the Connector architecture specification.

```

    ]]>
  </xsd:documentation>
</xsd:annotation>

```

```

</xsd:element>
<xsd:element name="credential-interface"
  type="j2ee:credential-interfaceType"/>
</xsd:sequence>
<xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>

```

```

<!-- ***** -->

```

```

<xsd:complexType name="config-property-nameType">

```

CODE EXAMPLE 17-1 Schema Definition for the Deployment Descriptor for a Resource Adapter

```

<xsd:annotation>
  <xsd:documentation>
    <![CDATA[

      The config-property-nameType contains the name of a
      configuration property.

      The connector architecture defines a set of well-defined
      properties all of type java.lang.String. These are as
      follows.

          ServerName
          PortNumber
          UserName
          Password
          ConnectionURL

      A resource adapter provider can extend this property set to
      include properties specific to the resource adapter and its
      underlying EIS.

      Possible values include
      ServerName
      PortNumber
      UserName
      Password
      ConnectionURL

      Example: <config-property-name>ServerName</config-property-name>

    ]]>
  </xsd:documentation>
</xsd:annotation>
<xsd:simpleContent>
  <xsd:restriction base="j2ee:xsdStringType"/>
</xsd:simpleContent>
</xsd:complexType>

<!-- ***** -->

<xsd:complexType name="config-property-typeType">
  <xsd:annotation>
    <xsd:documentation>
      <![CDATA[

        The config-property-typeType contains the fully
        qualified Java type of a configuration property.

```

CODE EXAMPLE 17-1 Schema Definition for the Deployment Descriptor for a Resource Adapter

The following are the legal values:

java.lang.Boolean, java.lang.String, java.lang.Integer,
java.lang.Double, java.lang.Byte, java.lang.Short,
java.lang.Long, java.lang.Float, java.lang.Character

Used in: config-property

Example:

```
<config-property-type>java.lang.String</config-property-type>
```

```
]]>
```

```
</xsd:documentation>
```

```
</xsd:annotation>
```

```
<xsd:simpleContent>
```

```
  <xsd:restriction base="j2ee:string">
```

```
    <xsd:enumeration value="java.lang.Boolean"/>
```

```
    <xsd:enumeration value="java.lang.String"/>
```

```
    <xsd:enumeration value="java.lang.Integer"/>
```

```
    <xsd:enumeration value="java.lang.Double"/>
```

```
    <xsd:enumeration value="java.lang.Byte"/>
```

```
    <xsd:enumeration value="java.lang.Short"/>
```

```
    <xsd:enumeration value="java.lang.Long"/>
```

```
    <xsd:enumeration value="java.lang.Float"/>
```

```
    <xsd:enumeration value="java.lang.Character"/>
```

```
  </xsd:restriction>
```

```
</xsd:simpleContent>
```

```
</xsd:complexType>
```

```
<!-- ***** -->
```

```
<xsd:complexType name="config-propertyType">
```

```
  <xsd:annotation>
```

```
    <xsd:documentation>
```

The config-propertyType contains a declaration of a single configuration property that may be used for providing configuration information.

The declaration consists of an optional description, name, type and an optional value of the configuration property. If the resource adapter provider does not specify a value than the deployer is responsible for providing a valid value for a configuration property.

Any bounds or well-defined values of properties should be described in the description element.

CODE EXAMPLE 17-1 Schema Definition for the Deployment Descriptor for a Resource Adapter

```

    </xsd:documentation>
  </xsd:annotation>
  <xsd:sequence>
    <xsd:element name="description"
      type="j2ee:descriptionType"
      minOccurs="0"
      maxOccurs="unbounded"/>
    <xsd:element name="config-property-name"
      type="j2ee:config-property-nameType"/>
    <xsd:element name="config-property-type"
      type="j2ee:config-property-typeType"/>
    <xsd:element name="config-property-value"
      type="j2ee:xsdStringType"
      minOccurs="0">
      <xsd:annotation>
        <xsd:documentation>
          <![CDATA[

            The element config-property-value contains the value
            of a configuration entry. Note, it is possible for a
            resource adapter deployer to override this
            configuration information during deployment.

            Example:
            <config-property-value>WombatServer</config-property-value>

          ]]>
        </xsd:documentation>
      </xsd:annotation>
    </xsd:element>
  </xsd:sequence>
  <xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>

<!-- ***** -->

<xsd:complexType name="connection-definitionType">
  <xsd:annotation>
    <xsd:documentation>

      The connection-definitionType defines a set of connection
      interfaces and classes pertaining to a particular connection
      type. This also includes configurable properties for
      ManagedConnectionFactory instances that may be produced out
      of this set.
    
```

CODE EXAMPLE 17-1 Schema Definition for the Deployment Descriptor for a Resource Adapter

```
</xsd:documentation>
</xsd:annotation>

<xsd:sequence>
  <xsd:element name="managedconnectionfactory-class"
    type="j2ee:fully-qualified-classType">
    <xsd:annotation>
      <xsd:documentation>
        <![CDATA[

          The element managedconnectionfactory-class specifies
          the fully qualified name of the Java class that
          implements the
          javax.resource.spi.ManagedConnectionFactory interface.
          This Java class is provided as part of resource
          adapter's implementation of connector architecture
          specified contracts. The implementation of this
          class is required to be a JavaBean.

          Example:
          <managedconnectionfactory-class>
            com.wombat.ManagedConnectionFactoryImpl
          </managedconnectionfactory-class>

        ]]>
      </xsd:documentation>
    </xsd:annotation>
  </xsd:element>
  <xsd:element name="config-property"
    type="j2ee:config-propertyType"
    minOccurs="0"
    maxOccurs="unbounded"/>
  <xsd:element name="connectionfactory-interface"
    type="j2ee:fully-qualified-classType">
    <xsd:annotation>
      <xsd:documentation>
        <![CDATA[

          The element connectionfactory-interface specifies
          the fully qualified name of the ConnectionFactory
          interface supported by the resource adapter.

          Example:
          <connectionfactory-interface>com.wombat.ConnectionFactory
          </connectionfactory-interface>

        ]]>
      </xsd:documentation>
    </xsd:annotation>
  </xsd:element>
</xsd:sequence>

OR
```

CODE EXAMPLE 17-1 Schema Definition for the Deployment Descriptor for a Resource Adapter

```
<connectionfactory-interface>javax.resource.cci.ConnectionFactory
</connectionfactory-interface>

]]>
</xsd:documentation>
</xsd:annotation>
</xsd:element>
```

```
<xsd:element name="connectionfactory-impl-class"
  type="j2ee:fully-qualified-classType">
<xsd:annotation>
  <xsd:documentation>
    <![CDATA[
```

The element `connectionfactory-impl-class` specifies the fully qualified name of the `ConnectionFactory` class that implements resource adapter specific `ConnectionFactory` interface.

Example:

```
<connectionfactory-impl-class>com.wombat.ConnectionFactoryImpl
</connectionfactory-impl-class>
```

```
]]>
</xsd:documentation>
</xsd:annotation>
</xsd:element>
<xsd:element name="connection-interface"
  type="j2ee:fully-qualified-classType">
<xsd:annotation>
  <xsd:documentation>
    <![CDATA[
```

The `connection-interface` element specifies the fully qualified name of the `Connection` interface supported by the resource adapter.

Example:

```
<connection-interface>javax.resource.cci.Connection
</connection-interface>
```

```
]]>
</xsd:documentation>
</xsd:annotation>
```

CODE EXAMPLE 17-1 Schema Definition for the Deployment Descriptor for a Resource Adapter

```
</xsd:element>
<xsd:element name="connection-impl-class"
  type="j2ee:fully-qualified-classType">
  <xsd:annotation>
    <xsd:documentation>
      <![CDATA[

        The connection-impl-classType specifies the fully
        qualified name of the Connection class that
        implements resource adapter specific Connection
        interface. It is used by the connection-impl-class
        elements.

        Example:

        <connection-impl-class>com.wombat.ConnectionImpl
        </connection-impl-class>

        ]]>
      </xsd:documentation>
    </xsd:annotation>
  </xsd:element>
</xsd:sequence>
<xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>

<!-- ***** -->

<xsd:complexType name="connectorType">
  <xsd:annotation>
    <xsd:documentation>

      The connectorType defines a resource adapter.

    </xsd:documentation>
  </xsd:annotation>

  <xsd:sequence>
    <xsd:group ref="j2ee:descriptionGroup"/>
    <xsd:element name="vendor-name"
      type="j2ee:xsdStringType">
    </xsd:element>
  </xsd:sequence>
  <xsd:annotation>
    <xsd:documentation>

      The element vendor-name specifies the name of
      resource adapter provider vendor.

    </xsd:documentation>
  </xsd:annotation>
</xsd:complexType>
```

CODE EXAMPLE 17-1 Schema Definition for the Deployment Descriptor for a Resource Adapter

```
</xsd:documentation>
</xsd:annotation>
</xsd:element>
<xsd:element name="eis-type"
  type="j2ee:xsdStringType">
<xsd:annotation>
  <xsd:documentation>
```

The element `eis-type` contains information about the type of the EIS. For example, the type of an EIS can be product name of EIS independent of any version info.

This helps in identifying EIS instances that can be used with this resource adapter.

```
</xsd:documentation>
</xsd:annotation>
</xsd:element>
<xsd:element name="resourceadapter-version"
  type="j2ee:xsdStringType">
<xsd:annotation>
  <xsd:documentation>
```

The element `resourceadapter-version` specifies a string-based version of the resource adapter from the resource adapter provider.

```
</xsd:documentation>
</xsd:annotation>

</xsd:element>
<xsd:element name="license"
  type="j2ee:licenseType"
  minOccurs="0"/>
<xsd:element name="resourceadapter"
  type="j2ee:resourceadapterType"/>
```

```
</xsd:sequence>
<xsd:attribute name="version"
  type="j2ee:dewey-versionType"
  fixed="1.5"
  use="required">
<xsd:annotation>
<xsd:documentation>
```

The version specifies the version of the

CODE EXAMPLE 17-1 Schema Definition for the Deployment Descriptor for a Resource Adapter

```

connector architecture specification that is
supported by this resource adapter. This information
enables deployer to configure the resource adapter to
support deployment and runtime requirements of the
corresponding connector architecture specification.

</xsd:documentation>
</xsd:annotation>
</xsd:attribute>
<xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>

<!-- ***** -->

<xsd:complexType name="credential-interfaceType">
  <xsd:annotation>
    <xsd:documentation>

The credential-interfaceType specifies the
interface that the resource adapter implementation
supports for the representation of the
credentials. This element(s) that use this type,
i.e. credential-interface, should be used by
application server to find out the Credential
interface it should use as part of the security
contract.

The possible values are:

javax.resource.spi.security.PasswordCredential
org.ietf.jgss.GSSCredential
javax.resource.spi.security.GenericCredential

    </xsd:documentation>
  </xsd:annotation>
  <xsd:simpleContent>
    <xsd:restriction base="j2ee:fully-qualified-classType">
      <xsd:enumeration
        value="javax.resource.spi.security.PasswordCredential"/>
      <xsd:enumeration
        value="org.ietf.jgss.GSSCredential"/>
      <xsd:enumeration
        value="javax.resource.spi.security.GenericCredential"/>
    </xsd:restriction>
  </xsd:simpleContent>
</xsd:complexType>

```

CODE EXAMPLE 17-1 Schema Definition for the Deployment Descriptor for a Resource Adapter

```
<!-- ***** -->

<xsd:complexType name="inbound-resourceadapterType">
  <xsd:annotation>
    <xsd:documentation>

      The inbound-resourceadapterType specifies information
      about an inbound resource adapter. This contains information
      specific to the implementation of the resource adapter
      library as specified through the messageadapter element.

    </xsd:documentation>
  </xsd:annotation>

  <xsd:sequence>
    <xsd:element name="messageadapter"
      type="j2ee:messageadapterType"
      minOccurs="0">
      <xsd:unique name="messagelistener-type-uniqueness">
        <xsd:annotation>
          <xsd:documentation>

            The messagelistener-type element content must be
            unique in the messageadapter. Several messagelisteners
            can not use the same messagelistener-type.

          </xsd:documentation>
        </xsd:annotation>
        <xsd:selector xpath="j2ee:messagelistener"/>
        <xsd:field xpath="j2ee:messagelistener-type"/>
      </xsd:unique>
    </xsd:element>
  </xsd:sequence>
  <xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>

<!-- ***** -->

<xsd:complexType name="licenseType">
  <xsd:annotation>
    <xsd:documentation>

      The licenseType specifies licensing requirements for the
      resource adapter module. This type specifies whether a
      license is required to deploy and use this resource adapter,
      and an optional description of the licensing terms
      (examples: duration of license, number of connection
```

CODE EXAMPLE 17-1 Schema Definition for the Deployment Descriptor for a Resource Adapter

restrictions). It is used by the license element.

```
</xsd:documentation>
</xsd:annotation>

<xsd:sequence>
  <xsd:element name="description"
    type="j2ee:descriptionType"
    minOccurs="0"
    maxOccurs="unbounded" />
  <xsd:element name="license-required"
    type="j2ee:true-falseType">
<xsd:annotation>
  <xsd:documentation>
```

The element license-required specifies whether a license is required to deploy and use the resource adapter. This element must be one of the following, "true" or "false".

```
</xsd:documentation>
</xsd:annotation>
</xsd:element>
</xsd:sequence>
<xsd:attribute name="id" type="xsd:ID" />
</xsd:complexType>

<!-- ***** -->
```

```
<xsd:complexType name="messageadapterType">
  <xsd:annotation>
    <xsd:documentation>
```

The messageadapterType specifies information about the messaging capabilities of the resource adapter. This contains information specific to the implementation of the resource adapter library as specified through the messagelistener element.

```
</xsd:documentation>
</xsd:annotation>

<xsd:sequence>
  <xsd:element name="messagelistener"
    type="j2ee:messagelistenerType"
    maxOccurs="unbounded" />
</xsd:sequence>
```


CODE EXAMPLE 17-1 Schema Definition for the Deployment Descriptor for a Resource Adapter

```

    <xsd:attribute name="id" type="xsd:ID" />
  </xsd:complexType>

<!-- ***** -->

  <xsd:complexType name="messagelistenerType">
    <xsd:annotation>
      <xsd:documentation>

        The messagelistenerType specifies information about a
        specific message listener supported by the messaging
        resource adapter. It contains information on the Java type
        of the message listener interface and an activation
        specification.

      </xsd:documentation>
    </xsd:annotation>

    <xsd:sequence>
      <xsd:element name="messagelistener-type"
        type="j2ee:fully-qualified-classType">
        <xsd:annotation>
          <xsd:documentation>
            <![CDATA[

              The element messagelistener-type specifies the fully
              qualified name of the Java type of a message
              listener interface.

              Example:

              <messagelistener-type>javax.jms.MessageListener
              </messagelistener-type>

            ]]>
          </xsd:documentation>
        </xsd:annotation>

        </xsd:element>
        <xsd:element name="activation-spec"
          type="j2ee:activation-specType" />
      </xsd:sequence>
      <xsd:attribute name="id" type="xsd:ID" />
    </xsd:complexType>

  <!-- ***** -->

```

CODE EXAMPLE 17-1 Schema Definition for the Deployment Descriptor for a Resource Adapter

```
<xsd:complexType name="outbound-resourceadapterType">
  <xsd:annotation>
    <xsd:documentation>
```

The outbound-resourceadapterType specifies information about an outbound resource adapter. The information includes fully qualified names of classes/interfaces required as part of the connector architecture specified contracts for connection management, level of transaction support provided, one or more authentication mechanisms supported and additional required security permissions.

If there is no authentication-mechanism specified as part of resource adapter element then the resource adapter does not support any standard security authentication mechanisms as part of security contract. The application server ignores the security part of the system contracts in this case.

```
    </xsd:documentation>
  </xsd:annotation>
```

```
<xsd:sequence>
  <xsd:element name="connection-definition"
    type="j2ee:connection-definitionType"
    maxOccurs="unbounded"/>
  <xsd:element name="transaction-support"
    type="j2ee:transaction-supportType"/>
  <xsd:element name="authentication-mechanism"
    type="j2ee:authentication-mechanismType"
    minOccurs="0"
    maxOccurs="unbounded"/>
  <xsd:element name="reauthentication-support"
    type="j2ee:true-falseType">
<xsd:annotation>
  <xsd:documentation>
```

The element reauthentication-support specifies whether the resource adapter implementation supports re-authentication of existing Managed- Connection instance. Note that this information is for the resource adapter implementation and not for the underlying EIS instance. This element must have either a "true" or "false" value.

```
    </xsd:documentation>
  </xsd:annotation>
```

CODE EXAMPLE 17-1 Schema Definition for the Deployment Descriptor for a Resource Adapter

```
</xsd:element>
</xsd:sequence>
<xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>

<!-- ***** -->

<xsd:complexType name="required-config-propertyType">
  <xsd:annotation>
    <xsd:documentation>
      <![CDATA[

        The required-config-propertyType contains a declaration
        of a single configuration property used for specifying a
        required configuration property name. It is used
        by required-config-property elements.

        Example:

        <required-config-property>Destination</required-config-property>

      ]]>
    </xsd:documentation>
  </xsd:annotation>

  <xsd:sequence>
    <xsd:element name="description"
      type="j2ee:descriptionType"
      minOccurs="0"
      maxOccurs="unbounded"/>
    <xsd:element name="config-property-name"
      type="j2ee:config-property-nameType"/>
  </xsd:sequence>
  <xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>

<!-- ***** -->

<xsd:complexType name="resourceadapterType">
  <xsd:annotation>
    <xsd:documentation>

      The resourceadapterType specifies information about the
      resource adapter. The information includes fully qualified
      resource adapter Java class name, configuration properties,
      information specific to the implementation of the resource
      adapter library as specified through the
```

CODE EXAMPLE 17-1 Schema Definition for the Deployment Descriptor for a Resource Adapter

outbound-resourceadapter and inbound-resourceadapter elements, and an optional set of administered objects.

```
</xsd:documentation>
</xsd:annotation>
```

```
<xsd:sequence>
  <xsd:element name="resourceadapter-class"
    type="j2ee:fully-qualified-classType"
    minOccurs="0">
```

```
<xsd:annotation>
  <xsd:documentation>
```

The element resourceadapter-class specifies the fully qualified name of a Java class that implements the javax.resource.spi.ResourceAdapter interface. This Java class is provided as part of resource adapter's implementation of connector architecture specified contracts. The implementation of this class is required to be a JavaBean.

```
</xsd:documentation>
</xsd:annotation>
```

```
</xsd:element>
<xsd:element name="config-property"
  type="j2ee:config-propertyType"
  minOccurs="0"
  maxOccurs="unbounded"/>
```

```
<xsd:element name="outbound-resourceadapter"
  type="j2ee:outbound-resourceadapterType"
  minOccurs="0">
<xsd:unique name="connectionfactory-interface-uniqueness">
  <xsd:annotation>
    <xsd:documentation>
```

The connectionfactory-interface element content must be unique in the outbound-resourceadapter. Multiple connection-definitions can not use the same connectionfactory-type.

```
</xsd:documentation>
</xsd:annotation>
<xsd:selector xpath="j2ee:connection-definition"/>
<xsd:field xpath="j2ee:connectionfactory-interface"/>
</xsd:unique>
</xsd:element>
```

CODE EXAMPLE 17-1 Schema Definition for the Deployment Descriptor for a Resource Adapter

```
<xsd:element name="inbound-resourceadapter"
  type="j2ee:inbound-resourceadapterType"
  minOccurs="0"/>
<xsd:element name="adminobject"
  type="j2ee:adminobjectType"
  minOccurs="0"
  maxOccurs="unbounded"/>
<xsd:element name="security-permission"
  type="j2ee:security-permissionType"
  minOccurs="0"
  maxOccurs="unbounded"/>
</xsd:sequence>
<xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>

<!-- ***** -->

<xsd:complexType name="security-permissionType">
  <xsd:annotation>
    <xsd:documentation>

      The security-permissionType specifies a security
      permission that is required by the resource adapter code.

      The security permission listed in the deployment descriptor
      are ones that are different from those required by the
      default permission set as specified in the connector
      specification. The optional description can mention specific
      reason that resource adapter requires a given security
      permission.

    </xsd:documentation>
  </xsd:annotation>

  <xsd:sequence>
    <xsd:element name="description"
      type="j2ee:descriptionType"
      minOccurs="0"
      maxOccurs="unbounded"/>
    <xsd:element name="security-permission-spec"
      type="j2ee:xsdStringType">
      <xsd:annotation>
        <xsd:documentation>

          The element security-permission-spec specifies a security
          permission based on the Security policy file
          syntax. Refer to the following URL for Sun's
```

CODE EXAMPLE 17-1 Schema Definition for the Deployment Descriptor for a Resource Adapter

```
        implementation of the security permission
        specification:

http://java.sun.com/products/jdk/1.4/docs/guide/security/PolicyFiles.html#File
Syntax
        </xsd:documentation>
    </xsd:annotation>
    </xsd:element>
</xsd:sequence>
<xsd:attribute name="id" type="xsd:ID"/>
</xsd:complexType>

<!-- ***** -->

<xsd:complexType name="transaction-supportType">
    <xsd:annotation>
        <xsd:documentation>

            The transaction-supportType specifies the level of
            transaction support provided by the resource adapter. It is
            used by transaction-support elements.

            The value must be one of the following:

                NoTransaction
                LocalTransaction
                XATransaction

        </xsd:documentation>
    </xsd:annotation>
    <xsd:simpleContent>
        <xsd:restriction base="j2ee:string">
            <xsd:enumeration value="NoTransaction"/>
            <xsd:enumeration value="LocalTransaction"/>
            <xsd:enumeration value="XATransaction"/>
        </xsd:restriction>
    </xsd:simpleContent>
</xsd:complexType>

</xsd:schema>
```

Runtime Environment

This chapter focuses on the Java portion of a resource adapter that executes within a Java compatible runtime environment. A Java runtime environment is provided by an application server and its containers.

The chapter specifies the Java APIs that a J2EE-compliant application server and its containers must make available to a resource adapter at runtime. A portable resource adapter can rely on these APIs to be available on all J2EE-compliant application servers.

The chapter also specifies programming restrictions imposed on a resource adapter. These restrictions enable an application server to enforce security and manage a runtime environment with multiple configured resource adapters.

18.1 Programming APIs

A resource adapter provider relies on a J2EE compliant application server to provide the following APIs:

- Java 2 SDK, Standard Edition, version 1.4 that includes the following as part of either the core platform or standard extensions: JavaIDL, JNDI Standard Extension, and RMI-IIOP.
- Required APIs for Java 2 SDK, Enterprise Edition, version 1.4 as specified in the J2EE platform specification (Section 8., “Java 2 Platform Enterprise Edition (J2EETM), Platform specification, version 1. 4:” on page F-1), version 1.4.
- Java Authentication and Authorization Service (JAAS) 1.0 that requires at least Java 2 SDK, Standard Edition, version 1.3 or the Java 2 Runtime Environment version 1.3.

18.2 Security Permissions

An application server must provide a set of security permissions for executing a resource adapter in a managed runtime environment. A resource adapter must be granted explicit permissions to access system resources.

Since the exact set of required security permissions for a resource adapter depends on the overall security policy for an operational environment and the implementation requirements of a resource adapter, the connector architecture does not define a fixed set of permissions.

The following permission set represents the default set of security permissions that a resource adapter should expect from an application server. These security permissions are described in detail in the Java 2 platform documentation. Refer to <http://java.sun.com/products/jdk/1.4/docs/guide/security/permissions.html>.

TABLE 18-1 Default Security Permission Set

Security Permission	Default Policy	Notes
java.security. AllPermission	deny	Extreme care should be taken before granting this permission to a resource adapter. This permission should only be granted if the resource adapter code is completely trusted and when it is prohibitively cumbersome to add necessary permissions to the security policy.
java.awt.AWTPermission	deny *	A resource adapter must not use AWT code to interact with display or input devices.
java.io.FilePermission	grant read and write <pathname> deny rest	A java.io.FilePermission represents access to a file or directory. A FilePermission consists of a pathname and a set of actions valid for that pathname. A resource adapter is granted permission to read/write files as specified by the pathname, which is specific to a configured operational environment. It is important to consider the implications of granting Write permission for <<ALL FILES>> because this grants the resource adapter permissions to write to the entire file system. This can allow a malicious resource adapter to mangle system binaries for the JVM environment.
java.net.NetPermission	deny *	

TABLE 18-1 Default Security Permission Set

Security Permission	Default Policy	Notes
java.util. PropertyPermission	grant read (allows System.getProper- ty to be called) deny rest	Granting code permission to access certain system properties (java.home) can potentially give malevolent code sensitive information about the system environment, such as the Java installation directory.
java.lang.reflect.ReflectPermission	deny *	
java.lang. RuntimePermission	deny *	By default, RuntimePermission is denied to the resource adapter code. A resource adapter should explicitly request LoadLibrary.{libraryName} } to link a dynamic library. The libraryName represents a specific library. A resource adapter that manages threads must explicitly request permission to modifyThread through its deployment descriptor. A resource adapter should never be granted exitVM permission in a managed application server environment.

TABLE 18-1 Default Security Permission Set

Security Permission	Default Policy	Notes
java.security. SecurityPermission	deny *	
java.net. SocketPermission	grant connect * deny rest	This represents permission to access a network via sockets. A <code>SocketPermission</code> consists of a host specification and a set of actions specifying ways to connect to that host. A resource adapter is granted permission to connect to any host as indicated by the wildcard *. A resource adapter may be granted permission to accept connections from other hosts via a “grant accept *”. This may be necessary for resource adapters that support inbound communication.
java.security. SerializablePermission	deny *	This ensures that a resource adapter cannot subclass <code>ObjectOutputStream</code> or <code>ObjectInputStream</code> to override the default serialization or deserialization of objects or to substitute one object for another during serialization or deserialization.

18.3 Requirements

A resource adapter provider must ensure that resource adapter code does not conflict with the default security permission set. By ensuring this, a resource adapter can be deployed and run in any application server without execution or manageability problems.

If a resource adapter needs security permissions other than those specified in the default set, it must describe such requirements in the XML deployment descriptor using the `security-permission` element.

A deployment descriptor-based specification of an extended permission set for a resource adapter allows the deployer to analyze the security implications of the extended permission set and make a deployment decision accordingly. An application server must be capable of deploying a resource adapter with the default permission set.

18.3.1 Example

The resource adapter implementation creates a `java.net.Socket` and retrieves the hostname using the `getHostName` method in `java.net.InetAddress`.

TABLE 18-2 Methods and Security Permissions required

Method	Security Manager Method Called	Permission
<code>java.net.Socket</code> <code>Socket(...)</code>	<code>checkConnect({host}, {port})</code>	<code>java.net.SocketPermission</code> <code>"{host}:{port}", "connect"</code>
<code>java.net.InetAddress</code> <code>s</code> <code>public String</code> <code>getHostName()</code>	<code>checkConnect({host}, -1)</code>	<code>java.net.SocketPermission</code> <code>"{host}", "resolve"</code>

The default `SocketPermission`, as specified in Table 18-1, is `grant connect` and `deny rest`. This means that if resource adapter uses the default permission set, the first method `Socket(...)` will be allowed while the second method `InetAddress.getHostName` is disallowed.

The resource adapter needs to explicitly request security permission for the `InetAddress.getHostName` method in the `security-permission-spec` element of its XML deployment descriptor. The following is an example of allowing additional security permissions:

```
<security-permission-spec>
  grant {
    permission java.net.SocketPermission *, "resolve";
  };
</security-permission-spec>
```

18.4 Privileged Code

A resource adapter runs in its own protection domain as identified by its code source and security permission set. For the resource adapter to be allowed to perform a secured action, such as writing a file, it must have been granted permission for that particular action.

Resource adapter code is considered system code which may require more security permissions than the calling application component code. For example, when an application component calls a resource adapter method to execute a function call on the underlying EIS instance, the resource adapter code may need more security permissions than allowed to the calling component, such as the ability to create a thread.

The Java 2 security architecture requires that whenever a system resource access or any secured action is attempted, all code traversed by the current execution thread up to that point must have the necessary permissions for the system resource access, unless some code on the thread has been marked as privileged. Refer to <http://java.sun.com/products/jdk/1.4/docs/guide/security/doprivileged.html>.

To support such scenarios, the resource adapter code should use the `privileged` code feature in the Java security architecture. This enables the resource adapter code to temporarily perform more secured actions than are available directly to the application code calling the resource adapter.

18.4.1 Example

A resource adapter from Wombat Inc. packaged in the `wombat.rar` file contains the following permission specification:

```
<security-permission>
  <security-permission-spec>
    grant {
      permission java.io.FilePermission
        "${user.home}${file.separator}trace{file.separator}-",
        "read,write,delete";
    };
  </security-permission-spec>
</security-permission>
```

During resource adapter deployment, the application server processes this security-permission-spec and grants the necessary permissions to the wombat.rar code base. This is an implementation-specific mechanism and not prescribed by the specification. As an example, the application server may append these permissions to the java.policy file or some implementation-specific policy file, and this may involve manual intervention.

```
// application code
...
WombatConnectionFactory wcf =
    (WombatConnectionFactory)
jndi.lookup("WombatConnectionFactory");
WombatConnection wc = wcf.getConnection(..);
doWork(wc); // calls into resource adapter code
// resource adapter implementation of WombatConnection
...
AccessController.doPrivileged(new PrivilegedAction() {
    public Object run() {
        // privileged code goes here, for example:
        File file = File.createNewFile();
        writeTraceInfoToFile(file);
        return null; // nothing to return
    }
});
```

In addition to specifying these required permissions, the resource adapter must also use `doPrivileged` blocks at strategic locations in its code to prevent the permission checking from reaching the application code or the application server code. The `doPrivileged` block allows the `AccessController` to temporarily grant the necessary permissions to the resource adapter code and to stop checking the rest of the call stack. This allows the resource adapter code to be unaffected by the calling application code's security permission restrictions.

Exceptions

This chapter specifies standard exceptions that identify error conditions which may occur as part of the connector architecture.

The connector architecture defines two classes of exceptions:

- **System Exceptions**—Indicate an unexpected error condition that occurs as part of an invocation of a method defined in the system contracts. For example, system exceptions are used to indicate transaction management-related errors. A system exception is targeted for handling by an application server or resource adapter, depending on who threw the exception, and may not be reported in its original form directly to an application component.
- **Application Exceptions**—Thrown when an application component accesses an EIS resource. For example, an application exception may indicate an error in the execution of a function on a target EIS. These exceptions are meant to be handled directly by an application component.

The connector architecture defines the `javax.resource.ResourceException` class as the root of the system exception hierarchy. The `ResourceException` class extends the `java.lang.Exception` class and is a checked exception.

The `javax.resource.ResourceException` is also the root of the application exception hierarchy for CCI. Application level exceptions are specified in more detail in the API documentation for CCI.

Note, an extended implementation of an exception type provided by a resource adapter may override the `getLocalizedMessage` method to provide a localized message.

19.1 ResourceException

A `ResourceException` provides the following information:

- A resource adapter-specific string describing the error. This string is a standard Java exception message and is available through the `getMessage` method.
- A resource adapter-specific error code that identifies the error condition represented by the `ResourceException`.
- A reference to another exception. Often a `ResourceException` results from a lower-level problem. If appropriate, a lower-level exception, such as `java.lang.Exception` or any derived exception type, may be linked to a `ResourceException` instance.

19.2 System Exceptions

The connector architecture requires that methods, as part of a system contract implementation, use the checked `ResourceException` and other standard exceptions derived from it to indicate system-level error conditions. Using checked exceptions leads to a strict enforcement of the contract for throwing and catching system exceptions and dealing with error conditions.

In addition, a method implementation may use `java.lang.RuntimeException` or any derived exception to indicate runtime error conditions of varying severity levels. Using unchecked exceptions to indicate important system-level error conditions is not recommended for an implementation of system contracts.

If a method needs to indicate a serious error condition that it does not want the caller to catch, the method should use `java.lang.Error` to indicate such conditions. A method is not required to declare in its throws clause any subclasses of `Error` that may be thrown but not caught during the execution of the method, since these errors are abnormal conditions that should never occur.

19.2.1 Exception Hierarchy

The `ResourceException` represents a generic form of exception. A derived exception represents a specific class of error conditions. This design enables the method invocation code to catch a class of error conditions based on the exception type and to handle error conditions appropriately.

The following exceptions are derived from `ResourceException` to indicate more specific classes of system error conditions:

- `javax.resource.spi.SecurityException`: A `SecurityException` indicates error conditions related to the security contract between an application server and resource adapter. The common error conditions represented by this exception are:

- Invalid security information, represented by a `Subject` instance, passed across the security contract. For example, credentials may have expired or be in an invalid format.
- Lack of support for a specific security mechanism in an EIS or resource adapter.
- Failure to create a connection to an EIS because of failed authentication or authorization.
- Failure to authenticate a resource principal to an EIS or failure to establish a secure association with an underlying EIS instance.
- Access control exception indicating that a requested access to an EIS resource or a request to create a new connection has been denied.
- `javax.resource.spi.LocalTransactionException`: A `LocalTransactionException` represents various error conditions related to the local transaction management contract. The JTA specification specifies the `javax.transaction.xa.XAException` class for exceptions related to an `XAResource`-based transaction management contract. The `LocalTransactionException` is used for the local transaction management contract to indicate the following types of error conditions:
 - Invalid transaction context when a transaction operation is executed. For example, calling the `LocalTransaction.commit` method without an active local transaction is an error condition.
 - Transaction is rolled back instead of being committed in the `LocalTransaction.-commit` method.
 - Attempt to start a local transaction from the same thread on a `ManagedConnection` instance that is already associated with an active local transaction.
 - All resource adapter or resource manager-specific error conditions related to local transaction management. Examples are violation of integrity constraints, deadlock detection, communication failure during transaction completion, or any retry requirement.
- `javax.resource.spi.ResourceAdapterInternalException`: This exception indicates all system-level error conditions related to a resource adapter. The common error conditions indicated by this exception type are:
 - Invalid configuration of the `ManagedConnectionFactory` for creating a new physical connection. An example is an invalid server name for a target EIS instance.
 - Failure to create a physical connection to a EIS instance due to a communication protocol error or a resource adapter implementation-specific error.
 - Error conditions internal to a resource adapter implementation.
- `javax.resource.spi.EISSystemException`: An `EISSystemException` is used to indicate any EIS-specific system-level error conditions. Examples of common error conditions are failure or inactivity of an EIS instance, communication failure, and an EIS-specific error during the creation of a physical connection.

- `javax.resource.spi.ApplicationServerInternalException`: This exception is thrown by an application server to indicate error conditions specific to an application server. Example error conditions are: errors related to an application server configuration or implementation of mechanisms internal to an application server, such as connection pooling and thread management.
- `javax.resource.spi.ResourceAllocationException`: This exception is thrown by an application server or resource adapter to indicate a failure to allocate system resources, such as threads and physical connections. An example is an error condition that results when an upper bound is reached for the maximum number of physical connections that can be managed by an application server-specific connection pool.
- `javax.resource.spi.IllegalStateException`: This exception is thrown from a method if the invoked code, either the resource adapter or the application server for system contracts, is in an illegal or inappropriate state for the method invocation.
- `javax.resource.NotSupportedException`: This exception is thrown to indicate that invoked code, either the resource adapter or the application server for system contracts, cannot execute an operation because the operation is not a supported feature. For example, if the transaction support level for a resource adapter is `NoTransaction`, an invocation of the `ManagedConnection.getXAResource` method throws a `NotSupportedException` exception.
- `javax.resource.spi.CommException`: This exception indicates errors related to failed or interrupted communication with an EIS instance. Examples of common error conditions represented by this exception type include communication protocol errors and invalidated connections due to server failure.
- `javax.resource.spi.InvalidPropertyException`: This exception is thrown to indicate invalid configuration property settings.
- `javax.resource.spi.UnavailableException`: This exception is thrown to indicate that a service is unavailable.

19.3 Work Exceptions

These exceptions are thrown by an application server to report error conditions related to the work management contract.

- `javax.resource.spi.work.WorkException`: A common base class for all Work processing related exceptions.
- `javax.resource.spi.work.WorkRejectedException`: This exception is thrown to indicate that a submitted `Work` instance has been rejected. The rejection may be due to internal factors or start timeout expiration.
- `javax.resource.spi.work.WorkCompletedException`: This exception is thrown to indicate that a submitted `Work` instance has completed with an exception.

19.4 Additional Exceptions

The JTA specification (Section 2., “Java Transaction API (JTA) specification, version 1.0.1B” on page F-1) specifies the `javax.transaction.xa.XAException` class for exceptions related to the `XAResource`-based transaction management contract.

Projected Items

The following are some of the features planned for future versions of this specification:

- **Security Inflow:** This would allow an EIS to propagate a security context to the application server to be used while accessing application components.
- **Common Client Interface:** The CCI may become required as part of a future version of the connector architecture. The CCI may also be extended to include support for XML, type mapping, and metadata facility.

Previous Version DTDs

This appendix contains Document Type Definitions (DTDs) for Deployment Descriptors from previous versions of the J2EE Connector Architecture specification. All J2EE products are required to support these DTDs as well as the XML Schema Definition specified in this version of the specification. This ensures that resource adapters written to previous versions of this specification can be deployed on products supporting the current version of this specification.

A.1 J2EE Connector Architecture 1.0 Resource Adapter XML DTD

This section specifies the XML DTD for the 1.0 resource adapter deployment descriptor. The comments in the DTD specify additional requirements for syntax and semantics that cannot be specified by the DTD mechanism.

A resource adapter, or an application server on behalf of a resource adapter, may specify additional deployment information beyond the standard deployment descriptor. The additional information should be stored in a separate file and should refer to the standard deployment descriptor.

A resource adapter is not allowed to add any non-standard information into a standard deployment descriptor.

CODE EXAMPLE A-1 Connector Architecture 1.0 Resource Adapter DTD

```
<!--  
This is the XML DTD for the Connector 1.0 deployment descriptor.  
All Connector 1.0 deployment descriptors must include a DOCTYPE of  
the following form:
```

CODE EXAMPLE A-1 Connector Architecture 1.0 Resource Adapter DTD

```
<!DOCTYPE connector PUBLIC "-//Sun Microsystems, Inc.//DTD
Connector 1.0//EN" "http://java.sun.com/dtd/connector_1_0.dtd">

-->

<!--
The following conventions apply to all J2EE deployment descriptor
elements unless indicated otherwise.

- In elements that contain PCDATA, leading and trailing whitespace
  in the data may be ignored.

- In elements whose value is an "enumerated type", the value is
  case sensitive.

- In elements that specify a pathname to a file within the same
  JAR file, relative filenames (i.e., those not starting with "/")
  are considered relative to the root of the JAR file's namespace.
  Absolute filenames (i.e., those starting with "/") also specify
  names in the root of the JAR file's namespace. In general,
relative
  names are preferred. The exception is .war files where absolute
  names are preferred for consistency with the Servlet API.
-->

<!--
The connector element is the root element of the deployment
descriptor
for the resource adapter. This element includes general information
- vendor
name, version, specification version supported, icon - about the
resource adapter module. It also includes information specific to the
implementation of the resource adapter library as specified through
the element resourceadapter.
-->
<!ELEMENT connector (display-name?, description?, icon?, vendor-
name,
spec-version, eis-type, version, license?, resourceadapter>

<!--
The element authentication-mechanism specifies an authentication
mechanism
supported by the resource adapter. Note that this support is for
the resource adapter and not for the underlying EIS instance. The
```


CODE EXAMPLE A-1 Connector Architecture 1.0 Resource Adapter DTD

optional description specifies any resource adapter specific requirement for the support of security contract and authentication mechanism.

Note that BasicPassword mechanism type should support the javax.resource.spi.security.PasswordCredential interface. The Kerbv5 mechanism type should support the javax.resource.spi.security.Generic-Credential interface.

Used in: resourceadapter

-->

```
<!ELEMENT authentication-mechanism (
description?, authentication-mechanism-type, credential-
interface)>
```

<!--

The element authentication-mechanism-type specifies type of an authentication mechanism.

The example values are:

```
<authentication-mechanism-type>BasicPassword
    </authentication-mechanism-type>
    <authentication-mechanism-type>Kervb5
    </authentication-mechanism-type>
```

Any additional security mechanisms are outside the scope of the Connector architecture specification.

Used in: authentication-mechanism

-->

```
<!ELEMENT authentication-mechanism-type (#PCDATA)>
```

<!--

The element config-property contains a declaration of a single configuration property for a ManagedConnectionFactory instance.

Each ManagedConnectionFactory instance creates connections to a specific EIS instance based on the properties configured on the ManagedConnectionFactory instance. The configurable properties are specified only once in the deployment descriptor, even though a resource adapter can be used to configure multiple ManagedConnection-Factory instances (that create connections to different instances of

CODE EXAMPLE A-1 Connector Architecture 1.0 Resource Adapter DTD

the same EIS).

The declaration consists of an optional description, name, type and an optional value of the configuration property. If the resource

adapter provider does not specify a value than the deployer is responsible for providing a valid value for a configuration property.

Any bounds or well-defined values of properties should be described in the description element.

Used in: resourceadapter

-->

```
<!ELEMENT config-property (description?, config-property-name,
config-property-type, config-property-value?)>
```

<!--

The element config-property-name contains the name of a configuration property.

The connector architecture defines a set of well-defined properties all of type java.lang.String. These are as follows:

```
<config-property-name>ServerName</config-property-name>
<config-property-name>PortNumber</config-property-name>
<config-property-name>UserName</config-property-name>
<config-property-name>Password</config-property-name>
<config-property-name>ConnectionURL</config-property-name>
```

A resource adapter provider can extend this property set to include properties specific to the resource adapter and its underlying EIS.

Used in: config-property

Example: <config-property-name>ServerName</config-property-name>

-->

```
<!ELEMENT config-property-name (#PCDATA)>
```

<!--

The element config-property-type contains the fully qualified Java type of a configuration property as required by ManagedConnectionFactory instance.

The following are the legal values of config-property-type:

```
java.lang.Boolean, java.lang.String, java.lang.Integer,
java.lang.Double, java.lang.Byte, java.lang.Short,
```

CODE EXAMPLE A-1 Connector Architecture 1.0 Resource Adapter DTD

```
java.lang.Long, java.lang.Float, java.lang.Character

Used in: config-property

Example: <config-property-type>java.lang.String</config-property-
type>
-->
<!ELEMENT config-property-type (#PCDATA)>

<!--
The element config-property-value contains the value of a
configuration
entry.

Used in: config-property

Example: <config-property-value>WombatServer</config-property-
value>
-->
<!ELEMENT config-property-value (#PCDATA)>

<!--
The element connection-impl-class specifies the fully-qualified
name of the Connection class that implements resource adapter
specific Connection interface.

Used in: resourceadapter

Example: <connection-impl-class>com.wombat.ConnectionImpl
        </connection-impl-class>
-->
<!ELEMENT connection-impl-class (#PCDATA)>

<!--
The element connection-interface specifies the fully-qualified
name of the Connection interface supported by the resource
adapter.

Used in: resourceadapter

Example: <connection-interface>javax.resource.cci.Connection
        </connection-interface>
-->
<!ELEMENT connection-interface (#PCDATA)>

<!--
```

CODE EXAMPLE A-1 Connector Architecture 1.0 Resource Adapter DTD

The element `connectionfactory-impl-class` specifies the fully-qualified name of the `ConnectionFactory` class that implements resource adapter specific `ConnectionFactory` interface.

Used in: `resourceadapter`

Example: `<connectionfactory-impl-class>com.wombat.ConnectionFactoryImpl</connectionfactory-impl-class>`

-->

`<!ELEMENT connectionfactory-impl-class (#PCDATA)>`

<!--

The element `connectionfactory-interface` specifies the fully-qualified name of the `ConnectionFactory` interface supported by the resource adapter.

Used in: `resourceadapter`

Example: `<connectionfactory-interface>com.wombat.ConnectionFactory</connectionfactory-interface>`

OR

`<connectionfactory-interface>javax.resource.cci.ConnectionFactory</connectionfactory-interface>`

-->

`<!ELEMENT connectionfactory-interface (#PCDATA)>`

<!--

The element `credential-interface` specifies the interface that the resource adapter implementation supports for the representation of the credentials. This element should be used by application server to find out the `Credential` interface it should use as part of the security contract.

The possible values are:

`<credential-interface>javax.resource.spi.security.PasswordCredential</credential-interface>`

`<credential-interface>javax.resource.spi.security.GenericCredential</credential-interface>`

CODE EXAMPLE A-1 Connector Architecture 1.0 Resource Adapter DTD

```
Used in: authentication-mechanism
-->
<!ELEMENT credential-interface (#PCDATA)>

<!--
The description element is used to provide text describing the
parent
element. The description element should include any information
that
the component file producer wants to provide to the consumer of the
component file (i.e., to the Deployer). Typically, the tools used
by
the component file consumer will display the description when
processing the parent element that contains the description.

Used in: authentication-mechanism, config-property, connector,
license,
security-permission
-->
<!ELEMENT description (#PCDATA)>

<!--
The display-name element contains a short name that is intended to
be
displayed by tools. The display name need not be unique.

Used in: connector

Example:

<display-name>Employee Self Service</display-name>
-->
<!ELEMENT display-name (#PCDATA)>

<!--
The element eis-type contains information about the type of the
EIS. For example, the type of an EIS can be product name of EIS
independent of any version info.

This helps in identifying EIS instances that can be used with
this resource adapter.

Used in: connector
-->
<!ELEMENT eis-type (#PCDATA)>

<!--
```

CODE EXAMPLE A-1 Connector Architecture 1.0 Resource Adapter DTD

The icon element contains a small icon and large icon element which specify the URIs for a small and a large GIF or JPEG icon image to represent the application in GUI.

Used in: connector

-->

<!ELEMENT icon (small-icon?, large-icon?)>

<!--

The large-icon element contains the name of a file containing a large

(32 x 32) icon image. The file name is a relative path within the component's jar file.

The image may be either in the JPEG or GIF format.

The icon can be used by tools.

Used in: icon

Example:

<large-icon>employee-service-icon32x32.jpg</large-icon>

-->

<!ELEMENT large-icon (#PCDATA)>

<!--

The element license specifies licensing requirements for the resource

adapter module. This element specifies whether a license is required

to deploy and use this resource adapter, and an optional description

of the licensing terms (examples: duration of license, number of connection restrictions).

Used in: connector

-->

<!ELEMENT license (description?, license-required)>

<!--

The element license-required specifies whether a license is required

to deploy and use the resource adapter. This element must be one of the following:

<license-required>true</license-required>

<license-required>>false</license-required>

CODE EXAMPLE A-1 Connector Architecture 1.0 Resource Adapter DTD

```
Used in: license
-->
<!ELEMENT license-required (#PCDATA)>

<!--
The element managedconnectionfactory-class specifies the fully
qualified
name of the Java class that implements the
javax.resource.spi.Managed-
ConnectionFactory interface. This Java class is provided as part of
resource adapter's implementation of connector architecture
specified
contracts.

Used in: resourceadapter

Example:
    <managedconnectionfactory-
class>com.wombat.ManagedConnectionFactoryImpl
    </managedconnectionfactory-class>
-->
<!ELEMENT managedconnectionfactory-class (#PCDATA)>

<!--
The element reauthentication-support specifies whether the
resource
adapter implementation supports re-authentication of existing
Managed-
Connection instance. Note that this information is for the resource
adapter implementation and not for the underlying EIS instance.

This element must be one of the following:
    <reauthentication-support>true</reauthentication-support>
    <reauthentication-support>false</reauthentication-support>
Used in: resourceadapter
-->
<!ELEMENT reauthentication-support (#PCDATA)>

<!--
The element resourceadapter specifies information about the
resource
adapter. The information includes fully-qualified names of
class/interfaces required as part of the connector architecture
specified contracts, level of transaction support provided,
configurable properties for ManagedConnectionFactory instances,
one or more authentication mechanisms supported and additional
```

CODE EXAMPLE A-1 Connector Architecture 1.0 Resource Adapter DTD

required security permissions.

If there is no authentication-mechanism specified as part of resource adapter element then the resource adapter does not support any standard security authentication mechanisms as part of security contract. The application server ignores the security part of the system contracts in this case.

Used in: connector

-->

```
<!ELEMENT resourceadapter (  
managedconnectionfactory-class, connectionfactory-interface,  
connectionfactory-impl-class, connection-interface,  
connection-impl-class, transaction-support, config-property*,  
authentication-mechanism*, reauthentication-support, security-  
permission*  
)>
```

<!--

The element security permission specifies a security permission that is required by the resource adapter code.

The security permission listed in the deployment descriptor are ones that are different from those required by the default permission set as specified in the connector specification. The optional description can mention specific reason that resource adapter requires a given security permission.

Used in: resourceadapter

-->

```
<!ELEMENT security-permission (description?, security-permission-  
spec)>
```

<!--

The element permission-spec specifies a security permission based on the Security policy file syntax. Refer the following URL for Sun's implementation of security permission specification:

<http://java.sun.com/products/jdk/1.3/docs/guide/security/PolicyFiles.html#FileSyntax>

Used in: security-permission

CODE EXAMPLE A-1 Connector Architecture 1.0 Resource Adapter DTD

```
-->
<!ELEMENT security-permission-spec (#PCDATA)>

<!--
The small-icon element contains the name of a file containing a
small
(16 x 16) icon image. The file name is a relative path within the
component's jar file.

The image may be either in the JPEG or GIF format.
The icon can be used by tools.

Used in: icon

Example:

<small-icon>employee-service-icon16x16.jpg</small-icon>
-->
<!ELEMENT small-icon (#PCDATA)>

<!--
The element spec-version specifies the version of the connector
architecture specification that is supported by this resource
adapter. This information enables deployer to configure the
resource
adapter to support deployment and runtime requirements of the
corresponding connector architecture specification.

Used in: connector

Example:
    <spec-version>1.0</spec-version>
-->
<!ELEMENT spec-version (#PCDATA)>

<!--
The transaction-support element specifies the level of transaction
support provided by the resource adapter.
The value of transaction-support must be one of the following:
<transaction-support>NoTransaction</transaction-support>
<transaction-support>LocalTransaction</transaction-support>
<transaction-support>XATransaction</transaction-support>

Used in: resourceadapter
-->
<!ELEMENT transaction-support (#PCDATA)>
```

CODE EXAMPLE A-1 Connector Architecture 1.0 Resource Adapter DTD

```
<!--
The element vendor-name specifies the name of resource adapter
provider
vendor.

Used in: connector

Example:
  <vendor-name>Wombat Corp.</vendor-name>
-->
<!ELEMENT vendor-name (#PCDATA)>

<!--
The element version specifies a string-based version of the
resource adapter from the resource adapter provider.

Used in: connector

Example:
  <version>1.0</version>
-->
<!ELEMENT version (#PCDATA)>

<!--
The ID mechanism is to allow tools that produce additional
deployment
information (i.e., information beyond the standard deployment
descriptor information) to store the non-standard information in a
separate file, and easily refer from these tool-specific files to
the
information in the standard deployment descriptor.

Tools are not allowed to add the non-standard information into the
standard deployment descriptor.
-->

<!ATTLIST authentication-mechanism id ID #IMPLIED>
<!ATTLIST authentication-mechanism-type id ID #IMPLIED>
<!ATTLIST config-property id ID #IMPLIED>
<!ATTLIST config-property-name id ID #IMPLIED>
<!ATTLIST config-property-type id ID #IMPLIED>
<!ATTLIST config-property-value id ID #IMPLIED>
<!ATTLIST connection-impl-class id ID #IMPLIED>
<!ATTLIST connection-interface id ID #IMPLIED>
<!ATTLIST connectionfactory-impl-class id ID #IMPLIED>
<!ATTLIST connectionfactory-interface id ID #IMPLIED>
<!ATTLIST connector id ID #IMPLIED>
```

CODE EXAMPLE A-1 Connector Architecture 1.0 Resource Adapter DTD

```
<!--ATTLIST credential-interface id ID #IMPLIED>
<!--ATTLIST description id ID #IMPLIED>
<!--ATTLIST display-name id ID #IMPLIED>
<!--ATTLIST eis-type id ID #IMPLIED>
<!--ATTLIST icon id ID #IMPLIED>
<!--ATTLIST large-icon id ID #IMPLIED>
<!--ATTLIST license id ID #IMPLIED>
<!--ATTLIST license-required id ID #IMPLIED>
<!--ATTLIST managedconnectionfactory-class id ID #IMPLIED>
<!--ATTLIST reauthentication-support id ID #IMPLIED>
<!--ATTLIST resourceadapter id ID #IMPLIED>
<!--ATTLIST security-permission id ID #IMPLIED>
<!--ATTLIST security-permission-spec id ID #IMPLIED>
<!--ATTLIST small-icon id ID #IMPLIED>
<!--ATTLIST spec-version id ID #IMPLIED>
<!--ATTLIST transaction-support id ID #IMPLIED>
<!--ATTLIST vendor-name id ID #IMPLIED>
<!--ATTLIST version id ID #IMPLIED>
```


Activation Configuration for Message Inflow to JMS Endpoints

B.1 Introduction

The following conventions are recommended to support message inflow to a JMS endpoint, a message-driven bean application, using a Connector API- based JMS resource adapter. The following describes the `ActivationSpec` JavaBean used by a JMS resource adapter, the mapping of an endpoint's activation configuration elements, specified in the endpoint deployment descriptor, to the `ActivationSpec` JavaBean properties, and how to map earlier versions of the endpoint's deployment descriptor information to the `ActivationSpec` JavaBean properties. The JMS `ActivationSpec` JavaBean properties should be standardized by the JMS community.

B.2 JMS `ActivationSpec` JavaBean

The `ActivationSpec` JavaBean is used by the endpoint deployer to provide configuration information when an endpoint is deployed and activated. JMS providers are strongly encouraged to include the following elements in their `ActivationSpec` JavaBean implementation.

The recommended properties for JMS `ActivationSpec` JavaBean are: *destination* (required), *destinationType* (required), *messageSelector*, *acknowledgeMode*, *subscriptionDurability*, *clientId*, *subscriptionName*.

The `ActivationSpec` JavaBean may contain more properties than the recommended set, so that an endpoint deployer may configure additional properties during endpoint deployment.

B.2.1 JMS `ActivationSpec` JavaBean Properties

The following is a description of the recommended set of JMS `ActivationSpec` JavaBean properties:

B.2.1.1 `destination`

In general, a destination is a logical link between a message producer and a consumer. A message producer produces messages addressed to a destination, and a message consumer consumes messages from a chosen destination.

A destination could be represented by a JMS provider in different ways. For example, it could be encapsulated as textual data (this may be specified by endpoint developer or assembler in the endpoint deployment descriptor) or as a private object available only at deployment time. Hence, it may not always be specified in an endpoint deployment descriptor (since deployment descriptor entries contain textual data only).

However, during endpoint deployment JMS `ActivationSpec` JavaBean requires a value for the *destination* property. While configuring the JMS `ActivationSpec` JavaBean, the endpoint deployer must provide a value for this property (if it is not already present in the endpoint deployment descriptor).

It is recommended that *destination* be specified by a JMS resource adapter provider as a property on the `ActivationSpec` JavaBean that requires a value. This can be specified in the JMS resource adapter deployment descriptor.

The *destination* property value may also be an object that implements the `javax.jms.Destination` interface. In such a case, the resource adapter should provide an administered object that implements the `javax.jms.Destination` interface.

B.2.1.2 `destinationType`

The *destinationType* identifies whether the specified *destination* is a JMS Queue or JMS Topic. The valid values are `javax.jms.Queue` or `javax.jms.Topic`.

This property requires a value to be specified. The endpoint deployer must provide a value for this property (if it is not already present in the endpoint deployment descriptor). It is recommended that *destinationType* be specified by a JMS resource adapter provider as a property on the `ActivationSpec` JavaBean that requires a value. This can be specified in the JMS resource adapter deployment descriptor.

B.2.1.3 `messageSelector`

The value for a *messageSelector* property is a string which is used to select a subset of the available messages. Its syntax is based on a subset of the SQL 92 conditional expression syntax and is described in detail in JMS specification (Section 11., “Java Message Service, version 1.1:” on page F-1). Specifying a value for the *messageSelector* property on the `ActivationSpec` JavaBean is optional.

B.2.1.4 `acknowledgeMode`

The *acknowledgeMode* indicates how a message received by an endpoint (message-driven bean) should be acknowledged. The EJB specification (Section 1., “Enterprise JavaBeans (EJB™) specification, version 2.1:” on page F-1) describes the restrictions for message acknowledgements under different transactional modes. The valid values are `Auto-acknowledge` or `Dups-ok-acknowledge`.

Specifying a value for the *acknowledgeMode* property on the `ActivationSpec` JavaBean is optional. If no value is specified, `Auto-acknowledge` is assumed.

B.2.1.5 `subscriptionDurability`

This property applies only to endpoints (message-driven beans) that receive messages published to a `JMS Topic`. It indicates whether messages should be retained by the message provider while the endpoint (message-driven bean) is not available, and subsequently delivered when the endpoint becomes available.

The valid values are `Durable` or `NonDurable`. If the value is set to `Durable`, messages will be retained and redelivered by the provider. If the value is set to `NonDurable`, messages will be not retained or redelivered by the provider.

Specifying a value for the *subscriptionDurability* property on the `ActivationSpec` JavaBean is optional. If no value is specified, `NonDurable` is assumed.

B.2.1.6 `clientId`

The value specified for a *clientId* property is a unique identifier that identifies a client (message consumer) within the context of a specific message provider. This is used by the message provider for durable subscriptions, in order to uniquely identify a client consumer and redeliver persisted messages upon client availability.

It is the responsibility of the endpoint deployer to specify a unique client identifier in the `JMS ActivationSpec` JavaBean for durable JMS Topic subscriptions, that uniquely identifies a client consumer within the context of a message provider.

Specifying a value for the *clientId* property on the `ActivationSpec` JavaBean is required for durable JMS Topic subscriptions.

B.2.1.7 `subscriptionName`

This applies only to endpoints (message-driven beans) that durably subscribe to a JMS Topic. The value for the *subscriptionName* property is a string that is combined with the client identifier to uniquely identify the message consumer within the context of a message provider.

It is the responsibility of the endpoint deployer to specify a subscription name in the `JMS ActivationSpec` JavaBean for durable JMS Topic subscriptions.

Specifying a value for the *subscriptionName* property on the `ActivationSpec` JavaBean is required for durable JMS Topic subscriptions.

B.2.2 JMS `ActivationSpec` JavaBean Property Values

The following is a recommended set of acceptable values for the various JMS `ActivationSpec` properties.

TABLE B-1 JMS `ActivationSpec` JavaBean Property Values

| JMS <code>ActivationSpec</code> property name | JMS <code>ActivationSpec</code> property values |
|---|--|
| <code>destinationType</code> | <code>javax.jms.Queue</code> or <code>javax.jms.Topic</code> |
| <code>acknowledgeMode</code> | Auto-acknowledge (default) or Dups-ok-acknowledge |
| <code>subscriptionDurability</code> | NonDurable (default) or Durable |

B.3 JMS Endpoint with EJB 2.1 Activation Configuration Elements

EJB 2.1 specification (Section 1., “Enterprise JavaBeans (EJB™) specification, version 2.1:” on page F-1) allows a JMS endpoint developer or assembler to specify activation configuration elements in the endpoint deployment descriptor.

The following is a recommended set of endpoint activation configuration elements in the endpoint deployment descriptor (EJB 2.1 style) for endpoints which implement the `javax.jms.MessageListener`. Further, the mapping between each recommended JMS endpoint activation configuration element to a corresponding JMS `ActivationSpec` JavaBean property is described.

TABLE B-2 Activation Configuration Elements (EJB 2.1) and `ActivationSpec` JavaBean Properties

| Activation configuration elements in endpoint deployment descriptor (EJB 2.1 style) | JMS <code>ActivationSpec</code> properties |
|---|--|
| <code>destination</code> | <code>destination</code> (required) |
| <code>destinationType</code> | <code>destinationType</code> (required) |
| <code>messageSelector</code> | <code>messageSelector</code> |
| <code>acknowledgeMode</code> | <code>acknowledgeMode</code> |
| <code>subscriptionDurability</code> | <code>subscriptionDurability</code> |
| <code>clientId</code> | <code>clientId</code> |
| <code>subscriptionName</code> | <code>subscriptionName</code> |

Note the one-to-one correspondence between the endpoint activation configuration elements and the JMS `ActivationSpec` JavaBean property names. This allows automatic merging of the activation configuration element values with the corresponding entries in the JMS `ActivationSpec` JavaBean, while configuring the JMS `ActivationSpec` JavaBean during endpoint deployment.

The JMS `ActivationSpec` JavaBean may contain more properties than the above recommended set, that an endpoint deployer may configure during endpoint deployment.

B.4 JMS Endpoint with EJB 2.0 Deployment Descriptor Elements

The EJB 2.0 specification provides endpoint deployment descriptor elements which allows a JMS endpoint developer or assembler to specify JMS specific configuration information in the endpoint deployment descriptor. These elements are contained in the `message-driven` element as specified in the EJB 2.0 specification.

The following is a recommended mapping between relevant EJB 2.0 endpoint deployment descriptor elements and JMS `ActivationSpec` JavaBean properties.

TABLE B-3 EJB 2.0 Deployment Descriptor Elements and JMS `ActivationSpec` JavaBean Properties

| EJB 2.0 deployment descriptor elements | JMS <code>ActivationSpec</code> properties |
|--|--|
| <code>destination-type</code> | <code>destinationType</code> |
| <code>message-selector</code> | <code>messageSelector</code> |
| <code>acknowledge-mode</code> | <code>acknowledgeMode</code> |
| <code>subscription-durability</code> | <code>subscriptionDurability</code> |

Caching Manager

The following section describes how the connector architecture supports caching.

This section serves as a brief introduction to the caching support in the connector architecture. A future version of the connector architecture will address this issue in detail.

C.1 Overview

The connector architecture provides a standard way of extending an application manager for plugging in caching managers. A caching manager may be provided by a third party vendor or a resource adapter provider.

A caching manager manages cached state for application components while they access EISs across transactions.

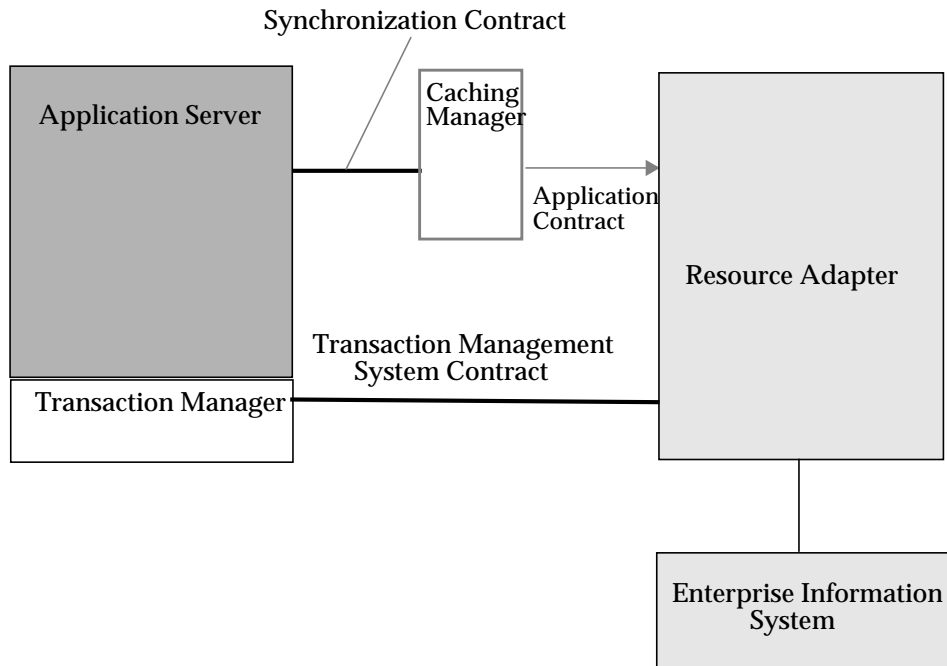
A caching manager is provided above a resource adapter. An application component may access a resource manager either through a caching manager (thereby maintaining a cached state across application requests) or directly through the resource adapter with no caching involved.

The `XAResource` based transaction management contract enables an external transaction manager to control and coordinate transactions across multiple resource managers. A caching manager (provided above the resource adapter) needs to be synchronized relative to the transaction coordination flow (defined by the JTA `XAResource` interface) on the underlying resource manager. This leads to a requirement for a synchronization contract between the application server and caching manager.

The connector architecture defines a standard synchronization contract between the application server and caching manager. The caching manager uses the synchronization notifications to manage its cached state and to flush it to the

resource adapter. The resource adapter then takes the responsibility of managing its recoverable units of work and participates in the transaction coordination protocol from the transaction manager.

FIGURE C-1 Synchronization Contract between Caching Manager and Application Server



The above diagram shows a caching manager layered above a resource adapter. The contract between caching manager and resource adapter is specific to a resource adapter.

C.2 Synchronization contract

Note – To support a caching manager as a standard extension to the application server, additional contracts between the application server and the caching manager are required. This version of the specification introduces only the synchronization contract.

This section specifies the synchronization contract between the application server and the caching manager.

C.2.1 Interface

Each caching manager implements the `javax.transaction.Synchronization` interface. A caching manager registers its `Synchronization` instance with the application server when it is configured with the application server.

The caching manager receives synchronization notifications only for transactions managed by an external transaction manager. In the case of transactions managed internally by a resource manager, the resource adapter and caching manager define their own implementation-specific mechanisms for synchronizing caches.

The `Synchronization.beforeCompletion` method is called prior to the start of the two-phase commit transaction completion process. This call executes in the same transaction context of the caller who initiated the transaction completion. The caching manager uses this notification to flush its cached state to the resource adapter.

The `Synchronization.afterCompletion` method is called after the transaction has completed. The status of transaction completion is passed in as a parameter. The caching manager uses this notification to do any cache cleanups if a rollback has occurred.

C.2.2 Implementation

The caching manager must support the `javax.transaction.Synchronization` interface. If the caching manager implements the `Synchronization` interface and registers it with the application server, then the application server must invoke the `beforeCompletion` and `afterCompletion` notifications.

The application server is responsible for ensuring that synchronization notifications are delivered first to the application components (that have expressed interest in receiving synchronization notification through their respective application component and container-specific mechanisms) and then to the caching managers that implement the `Synchronization` interface.

Security Scenarios

The following section describes various scenarios for EIS integration. These scenarios focus on security aspects of the connector architecture.

Note that these scenarios establish the requirements to be addressed by the connector architecture. Chapter 8, “Security Architecture” and Chapter 9, “Security Contract” specify the requirements that are supported in this version of the specification.

A J2EE application is a multi-tier, web-enabled application that accesses EISs. It consists of one or more application components—EJBs, JSPs, servlets—which are deployed on containers. These containers can be one of the following:

- Web containers that host JSP, servlets, and static HTML pages
- EJB containers that host EJB components
- Application client containers that host standalone application clients

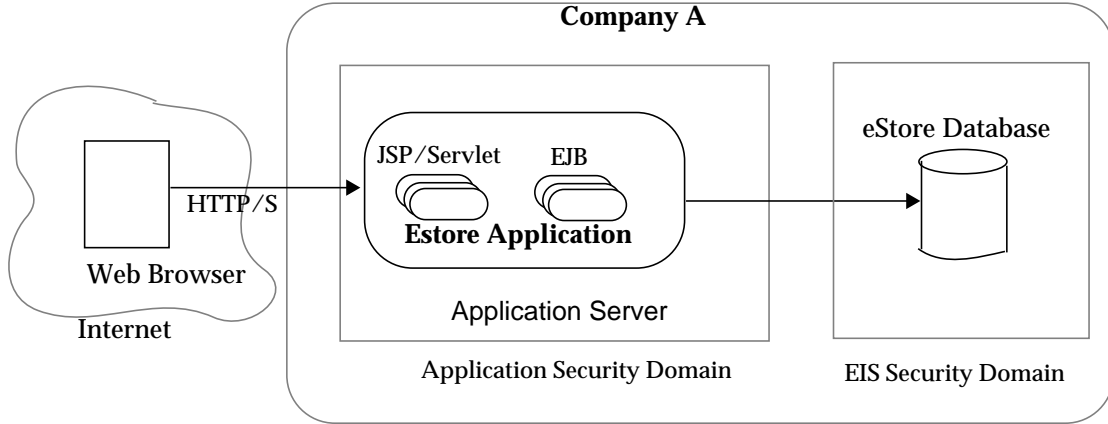
In the following scenarios, the description of the architecture and security environments are illustrative in scope.

D.1 EStore Application

Company A has an eStore application based on the J2EE platform. The eStore application is composed of EJBs and JSP/servlets; together they collaborate to provide the overall functionality of the application. The application also utilizes an eStore database to store data related to product catalog;, shopping carts; customer registration and profiles; transaction status and records; and order status.

The architecture of this application is illustrated in the following diagram:

FIGURE D-1 Illustrative Architecture of an Estore Application



D.1.1 Scenario

A customer, using a web browser, initiates an e-commerce transaction with the eStore application. The e-commerce transaction consists of a series of customer actions. The customer:

- Browses the catalog
- Makes a selection of products
- Puts the selected products into a shopping cart
- Enters her user name and password to initiate a secure transaction
- Fills in order-related information
- And, finally, places an order

In this scenario, the eStore application stores all persistent information about customers and their transactions in a database.

D.1.2 Security Environment

To support the above interaction scenario, the system administrator configures a unique security domain (with specific security technology and security policies) for the eStore application. A firewall protects this security domain from unauthorized Internet access.

The security domain configuration for the eStore application includes secure web access to the eStore application. Secure web access is set up based on the requirements specified in the J2EE specification. Note that the focus of this section is security related to EIS integration, not on web access security. As a result, this description ignores web access security.

The system administrator sets up a database to manage persistent data for the eStore application. In terms of security, the database system is configured with an independent security domain. This domain has its own set of user accounts, plus its own security policies and mechanisms for authentication and authorization.

The system administrator (or database administrator DBA) creates a unique database account (called `EStoreUser`) to handle database transactions; the database transactions correspond to different customer-driven interactions with the eStore application. He also sets up an additional database account (called `EStoreAdministrator`) to manage the database on behalf of the eStore administrator. This administrative account has a higher level of access privileges.

To facilitate better scaling of the eStore application, the system administrator may choose to set the load balancing of database operations across multiple databases. He may also partition persistent data and transactions across multiple database accounts, based on various performance optimization criteria. These areas are out of the scope for this document.

This scenario deals only with the simple case of a single database and a single user account to handle all database transactions.

D.1.3 Deployment


Note – This document does not address how principal delegation happens between the web and EJB containers. When an EJB instance acquires an EIS connection, a caller principal is associated with the EJB instance. This document does not address determining which caller principal is associated with the EJB instance.

During the deployment of the eStore application, the deployer sets up access control for all authenticated customer accounts—the customer accounts that are driving e-commerce transactions over the web—based on a single role `eStoreUserRole`.

The deployer configures the resource adapter with the security information that is required for the creation of database connections. This security information is the database account `EStoreUser` and its password.

The deployer sets up the resource principal for accessing the database system as illustrated in the FIGURE D-2:

FIGURE D-2 Resource Principal for Estore Application Scenario



The deployment configuration ensures that all database access is always performed under the security context of the database account `EStoreUser`.

All authenticated customers (referred to as `Initiating Principal`) map to a single `EStoreUser` database account. The eStore application uses an implementation-specific mechanism to tie database transactions (performed under a single database account) to the unique identity (social security number or eStore account ID) of the initiating principal. To ensure that database access has been properly authorized, the eStore application also performs access control based on the role of the initiating principal. Because all initiating principals map to a single role, this is in effect a simple case.

This scenario describes an n-to-1 mapping. However, depending on the requirements of an application, the deployer can set the principal mapping to be different from an n-to-1 mapping. For example, the deployer can map each role to a single resource principal, where a role corresponds to an initiating principal. This results in a $[m \text{ principals and } n \text{ roles}]$ to $[p \text{ resource principals}]$ mapping. When doing such principal mapping, the deployer has to ensure not to compromise the access rights of the mapped principals. An illustrative example is:

- User is in administrator role: Principal `EISadmin`
- User is in manager role: Principal `EISmanager`
- User is in employee role: Principal `EISemployee`

D.2 Employee Self Service Application

Company B has developed and deployed an employee self-service (ESS) application based on the J2EE platform. This application supports a web interface to the existing Human Resources (HR) applications, which are supported by the ERP system from Vendor X. The ESS application also provides additional business processes customized to the needs of Company B.

The application tier is composed of EJBs and JSPs that provide the customization of the business processes and support a company-standardized web interface. The ESS application enables an employee (under the roles of Manager, HR manager, and Employee) to perform various HR functions, including personal information management, payroll management, compensation management, benefits administration, travel management, and HR cost planning.

D.2.1 Architecture

The IS department of Company B has deployed its HR ESS application and ERP system in a secure environment on a single physical location. Any access to the HR application is permitted Only legal employees of the organization are permitted access to the HR application. Access is based on the employee's roles and access privileges. In addition, access to the application can only be from within the organization-wide intranet. See FIGURE D-3.

D.2.2 Security Environment

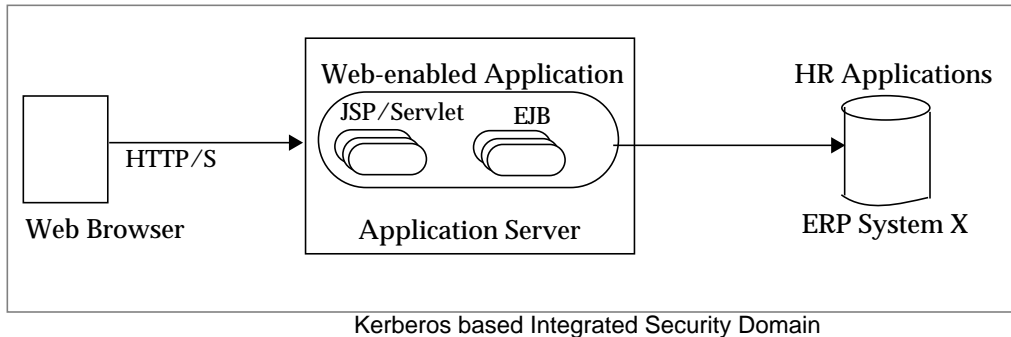
To support the various interaction scenarios related to the ESS application, the system administrator sets up an end-to-end Kerberos-based security domain for this application environment.

Note – The Security policies and mechanisms that are required to achieve this single security domain are technology dependent. Refer to Kerberos V5 specification for more details.

The system administrator configures the security environment to support single sign-on; the user logs on only once and can then access all the services provided by the ESS application and its underlying ERP system. Single sign-on is achieved through the security mechanism and policies specific to the underlying security technology, which in this case is Kerberos.

The ERP system administrator configures all legal employees as valid user accounts in the ERP system. He also must set up various roles (Manager, HRManager, and Employee), default passwords, and access privileges. This security information is kept synchronized with the enterprise-wide directory service, which is used by Kerberos to perform the initial authentication of end-users.

FIGURE D-3 Illustrative Architecture of an Employee Self-service Application
Company B



D.2.3 Deployment

During deployment of the ESS application, the deployer sets a default delegation policy of client impersonation for EIS sign-on. In this case, the application server and ERP system know that it is the initiating principal accessing their respective services and they perform access control based on this knowledge. See FIGURE D-4.

FIGURE D-4 Principal Mapping



In this scenario, both the initiating principal and the resource principal refer to the same principal. This common principal is authenticated using Kerberos and its Kerberos credentials are valid in the security domains of both the application and the ERP system.

The deployer sets up access control for all authenticated employees (initiating principal) based on the configured roles—Manager, HR Manager, and Employee.

If the ERP system does not support Kerberos, then an alternate scenario is utilized. The deployer or application server administrator sets up an automatic mapping of Kerberos credentials (for the initiating principal) to valid credentials (for the same principal) in the security domain of the ERP system. Note that when the ERP system does support Kerberos, the application server performs no credentials mapping.

D.2.4 Scenario

An employee initiates an initial login to his client desktop. He enters his username and password. As part of this initial login, the employee (called initiating principal C) gets authenticated with Kerberos KDC. [Refer to the details for Kerberos KDC authentication in the Kerberos v5 specification.]

After a successful login, the employee starts using his desktop environment. He directs his web browser to the URL for the ESS application deployed on the application server. At this point, the initiating principal C authenticates itself to the application server and establishes a session key with the application server.

The ESS application is set up to impersonate initiating principal C when accessing the ERP system, which is running on another server. Though the application server directly connects to the ERP system, access to the ERP system is requested on behalf of the initiating principal. For this to work, principal C needs to delegate its identity and Kerberos credential to the application server and allow the application server to make requests to the ERP system on C's behalf.

D.3 Integrated Purchasing Application

Company C has an integrated purchasing application that enables an employee to use a web-based interface to perform multiple purchasing transactions. An employee can manage the entire procurement process, from creating a purchase requisition through invoice approval. The purchasing application also integrates with the enterprise's existing financial applications so that the accounting and financial aspects of the procurement business processes can be tracked.

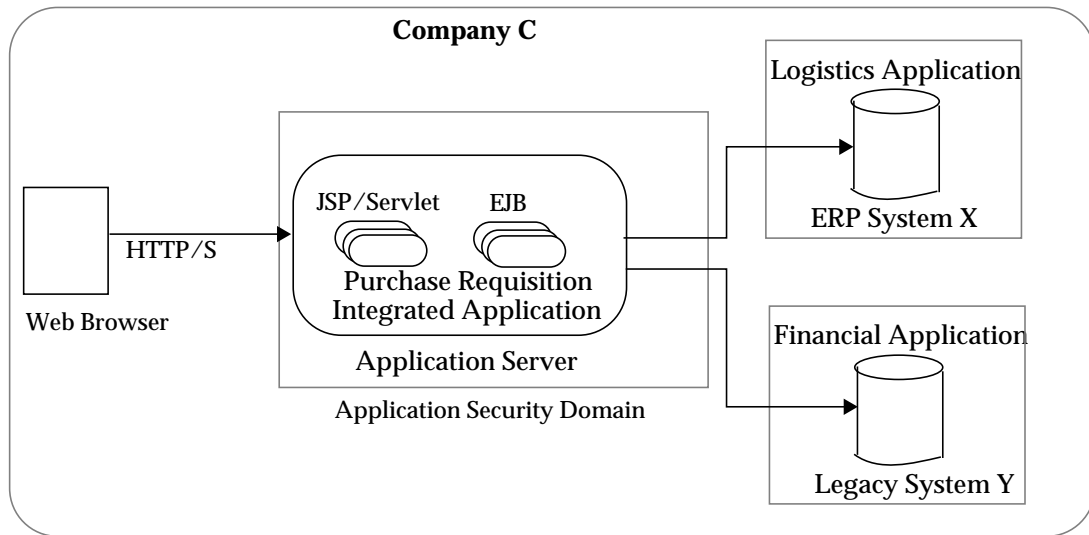
D.3.1 Architecture

FIGURE D-5 illustrates an architecture for this purchasing application. The application has been developed and deployed based on the J2EE platform and is composed of EJBs and JSPs. The EJB components provide the integration across the different applications—the logistics application from a separate vendor (this application

provides integrated purchasing and inventory management functions) and the financial accounting applications (the applications supported by the legacy system from vendor Y).

Company B is a huge decentralized enterprise; its business units and departments are geographically distributed. In this scenario, different IS departments manage ERP system X and legacy system Y. In addition, ERP system X and legacy system Y have been deployed at secured data centers in different geographic locations. Lastly, the integrated purchasing application has been deployed at a geographic location different from both ERP system X and legacy system Y.

FIGURE D-5 Illustrative Architecture of an Integrated Purchasing Application



D.3.2 Security Environment

ERP system X and legacy system Y are also in different security domains; they use different security technologies and have their own specific security policies and mechanisms. The integrated purchasing application is deployed in a security domain that is different from both that of ERP system X and legacy system Y.

To support the various interaction scenarios for this integrated purchasing application, the ERP system administrator creates a unique account `LogisticsAppUser` in the ERP system. He sets up the password and specific access rights for this account. This user account is allowed access only to the logistics business processes that are used by the integrated purchasing application.

Likewise, the system administrator for the legacy system creates a unique account `FinancialAppUser`. He also sets up the password and specific access rights for this account.

The application server administrator, as part of the operational environment of the application server, configures the access to an organization-wide directory. This directory contains security information (name, password, role, and access rights) for all the employees in the organization. It is used for authentication and authorization of employees accessing the purchasing application.

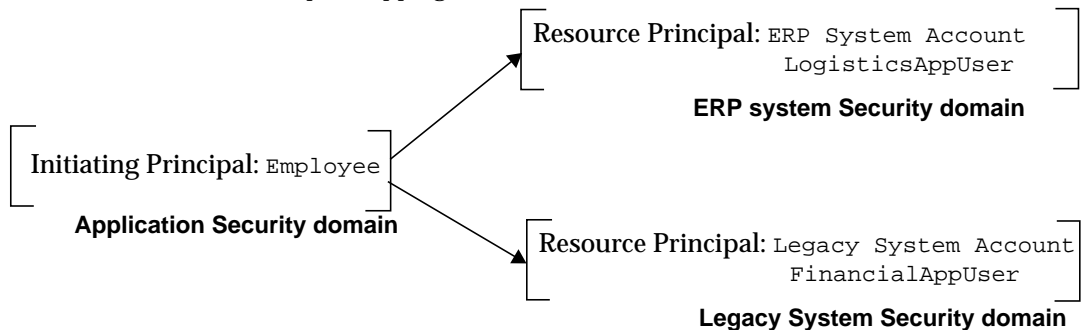
Due to their physical separation in this scenario, EISs X and Y are accessed over either a secure private network or over the Internet. This requires that a secure association be established between the application server and the EISs. A secure association allows a component on the application server to communicate securely with an EIS.

D.3.3 Deployment

During the deployment of this application, the deployer configures the security information (that is, the user account `LogisticsAppUser` and its password) required to create connections to the ERP system. This configuration is done using the resource adapter for ERP system X. The deployer also configures the security information (that is, user account `FinancialAppUser` and its password) required to create connections to the legacy system Y.

The deployer configures security information in the application server to achieve the principal mapping shown in FIGURE D-6.

FIGURE D-6 Principal Mapping



This principal mapping ensures that all connections to the ERP system are established under the security context of `LogisticsAppUser`, the resource principal for the ERP system security domain. Similarly, all connections to legacy system Y are established under the security context of the `FinancialAppUser`.

The application server does this principal mapping for all authenticated initiating principals (that is, employees accessing the integrated purchasing application) when the application connects to either the ERP system or the legacy system.

JAAS Based Security Architecture

This chapter extends the security architecture specified in Chapter 8, “Security Architecture” and Chapter 9, “Security Contract” to include support for JAAS-based pluggable authentication. The chapter refers to the following documents:

- White Paper on User Authentication and Authorization in Java platform:
<http://java.sun.com/security/jaas/doc/jaas.html>
- JAAS 1.0 documentation

E.1 Java Authentication and Authorization Service (JAAS)

JAAS provides a standard Java framework and programming interface that enables applications to authenticate and enforce access controls upon users. JAAS is divided into two parts based on the security services that it provides:

- **Pluggable Authentication:** This part of the JAAS framework allows a system administrator to plug in the appropriate authentication services to meet the security requirements of an application environment. There is no need to modify or recompile an existing application to support new or different authentication services.
- **Authorization:** Once authentication has successfully completed, JAAS provides the ability to enforce access controls based upon the principals associated with an authenticated subject. The JAAS principal-based access controls (access controls based on who runs code) supplement the existing Java 2 code source-based access controls (access controls based on where code came from and who signed it).

E.2 Requirements

The connector security architecture uses JAAS in two ways:

- **Security Contract:** The connector security architecture uses the JAAS `Subject` class as part of the security contract between an application server and a resource adapter. Use of JAAS interfaces enables the security contract to remain independent of specific security technologies or mechanisms. The security contract has been specified in Section 9.2, “Requirements” on page 9-15.
- **JAAS Pluggable Authentication framework:** This framework lets an application server and its underlying authentication services remain independent from each other. When additional EISs and new authentication services are required (or are upgraded), they can be plugged in an application server without requiring modifications to the application server.

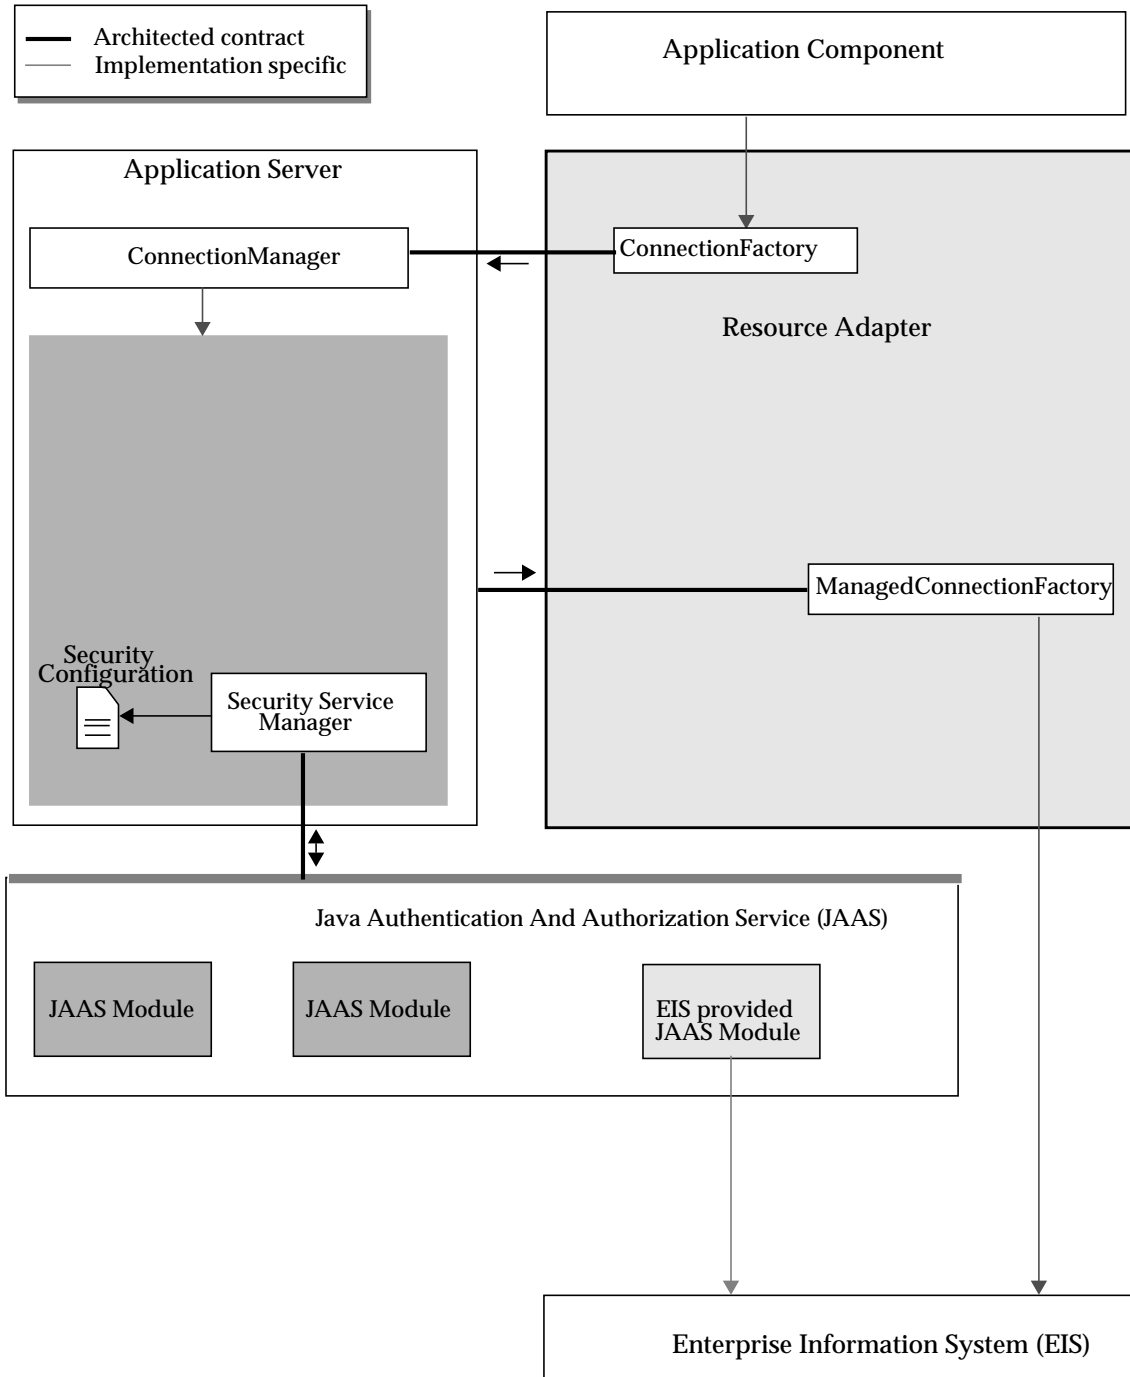
The connector architecture requires that the application server and the resource adapter must support the JAAS `Subject` class as part of the security contract. However, it recommends (but does not mandate) that an application server use the JAAS pluggable authentication framework.

The connector architecture does not require support for the authorization portion of the JAAS framework.

E.3 Security Architecture

The following section specifies the JAAS based security architecture. The security architecture addresses how JAAS may be used by an application server to support authentication requirements of heterogeneous EISs.

FIGURE E-1 Security Architecture.



E.3.1 JAAS Modules

The connector architecture recommends (but does not mandate) that an application server support *platform-wide* JAAS modules (also called authentication modules) for authentication mechanisms that are common across multiple EISs. The implementation of these JAAS modules is typically specific to an application server. However, these modules may be developed to be reusable across application servers.

A resource adapter provider can provide a resource adapter-specific *custom* implementation of a JAAS module. The connector architecture recommends that a resource adapter provider provide a *custom* JAAS module when the underlying EIS supports an authentication mechanism that is EIS specific and is not supported by an application server.

A *custom* JAAS module can be packaged together with a resource adapter and can be pluggable into an application server using the JAAS architecture.

The JAAS specification (Section 7., “Java Authentication and Authorization Service, version 1.0:” on page F-1) specifies requirements for developing and configuring JAAS modules.

E.3.2 Illustrative Examples: JAAS Module

It is not a goal of the connector architecture to specify a standard architecture for JAAS modules. The following are illustrative examples of JAAS modules used typically in the JAAS-based security architecture:

E.3.2.1 Principal Mapping Module

The application server invokes the principal mapping module passing in the `Subject` instance corresponding to the caller/initiating principal. The JAAS specification specifies the interfaces/classes and mechanisms involved in the invocation of a JAAS module.

The principal mapping module maps a caller/initiating principal to a valid resource principal and returns the mapped resource principal as part of a `Subject` instance. The authentication data (example, password) for the mapped resource principal is added to the `Subject`'s credentials. The authentication data is used later to authenticate the resource principal to the underlying EIS.

A special case of the principal mapping module takes a `null` `Subject` as an input parameter and forms a `Subject` instance with a valid resource principal and authentication data. This is the case of default principal mapping.

The principal mapping module achieves its mapping functionality by using security information configured in the application server or an enterprise directory.

The principal mapping module does not authenticate a resource principal and is configured to perform only principal mapping. The authentication of a mapped resource principal is performed separately by an authentication mechanism-specific JAAS module.

E.3.2.2 Credential Mapping Module

The credential mapping module automatically maps credentials from one authentication domain to those in a different target authentication domain. For example, an application server can provide a module that maps the public key certificate-based credential associated with a principal to a Kerberos credential.

The credentials mapping module can use the JAAS callback mechanism (note that this involves no user-interface based interaction) to get authentication data from the application server. The authentication data is used to authenticate the principal to the target authentication domain during the credentials mapping. This module can also use an enterprise directory to get security information or pre-configured mapped credentials.

E.3.2.3 Kerberos Module

This type of JAAS module supports Kerberos-based authentication for a principal. A sample Kerberos module supports:

- Getting a TGT (ticket granting ticket) to the Kerberos server in the local domain. The TGT is created by the KDC. The TGT is placed on the credentials structure for a principal.
- Delegation of authentication based on either a forwardable or proxy mechanism as specified in the Kerberos specification.

0.1 Generic Security Service API: GSS-API

The GSS-API is a standard API that provides security services to caller applications in a generic fashion. These security services include authentication, authorization, principal delegation, secure association establishment, per-message confidentiality, and integrity. These services can be supported by a wide range of security mechanisms and technologies. However, an application using GSS-API accesses these services in a generic mechanism-independent fashion and achieves source-level portability.

In the context of the connector architecture, a resource adapter uses GSS-API to establish a secure association with the underlying EIS. The use of the GSS mechanism by a resource adapter is typical in the following scenarios:

- The EIS supports Kerberos as a third-party authentication service and uses GSS-API as a generic API for accessing security services.
- The resource adapter and EIS need data integrity and confidentiality services during their communication over insecure links.

The GSS-API has been implemented over a range of security mechanisms, including Kerberos V5. See Section 6., “Java Specification Request: Generic Security Service API (GSS-API), Java bindings:” on page F-1 for a Java binding of GSS-API.

Note – The connector architecture does not require a resource adapter to use GSS-API.

E.4 Security Configuration

During deployment of a resource adapter, the deployer is responsible for configuring JAAS modules in the operational environment. The configuration of JAAS modules is based on the security requirements specified by a resource adapter in its deployment descriptor. Refer to Section 17.5.4, “Requirements” on page 17-23.

The element `authentication-mechanism` in the deployment descriptor specifies an authentication mechanism supported by a resource adapter. The standard types of authentication mechanisms are: `BasicPassword` and `Kerbyv5`. For example, if a resource adapter specifies support for `kerbyv5` authentication mechanism, the deployer configures a Kerberos JAAS module in the operational environment.

E.4.1 JAAS Configuration

The deployer sets up the configuration of JAAS modules based on the JAAS-specified mechanism. Refer to `javax.security.auth.login.Configuration` specification for more details. The JAAS configuration includes the following information on a per resource adapter basis:

- One or more authentication modules used to authenticate a resource principal.
- The order in which authentication modules need to be invoked during a stacked authentication.
- The flag value controlling authentication semantics if stacked modules are invoked.

The format for the above configuration is specific to an application server implementation.

E.5 Scenarios

The following section illustrates security scenarios for JAAS based security architecture.

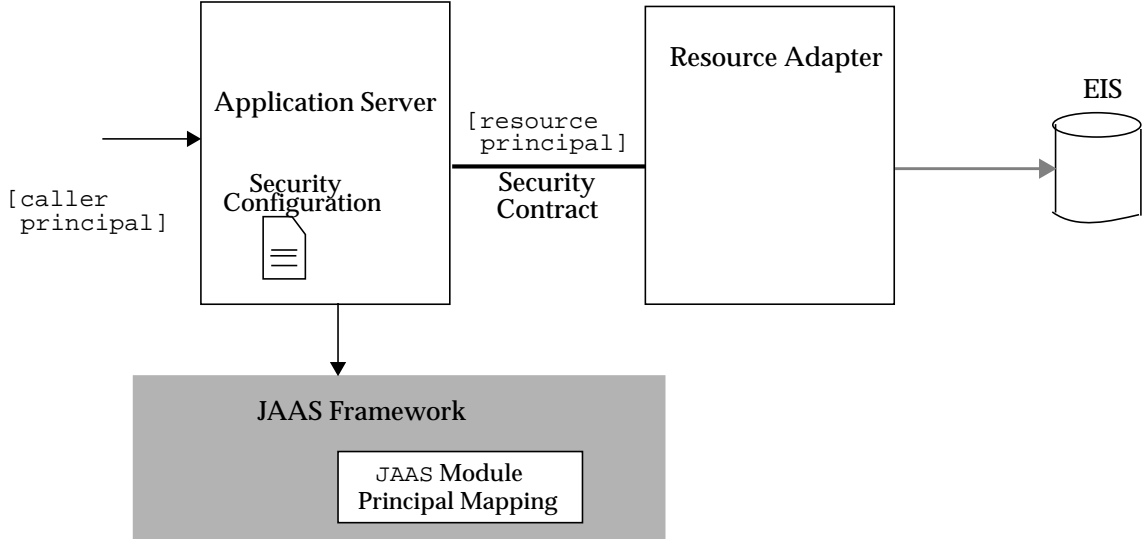
E.5.1 Scenario: Resource Adapter Managed Authentication

This scenario enables the connector architecture to support EIS specific username and password-based authentication. It involves the following steps:

- The application component invokes connection request method on the resource adapter without passing in any security arguments. The resource adapter passes the connection request to the application server.
- During the deployment of the resource adapter, the application server is configured to use a principal mapping module. This principal mapping module takes a `Subject` instance with the caller principal and returns a `Subject` instance with a valid resource principal and `PasswordCredential` instance. The `PasswordCredential` has the password for authentication of the resource principal.

- The application server calls `LoginContext.login` method. On a successful return from the principal mapping module, the application server gets a `Subject` instance that has the mapped resource principal with a valid `PasswordCredential`.

FIGURE E-2 Resource Adapter-Managed Authentication

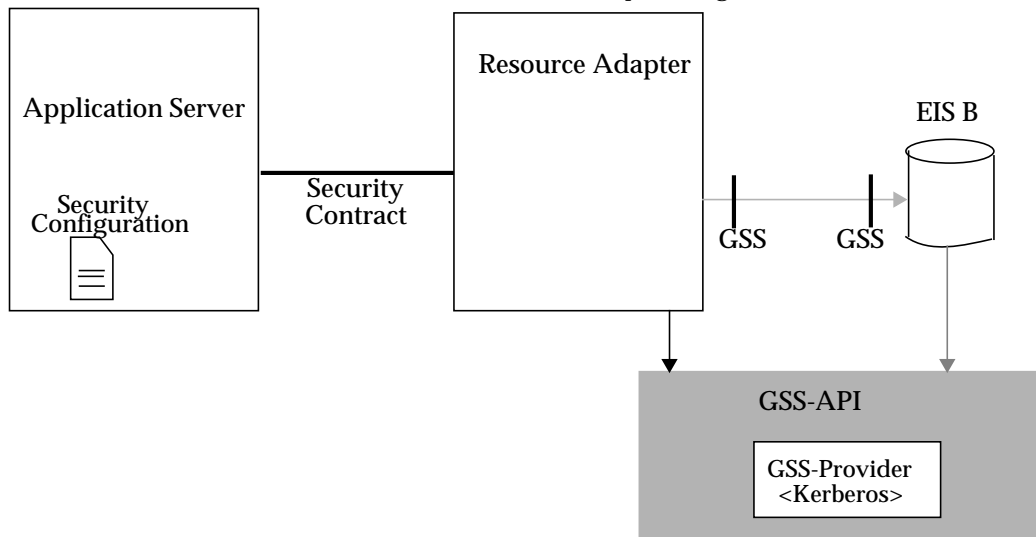


- The application server invokes the method `ManagedConnectionFactory.createManagedConnection` passing in a non-null `Subject` instance. The `Subject` instance carries the resource principal and its corresponding `PasswordCredential`, which holds the user name and password.
- The resource adapter extracts the user name and password from the `PasswordCredential` instance. The resource adapter uses the getter methods (`getPrivateCredentials` method) defined on the `Subject` interface to extract the `PasswordCredential` instance.
- The resource adapter uses username and password information (extracted from the `PasswordCredential` instance) to authenticate the resource principal to the EIS. The authentication happens during the creation of the connection through an authentication mechanism specific to the underlying EIS.

E.5.2 Scenario: Kerberos and Principal Delegation

The scenario in FIGURE E-3 involves the following steps:

FIGURE E-3 Kerberos Authentication with Principal Delegation



- The initiating principal has already authenticated itself to the application server using Kerberos. The initiating principal has a service ticket for the application server and a TGT (ticket granting ticket issued by the KDC) as part of its Kerberos based credentials.
- In this scenario, the application server is configured to impersonate the initiating principal when connecting to the EIS instance. So even though application server is directly connecting to the EIS, access to the EIS is being requested on behalf of the initiating principal. The initiating principal needs to pass its identity to the application server and allow the application server to make requests to the EIS on behalf of the initiating principal. The above is achieved through delegation of authentication.
- The application server calls the method `ManagedConnectionFactory.createManaged-Connection` by passing in a `Subject` instance with the initiating principal and its Kerberos credentials. The credentials contain a Kerberos TGT and are represented through the `GSSCredential` interface.
- The resource adapter extracts the resource principal and its Kerberos credentials from the `Subject` instance.
- The resource adapter creates a new physical connection to the EIS.
- If the resource adapter and EIS support GSS-API for establishing a secure association, the resource adapter uses the Kerberos credentials based on the GSS mechanism as follows. For details, see GSS-API specification:

- resource adapter calls `GSS_Acquire_cred` method to acquire `cred_handle` in order to reference the credentials for establishing the shared security context.
- resource adapter calls the `GSS_Init_sec_context` method. The method `GSS_Init_sec_context` yields a service ticket to the requested EIS service with the corresponding session key.

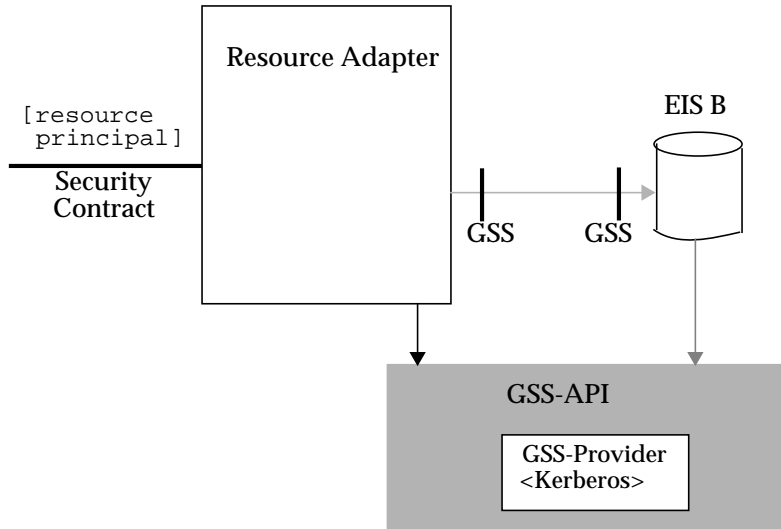
Note – The mechanism and representation through which Kerberos credentials are shared across the underlying JAAS module and GSS provider is beyond the scope of the connector architecture.

- After success, `GSS_Init_sec_context` builds a specific Kerberos-formatted message and returns it as an output token. The resource adapter sends the output token to the EIS instance.
- EIS service passes the received token to the `GSS_Accept_sec_context` method.
- Resource adapter and EIS now hold the shared security context (so have established a secure association) in the form of a session key associated with the service ticket. They can now use the session key in the subsequent per-message methods: `GSS_GetMIC`, `GSS_VerifyMIC`, `GSS_Wrap`, `GSS_Unwrap`.
- If the resource adapter and EIS fail to establish a secure association, the resource adapter cannot use the physical connection as a valid connection to the EIS instance. The resource adapter returns a security exception on the `createManagedConnection` method.

E.5.3 Scenario: GSS-API

If an EIS supports the GSS mechanism, a resource adapter may (but is not required to) use GSS-API to set up a secure association with the EIS instance. (See FIGURE E-4.) The section Generic Security Service API: GSS-API on page 6 gives a brief overview of GSS-API.

FIGURE E-4 GSS-API use by Resource Adapter

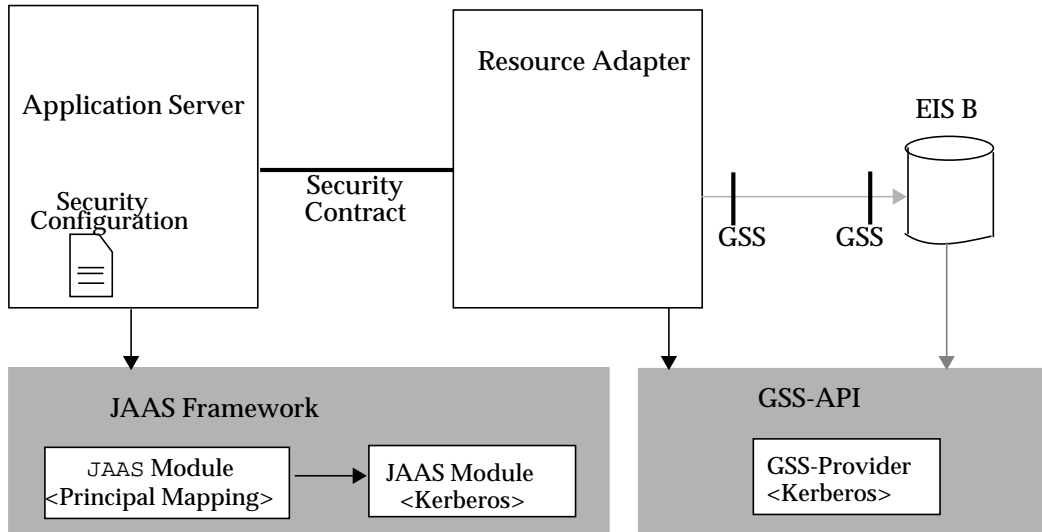


A formal specification of the use of GSS-API by a resource adapter is beyond the scope of the connector architecture. However, GSS-API has been mentioned as a possible implementation option for a resource adapter that has the GSS mechanism supported by its underlying EIS.

E.5.4 Scenario: Kerberos Authentication After Principal Mapping

The scenario depicted in FIGURE E-5 involves the following steps:

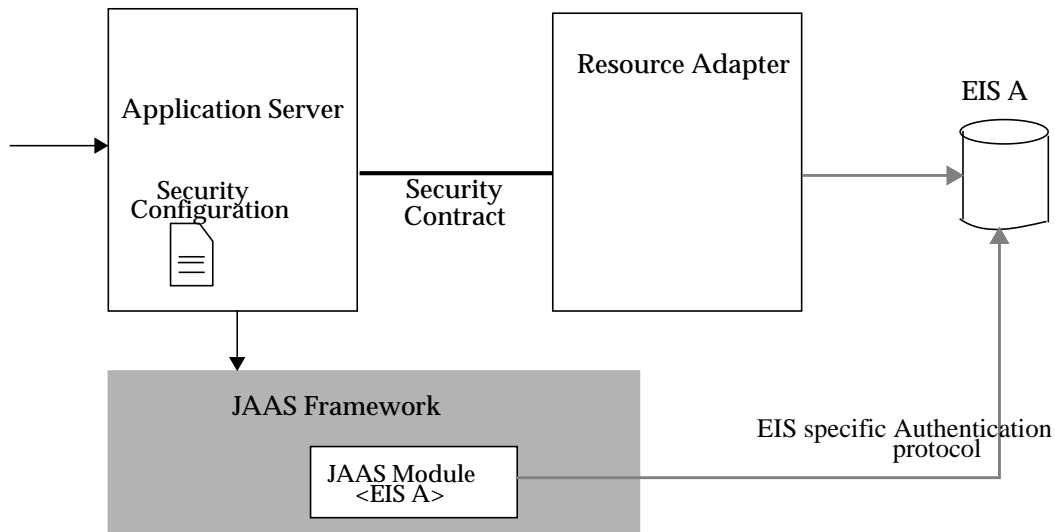
FIGURE E-5 Kerberos Authentication After Principal Mapping



- The application server is configured to use the principal mapping module and Kerberos module. The two authentication modules are stacked together with the principal mapping module first.
- The application server creates a `LoginContext` instance by passing in the `Subject` instance for the caller principal and a `CallbackHandler` instance. Next, the application server calls the `login` method on the `LoginContext` instance.
- The principal mapping module takes a `Subject` instance with caller principal and returns a `Subject` instance with a valid resource principal and Kerberos-based authentication data. The principal mapping module does not authenticate the resource principal; it does only principal mapping to find the mapped resource principal and its authentication data.
- Next, the Kerberos module (called after the principal mapping module) uses the resource principal and its authentication data to authenticate the resource principal. The Kerberos module leads to a valid TGT for the Kerberos domain supported by the EIS. The TGT is contained in the Kerberos credentials represented through the `GSSCredential` interface.
- The application server calls the method `ManagedConnectionFactory.createManagedConnection` passing in a `Subject` instance with the resource principal and its Kerberos credentials.
- The remaining steps are the same as in the previous scenario, Section E.5.2, “Scenario: Kerberos and Principal Delegation” on page E-10

E.5.5 Scenario: EIS-Specific Authentication

FIGURE E-6 Authentication Through EIS-Specific JAAS Module



The scenario in FIGURE E-6 involves the following steps:

- During the configuration of a resource adapter, the application server is configured to use an EIS-specific JAAS module for authentication to the underlying EIS.

The configured JAAS module supports an authentication mechanism specific to the EIS. The application server is responsible for managing the authentication data and JAAS configuration.

- The application server gets a request from the application component to create a new physical connection to the EIS. Creating a new physical connection requires the resource principal to authenticate itself to the underlying EIS instance.
- The application server initiates the authentication of the resource principal. It creates a `LoginContext` instance by passing in the `Subject` instance and a `CallbackHandler` instance. Next, the application server calls the `login` method on the `LoginContext` instance.
- The JAAS module authenticates the resource principal to the underlying EIS. It uses the callback handler provided by the application server to get the authentication data.
- The application server invokes the method `ManagedConnectionFactory.createManagedConnection` passing in the `Subject` instance with the authenticated resource principal and its credential.

- The resource adapter extracts the credential (associated with the `Subject` instance) for the resource principal using the getter methods defined on the `Subject` interface. The resource adapter uses this credential to create a connection to the underlying EIS.

In this scenario, authenticating a resource principal (initiated by the application server and performed by the JAAS module) is separate from creating a connection to the EIS. The resource adapter uses the credential of the resource principal to create a connection to the EIS. This connection creation can involve further authentication.

After successfully creating a connection to the EIS, the resource adapter returns the newly created connection from the method

`ManagedConnectionFactory.createManagedConnection`.

Related Documents

1. Enterprise JavaBeans (EJB™) specification, version 2.1:
<http://java.sun.com/products/ejb/>
2. Java Transaction API (JTA) specification, version 1.0.1B
<http://java.sun.com/products/jta/>
3. JDBC API specification, version 3.0
<http://java.sun.com/products/jdbc/>
4. X/Open CAE Specification -- Distributed Transaction Processing: the XA specification, X/Open document
5. RFC: Generic Security Service API (GSS-API) specification, version 2:
<http://www.ietf.org/rfc/rfc2078.txt>
6. Java Specification Request: Generic Security Service API (GSS-API), Java bindings:
http://java.sun.com/aboutJava/communityprocess/jsr/jsr_072_gss.html
7. Java Authentication and Authorization Service, version 1.0:
<http://java.sun.com/products/jaas/>
8. Java 2 Platform Enterprise Edition (J2EE™), Platform specification, version 1. 4:
<http://java.sun.com/j2ee/>
9. Java Server Pages (JSP) specification, version 1.3:
<http://java.sun.com/products/jsp/>
10. Java Servlet specification, version 2.4:
<http://java.sun.com/products/servlet/>
11. Java Message Service, version 1.1:
<http://java.sun.com/products/jms/>
12. Java 2 Platform, Standard Edition (J2SE™), API specification, version 1.4:

<http://java.sun.com/j2se/1.4/docs/api/>

13. J2EE™ Deployment API Specification

<http://java.sun.com/j2ee/tools/deployment/index.html>

14. JavaBeans™ Specification 1.01 Final Release

<http://java.sun.com/products/javabeans>

Change History for Version 1.0

G.1 Version 0.9

- Editorial run through the document
- Added section 1.4 on relationship between JDBC and Connector architecture
- Added scenario on B2B in the chapter 4
- Added `java.io.Serializable` to the code specification of interfaces that are required to support `Serializable` interface
- Added clarifications in the chapter 5 based on the expert comments. The changes are marked by change bars.
- Added `equals` and `hashCode` methods to interface `ConnectionRequestInfo`
- Added section 6.8 on Connection Association
- Added clarifications to the chapter 7. Did minor restructuring of the chapter based on review comments. The changes are marked by change bars.
- Added clarifications to the chapter 8 based on expert comments
- Changed few details and added clarifications in the chapter 9 based on the review comments. The changes are marked by change bars.
- Added more description for packaging and deployment in the chapter 10
- Clarified version dependencies in the chapter 11
- Introduced interface `javax.resource.Referenceable` for the standard `setReference` method
- Removed scenarios on Credentials Mapping and Single sign-on from Appendix C. Updated scenario C.6.2 to refer GSS-API.

G.2 Version 1.0 - Public Draft 1

- Removed definition of "Connector" from 2.1. The term Connector is now used broadly refer to the Connector architecture, while resource adapter refers to the system library.
- Added requirement for ConnectionEventListener to 6.9.2: Application Server
- Added connection handle property to the ConnectionEvent, section 5.5.7
- Introduced getResultSetInfo method in the Connection interface
- Added "Administered Object" in the section 9.6.2
- Added more details to section "Auto Commit" in 9.5.2
- Introduced separate interface for ResultSetInfo in the section 9.10.3
- Changed specification of element config-property-type in section 10.6
- Added an example to illustrate security permission specification in the section 11.3
- Added CCI related information to Projected Items, chapter 12

G.3 Version 1.0 - Public Draft 2

- Section 5.5.1: Change based on introduction of ConnectionSpec interface
- Section 5.5.1: Added clarification to ConnectionRequestInfo section
- Section 5.5.4: Added clarification to section on "Cleanup of ManagedConnection"
- Section 5.5.6: Added clarification to paragraph after the interface for ConnectionEventListener
- Section 5.9.1: Added clarification to description of the scenario
- Section 6.8: Moved earlier section "Details on Local transaction" ahead of connection sharing section and renamed it "Scenarios: Local Transaction Management". No change in any content.
- Section 6.9: Added more details on connection sharing based on the changes in EJB 2.0 and J2EE 1.3 platform specification.
- Section 6.10: Added this section to clarify local transaction optimization. This is based on changes in EJB 2.0 and J2EE 1.3 platform specification.
- Section 6.11: Made a new section on "Scenarios: Connection sharing". No change in content.
- Section 6.12: Added clarifications and requirements in the section on "Connection Association"
- Section 6.13.2: Moved requirements on connection sharing to section 6.9
- Section 7.4.2: Code sample changed to reflect ConnectionSpec usage
- Section 9.5.1: Changed getConnection(Map) to getConnection(ConnectionSpec) and added clarifications.
- Section 9.5.2: Introduced a section on ConnectionSpec

- Section 9.7.2: Added methods to ResourceAdapterMetaData interface. Added description of these methods.
- Section 9.9.1: Record, MappedRecord and IndexedRecord now extend Serializable interface.
- Section 9.10: Added note on JDBC semantics in relation to CCI ResultSet
- Section 9.10.3: Added note on ResultSetInfo implementation requirements
- Section 10.6: Change to auth-mechanism specification in DTD. Removed + from credential-interface.
- Figure 29: Added clarifications for the diagram
- Section 8.3: Clarified security contract requirements for the application server
- Section 9.5.1: Moved method getRecordFactory from Interaction to ConnectionFactory. Note that it is not necessary to have an active connection to create generic record instances.

G.4 Version 1.0 - Proposed Final Draft 2

- Reviewed requirements in terms of compliance testing. marked with change bars in the document
- Fixed documentation errors
- System Contracts:
 - Section 5.5.4: Clarified requirements for the method `matchManagedConnections` on `ManagedConnectionFactory` interface
 - Section 6.9: Made requirements for connection sharing consistent with J2EE 1.3 platform specification
 - Section 6.10: Added specification of requirements for different transaction scenarios. Added illustrative scenarios
 - Section 6.11: Removed a transaction scenario that illustrated connection sharing
 - Section 6.11: Clarified requirements for connection association
- Common Client Interface:
 - Removed `setLogWriter`, `getLogWriter`, `setTimeout`, `getTimeout` methods from `ConnectionFactory` interface
 - Added description for exceptions in Java docs for the CCI interfaces. Note that no new exception has been introduced
 - `ConnectionFactory` implementation class required to provide a default constructor
 - Added clarifications; marked by change bars
- Deployment and Packaging:
 - Section 10.2: Clarified requirements for packaging and deployment of a resource adapter
- DTD changes based on a review of DTDs for various J2EE specifications:
 - Ordered elements alphabetically except the root element
 - `<display-name>` changed to optional in `<connector>` element

- Used common elements from other DTDs: <description>, <small-icon>, <large-icon>
- Used common header comments across all J2EE DTDs
- <auth-mechanism> changed to <authentication-mechanism>
- <auth-mech-type> changed to <authentication-mechanism-type>
- Added `java.lang.Character` to <config-property-type>:
- Changed defined values in <authentication-mechanism-type>: `basic-password` to `BasicPassword`, `kerbv5` to `Kervb5`
- Changed defined values in <transaction-support> element: `no_transaction` to `NoTransaction`, `local_transaction` to `LocalTransaction`, `xa_transaction` to `XATransaction`

G.5 Version 1.0 - Final Release

- Clarification on reauthentication in the section 8.2.7
- Change in auto-commit in section 9.5.3. Removed `set/getAutoCommit` methods from the `Connection` interface

Change History for Version 1.5

H.1 Public Review Draft

- some minor clarifications to sections 5.3.1, 5.3.4, 5.3.5
- some minor clarifications to section 10.3
- added chapter 11: Inbound communication
- added chapter 13: EJB invocation
- section 12.3
- added method release on Endpoint interface
- added method getXAResources on ResourceAdapter interface
- updated section 12.5.1 to illustrate XAResource implementation
- added section 12.5.2: Message redelivery upon crash recovery
- section 12.5.6: added an application server requirement to set thread context class loader
- section 12.7: updated the JMS use case to comply with the Connector 1.5 XSD
- updated section 16.2: application server requirements
- some minor clarifications to sections 17.4.1, 17.4.2
- section 17.6 Converted Connector 1.5 DTD to Connector 1.5 XSD
- added Appendix B: Activation Configuration for Message Inflow to JMS endpoints

H.2 Proposed Final Draft

- Section 7.9: connection sharing requirements have been updated.
- Section 9.2.3: `GenericCredential` interface has been deprecated.
- Section 9.3: added a resource adapter requirement to throw a `SecurityException` if the authentication information provided is erroneous.

- Section 10.3.1: clarified use of synchronized blocks within a `Work` implementation.
- Section 12.3: a new method parameter `java.lang.reflect.Method` has been added to `beforeDelivery` and `isDeliveryTransacted` methods.
- Section 12.4.2: added separate sections to describe `ActivationSpec` `JavaBean` and `Administered` objects.
- Section 12.4.9: clarified multiple endpoint activations
- Section 12.4.7: this new section describes the structure of a message listener interface.
- Section 12.5.2: clarified crash recovery processing
- Section 12.7.2: updated the sample code to use JMS 1.1 style generic APIs.
- Section 16.2: removed application client container requirements for supporting connector contracts.
- Section 17.2: added description of RAR file contents.
- Section 17.6: updated Connector 1.5 XSD - `resourceadapter-class` element has been made optional, and `config-property` element has been added to `adminobject` element.
- Section 18.4: added an example to illustrate usage of `doPrivileged` blocks in resource adapter code.
- Section 19.2: added `InvalidPropertyException` and `UnavailableException` to the exception list.
- Appendix B: added descriptions to the various `ActivationSpec` `JavaBean` properties.
- Appendix F: updated links to related documents.

H.3 Proposed Final Draft 2

- Section 6.5.2: clarifications on `ConnectionManager` requirements.
- Section 7.9.1: this new section describes how a resource adapter may detect and report sharing violations.
- Section 7.13.1: added description to clarify the auto-commit behavior of transactional resource adapters.
- Section 7.14: this new section describes two optional connection optimizations.
- Section 12.5.6: clarified that the message delivery must not rely on the class loader that is associated with the `java.lang.reflect.Method` object parameter.
- Section 12.5.8: clarified the usage of source managed transactions.
- Section 14.4.4: clarified the requirements on transaction inflow contracts.
- Section 17.6: XSD modification: The contents of the `connectionfactory-interface` element must be unique in the `outbound-resourceadapter`.

H.4 Final Release

- Section 5.3.1: clarification on the behavior of `start` method call.
- Section 5.3.2: `ResourceAdapterAssociation` interface.
- Section 5.3.4.1, 5.3.4.2: clarifications on resource adapter shutdown.
- Section 6.4.1: fifth bullet item, clarification on access to application JNDI context.
- Section 6.5.3: `ValidatingManagedConnectionFactory` interface.
- Section 6.5.4: clarifications on connection sharing and multiple connection handles.
- Section 6.5.6: clarifications on event listener mechanism.
- Section 7.7.2: clarifications on local transaction events and implicit transaction begin.
- Section 7.9: `SharingViolationException`, and clarification on non-transactional resource adapters.
- Section 7.9.1: clarifications on sharing violation detection.
- Section 9.1.8.2: clarifications on security options A, B and C.
- Section 10.3: numeric string values for `WorkException` error code values.
- Section 10.3.1: clarification on access to JNDI context from `Work` objects.
- Section 10.3.7: `createTimer` method throws `UnavailableException`.
- Section 12.4.2.2: clarifications on required configuration properties on `ActivationSpec` `JavaBean`, and `JMS ActivationSpec` `JavaBean`.
- Section 12.4.2.4: clarifications on configuring administered objects.
- Section 12.3: changes to `endpointActivation`, `beforeDelivery`, `afterDelivery` and `isDeliveryTransacted` method signatures.
- Section 12.4.4: clarifications on message delivery during endpoint activation.
- Section 12.4.4: clarifications on `endpointDeactivation` method.
- Section 12.4.6: clarifications on contract between deployment tool and application server.
- Section 12.5.6: description of `IllegalStateException` thrown by `beforeDelivery` and `afterDelivery` method calls.
- Section 12.5.8: clarifications on use of transaction inflow.
- Section 15.8: service endpoint `javax.resource.cci.MessageListener` interface.
- Section 16.2: Table 16-1, updates to column corresponding to lifecycle management.
- Section 16.3: `JavaBean` requirements.
- Section 16.4: Equality constrains.
- Section 17.5.3: clarifications on `getReference` method.
- All exception classes have been provided constructors that have a `cause` param.
- Section 17.6: XSD modifications: `vendor-name` has been assigned a string type, `security-permission` has been moved one level up such that the permissions apply to the whole resource adapter code, `deployment-extension` element has been removed, indentations have been reformatted to follow standard XML conventions.

