

THE PRENTICE HALL SERVICE TECHNOLOGY SERIES FROM THOMAS ERL



*"An outstanding depth-and-breadth resource for IT architects and Java professionals to understand and apply the marriage of SOA and modern Java."*

—Antonio Bruno, Enterprise Architecture and Strategy, digitalStrom

*"Provides clarity on abstract concepts and is filled with concrete examples of implementing SOA principles in Java environments."*

—Sanjay Singh, Certified SOA Architect

*"A great self-contained book on SOA using flexible Java implementations..."*

—Roger Stoffers, Hewlett-Packard

*"...provides a holistic, comprehensive view on leveraging SOA principles and architecture for building and deploying performant Java services."*

—Suzanne D'Souza, KBACE Technologies


# SOA with Java

## Realizing Service-Oriented Architecture with Java Technologies

Co-Authored and Edited by Thomas Erl, World's Top-Selling SOA Author  
Co-Authored by Andre Tost, Satadru Roy, Philip Thomas

Contributions by Raj Balasubramanian, David Chou, Thomas Plunkett  
Foreword by Dr. Mark Little, Vice President of Engineering, Red Hat

 PRENTICE  
HALL

ServiceTech  
 PRESS

FREE SAMPLE CHAPTER

SHARE WITH OTHERS



# SOA with Java

*Realizing Service-Oriented  
with Java Technologies*

Thomas Erl, Andre Tost,  
Satadru Roy, and Philip Thomas



PRENTICE  
HALL

PRENTICE HALL

UPPER SADDLE RIVER, NJ • BOSTON • INDIANAPOLIS • SAN FRANCISCO

NEW YORK • TORONTO • MONTREAL • LONDON • MUNICH • PARIS • MADRID

CAPE TOWN • SYDNEY • TOKYO • SINGAPORE • MEXICO CITY



Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and the publisher was aware of a trademark claim, the designations have been printed with initial capital letters or in all capitals.

The authors and publisher have taken care in the preparation of this book, but make no expressed or implied warranty of any kind and assume no responsibility for errors or omissions. No liability is assumed for incidental or consequential damages in connection with or arising out of the use of the information or programs contained herein.

For information about buying this title in bulk quantities, or for special sales opportunities (which may include electronic versions; custom cover designs; and content particular to your business, training goals, marketing focus, or branding interests), please contact our corporate sales department at [corpsales@pearsoned.com](mailto:corpsales@pearsoned.com) or (800) 382-3419.

For government sales inquiries, please contact [governmentsales@pearsoned.com](mailto:governmentsales@pearsoned.com).

For questions about sales outside the U.S., please contact [international@pearsoned.com](mailto:international@pearsoned.com).

Visit us on the Web: [informit.com/ph](http://informit.com/ph)

Library of Congress Control Number: 2014939738

Copyright © 2014 Arcitura Education Inc.

All rights reserved. Printed in the United States of America. This publication is protected by copyright, and permission must be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise. To obtain permission to use material from this work, please submit a written request to Pearson Education, Inc., Permissions Department, One Lake Street, Upper Saddle River, New Jersey 07458, or you may fax your request to (201) 236-3290.

ISBN-13: 978-0-13-385903-4

ISBN-10: 0-13-385903-7

Text printed in the United States on recycled paper at Courier in Westford, Massachusetts.

First printing: June 2014

**Editor-in-Chief**

Mark Taub

**Senior Acquisitions Editor**

Trina MacDonald

**Development Editor**

Teejay Keepence

**Managing Editor**

Kristy Hart

**Copy Editors**

Teejay Keepence

Maria Lee

**Senior Indexer**

Cheryl Lenser

**Proofreader**

Maria Lee

**Publishing Coordinator**

Olivia Basegio

**Cover Designer**

Thomas Erl

**Composer**

Jasper Paladino

**Photos**

Thomas Erl

**Graphics**

Jasper Paladino

*To my family and friends.*

—Thomas Erl

*Thanks go to my wife and my boys for their  
continued understanding, patience and support.*

—Andre Tost

*I'd like to express my sincere thanks to Thomas Erl for giving me the  
opportunity to work on this book. It has been a privilege to be associated  
with him and along the way, I have learned a lot about service-orientation.*

*I'd also like to thank Mark Taub at Pearson for supplying technical  
books supporting my work.*

—Satadru Roy

*To my wife, Sarah, for her tireless support.*

—Philip Thomas

*This page intentionally left blank*

# Contents at a Glance

Foreword .....	xix
CHAPTER 1: Introduction .....	1
CHAPTER 2: Case Study Examples .....	11
<b>PART I: FUNDAMENTALS</b>	
CHAPTER 3: Fundamental SOA Concepts .....	21
CHAPTER 4: Basic Java Distributed Technologies .....	39
CHAPTER 5: Web-Based Service Technologies .....	81
CHAPTER 6: Building Web-Based Services with Java .....	111
<b>PART II: SERVICES</b>	
CHAPTER 7: Service-Oriented Principles with Java Web-Based Services .....	139
CHAPTER 8: Utility Services with Java .....	211
CHAPTER 9: Entity Services with Java .....	261
<b>PART III: SERVICE COMPOSITION AND INFRASTRUCTURE</b>	
CHAPTER 10: Task Services with Java .....	307
CHAPTER 11: Service Composition with Java .....	335
CHAPTER 12: ESB as SOA Infrastructure .....	379
<b>PART IV: APPENDICES</b>	
APPENDIX A: Case Study Conclusion .....	405
APPENDIX B: Service-Oriented Principles Reference .....	409
APPENDIX C: SOA Design Patterns Reference .....	425
APPENDIX D: The Annotated SOA Manifesto .....	519
About the Authors .....	533
About the Foreword Contributor .....	535
About the Contributors .....	537
Index .....	539

*This page intentionally left blank*

# Contents

<b>Foreword</b> .....	<b>xix</b>
-----------------------	------------

<b>CHAPTER 1: Introduction</b> .....	<b>1</b>
--------------------------------------	----------

1.1 About This Book .....	2
Objectives of This Book .....	2
Who This Book Is For .....	2
What This Book Does Not Cover .....	3
1.2 Prerequisite Reading .....	3
How This Book Is Organized .....	4
Part I: Fundamentals .....	4
<i>Chapter 3: Fundamental SOA Concepts</i> .....	5
<i>Chapter 4: Basic Java Distributed Technologies</i> .....	5
<i>Chapter 5: Web-Based Service Technologies</i> .....	5
<i>Chapter 6: Building Web-Based Services with Java</i> .....	5
Part II: Services .....	5
<i>Chapter 7: Service-Oriented Principles with Java Web-Based Services</i> .....	5
<i>Chapter 8: Utility Services with Java</i> .....	5
<i>Chapter 9: Entity Services with Java</i> .....	6
Part III: Service Composition and Infrastructure .....	6
<i>Chapter 10: Task Services with Java</i> .....	6
<i>Chapter 11: Service Composition with Java</i> .....	6
<i>Chapter 12: ESB as SOA Infrastructure</i> .....	6
Part IV: Appendices .....	6
<i>Appendix A: Case Study Conclusion</i> .....	6
<i>Appendix B: Service-Oriented Principles Reference</i> .....	6
<i>Appendix C: SOA Design Patterns Reference</i> .....	7
<i>Appendix D: The Annotated SOA Manifesto</i> .....	7
1.3 How Principles and Patterns Are Used in This Book .....	7
1.4 Symbols and Figures .....	7
1.5 Additional Information .....	7
Updates, Errata, and Resources ( <a href="http://www.servicetechbooks.com">www.servicetechbooks.com</a> ) ..	8
Referenced Specifications ( <a href="http://www.servicetechspecs.com">www.servicetechspecs.com</a> ) ..	8



The Service Technology Magazine ([www.servicetechmag.com](http://www.servicetechmag.com)) . . . 8

Service-Orientation ([www.serviceorientation.com](http://www.serviceorientation.com)) . . . . . 8

What Is REST? ([www.whatisrest.com](http://www.whatisrest.com)) . . . . . 8

What Is Cloud? ([www.whatiscloud.com](http://www.whatiscloud.com)) . . . . . 8

SOA and Cloud Computing Design Patterns  
 ([www.soapatterns.org](http://www.soapatterns.org), [www.cloudpatterns.org](http://www.cloudpatterns.org)) . . . . . 8

SOA Certified (SOACP) Professional ([www.soaschool.com](http://www.soaschool.com)) . . . . . 9

Cloud Certified Professional (CCP) ([www.cloudschool.com](http://www.cloudschool.com)). . . . . 9

Big Data Science Certified Professional (BDSCP)  
 ([www.bigdatascienceschool.com](http://www.bigdatascienceschool.com)) . . . . . 9

Notification Service . . . . . 9

**CHAPTER 2: Case Study Examples . . . . . 11**

2.1 How Case Study Examples Are Used . . . . . 12

    Style Characteristics . . . . . 12

    Relationship to Abstract Content . . . . . 12

    Code Samples . . . . . 12

2.2 Case Study Background: NovoBank . . . . . 13

    Technical Infrastructure . . . . . 13

    Automation Solutions . . . . . 13

    Business Obstacles and Goals . . . . . 14

    Future IT Roadmap . . . . . 15

        1. *Build Reusable Business Services* . . . . . 15

        2. *Consolidate Information* . . . . . 16

        3. *Improve Channel Experience* . . . . . 16

        4. *Build Services Infrastructure* . . . . . 16

2.3 Case Study Background: SmartCredit Co. . . . . 16

    Technical Infrastructure . . . . . 17

    Automation Solutions . . . . . 17

    Business Goals . . . . . 17

    Future IT Roadmap . . . . . 18

**PART I: FUNDAMENTALS**

**CHAPTER 3: Fundamental SOA Concepts . . . . .21**

- 3.1 Basic Terminology and Concepts . . . . . 22
  - Service-Oriented Computing . . . . . 22
  - Service-Oriented Architecture (SOA) . . . . . 27
  - SOA Manifesto . . . . . 28
  - Services . . . . . 29
  - Cloud Computing . . . . . 29
  - IT Resources . . . . . 30
  - Service Models. . . . . 30
    - Agnostic Logic and Non-Agnostic Logic* . . . . . 31
  - Service Inventory . . . . . 32
  - Service Portfolio . . . . . 32
  - Service Candidate . . . . . 33
  - Service Contract. . . . . 33
  - Service-Related Granularity . . . . . 34
  - Service Profiles. . . . . 35
  - SOA Design Patterns . . . . . 36
- 3.2 Further Reading . . . . . 38

**CHAPTER 4: Basic Java Distributed Technologies . . . . .39**

- 4.1 Java Distributed Computing Basics . . . . . 40
  - Java SE Architecture . . . . . 41
  - Java EE Architecture . . . . . 41
    - The Beginning of Java EE.* . . . . . 41
  - Application Packaging and Deployment . . . . . 44
    - Deployment Descriptors* . . . . . 44
  - Java EE Architectural Tiers. . . . . 45
- 4.2 Java Distributed Technologies and APIs. . . . . 46
  - Java SE APIs . . . . . 46
    - RMI* . . . . . 46
    - RMI / IIOP* . . . . . 47
    - JNDI* . . . . . 48
    - JDBC.* . . . . 49

Java EE APIs . . . . .	49
<i>Contexts and Dependency Injection</i> . . . . .	49
<i>JTA</i> . . . . .	50
<i>Java EE Connector Architecture</i> . . . . .	51
<i>EJB</i> . . . . .	51
<i>Session EJBs</i> . . . . .	52
<i>Persistence Entities</i> . . . . .	53
<i>Service-Oriented Principles and the EJB Model</i> . . . . .	55
<i>JMS</i> . . . . .	56
<i>Message-Driven Beans</i> . . . . .	58
<i>Security in Java EE</i> . . . . .	58
4.3 XML Standards and Java APIs . . . . .	59
XML . . . . .	59
XML Schema Definition . . . . .	61
XSLT . . . . .	63
JAXP. . . . .	63
JAXB . . . . .	64
4.4 Building Services with Java Components . . . . .	64
Components as Services. . . . .	65
Application Protocols . . . . .	65
Service Contracts. . . . .	68
<i>Location</i> . . . . .	68
<i>Operations</i> . . . . .	69
<i>Messages</i> . . . . .	69
<i>Further Considerations</i> . . . . .	70
Components as Services and Service-Oriented . . . . .	71
<i>Standardized Service Contract</i> . . . . .	71
<i>Service Loose Coupling</i> . . . . .	71
<i>Service Abstraction</i> . . . . .	72
<i>Service Discoverability</i> . . . . .	72
4.5 Java Vendor Platforms . . . . .	74
GlassFish Enterprise Server . . . . .	74
IBM WebSphere Application Server. . . . .	75
IBM WebSphere Application Server Community Edition . . . . .	78
Oracle WebLogic Server . . . . .	79

**CHAPTER 5: Web-Based Service Technologies . . . . . 81**

- 5.1 SOAP-Based Web Services . . . . . 82
  - Extensibility of Web Services Standards (WS-\*) . . . . . 88
    - WS-Addressing* . . . . . 89
    - SOAP with Attachments (SwA)* . . . . . 90
    - WS-ReliableMessaging* . . . . . 91
    - WS-Transaction* . . . . . 92
    - WS-Security* . . . . . 93
    - WS-Policy* . . . . . 94
    - Web Services Distributed Management* . . . . . 95
  - Common Web Services Middleware . . . . . 95
    - Enterprise Service Bus (ESB)* . . . . . 95
    - Orchestration* . . . . . 97
    - Management and Monitoring* . . . . . 99
    - Registries and Repositories* . . . . . 99
    - Service Construction and Assembly* . . . . . 100
- 5.2 REST Services . . . . . 101
  - HTTP Response Codes . . . . . 102
  - Resources and Addresses . . . . . 103
    - Service Request* . . . . . 104
    - Service Response* . . . . . 104
    - Service Request* . . . . . 104
    - Service Response* . . . . . 104
  - HTTP Methods . . . . . 106
  - Resource Representations . . . . . 108
  - The ACCEPT Header . . . . . 109

**CHAPTER 6: Building Web-Based Services with Java . . . . 111**

- 6.1 JAX-WS . . . . . 112
  - SAAJ . . . . . 115
  - Handlers . . . . . 118
  - Web Services Engines and Toolkits . . . . . 119
  - JAXR . . . . . 120
- 6.2 Java Implementations of WS-\* Standards . . . . . 122
  - Advanced Web Services Standards and Frameworks . . . . . 122
  - Service Component Architecture . . . . . 123
  - Spring-WS . . . . . 124

- 6.3 JAX-RS . . . . . 124
  - Implementing JAX-RS . . . . . 125
  - Implementing REST Services. . . . . 129
  - Scalability . . . . . 130
    - Statelessness* . . . . . 131
    - Uniform Contract* . . . . . 131
    - Cacheability* . . . . . 131
    - Addressability* . . . . . 132
  - Security . . . . . 132
  - REST Service Support . . . . . 134

**PART II: SERVICES**

**CHAPTER 7: Service-Oriented Principles with Java Web-Based Services . . . . . 139**

- 7.1 Service Reusability . . . . . 140
  - Agnostic Functional Contexts. . . . . 140
  - Highly Generic Service Logic . . . . . 141
  - Generic and Extensible Service Contracts . . . . . 144
  - Concurrent Access to Service Logic . . . . . 145
- 7.2 Standardized Service Contract . . . . . 151
  - Top-Down vs. Bottom-Up. . . . . 151
  - Mapping Between Java and WSDL . . . . . 152
  - Wrapped Document/Literal Contracts . . . . . 153
  - Implicit and Explicit Headers . . . . . 154
    - Explicit Headers* . . . . . 154
    - Implicit Headers* . . . . . 157
    - No Headers*. . . . . 159
  - Data Mapping with REST . . . . . 159
    - Conversion Between JSON and POJOs*. . . . . 161
    - Binary Data in Web Services*. . . . . 165
    - Binary Data in REST Services* . . . . . 170
  - Use of Industry Standards . . . . . 175

- 7.3 Service Loose Coupling . . . . . 176
  - Separation of Contract and Implementation . . . . . 177
  - Independent Functional Contexts . . . . . 179
  - Service Consumer Coupling . . . . . 180
- 7.4 Service Abstraction . . . . . 184
  - Abstracting Technology Details . . . . . 185
  - Hiding Service Details . . . . . 185
  - Document Constraints . . . . . 188
- 7.5 Service Composability . . . . . 189
  - Runtime Environment Efficiency . . . . . 190
  - Service Contract Flexibility . . . . . 192
  - Standards-Based Runtime . . . . . 193
- 7.6 Service Autonomy . . . . . 194
  - Well-Defined Functional Boundary . . . . . 194
  - Runtime Environment Control . . . . . 195
  - High Concurrency . . . . . 196
- 7.7 Service Statelessness . . . . . 197
  - Orchestration Infrastructure . . . . . 198
  - Session State . . . . . 198
  - Storing State . . . . . 199
- 7.8 Service Discoverability . . . . . 204
  - Design-Time Discoverability . . . . . 204
  - Runtime Discoverability . . . . . 205
  - Service Registries . . . . . 208

**CHAPTER 8: Utility Services with Java . . . . . 211**

- 8.1 Inside the Java Utility Service . . . . . 212
  - Architectural Considerations . . . . . 212
  - Utility Service Taxonomy . . . . . 220
- 8.2 Utility Service Design and Implementation . . . . . 221
  - Utility Service Design . . . . . 221
  - Utility Services and Java Editions . . . . . 226
    - Utility Services in Java SE* . . . . . 226
    - Utility Services in Java EE* . . . . . 227

Utility Services and Open-Source Frameworks . . . . .	229
<i>Spring Framework</i> . . . . .	229
<i>Transaction Management</i> . . . . .	229
<i>Data Access Objects</i> . . . . .	230
<i>Object-Relational Mapping</i> . . . . .	230
<i>JMS</i> . . . . .	230
<i>JMX</i> . . . . .	230
<i>JCA</i> . . . . .	231
<i>Spring MVC</i> . . . . .	231
<i>Hibernate</i> . . . . .	231
<i>Commons Logging and Log4J</i> . . . . .	231
Utility Services as Web-Based Services . . . . .	231
<i>Sending XML Data as a String</i> . . . . .	232
<i>Utilizing &lt;xsd:any/&gt;</i> . . . . .	233
<i>Provider-Style Web Service Logic in JAX-WS</i> . . . . .	234
<i>Building REST Utility Services</i> . . . . .	236
Testing Considerations . . . . .	238
Packaging Considerations . . . . .	239
8.3 Utility Service Types . . . . .	240
Omni Utility Services . . . . .	240
<i>Design Considerations</i> . . . . .	241
<i>Service Implementation</i> . . . . .	241
<i>Service Consumption</i> . . . . .	241
Resource Utility Services . . . . .	248
<i>Persistence/Data Access Resources</i> . . . . .	248
<i>Messaging Resources</i> . . . . .	248
<i>Transaction Resources</i> . . . . .	249
<i>Design Considerations</i> . . . . .	249
<i>Service Implementation</i> . . . . .	251
<i>Service Consumption</i> . . . . .	251
Micro-Utility Services . . . . .	253
<i>Design Considerations</i> . . . . .	253
<i>Service Implementation</i> . . . . .	253
<i>Service Consumption</i> . . . . .	254
Wrapper Utility Services . . . . .	257
<i>Design Considerations</i> . . . . .	258
<i>Service Implementation</i> . . . . .	259
<i>Service Consumption</i> . . . . .	259

**CHAPTER 9: Entity Services with Java . . . . .261**

- 9.1 Inside the Java Entity Service . . . . . 262
  - Architectural Considerations . . . . . 263
  - Domain Entities vs. Message Entities. . . . . 265
  - Data Aggregation . . . . . 266
  - Data Access Modes. . . . . 267
  - Change Notifications . . . . . 268
- 9.2 Java Entity Service Design and Implementation. . . . . 270
  - Entity Service Design . . . . . 270
    - Designing Domain Entities and Message Entities* . . . . . 271
    - Designing Stateless Entity Services* . . . . . 272
    - Designing Business-Relevant Entity Services* . . . . . 273
    - Designing Generic Entity Services* . . . . . 273
    - Designing Aggregating Entity Services* . . . . . 275
  - Entity Service Implementation . . . . . 278
    - Java Editions* . . . . . 278
  - Entity Services as Web-Based Services . . . . . 282
  - Entity Web Services Using SOAP . . . . . 283
  - REST Entity Services . . . . . 291
    - Read-Only and Read-Write Resources*. . . . . 292
    - Resource Granularity*. . . . . 292
    - Resource Creation and Location*. . . . . 292
    - Request Message*. . . . . 292
    - Response Message* . . . . . 293
    - Resource Relationships* . . . . . 294
      - Request Message* . . . . . 294
      - Response Message* . . . . . 294
    - Resource Collections* . . . . . 295
      - Request Message* . . . . . 295
      - Response Message* . . . . . 295
    - Aggregate Entities*. . . . . 297
      - Request Message* . . . . . 298
      - Response Message* . . . . . 298
    - Open-Source Frameworks* . . . . . 302
  - Testing Considerations. . . . . 302
  - Java Packaging Considerations. . . . . 303



**PART III: SERVICE COMPOSITION AND INFRASTRUCTURE**

**CHAPTER 10: Task Services with Java . . . . .307**

- 10.1 Inside a Task Service. . . . . 308
  - Performance Considerations . . . . . 315
- 10.2 Building Task Services . . . . . 316
  - Implementation Considerations . . . . . 316
  - Web-Based Task Services . . . . . 320
    - Task Services with SOAP and WSDL* . . . . . 320
    - Task Services with REST* . . . . . 324
  - Testing Considerations . . . . . 332
  - Packaging Considerations . . . . . 334

**CHAPTER 11: Service Composition with Java . . . . .335**

- 11.1 Inside Service Compositions . . . . . 336
  - Service Composition Roles . . . . . 336
  - Compositions and MEPs . . . . . 337
  - Synchronous and Asynchronous Invocation . . . . . 338
  - Service Level Agreements (SLAs) . . . . . 339
- 11.2 Java Service Composition Design and Implementation . . . . . 340
  - Composition Logic: Coding vs. Orchestration . . . . . 340
  - REST Service Composition Considerations . . . . . 341
  - Composition Member Endpoints . . . . . 344
  - Error Handling . . . . . 345
  - Schema Type Reuse . . . . . 353
  - Web-Based Services vs. Java Components . . . . . 359
  - Packaging, Testing and Deploying Composed Services. . . . . 362
- 11.3 Service and Service Composition Performance Guidelines . . . . . 368
  - Measuring Performance. . . . . 368
  - Testing Performance . . . . . 370
  - Caching . . . . . 371
    - Data Grids*. . . . . 371
    - REST Caching*. . . . . 372

Scaling Out Services with State . . . . .374  
 Handling Failures . . . . .375  
 Parsing and Marshaling . . . . .376

**CHAPTER 12: ESB as SOA Infrastructure . . . . .379**

12.1 Basic Traditional Messaging Frameworks . . . . .380  
 RPC vs. Messaging . . . . .381  
   *Technology Coupling* . . . . .382  
   *Spatial Coupling* . . . . .382  
   *Temporal Coupling* . . . . .382  
   *Message Producers and Message Consumers* . . . . .385  
 12.2 Basic Service Messaging Frameworks . . . . .389  
 Basic Service Message Processing without ESBs. . . . . 389  
   *Message Routing without an ESB.* . . . . .390  
   *Message Transformation without an ESB* . . . . .391  
 Basic Service Message Processing with ESBs . . . . . 392  
   *Message Routing with an ESB* . . . . .392  
   *Message Transformation with an ESB.* . . . . .392  
 12.3 Common ESB Features Relevant to SOA . . . . . 397  
 Service Lookup and Invocation . . . . .397  
 Service Processing . . . . . 399  
 Service Composition Support . . . . .401  
 REST API Management Support . . . . . 402

**PART IV: APPENDICES**

**APPENDIX A: Case Study Conclusion . . . . .405**

A.1 NovoBank . . . . .406  
 A.2 SmartCredit Co. . . . .407

**APPENDIX B: Service-Oriented Principles Reference . . . . .409**

- APPENDIX C: SOA Design Patterns Reference . . . . .425**
  
- APPENDIX D: The Annotated SOA Manifesto. . . . .519**
  
- About the Authors . . . . .533**
  - Thomas Erl. . . . .533
  - Andre Tost . . . . .533
  - Satadru Roy . . . . .534
  - Philip Thomas . . . . .534
  
- About the Foreword Contributor . . . . .535**
  - Mark Little. . . . .535
  
- About the Contributors . . . . .537**
  - Raj Balasubramanian. . . . .537
  - David Chou . . . . .537
  - Thomas Plunkett . . . . .538
  
- Index. . . . .539**

# Foreword by Mark Little

Over the past decade SOA has moved out of the hype and into mainstream development, becoming a standard part of any architect's repertoire. Although its roots go back beyond CORBA, SOA came to the forefront with the advent of two technology waves: Web Services and the Enterprise Service Bus (ESB). As REST picked up adoption beyond the Web, so it too became useful within the SOA paradigm. Of course, with the rise of cloud and mobile industries, SOA's applicability has continued to grow and we are seeing more and more business-critical deployments that are examples of good SOA implementations.

In parallel to the evolution and adoption of SOA, the Java language and platform have grown in maturity and dominance within the industry. In fact, the first ESB was written in Java and so are the most popular Web service stacks, such as Apache CXF. As relevant standards have evolved, such as Java Enterprise Edition, so too have they embraced more and more service-orientation principles, offering new APIs and development approaches. It is therefore no surprise that SOA and Java appear together more often than not in real-world deployments across a range of industry sectors.

Building good service-oriented applications is often a challenge, not just for individual developers but also at the organizational level. Even with a range of interfaces, frameworks, and other tools aimed at making SOA development easier, it is often difficult to know which to use and when. In terms of open standards, Java offers the richest suite of such aids, but at the same time can complicate a developer's life by their sheer number.

As Shakespeare once said: “We suffer a lot the few things we lack and we enjoy too little the many things we have.”

The authors of this book have created a great resource for anyone new to Java to learn how to translate SOA practices into reality. Likewise, it’s a must-have book for Java developers to see how they can use their Java experiences when developing service-oriented applications. Rather than just throw out example code for this or that framework in isolation, the authors have pulled together some wonderful case studies. These case studies are used throughout the book and help put everything into perspective.

It’s nice to see that the authors have also made this book self-contained. Far too often we are introduced to books, particularly those in a series, that can only be read if you’ve also purchased other books that came before and perhaps even some that come afterwards. Not so here. We have introductory chapters that spend a lot of time ensuring that even the most novice reader has all of the information they need at their fingertips. Very refreshing, especially when you consider that some of these chapters could be books in their own right! So whether you are a Java expert who needs to learn about SOA, or an SOA expert who needs to learn about Java, this book has all of the details you need. The Java background text covers the latest Java EE 7 specification, including topics such as JTA/JCA, EJB3, CDI, JPA, and JMS.

One of the dominant SOA implementation approaches today remains SOAP-based Web services. The WS-\* standards are the most widely adopted cross-vendor standards relevant to SOA and are used extensively in enterprises both inside and outside of the cloud. This book pulls together the often confusing plethora of standards and implementation approaches in a coherent manner, offering the reader an easier on-ramp to this critical area than is available elsewhere. In terms of how Java interfaces with Web-based services, the authors cover standards such as JAX-WS, JAX-RS and their Spring equivalents. Copious amounts of code are used to help illustrate critical points such as how to get the best throughput or security for your applications. The chapter called Service-Oriented Principles with Java Web-Based Services is one of the best of its kind, covering details such as WSDL, data mapping with REST and how to achieve loose coupling.

What also helps to positively differentiate this book from others in this area are chapters on how to build specific types of services in Java following service-orientation principles and getting them to perform well at the same time: this isn’t a book that is just heavy on theory, but it more than complements any theory with a very heavy practical aspect, making it ideal for architects and developers alike. If you’re looking for how to

build individual services or, more likely as your applications grow, composite services in a way that does not break service-orientation principles, then the authors have created a book that you need to have on your shelf, well read, or next to your computer. The code in chapters such as Service Composition with Java is presented in such a way and with the right context that you should be able to use much of it as a template for you and your teams.

Another technology which is often used within SOA deployments, but also just as often misunderstood, is the Enterprise Service Bus. Like REST, some people love the ESB whilst others believe they bring nothing of value. This book helps to shed some light on the ESB and inject some much needed reality to the debate, allowing you to select the right tool for the right job. ESBs aren't a global panacea for everything SOA-related. But neither are they something which should be ignored, and the authors help to make this clear in a succinct manner.

Whether you're a Java developer looking for a book to explain SOA or an SOA expert looking for something to help you turn those ideas into a Java reality, this is the book for you. It also doesn't matter whether you are a software developer, an architect, working in a team, or by yourself, this book works on many levels and is relevant to you as you move through the SOA design and development phases.

Dr. Mark Little, Red Hat

*This page intentionally left blank*

# Acknowledgments

- Claude Baillargeon, La Capitale
- Chris Barran
- Antonio Bruno, UBS Wealth Management
- Jason Burrows, Pacific Blue Cross
- David Chou, Microsoft
- Leonid Felikson
- Olaf Heimburger, Oracle
- Khanderao Kand, Guavus
- Anish Karmarkar, Oracle
- Kris Kothumb
- Dennis Lamarre, System Science Corporation
- Briana Lee
- Beth Liang, HeathWallace
- Karl Lopes
- Anil Luthra
- Damian Maschek, Deutsche Bahn System GmbH
- Veeru Mehta
- Paulo Merson
- Prakash Narayan
- Songlin Qiu
- Antony Reynolds, Oracle
- Rich Rosen
- Sam Rostam
- Sanjay Singh



- Dennis Sosnoski
- Kanu Tripathi
- Sameer Tyagi
- Clemens Utschig, Boehringer Ingelheim Pharma
- Matt Wright, Oracle
- Kareem Yusuf PhD, IBM

# Chapter 1



## Introduction

- 1.1 About This Book
- 1.2 Prerequisite Reading
- 1.3 How Principles and Patterns Are Used in This Book
- 1.4 Symbols and Figures
- 1.5 Additional Information

## **1.1 About This Book**

The Java platform has evolved significantly over the past decade and has become a preferred platform for building Web-based enterprise applications. The service-orientation paradigm introduces a number of requirements and design principles that formalize the application of Java platforms and technologies in support of the strategic goals of service-oriented computing. This book explores service-oriented solution design and implementation through the application of techniques and best practices that utilize Java technology advances.

### **Objectives of This Book**

This book was written with the following primary goals in mind:

- to provide coverage of contemporary Java distributed technologies and modern service technologies
- to document the application of service-orientation principles to the Java technology platform
- to explore the application of SOA design patterns to various Java technologies and solutions built with these technologies
- to provide coverage of Java infrastructure extensions relevant to service-oriented solutions

### **Who This Book Is For**

This book can be used as a tutorial and reference text intended for the following types of readers:

- experienced Java developers and architects who want to learn how to apply service-orientation principles and select SOA design patterns in order to create shared services and service-oriented solutions
- enterprise architects that want to learn more about how to position and establish Java platforms and technologies within an IT enterprise that has or will undergo SOA adoption
- developers who want to build solutions using modern Java service technologies

## What This Book Does Not Cover

This is not a “how-to” book for Java. Although the chapters in Part I contain a great deal of introductory coverage of Java distributed technologies, the overall purpose of this book is to explore the intersection of Java technologies and the application of service-orientation principles and SOA design patterns. You should not proceed to read this book if you are new to Java.

## 1.2 Prerequisite Reading

This book assumes you have a basic knowledge of:

- Java technologies, APIs, and standards
- fundamental XML concepts
- fundamental service-orientation

If you have not yet worked with XML, you can read some of the brief tutorials published at [www.servicetechspecs.com](http://www.servicetechspecs.com). If you are new to SOA, you can get a basic understanding of service-oriented computing, service-orientation, and related design patterns by studying the content at the following Web sites:

- [www.serviceorientation.com](http://www.serviceorientation.com)
- [www.soapatterns.org](http://www.soapatterns.org)
- [www.soa-manifesto.com](http://www.soa-manifesto.com)

To ensure that you have a clear understanding of key terms used and referenced in the upcoming chapters, you can also visit the online master glossary for this book series at [www.soaglossary.com](http://www.soaglossary.com) to look up definitions for terms that may not be fully described in this book.

Here are some recommendations for additional books that elaborate on key topics covered by this title:

- *SOA Principles of Service Design* – A comprehensive documentation of the service-orientation design paradigm with full descriptions of all of the principles referenced in this book. Concise profiles of these principles are provided in Appendix B.

- *SOA Design Patterns* – This is the official SOA design pattern catalog containing descriptions and examples for most of the patterns referenced in this book. You can also look up concise descriptions for these patterns at [www.soapatterns.org](http://www.soapatterns.org) and in Appendix C.
- *Web Service Contract Design & Versioning for SOA* – Any content pertaining to WSDL, SOAP, and XML Schema development and design, development, and versioning will be aided by the detailed coverage in this title.
- *SOA with REST: Principles, Patterns & Constraints for Building Enterprise Solutions with REST* – REST service development and implementation is covered from various perspectives in the upcoming chapters. This book provides detailed coverage of associated REST technology, architecture, and design constraints.
- *Service-Oriented Architecture: Concepts, Technology, and Design* – The coverage of service-oriented analysis and design processes in this title supplements the technology-centric focus of this book with methodology-related topics.

For more information, visit [www.servicetechbooks.com](http://www.servicetechbooks.com).

## How This Book Is Organized

This book begins with Chapters 1 and 2 providing introductory content and case study background information respectively. All subsequent chapters are grouped into the following parts:

- Part I: Fundamentals
- Part II: Services
- Part III: Service Composition and Infrastructure
- Part IV: Appendices

## Part I: Fundamentals

The four chapters in this part cover introductory topics related to SOA, service-orientation, Web-based services technology, and Java. Note that these chapters do not contain case study content.

*Chapter 3: Fundamental SOA Concepts*

This chapter provides an overview of key terms and concepts associated with SOA, service-orientation, and service-oriented computing.

*Chapter 4: Basic Java Distributed Technologies*

With an emphasis on Java EE, this chapter introduces Java technologies and APIs relevant to this book. Open-source and commercial Java vendor platforms are further referenced.

*Chapter 5: Web-Based Service Technologies*

This chapter provides a concise overview of the industry standards and conventions that form the basis for SOAP-based Web services and REST services. A series of examples is provided for each service implementation medium.

*Chapter 6: Building Web-Based Services with Java*

This chapter focuses on the JAX-WS standard for SOAP-based Web services and the JAX-RS standard for REST services, and provides numerous examples that demonstrate basic design and implementation.

**Part II: Services**

This part of the book contains a set of chapters that highlight key Java design considerations introduced by service-orientation principles, and further explores their application to the creation of Java-based utility and entity services.

*Chapter 7: Service-Orientation Principles with Java Web-Based Services*

This chapter documents each of the eight service-orientation design principles as they affect and relate to the use of Java-based technology for both SOAP-based Web services and REST services.

*Chapter 8: Utility Services with Java*

The unique design considerations that come with building utility services with Java are addressed in a series of sections that explore architecture, design, and implementation, further supported by a number of examples.

### *Chapter 9: Entity Services with Java*

This chapter discusses the development of entity services with Java with an emphasis on service contract and message standardization, data access, and processing.

## **Part III: Service Composition and Infrastructure**

The final set of chapters build upon Parts I and II to tackle the design and implementation of entire Java-based service-oriented solutions. Topics covered range from service composition to ESB-related infrastructure considerations.

### *Chapter 10: Task Services with Java*

This chapter discusses the characteristics, design, and implementation considerations of building task services with Java and further addresses preliminary service composition issues in preparation for Chapter 11.

### *Chapter 11: Service Composition with Java*

This chapter covers a spectrum of topics that pertain to and further expand upon many of the areas already covered in preceding chapters to focus on the design, implementation, and performance of services aggregated into Java-based service compositions.

### *Chapter 12: ESB as SOA Infrastructure*

This chapter highlights the use of ESB-based middleware and infrastructure in relation to service-oriented solution and service composition architecture and implementation.

## **Part IV: Appendices**

### *Appendix A: Case Study Conclusion*

This appendix provides a conclusion of the case study storylines first established in Chapter 2.

### *Appendix B: Service-Oriented Principles Reference*

This appendix provides the profile tables (originally from *SOA Principles of Service Design*) for the service-orientation design principles referenced in this book.

### *Appendix C: SOA Design Patterns Reference*

This appendix provides the profile tables from the design patterns documented in the books *SOA Design Patterns* and *SOA with REST: Principles, Patterns & Constraints for Building Enterprise Solutions with REST*.

### *Appendix D: The Annotated SOA Manifesto*

This appendix provides the annotated version of the SOA Manifesto declaration, which is also published at [www.soa-manifesto.com](http://www.soa-manifesto.com).

## **1.3 How Principles and Patterns Are Used in This Book**

As part of its exploration of service-orientation in relation to Java technologies and platforms, this book references and uses established design principles and patterns throughout its chapters.

Service-orientation design principles and SOA design patterns are always capitalized, as follows:

*“...the Service Loose Coupling principle is supported via the application of the Decoupled Contract pattern...”*

Note, as also demonstrated in this sample statement, a principle or a pattern can be referenced with or without being qualified. In other words, the statement *“...when the Decoupled Contract pattern is applied...”* has the same meaning as *“...when Decoupled Contract is applied...”*

## **1.4 Symbols and Figures**

This book contains a series diagrams that are referred to as *figures*. The primary symbols used throughout all figures are individually described in the symbol legend located on the inside of the front cover.

Symbol legend posters can also be downloaded from [www.arcitura.com/notation](http://www.arcitura.com/notation).

## **1.5 Additional Information**

These sections provide supplementary information and resources for the *Prentice Hall Service Technology Series from Thomas Erl*.



**Updates, Errata, and Resources ([www.servicetechbooks.com](http://www.servicetechbooks.com))**

Information about other series titles and various supporting resources can be found at the official book series Web site: [www.servicetechbooks.com](http://www.servicetechbooks.com). You are encouraged to visit this site regularly to check for content changes and corrections.

**Referenced Specifications ([www.servicetechspecs.com](http://www.servicetechspecs.com))**

This site provides a central portal to the original specification documents created and maintained by primary industry standards organizations.

**The Service Technology Magazine ([www.servicetechmag.com](http://www.servicetechmag.com))**

The Service Technology Magazine is a monthly publication provided by Arcitura Education Inc. and Prentice Hall and is officially associated with the *Prentice Hall Service Technology Series from Thomas Erl*. The Service Technology Magazine is dedicated to publishing specialized articles, case studies, and papers by industry experts and professionals.

**Service-Orientation ([www.serviceorientation.com](http://www.serviceorientation.com))**

This site provides papers, book excerpts, and various content dedicated to describing and defining the service-orientation paradigm, associated principles, and the service-oriented technology architectural model, including online access to the service-orientation principle profiles published in Appendix B.

**What Is REST? ([www.whatisrest.com](http://www.whatisrest.com))**

This Web site provides a concise overview of REST architecture and the official constraints that are referenced throughout this book.

**What Is Cloud? ([www.whatiscloud.com](http://www.whatiscloud.com))**

A reference site dedicated to fundamental cloud computing topics.

**SOA and Cloud Computing Design Patterns ([www.soapatterns.org](http://www.soapatterns.org), [www.cloudpatterns.org](http://www.cloudpatterns.org))**

The official SOA and cloud computing design pattern catalogs are published on these sites. The [www.soapatterns.org](http://www.soapatterns.org) site provides online access to the pattern profiles published in Appendix C.

**SOA Certified Professional (SOACP) ([www.soaschool.com](http://www.soaschool.com))**

The official site for the SOA Certified Professional (SOACP) curriculum dedicated to specialized areas of service-oriented architecture and service-orientation, including analysis, architecture, governance, security, development, and quality assurance.

**Cloud Certified Professional (CCP) ([www.cloudschool.com](http://www.cloudschool.com))**

The official site for the Cloud Certified Professional (CCP) curriculum dedicated to specialized areas of cloud computing, including technology, architecture, governance, security, capacity, virtualization, and storage.

**Big Data Science Certified Professional (BDSCP)  
([www.bigdatascienceschool.com](http://www.bigdatascienceschool.com))**

The official site for the Big Data Science Certified Professional (BDSCP) curriculum dedicated to specialized areas of data science, analytics, data analysis, machine learning, and Big Data solution engineering.

**Notification Service**

To be automatically notified of new book releases in this series, new supplementary content for this title, or key changes to the aforementioned resource sites, use the notification form at [www.servicetechbooks.com](http://www.servicetechbooks.com) or send a blank e-mail to [notify@arcitura.com](mailto:notify@arcitura.com).

*This page intentionally left blank*

*This page intentionally left blank*



## Service-Oriented Principles with Java Web-Based Services

- 7.1 Service Reusability
- 7.2 Standardized Service Contract
- 7.3 Service Loose Coupling
- 7.4 Service Abstraction
- 7.5 Service Composability
- 7.6 Service Autonomy
- 7.7 Service Statelessness
- 7.8 Service Discoverability

**B**uilding services for service-oriented solutions requires the application of the service-orientation paradigm whose established design principles drive many Java service contract and implementation decisions. In certain cases, the programming language and runtime environment used for services can also be influenced by these guiding principles. This chapter visits each of the eight service-orientation principles in depth to highlight considerations specific to design and development with Java.

#### NOTE

The service-orientation principles are formally documented in the series title *SOA Principles of Service Design*. Concise profiles of the principles are also available in Appendix B and at [www.serviceorientation.com](http://www.serviceorientation.com).

## 7.1 Service Reusability

The following are common design characteristics associated with reusable services:

- The service is defined by an agnostic functional context.
- The service logic is highly generic.
- The service has a generic and extensible contract.
- The service logic can be accessed concurrently.

Let's take a closer look at each of these characteristics in relation to Java.

### Agnostic Functional Contexts

Ensuring that the logic encapsulated by a service is agnostic to any particular functional context allows for the building of service interfaces independent of one particular business process or functional domain. The term "context" refers to the service's functional scope. An agnostic functional context is not specific to any one purpose and is therefore considered multi-purpose. A non-agnostic functional context, on the other hand, is intentionally single-purpose.

A checkpoint that can be part of a regular code review or service interface quality gate is to look at the imports in a Java interface or implementation class for a service. Java

interfaces and classes are often structured according to the applicable business domain, and focusing on the list of imported packages can help identify dependencies in the code. Returning to the simplified order management example, the Java service interface for the Credit Check service is seen in Example 7.1.

```
import com.acme.businessobjects.om.Order;
public interface CreditCheck {
    public boolean hasGoodCredit(Order order);
}
```

---

### Example 7.1

The `import` statement indicates that the service logic depends on the functional context of order management. Such a dependency cannot always be avoided if a service is developed specifically for a particular business domain at the cost of its reusability. Utility services are generally agnostic and reusable, as explained in Chapter 8.

## Highly Generic Service Logic

Generic Java service logic refers to logic independent of its service contract. In the Java world, this means that a Java service interface is created with no mapping for the data types referenced in the service contract.

The `javax.xml.ws.Provider` interface avoids dependency on the service contract when using the JAX-WS programming model for SOAP-based Web services. An incoming message can be received by the service as a SAAJ `javax.xml.soap.SOAPMessage` with the `Provider` interface, which allows the entire message to be parsed or navigated as a DOM tree, as seen in Example 7.2.

```
@ServiceMode(value=Service.Mode.MESSAGE)
@WebServiceProvider()
public class GenericProviderImpl implements
    Provider<javax.xml.soap.SOAPMessage> {
    public SOAPMessage invoke(SOAPMessage message) {
        // read and process SOAPMessage...
    }
}
```

---

### Example 7.2

For the same SOAP-based Web service, the request message can be alternatively read as a `javax.xml.transform.Source`. As shown in Example 7.3, the message can be treated

as a plain XML document with no relationship to SOAP. Only the payload of the request message can be retrieved. Developers can ignore the SOAP envelope or SOAP header to focus on the content in the body of the message.

```
@ServiceMode(value=Service.Mode.PAYLOAD)
@WebServiceProvider()
public class GenericProviderImpl implements
    Provider<javax.xml.transform.Source> {
    public Source invoke(Source source) {
        // read and process SOAPMessage...
    }
}
```

---

### Example 7.3

In both Examples 7.2 and 7.3, the response data returns in the same way the request data was received. If the request is received in an object of type `SOAPMessage`, then a new `SOAPMessage` object must be built for the response. Correspondingly, a new `Source` object must be returned if `Source` is used.

Generic Java types can capture the appropriate MIME media type or resource representation format produced or consumed by a target resource when building a REST service. For text-based request/response entities, Java `String`, `char[]`, and the character-based `java.io.Reader` or `Writer` interfaces can be used in the resource methods. For completely generic entity representations, which can include binary content, a `java.io.InputStream`, `OutputStream`, or a raw stream of bytes can be used as `byte[]`. For XML-based resource representations, the `javax.xml.transform.Source` type can be used to handle XML documents at a higher level than a raw stream.

As seen in Example 7.4, a slightly reworked customer resource example of the REST service from the Chapter 6 uses an `InputStream`. The contents of an entity are extracted in the incoming request to keep the service contract generic.

```
@Post
@Consumes("application/xml")
public void createCustomer(
    InputStream in){
    //extract customer from request
    Customer customer = extractCustomer(in);
    //create customer in system;
}
```

---

### Example 7.4



Similarly, `javax.xml.transform.Source` can extract the customer information from the incoming request. JAX-RS relies on a bottom-up service design philosophy for building REST APIs except when using WADL. Using generic types, such as `java.lang.Object` or `byte[]`, on a JAX-RS resource interface should be sufficient to keep the service contract generic. However, consider what corresponding data types will be used in the WSDL for SOAP-based Web services.

Avoid the use of concrete data types on the Java service interface. The payload of a message is cast into a Java object, such as a `byte` array or `string`. The service contract, such as the WSDL for a Web service, must match the generic type, such as `java.lang.Object` maps to `xsd:anyType`, `byte[]`, which maps to `xsd:hexBinary`, and `java.lang.String` maps to `xsd:string`. The matching generic types require specific code to be developed in both the service consumer and service for the data to be inserted into the request/response messages.

In Example 7.5, the public class employs a `byte` array on its interface to hide the details of the data processed by the service.

```
@WebService
public class OrderProcessing {
    public void processOrder( Order order,
        byte[] additionalData) {
        // process request...
    }
}
```

---

### Example 7.5

Supertypes in the service interface can aid in generalizing a service contract. For example, a service returns detailed account information for a bank's customer. When creating a data model for the different types of information provided by the different types of accounts, take advantage of inheritance in XML Schema. A superclass called `Account` can be created in Java, with a number of subclasses defined for each type of account, such as checking, savings, loans, and mortgages. A return type of `Account` which includes all of the different types of accounts can be specified in the service interface.

The considerations for supertypes are the same for both SOAP and REST services. As in both cases, the XML Java marshaling/unmarshaling is handled by JAXB for XML and MOXy or JSON-P for JSON. MOXy is a framework for marshaling/unmarshaling between JSON and Java objects. JSON-P (Java API for JSON Processing) supports low-level JSON parsing in Java EE 7.

A generic service implementation can serve multiple types of request/response messages, which generally increases reuse opportunities. However, the service logic must be implemented to handle different types of messages. Type-specific code uses language-specific data types. Tooling can generate code that automatically parses messages into the right Java objects, although at the cost of increased coupling. If generic types are used, the processing of incoming and outgoing messages is left to the service implementer. However, generic types offer greater flexibility in terms of type-independence and loose coupling.

Generalizing service logic also applies to the service consumer. For example, JAX-WS defines a generic service invocation API using the `javax.xml.ws.Dispatch` interface. Services can be invoked with unknown interfaces when the service consumer code is developed. Similar to how the use of the `Provider` interface supports handling requests from different types of service consumers, the use of the `Dispatch` API allows a service consumer to interact with different types of services. For REST service clients, if a JAX-RS implementation is used, all the generic Java types the JAX-RS implementation supports can be used to build requests and consume responses. However, generic service logic requires the client code to handle different types of messages and have some knowledge about the message formats expected.

### Generic and Extensible Service Contracts

Service logic can be made generic for reuse across a wide range of scenarios and adapted to changes in its environment, such as by changing and evolving services or service consumers. A service contract can be made generic by restricting the dependencies on data types referred to in the service contract and limiting the services composed inside the service logic to a minimum. When translated into Java-specific terms, reduce or eliminate the number of business-domain-specific classes in the service implementation.

Creating a generic service contract means applying generic types like `string`, `hexBinary`, or `anyType` in the schema type definitions for a Web service. Alternatively, message formats can be defined in a service contract with schema inheritance to use common supertypes, and the runtime allowed to determine which concrete subtype is used. Generic types are not only true for the top-level elements in a message but can also be used within a type definition. In Example 7.6, the schema describes a `Customer` type with a number of well-defined fields and a generic part.

```
<xs:complexType name="Customer">
  <xs:sequence>
    <xs:element name="accounts" type="ns:Account"
      nillable="true" maxOccurs="unbounded"
      minOccurs="0"/>
    <xs:element name="address" type="ns:Address"
      minOccurs="0"/>
    <xs:element name="customerId" type="xs:string"
      minOccurs="0"/>
    <xs:element name="name" type="ns:Name" minOccurs="0"/>
    <xs:element name="orderHistory" type="ns:OrderArray"
      nillable="true" maxOccurs="unbounded"
      minOccurs="0"/>
    <xs:any/>
  </xs:sequence>
</xs:complexType>
```

---

**Example 7.6**

When used in a service contract, the schema in Example 7.6 allows the service to process messages that have a fixed part at the beginning and a variable part at the end, which is represented by the `<xs:any>` element. SOAP-based Web services can use the Dispatch APIs in the service logic and Provider APIs in the service consumer logic without affecting how the service contract is built.

Service consumer logic can be implemented generically even if a detailed and specific service contract is provided. For REST services, the same considerations hold true for resource representations. XML-based representations can use highly specific types while the JAX-RS resource class can leverage generic Java types. The programmer then becomes responsible for performing the type mapping between XML and Java in the service logic.

**Concurrent Access to Service Logic**

A particular instance of a shared service will almost always be used by multiple service consumers simultaneously at runtime. How the service is deployed and the characteristics of the service's runtime environment as influences on service reuse are significant design considerations.

For example, each service request starts a new process that executes the service logic for the request, which ensures that the processing of one request does not affect the processing of another request. Each execution is completely independent of other executions and creates a new thread in the process. However, starting a new process is a

relatively expensive operation in terms of system resources and execution time in most runtime environments. Sharing services in this way is inefficient.

All service requests executed within the same process share all the resources assigned to that process, which provides a lightweight method of serving multiple concurrent requests to the same service. Starting a new thread is an inexpensive operation on most systems. Most, if not all, Web service engines work this way. Implementing services in a multithreaded environment requires adherence to the basic rules of concurrent Java programming.

### CASE STUDY EXAMPLE

As part of an initiative to comply with recently introduced legal obligations, the NovoBank IT team decides to create a Document Manager service that stores documents for auditing purposes. The service supports storing and retrieving XML documents. To maximize the reusability of this service, the NovoBank IT team creates a generic service contract and flexible service implementation to handle different types of documents, which can be extended over time.

Initially, the service will have the operations of `store` and `retrieve`. The service contract does not imply any structure of the documents stored, as shown in Example 7.7.

```
<definitions targetNamespace="http://utility.services.novobank.com/"
  name="DocumentManagerService" xmlns:tns="http://utility.services.
  novobank.com/" xmlns:xsd="http://www.w3.org/2001/XMLSchema"
  xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/"
  xmlns="http://schemas.xmlsoap.org/wsdl/">
  <types>
    <xs:schema version="1.0" targetNamespace="http://utility.
      services.novobank.com/" xmlns:xs="http://www.w3.org/2001/
     /XMLSchema">
      <xs:element name="retrieve" type="ns1:retrieve"
        xmlns:ns1="http://utility.services.novobank.com/" />
      <xs:complexType name="retrieve">
        <xs:sequence>
          <xs:element name="arg0" type="xs:string" minOccurs="0" />
        </xs:sequence>
      </xs:complexType>
      <xs:element name="retrieveResponse" type="ns2:retrieveResponse"
        xmlns:ns2="http://utility.services.novobank.com/" />
    </xs:schema>
  </types>

```

```

    <xs:complexType name="retrieveResponse">
      <xs:sequence>
        <xs:any processContents="skip"/>
      </xs:sequence>
    </xs:complexType>
  <xs:element name="store" type="ns3:store"
    xmlns:ns3="http://utility.services.novobank.com/" />
  <xs:complexType name="store">
    <xs:sequence>
      <xs:any processContents="skip"/>
    </xs:sequence>
  </xs:complexType>
  <xs:element name="storeResponse" type="ns4:storeResponse"
    xmlns:ns4="http://utility.services.novobank.com/" />
  <xs:complexType name="storeResponse">
  </xs:schema>
</types>
<message name="store">
  <part name="parameters" element="tns:store"/>
</message>
<message name="storeResponse">
  <part name="parameters" element="tns:storeResponse"/>
</message>
<message name="retrieve">
  <part name="parameters" element="tns:retrieve"/>
</message>
<message name="retrieveResponse">
  <part name="parameters" element="tns:retrieveResponse"/>
</message>
<portType name="DocumentManager">
  <operation name="store">
    <input message="tns:store"/>
    <output message="tns:storeResponse"/>
  </operation>
  <operation name="retrieve">
    <input message="tns:retrieve"/>
    <output message="tns:retrieveResponse"/>
  </operation>
</portType>
<binding name="DocumentManagerPortBinding"
  type="tns:DocumentManager">
  ...
</binding>
<service name="DocumentManagerService">
  ...

```

```
</service>
</definitions>
```

**Example 7.7**

NovoBank's service contract for the Document Manager service is deliberately generic.

Example 7.7 shows that the content of the stored messages is represented by an `<xs:any/>` element, which means that any well-formed XML can be inserted in the content. The `<xs:any/>` element allows the Document Manager service to be reused across many different types of messages and documents.

The development team decides to create a flexible implementation of the service which is independent of the specific type of document being sent or retrieved. The service must extend without affecting the existing implementation, preparing support for more specific processing of specific document types. A flexible implementation is achieved with a handler and factory. The implementation leverages the `javax.xml.ws.Provider` interface and delegates the processing of each message to a handler. The handler instance is then retrieved via a factory.

Example 7.8 shows the implementation class for the Document Manager service in the detailed source code.

```
package com.novobank.services.utility;
import javax.xml.transform.Source;
import javax.xml.ws.Provider;
import javax.xml.ws.Service;
import javax.xml.ws.ServiceMode;
import javax.xml.ws.WebServiceProvider;

@ServiceMode(value=Service.Mode.PAYLOAD)
@WebServiceProvider
public class DocumentManager implements Provider<Source> {
    public Source invoke(Source source) {
        DocumentHandler handler =
            DocumentHandlerFactory.instance().getHandler(source);
        handler.process(source);
        return null;
    }
}
```

**Example 7.8**

The implementation of the Document Manager service utilizes a factory to retrieve a handler.

This class implements the `Provider` interface, so that any invocation of the service is directed to the `invoke()` method. The `@ServiceMode` annotation indicates that only the actual content of the SOAP body should be passed into the implementation.

In the implementation of the `invoke()` method, a single instance of a class called `DocumentHandlerFactory` is used to retrieve a `DocumentHandler` for this message. Example 7.9 shows the source code for the `DocumentHandler` interface.

```
package com.novobank.services.utility;
import javax.xml.transform.Source;
public interface DocumentHandler {
    public void process(Source source);
}
```

---

**Example 7.9**

The source code for the `DocumentHandler` interface

The `DocumentHandler` interface only defines the `process()` method, which processes the message. Different implementations of the interface that process messages in different ways can exist. The message content is passed into the `process()` method as a stream of type `javax.xml.transform.Source`. When revisiting the source code for the Document Manager service in Example 7.2, the message returned by the factory can be seen to be passed to the handler.

A part of the source code for the factory is presented in Example 7.10.

```
package com.novobank.services.utility;
import javax.xml.transform.Source;
public class DocumentHandlerFactory {
    protected static DocumentHandlerFactory theInstance = new
        DocumentHandlerFactory();
    public static DocumentHandlerFactory instance() {
        return theInstance;
    }
    protected DocumentHandlerFactory() {}
    public DocumentHandler getHandler(Source source) {
        DocumentHandler handler = null;
        // the code where the message is parsed and
        // the appropriate handler is retrieved would be here.
        return handler;
    }
}
```

```
}  
}
```

**Example 7.10**

The `DocumentHandlerFactory` source code

The message is parsed to the point where an appropriate handler can be found. Handlers are registered via several mechanisms that can include hardcoding into the factory class, retrieval from a file, or lookup in a registry. The concrete `DocumentHandler` implementation chosen for a particular message is often based on the root element of the passed message. The root element typically provides adequate indication of the nature of the message.

Using the factory mechanism, where the service implementation class calls a factory to retrieve a handler to process the message, allows new handlers to be later added without affecting the existing code. The combination of a generic service definition using the `<xs:any/>` element, a flexible service implementation using the `Provider` interface, and a factory to delegate the processing of the message ensures maximum reusability of the service across a variety of environments and domains.

---

### SUMMARY OF KEY POINTS

---

- For SOAP-based Web services, the JAX-WS standard offers ways of implementing service logic as generic and therefore capable of handling different types of messages.
  - For REST services, Java generic types can be used in JAX-RS-based resource implementations. The request/response entities are treated as a raw sequence of bytes or characters.
  - Using generic data types for domain entities or resource representations allows the service to be reused across a greater number of potential service consumers.
-



## 7.2 Standardized Service Contract

A foundational criterion in service-orientation is that a service have a well-defined and standardized contract. When building Web services with SOAP and WS-\*, portable machine-readable service contracts are mandatory between different platforms as WSDL documents and associated XML schema artifacts describing the service data model. The finite set of widely used HTTP verbs for REST services form an implicit service contract. However, describing the entity representations for capture in a portable and machine-independent format is the same as SOAP and WS-\*.

For REST services, capturing and communicating various aspects of resources can be necessary, such as the set of resources, relationships between resources, HTTP verbs allowed on resources, and supported resource representation formats. Standards, such as WADL, can be used to satisfy the mandatory requirements. Having a standards-based service contract exist separate from the service logic, with service data entities described in a platform-neutral and technology-neutral format, constitutes a service by common definition. Even the self-describing contract of HTTP verbs for a REST service establishes a standards-based service contract. Recall the standards used for service contracts, such as WSDL/WADL and XML Schema, from Chapter 5.

### Top-Down vs. Bottom-Up

Ensuring that services are business-aligned and not strictly IT-driven is necessary when identifying services for a service portfolio in an SOA. Services are derived from a decomposition of a company's core business processes and a collection of key business entities. For a top-down approach, services are identified and interfaces are designed by creating appropriate schema artifacts to model either the operating data and WSDL-based service contracts, or model REST resources and resource methods. The completed service interface is implemented in code.

However, enterprises can have irreplaceable mission-critical applications in place. Therefore, another aspect of finding services is assessing existing applications and components to be refactored as services for a bottom-up approach. This includes creating standard service contracts, such as WSDL definitions or REST resource models, for the existing components.

Tooling provides support for both approaches in a Java world. For SOAP-based Web services, tools play a more prominent role than in Java-based REST services. JAX-WS defines the `wsimport` tool, which takes an existing WSDL definition as input to generate Java skeletons. These skeletons can be used as the starting point for implementing the

actual service logic. Similarly, the `wsgen` tool generates WSDL from existing Java code. The mapping between WSDL/XML schema and Java is an important function associated with the `wsimport` tool.

Machine-readable contracts are also necessary for REST services. JAX-RS, if WADL is not used, starts with a resource model to implement the resources in Java. Consider the contract as a logical collection of the resource model, with the supported resource methods, resource representations, and any hyperlinks embedded in the representations allowing navigability between resources. If WADL is used, tools like `wadl2java` can generate code artifacts. Initiatives exist to help generate WADL from annotated JAX-RS classes for a bottom-up approach, although these recent developments can have limited usefulness.

Some SOA projects will employ both a bottom-up and a top-down approach to identify and design services and service contracts, which often results in a meet-in-the-middle approach. Service definitions and Java interfaces are tuned and adjusted until a good match is found.

Sometimes an XML schema definition developed as part of the service design cannot map well into Java code. Conversely, existing Java code may not easily map into an XML schema. Java code that does not precisely map to a service interface designed as part of a top-down approach can exist. In this case, the Service Façade pattern can be applied to insert a thin service wrapper to satisfy the service interface and adapt incoming and outgoing data to the format supported by the existing Java code.

## Mapping Between Java and WSDL

WSDL is the dominant method of expressing the contract of a Java component. While typically related to Web services, the language can also be utilized for other types of services. Formalization and standardization of the relationship between Java and WSDL has made this possible, such as the work completed on the JAX-RPC standard.

The JAX-RPC standard initially defined basic tasks, such as “a service portType is mapped to a Java interface” and “an operation is mapped to a method.” However, formalizing allows the definitions described in the JAX-RPC standard to define how an existing Java component (a class or interface) can generate a WSDL definition, and vice versa. Consequently, most contemporary Java IDEs support generating one from the other without requiring any manual work.

JAX-WS, the successor standard for JAX-RPC, builds on top of its predecessor's definitions and delegates all the issues of mapping between Java and XML to the JAXB specification (as discussed in Chapter 6). These sections serve to highlight some of the issues raised when creating standard service contracts from Java or creating Java skeletons from existing service contracts. The majority of the details explained in the next section apply specifically to Web services.

### Wrapped Document/Literal Contracts

The WSDL standard identifies a variety of styles for transmitting information between a service consumer and service. Most of the styles are specific for the chosen message and network protocol, and specified in a section of the WSDL definition called the binding. A common binding found in a WSDL definition uses SOAP as the message protocol and HTTP as the network transport. Assume that SOAP/HTTP is the protocol used for the services presented as examples.

The `portType` is a binding-neutral part of a service definition in WSDL that describes the messages that travel in and out of a service. Reusable across multiple protocols, the `portType` is not bound to the use of a Web service. Any service, even if invoked locally, can be described by a WSDL `portType`, which allows service interfaces to be defined in a language-neutral fashion regardless of whether the service logic will be implemented in Java or another language.

As discussed in Chapter 5, the WSDL binding information defines the message format and protocol details for Web services. For SOAP-based bindings, two key attributes known as the encoding style and the invocation style determine how messages are encoded and how services are invoked.

The wrapped document/literal style supported by default in all Java environments for services dictates that an exchange should be literal. Literal means that no encoding happens in the message, so the payload of the message is a literal instantiation of the schema descriptions in the `<types>` element of the WSDL. The invocation style is document. Document means that the runtime environment should generate a direct copy of the input and output messages as defined in the `portType` and not just an arbitrary part of the message. Wrapped means that the payload of the message includes a wrapper element with the same name as the operation invoked.

In order to understand how the WSDL standard relates to Java, let's review Example 7.11 to expose the following class as a service and create a standardized contract.

```
package pack;
import javax.jws.*;
@WebService
public class Echo {
    public String echo(String msg) {
        return msg;
    }
}
```

---

### Example 7.11

Using the wrapped document/literal style implements a wrapper element called "echo" after the `echo()` method in the public class. Echo is included in the XML schema associated with this service. An excerpt in the resulting schema is provided in Example 7.12.

```
...
<xs:element name="echo">
  <xs:complexType>
    <xs:sequence>
      <xs:element name="arg0" type="xs:string" minOccurs="0"/>
    </xs:sequence>
  </xs:complexType>
</xs:element>
...
```

---

### Example 7.12

Wrapping a message in one additional element named after the operation is prevalent and the default in any commonly used tool. Note that naming the global element after the operation is common practice and not required by the specification.

## Implicit and Explicit Headers

Transferring information as part of the `<Header>` portion of the SOAP message, to be added to the WSDL definition, is another important part of the binding information for SOAP. This section discusses how to bind the information with explicit, implicit, or no headers.

### *Explicit Headers*

Header data is part of the messages referenced in the `portType` of the service, which is often called an explicit header. The header definition in the SOAP binding refers to a message part either included in the input message or the output message of an operation.

In Example 7.13, assume an Echo service takes a `string` as input and returns that `string` as the response. A `timestamp` must also be added into the header of the SOAP request message, indicating the time at which the request was sent.

```
<definitions targetNamespace="http://pack/" name="EchoService"
  xmlns:tns="http://pack/" xmlns:xs="http://www.w3.org/2001/
  XMLSchema" xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/"
  xmlns="http://schemas.xmlsoap.org/wsdl/"
  <types>
    <xs:schema targetNamespace="http://pack/">
      <xs:element name="echo">
        <xs:complexType>
          <xs:sequence>
            <xs:element name="arg0" type="xs:string" minOccurs="0"/>
          </xs:sequence>
        </xs:complexType>
      </xs:element>
      <xs:element name="echoResponse">
        <xs:complexType>
          <xs:sequence>
            <xs:element name="return" type="xs:string" minOccurs="0"/>
          </xs:sequence>
        </xs:complexType>
      </xs:element>
      <xs:element name="timestamp" type="xs:dateTime"/>
    </xs:schema>
  </types>
  <message name="echo">
    <part name="parameters" element="tns:echo"/>
    <part name="timestamp" element="tns:timestamp"/>
  </message>
  <message name="echoResponse">
    <part name="parameters" element="tns:echoResponse"/>
  </message>
  <portType name="Echo">
    <operation name="echo">
      <input message="tns:echo"/>
      <output message="tns:echoResponse"/>
    </operation>
  </portType>
  <binding name="EchoPortBinding" type="tns:Echo">
    <soap:binding transport="http://schemas.xmlsoap.org/ soap/http"
      style="document"/>
    <operation name="echo">
      <soap:operation soapAction=""/>
      <input>
        <soap:body parts="parameters" use="literal"/>
      </input>
    </operation>
  </binding>
</definitions>
```

```

        <soap:header message="tns:echo" part="timestamp"
            use="literal"/>
    </input>
    <output>
        <soap:body use="literal"/>
    </output>
</operation>
</binding>
...
</definitions>

```

---

**Example 7.13**

The Echo service WSDL definition with an explicit header contains a timestamp.

Example 7.13 contains an extract of the respective WSDL definition for an Echo service that shows:

- an additional element in the schema, called "timestamp" of type `xs:dateTime`
- an additional part in the input message definition, which refers to the `timestamp` element
- an additional definition for the header in the SOAP binding, which indicates that the `timestamp` element should be carried in the SOAP header of the request message

The header binding shown in Example 7.14 refers to a part also included in the portType of the service, the input message, and the Java service interface generated by the JAX-WS `wsimport` tool. Note that the import statements are left out.

```

@WebService(name = "Echo", targetNamespace = "http://pack/")
@SOAPBinding(parameterStyle = ParameterStyle.BARE)
public interface Echo {
    @WebMethod
    @WebResult(name = "echoResponse", targetNamespace = "http://pack/",
        partName = "parameters")
    public EchoResponse echo(
        @WebParam(name = "echo", targetNamespace = "http://pack/",
            partName = "parameters")
        Echo_Type parameters,
        @WebParam(name = "timestamp", targetNamespace = "http://pack/",
            header = true, partName = "timestamp")
        XMLGregorianCalendar timestamp);
}

```

---

**Example 7.14**

The service interface includes a parameter for the explicit header and indicates two parameters: one that contains the string wrapped into the `Echo_Type` class and another that carries the `timestamp` element. Note that nothing in the interface indicates that the `timestamp` will go into the SOAP header, as this information is only contained in the WSDL definition.

### *Implicit Headers*

Assume that the header data is not part of the `portType` but instead uses a message part unused in any input or output message, known as an implicit header. The header information is not included in the `portType` of the service or in the Java interface. Example 7.15 shows that the WSDL for the Echo service has been changed to include an explicit header.

```
<definitions targetNamespace="http://pack/" name="EchoService"
  xmlns:tns="http://pack/" xmlns:xs="http://www.w3.org/2001/
  XMLSchema" xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/"
  xmlns="http://schemas.xmlsoap.org/wsdl/"
  ... (this part is as before) ...
  <message name="echo">
    <part name="parameters" element="tns:echo"/>
  </message>
  <message name="echoResponse">
    <part name="parameters" element="tns:echoResponse"/>
  </message>
  <message name="header">
    <part name="timestamp" element="tns:timestamp"/>
  </message>
  <portType name="Echo">
    <operation name="echo">
      <input message="tns:echo"/>
      <output message="tns:echoResponse"/>
    </operation>
  </portType>
  <binding name="EchoPortBinding" type="tns:Echo">
    <soap:binding transport="http://schemas.xmlsoap.org/soap/http"
      style="document"/>
    <operation name="echo">
      <soap:operation soapAction=""/>
      <input>
        <soap:body parts="parameters" use="literal"/>
        <soap:header message="tns:header" part="timestamp"
          use="literal"/>
      </input>
```

```

    <output>
      <soap:body use="literal"/>
    </output>
  </operation>
</binding>
...
</definitions>

```

---

**Example 7.15**

A WSDL definition contains an implicit header for an Echo service.

The WSDL definition presented in Example 7.15 is not that much different from Example 7.13, with a separate message defined for the header. The separate message has a significant impact on the Java interface seen in Example 7.16, where the import statements have been omitted again.

```

@WebService(name = "Echo", targetNamespace = "http://pack/")
public interface Echo {
    @WebMethod
    @WebResult(targetNamespace = "")
    @RequestWrapper(localName = "echo",
        targetNamespace = "http://pack/", className = "pack.Echo_Type")
    @ResponseWrapper(localName = "echoResponse",
        targetNamespace = "http://pack/", className = "pack.EchoResponse")
    public String echo(
        @WebParam(name = "arg0", targetNamespace = "")
        String arg0);
}

```

---

**Example 7.16**

The service interface does not include the implicit header.

The interface in Example 7.16 takes a simple `String` parameter, and does not refer to the `timestamp` element or use the `Echo_Type` class to wrap the input message. The implicit header definition requires extra work on the service implementer by the service client developer to ensure the appropriate header information is inserted into the message. The implicit header definition cannot simply be passed to the service proxy as a parameter. In both cases, JAX-WS handlers can be leveraged to process the SOAP header, or intermediaries inserted between service consumer and service can manage all header information, such as part of an ESB.

The header portion of a service message should only contain contextual information, omitting any business connotation. The implementations of the service consumer and



the service should only deal with business logic and not with infrastructure-level information. The use of implicit headers is common, although extra code to generate and process the headers must be written.

### *No Headers*

A final option is to put no header information in the WSDL definition, which can appear to leave the contract incomplete but is actually preferable. Headers typically contain information independent from the business payload being exchanged between services. Recall a timestamp that had been inserted into the SOAP message presented in Examples 7.13 and 7.15. Inserting a timestamp might be a defined company policy across all services, and a common method for doing so can be established. Adding this detail to each WSDL definition is not required and creates an unnecessary dependency between the business-relevant service contract and technical cross-service policy.

## **Data Mapping with REST**

XML schemas can be used to represent service data elements, with JAXB and JAX-WS generating the mapped Java classes and Web service artifacts for SOAP-style Web service implementations. For REST services, the JAX-RS service implementations are similar. When the convenience of code generation is needed, JAXB annotated POJOs can be used as the service entities in JAX-RS resource classes. Behind the scenes, the JAX-RS runtime will marshal the Java objects to the appropriate MIME-type representations for the entities, such as `application/xml`. A `customer` object annotated with JAXB annotations is shown in Example 7.17.

```
@XmlElement(name="customer")
@XmlAccessorType(XmlAccessType.FIELD)
public class Customer {
    private int id;
    private String name;

    public Customer() {}
    //...other attributes omitted
    //...getters and setters for id and name
    //...
}
```

---

### **Example 7.17**

The resource methods in the JAX-RS implementation that produce or consume customer information in the form of XML can be seen in Example 7.18.

```
//...retrieve customer and return xml representation
@Get
@Path("id")
@Produces("application/xml")
public Customer getCustomer(
    @PathParam("id") Long id){
    Customer cust = findCustomer(id);
    return cust;
}
//...create customer with an xml input
@POST
@Consumes("application/xml")
public void createCustomer(
    Customer cust){
    //...create customer in the system
    Customer customer = createCustomer(cust);
    //...
}
```

---

#### Example 7.18

The JAX-RS implementation automatically handles the marshaling and unmarshaling of the XML-based resource representation to and from JAXB objects. Generic types, such as a `javax.xml.transform.Source`, can be handled in the resource class to keep the resource class independent of any specific types defined in the domain, such as a `customer` type. However, extra work is required to handle the extraction of the customer information from the `Source` object seen in Example 7.19.

```
@PUT
@Path("id")
@Consumes("application/xml")
public void updateCustomer(@PathParam("id") Long id,
    javax.xml.transform.Source cust) {
    // do all the hard work to
    // extract customer info in
    // extractCustomer()
    updateCustomer(id, extractCustomer(cust));
}
```

---

#### Example 7.19

JAX-RS supports alternate MIME types, such as JSON. Just as JAXB handles the binding of XML to and from Java objects, numerous frameworks exist that handle mapping JSON representations to Java objects and vice versa. Some commonly used frameworks for mapping between JSON and Java are MOXy, Jackson, and Jettison.

In the Jersey implementation of JAX-RS 2.0, the default mechanism for binding JSON data to Java objects leverages the MOXy framework. When the Jersey runtime is configured to use MOXy, the runtime will automatically perform binding between Java objects (POJOs or JAXB-based) and a JSON representation. Jackson or Jettison can also perform similar binding with appropriate runtime configuration. A low-level JSON parsing approach can be achieved with the newly introduced JSON-P API (Java API for JSON Processing) in Java EE 7. JSON-P should not be confused with JSONP (JSON with Padding), which is a JavaScript communication technique used to avoid certain types of browser restrictions.

#### *Conversion Between JSON and POJOs*

Given the same customer representation as illustrated in Example 7.20, no special code is required to handle an incoming JSON document or return a JSON-based representation.

```
//...
@GET
@Path("id")
@Produces("application/json")
public Customer getCustomer(
    @PathParam("id") Long id){
    return findCustomer(id);
}
```

---

#### **Example 7.20**

The returned customer representation would be a JSON string, such as `{"name": "John Doe", "id": "1234" ... }`.

The same JAXB objects can be used for handling JSON media types that would normally be used for XML representation. The addition of another MIME type in the `@Produces` can be seen in Example 7.21.

```
@GET
@Path("id")
@Produces("application/json", "application/xml")
```

```
public Customer getCustomer(  
//...
```

---

### Example 7.21

The JAX-RS runtime returns an appropriate representation (XML or JSON) that is determined by the client's preference. In spite of the convenience offered by JSON binding frameworks like MOXy or Jackson, greater control over the processing of the JSON input and output can be a requirement, as opposed to letting the JAX-RS runtime perform an automatic binding between JSON and Java objects.

For example, REST service operations must consume or produce only selective parts of large JSON documents, as converting the whole JSON document to a complete Java object graph can cause significant resource overheads. In such cases, a JSON parsing mechanism based on a streaming approach can be more suitable. JSON-P APIs allow a developer complete control over how JSON documents are processed. JSON-P supports two programming models, the JSON-P object model and the JSON-P streaming model.

The JSON-P object model creates a tree of Java objects representing a JSON document. The `JsonObject` class offers methods to add objects, arrays, and other primitive attributes to build a JSON document, while a `JsonValue` class allows attributes to be extracted from the Java object representing the JSON document. Despite the advantage of convenience, processing large documents with the object model can create substantial memory overheads, as maintaining a large tree of Java objects imposes significant demands on the Java heap memory. This API is similar to the Java DOM API for XML parsing (`javax.xml.parsers.DocumentBuilder`).

In comparison, the JSON-P streaming model uses an event parser that reads or writes JSON data one element at a time. The `JsonParser` can read a JSON document as a stream containing a sequence of events, offer hooks for intercepting events, and perform appropriate actions. The streaming model helps avoid reading the entire document into memory and offers substantial performance benefits. The API is similar to the StAX Iterator APIs for processing XML documents in a streaming fashion (`javax.xml.stream.XMLStreamReader`). The `JsonGenerator` class is used to write JSON documents in a streaming fashion similar to `javax.xml.stream.XMLStreamWriter` in StAX API.

JSON-P does not offer binding between JSON and Java. Frameworks, such as MOXy or Jackson, are similar to JAXB in how they will need to be leveraged to perform conversion between Java objects and JSON.

**CASE STUDY EXAMPLE**

SmartCredit launches an aggressive expansion campaign with the intention of offering premium credit cards with cashback offers to high-value customer prospects across a retail chain's locations. After signing an agreement with the retail chain, SmartCredit obtains prospect data from all the retail stores containing prospect names, e-mail addresses, and net transaction values at the end of every month. An internal Prospect Analyzer application will process the data to target prospects with high monthly transaction values and send out e-mails with new premium credit card offers.

The retail chain's IT department sends customer data to SmartCredit in large JSON documents containing prospect information. However, the SmartCredit Prospect Analyzer service is only interested in prospects that spend in excess of 2,000 dollars during the month. A fragment of a typical monthly JSON extract from the retail stores is provided in Example 7.22.

```
"txnsummary":{
  "date":"2014-01-31T23:30:00-0800",
  "store":"Fashion Trends #132",
  "txn": [
    {
      "type":"cash",
      "amount":235.50,
      "e-mail":null
    },
    {
      "type":"credit",
      "amount":3565.00,
      "e-mail":"jane@doe.com"
    }
  ]
}
```

**Example 7.22**

SmartCredit IT accepts the prospect and transaction data through an HTTP POST from the retail stores at the end of every month. A REST API that consumes this JSON data and extracts the prospects for marketing campaigns is built. After reviewing the size of the monthly feed, a SmartCredit architect quickly realizes that memory limitations will prevent a typical JSON-Java binding approach from working for

such large payloads. In addition, SmartCredit is only interested in processing selective parts of the payload, such as credit card transactions with amounts greater than 2,000 dollars.

Converting the entire JSON data into Java objects is a waste of time, memory, and resources. The JSON-P streaming API is a suitable option for allowing selective processing of only the data sections meeting the given criteria. A simplified version of the final resource class is illustrated in Example 7.23.

```
import javax.ws.rs.Consumes;
import javax.ws.rs.Path;
import javax.ws.rs.POST;
import javax.ws.rs.core.MediaType;
import java.io.Reader;
import javax.json.Json;
import javax.json.streaming.JsonParser;
import javax.json.streaming.JsonParser.Event
@Path("upload")
@Consumes(MediaType.APPLICATION_JSON)
public class ProspectFilteringResource {
    private final double THRESHOLD = 2000.00;
    @POST
    public void filter(final Reader transactions) {
        JsonParser parser = Json.createParser(transactions);
        Event event = null;
        while(parser.hasNext()) {
            event = parser.next();
            if(event == Event.KEY_NAME && "type".equals(parser.getString()))
            {
                event = parser.next(); //advance to Event.VALUE_STRING for
                the actual value of "type"
                if("credit".equals(parser.getString())) {
                    parser.next(); //Event.KEY_NAME for "amount"
                    event = parser.next(); //Event.VALUE_NUMBER for amount
                    value
                    if(parser.getBigDecimal().doubleValue() > THRESHOLD) {
                        parser.next(); //advance to Event.KEY_NAME for "e-mail"
                        parser.next(); //advance to Event.VALUE_STRING for
                        e-mail info
                        String e-mail = parser.getString();
                        addToCampaignList(e-mail);
                    }
                }
            }
        }
    }
}
```

```
    }  
  }  
  private void addToCampaignList(String e-mail) {  
    // actual logic of adding e-mail to campaign list  
  }  
}
```

**Example 7.23**

The JSON-P streaming API can parse a large JSON document selectively.

The code uses the streaming API to advance the parser to consume only specific events and avoids reading the entire JSON data structure into memory, which would have been the case using the standard JSON-Java binding approach. One drawback to the JSON-P approach is the cumbersome maneuvering of the event stream in the application code, although such trade-offs are often necessary in real-world usage.

### *Binary Data in Web Services*

Candidates looking to utilize a service-oriented solution often require the transfer of large amounts of data kept in some binary format, such as a JPEG image file. The Java language offers support for handling binary data in its JavaBeans Activation Framework (JAF) as well as classes and interfaces in other packages, which depend on the format. For example, the `java.awt.Image` class hierarchy supports image formats, such as JPEG. The JAXB specification defines how to map certain Java types to XML schema, such as the mapping of a `java.awt.Image` type to `base64Binary`.

Binary formats without a special Java type associated can use the `javax.activation.DataHandler` class defined in JAF. However, `byte[]` is the most generic way of representing binary data in Java. JAXB defines that a `byte[]` is mapped to `base64Binary` and `hexBinary`.

When generating a WSDL contract from a Java interface, binary data types are mapped using JAXB mapping rules. Generating the reverse is not as straightforward. An XML schema element of type `base64Binary` can map into multiple different Java types. By default, `byte[]` is used. Indicating that an element declared as `base64Binary` in the schema should be mapped to `java.awt.Image` in the Java service implementation can be performed in a number of ways. Binary data can be transferred in a SOAP message or inline, which means binary data is encoded into text and sent like any other data in the message.

Transferring the data inline maintains interoperability between different vendor environments, but is ineffective when dealing with large pieces of binary data, such as a CAD drawing of an airplane. The text encoding increases the size of the message by an average of 25%. The MTOM is an alternative to text encoding, which describes how parts of a message can be transferred in separate parts of a MIME-encoded multipart message. The binary content is removed from the SOAP payload, given a MIME type, and transferred separately.

JAXB provides support for MTOM, which must be explicitly enabled for the JAX-WS runtime. JAX-WS plugs into this support when building and parsing messages with attachments. An element definition can be annotated in an XML schema document with two specific attributes indicating which MIME type to give the element. When using MTOM, the `contentType` and `expectedContentTypes` attributes demonstrate how to MIME-encode the element and determine which Java type the element is mapped to in the Java service interface.

Nothing in the schema or in any of the Java code indicates whether the binary data is transferred as an attachment using MTOM or inserted directly into the SOAP envelope. In JAX-WS, distinguishing the difference is defined either by a setting on the service configuration file or programmatically. The following case study example illustrates a SOAP-based Web service managing attachments via MTOM.

### CASE STUDY EXAMPLE

NovoBank offers a remote method of opening an account from its Web site for customers to download a form, fill it out, and mail or fax it to a branch office. Alternatively, customers can fill out the form in branch and provide the completed form to a bank employee. The forms are scanned in at the branch office for further processing by NovoBank's back-end system before being archived.

To reduce processing time, the bank wants to offer customers and employees a Web application that accepts an uploaded binary image of the form to send to the bank. Internally, the Open Account service uses a Web service that takes the binary image as a parameter. To simplify the implementation of the associated service logic in Java, the service contract uses the `expectedContentType` attribute in its schema to indicate that the scanned image is formatted as a JPEG document. The resulting WSDL definition is shown in Example 7.24 with parts of the WSDL document omitted for brevity.



```

<definitions targetNamespace= "http://personalbanking.services.
  novobank.com/" name="AccountService"
  xmlns:tns="http://personal banking.services.novobank.com/"
  xmlns:xsd="http://www.w3.org/ 2001/XMLSchema"
  xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/"
  xmlns=http://schemas.xmlsoap.org/wsdl/
  xmlns:xmime="http://www.w3.org/2005/05/xmlmime">
  <types>
    <xsd:schema>
      <xs:element name="openAccount" type="tns:openAccount"
        xmlns:tns="http://personalbanking.services.novobank.com/"/>
      <xsd:complexType name="openAccount">
        <xsd:sequence>
          <xsd:element name="arg0" type="xsd:base64Binary"
            xmime:expectedContetTypes="image/jpeg" minOccurs="0"/>
        </xsd:sequence>
      </xsd:complexType>
      ...
    </xsd:schema>
  </types>
  <message name="openAccount">
    <part name="parameters" element="tns:openAccount"/>
  </message>
  <message name="openAccountResponse">
    <part name="parameters" element="tns:openAccountResponse"/>
  </message>
  <portType name="Account">
    <operation name="openAccount">
      <input message="tns:openAccount"/>
      <output message="tns:openAccountResponse"/>
    </operation>
  </portType>
</definitions>

```

**Example 7.24**

The WSDL for NovoBank's Open Account service with an MTOM content type definition will now accept the JPEG format.

An element of type `base64Binary` will be mapped to a `byte[]` in the Java interface. However, the additional annotation of the `parameter` element, using the `expectedContentTypes` attribute, leads to the following Java interface presented in Example 7.25.

```

package com.novobank.services.personalbanking;
import java.awt.Image;
// other imports omitted
@WebService(name = "Account", targetNamespace = "http://
  personalbanking.services.novobank.com/")
public interface Account {

    @WebMethod
    @WebResult(targetNamespace = "")
    @RequestWrapper(localName = "openAccount",
        targetNamespace = "http://personalbanking.services.novobank.
        com/", className = "com.novobank.services.personalbanking.
        OpenAccount")
    @ResponseWrapper(localName = "openAccountResponse",
        targetNamespace = "http://personalbanking.services.novobank.
        com/", className = "com.novobank.services.personal banking.
        OpenAccountResponse")
    public String openAccount(
        @WebParam(name = "arg0", targetNamespace = "")
        Image arg0);
}

```

**Example 7.25**

The content type is mapped to the appropriate Java type in the service interface.

Note how the `parameter` passed to the service implementation is mapped to the `java.awt.Image` type. Defining whether MTOM is used to transfer the form image as an attachment can be performed programmatically using JAX-WS, or in the client or the endpoint configuration for the service endpoint. A sample client for the new Open Account service is shown in Example 7.26.

```

package com.novobank.services.personalbanking.client;
import java.awt.Image;
import java.awt.Toolkit;
import javax.xml.ws.BindingProvider;
import javax.xml.ws.soap.SOAPBinding;

public class AccountServiceClient {
    public static void main(String[] args) {
        Image image = Toolkit.getDefaultToolkit().getImage("c:\\temp\\
        java.jpg");
        Account serviceProxy = new AccountService().getAccountPort();
        SOAPBinding binding = (SOAPBinding)((BindingProvider)
        serviceProxy).getBinding();
    }
}

```

```

binding.setMTOMEnabled(true);
String accountNumber = serviceProxy.openAccount(image);
System.out.println("returned account number is
    "+accountNumber);
}
}

```

**Example 7.26**

A Java client enables MTOM transport.

After deploying the service and running the test client listed in Example 7.26, the SOAP message sent can be reviewed by executing the client in Example 7.27.

```

POST /attachment/account HTTP/1.1
Content-Length: 13897
SOAPAction: ""
Content-Type: multipart/related; type="application/xop+xml";
boundary="-----_Part_0_14949315.1177991007796"; start-info="text/
xml"
Accept: text/xml, application/xop+xml, text/html, image/gif, image/
jpeg, *, q=.2, */*; q=.2
User-Agent: Java/1.5.0_11
Host: 127.0.0.1:8081
Connection: keep-alive

-----_Part_0_14949315.1177991007796
Content-Type: application/xop+xml; type="text/xml"; charset=utf-8

<?xml version="1.0" ?><soapenv:Envelope xmlns:soapenv=http://
schemas.xmlsoap.org/soap/envelope/ xmlns:xsd=http://www.w3.org/2001/
XMLSchema xmlns:ns1="http://personalbanking.services.novobank.
com/">/"><soapenv:Body><ns1:openAccount><arg0><xop:Include
xmlns:xop="http://www.w3.org/2004/08/xop/include" href="cid:f17b4f2b-
db2c-4bc5-96d1-2d4857aaa5b8@example.jaxws.sun.com"></xop:Include></
arg0></ns1:openAccount></soapenv:Body></soapenv:Envelope>

-----_Part_0_14949315.1177991007796
Content-Type: application/octet-stream
Content-ID: <f17b4f2b-db2c-4bc5-96d1-2d4857aaa5b8@example.jaxws.sun.
com>
Content-transfer-encoding: binary KÝÝ ÔF-;úPr*îÉu1_) Šlání}K²£•Ñ\
êKQZ_Ý4İ;É)*G'~ªohBíýfÜÈì|zùif_Ö(Ö &_Èõmáoâ``\Ó( ò¹âqç°È_W¶èÀh";_
ĐÚ'z0"ç5WÄ" üv|DÜî7IÚù_é6³ÈÚ1•leÉä|>BîöWÈ"ý|>|i"i™Åã]ÓÉÝ9f_"7Üý¶j9{
ßÁÉ?w("_86<üfÄ_?_:ŽšòÔŠvM|éÈè4_ÝÈ"=fÑ]™EýÖ×Es~Hýã÷©?7<â""r{<â³

```

```
Rè<...R×©Ô]z_íµÙ4_ÿBîùš_?ÿPé3ê
... the rest of the data omitted ...
-----=_Part_0_14949315.1177991007796--
```

**Example 7.27**

A SOAP message with an MTOM-compliant binary part

The request message is divided into the SOAP message, including reference to the binary data, and the binary data sent as an attachment. The message can be handled more efficiently and does not impact the processing of the remaining XML information. Again, whether the data is sent as an attachment or not is not defined in the service contract (the WSDL definition).

*Binary Data in REST Services*

JAX-RS supports the conversion of resource representations into Java types, such as `byte[]` or `java.io.InputStream`, in resource methods. To produce or consume binary data from resource methods, JAX-RS runtime handles the conversion through a number of built-in content handlers that map MIME-type representations to `byte[]`, `java.io.InputStream`, and `java.io.File`. If binary content must be handled as a raw stream of bytes in a resource method, the corresponding resource method is seen in Example 7.28.

```
@Get
@Produces("image/jpg")
public byte[] getPhoto() {
    java.io.Image img = getImage();
    return new java.io.File(img);
}

@Post
@Consumes("application/octet-stream")
public void processFile(byte[] bytes) {
    //...process raw bytes
}
```

**Example 7.28**

**NOTE**

Indicate the MIME type through an appropriate `@Produces` annotation, such as a JPG image, so that the runtime can set the right content-type for the returned resource representation.

The `java.io.File` type can help process large binary content in a resource method, such as an attachment containing medical images, as seen in Example 7.29.

```
@POST
@Consumes("image/*")
public void processImage(File file) {
    //...process image file
}
```

**Example 7.29**

The JAX-RS runtime streams the contents to a temporary file on the disk to avoid storing the entire content in memory. For complete control over content handling, JAX-RS offers several low-level utilities which can be useful for custom marshaling/unmarshaling of various content types. The `javax.ws.rs.ext.MessageBodyReader` and `javax.ws.rs.ext.MessageBodyWriter` interfaces can be implemented by developers to convert streams to Java types and vice versa.

Back-and-forth conversion is useful when mapping custom MIME types to the domain-specific Java types. The classes that handle the mapping are annotated with the `@Provider` annotation and generically referred to as JAX-RS entity providers. Entity providers are used by the JAX-RS runtime to perform custom mapping.

A special case of `javax.ws.rs.ext.MessageBodyWriter` is the `javax.ws.rs.core.StreamingOutput` callback interface, which is a wrapper around a `java.io.OutputStream`. JAX-RS does not allow direct writes to an `OutputStream`. The callback interface exposes a `write` method allowing developers to customize the streaming of the response entity. Example 7.30 demonstrates the `gzip` format used to compress a response entity.

```
import javax.ws.rs.core.StreamingOutput;
...
@GET
@Produces("application/gzip-compressed")
public StreamingOutput getCompressedEntity() {
```

```
return new StreamingOutput() {
    public void write(OutputStream out)
        throws IOException, WebApplicationException {
        try {
            ...
            GZipOutputStream gz =
                new GZipOutputStream(out);
            //...get array of bytes
            // to write to zipped stream
            byte[] buf = getBytes();
            gz.write(buf, 0, buf.length);
            gz.finish();
            gz.close();
            ...
        } catch(Exception e) { ... }
    }
};
}
```

---

**Example 7.30**

For mixed content containing both text and binary payload, entity providers can be used to perform custom marshaling, while multipart MIME representations are suitable for dealing with mixed payloads. JAX-RS standards do not mandate support for handling mixed MIME multipart content types apart from multipart FORM data (multipart/form-data), which is useful for HTML FORM posts but has limited use in a system-to-system interaction context. Various JAX-RS implementations, such as Jersey and RESTEasy, provide support for mixed multipart data. Handling mixed content with binary data is common, as seen in the following case study example for SmartCredit's Submit Credit Application service.

**CASE STUDY EXAMPLE**

SmartCredit is building a REST service known as the Submit Credit Application service that is intended for service consumers to submit credit card applications. Apart from basic information such as customer details, the supporting information in the form of various collaterals, such as mortgage papers and loan approvals, must be scanned as images and attached to the application.

The SmartCredit application development team considered using Base64 encoding, but moved onto other alternatives after realizing a substantial size bloat would

result. The development team decides on the mixed multipart representation for the application data. The multipart application data will have the customer information in an XML format as the first part, and a series of images in the subsequent parts.

A sample multipart application request over HTTP is shown in Example 7.31.

```
POST /creditapps/ HTTP/1.1
Host: smartcredit.com
Content-Type: multipart/mixed; boundary=xyzw

--xyzw
Content-Id: <abcdefgh-1>
Content-Type: application/xml
<customer>
  <name>John Doe</name>
  <Address>...</Address>
  ...
--xyzw
Content-Id: <abcdefgh-2>
Content-Type: image/jpg
...
&_Èômâoâ``\Ó\ ( ò³âq °Ê _W¶èÁh";_ĐÚ´z0"ç5WÃ"_üv|DÜî7IÚù_
é6³ÈU1•lèÉäl>BîöWÈ"ý|>i"i™Ää]ÓÉÝ9f_"7Üý¶j9{ßáÉ?w(_"86<úfÃ_´?_:ŽšòÔšv
M|éÈè4_ŸÊ"=fÑ]™EýÖ×Ès~Hý...rest of the binary data goes here
--xyzw-
```

### Example 7.31

The SmartCredit service development team considered using a custom JAX-RS Entity Provider to handle the mixed multipart data, but realized the reference JAX-RS implementation Jersey already provides support for mixed multipart data through an add-on called `jersey-media-multipart`. The key classes leveraged in this implementation include:

- The `org.glassfish.jersey.media.multipart.BodyPartEntity` represents the entity of a part when a MIME Multipart entity is received and parsed.
- The `org.glassfish.jersey.media.multipart.BodyPart` is a mutable model representing a body part nested inside a MIME Multipart entity.

The resource class method that handles the submitted application can be seen in Example 7.32.

```
import org.glassfish.jersey.media.multipart.MultiPart;
import org.glassfish.jersey.media.multipart.BodyPart;
import org.glassfish.jersey.media.multipart.BodyPartEntity;
import javax.ws.rs.core.Response;
...

import com.smartcredit.domain.Customer;
...

@Path("/creditapps")
public class CreditAppResource {

    @POST
    @Consumes("multipart/mixed")
    public Response post(MultiPart multiPart) {
        // First part contains a Customer object
        Customer customer =
            multiPart.getBodyParts().get(0).
                getEntityAs(Customer.class);

        // process customer information
        processCustomer(customer);

        // get the second part which is a scanned image
        BodyPartEntity bpe =
            (BodyPartEntity) multiPart.getBodyParts().
                get(1).getEntity();
        try {
            InputStream source = bpe.getInputStream();
            //process scanned image
        }

        // Similarly, process other images in the multipart
        // content, if any...

        //If everything was fine, return Response 200
        return Response.status(Response.Status.OK).build();

        //else if there were errors...
        return Response.status(Response.Status.BAD_REQUEST).
            build();
    }
}
```

---

**Example 7.32**



In the code fragment, the `@Consumes` annotation indicates that a resource representation of multipart/mixed is expected. In this case, the payload contains customer data in XML and one or more scanned images. The following Jersey utilities perform different steps in managing the mixed multipart data:

- `com.sun.jersey.multipart.MultiPart.getBodyParts()` returns a list of `com.sun.jersey.multipart.BodyParts`.
- `BodyPart.getEntityAs(Class<T> cls)` returns the entity converted to the passed-in class type.
- `com.smartcredit.domain.Customer` is a regular JAXB-annotated Java class. Since the first entity in the multipart message is known to be the customer entity, the `BodyPart.getEntityAs(Customer.class)` method is used to unmarshal the XML entity body into a JAXB `customer` object.
- `BodyPart.getEntity()` returns the entity object to be unmarshaled from a request. The entity object is known to be a `BodyPartEntity` and is cast accordingly.
- `BodyPartEntity.getInputStream()` returns the raw contents, which in this case are the contents of the scanned image.

Note that by using the MIME multipart utilities in Jersey, the development team is able to avoid writing much of the plumbing code that would otherwise be necessary to deal with multipart MIME messages.

## Use of Industry Standards

The use of industry standards in developing service contracts builds on the IT-specific standards and seldom offers challenges when using Java as the language and runtime environment. Many industries have established data formats that ensure interoperability between business partners, suppliers, and between systems within an enterprise, such as the ACORD standard for insurance, HL7 for the healthcare industry, or SWIFT for banking.

Used primarily to exchange information between companies or independent units within an enterprise, industry standards are prime candidates for use as part of the service contract. Industry standards are generally expressed as XML schema definitions,

which can be directly referenced in a service contract or serve as the basis for a specific schema.

Before the advent of JAXB 2.0 which supports the full set of XML schema constructs, a common issue was the inability to map all of the elements used in an industry schema into Java because such mapping was not defined. JAXB 2.x nearly resolves this issue, but cases still occur where a large and complex industry standard schema cannot be handled by a data binding tool like JAXB. Using an industry standard unchanged in a service contract can be tedious and lead to the generation of hundreds of Java classes mapping all of the complex types defined in the schema.

### SUMMARY OF KEY POINTS

---

- SOAP-based Web services can be developed top-down, bottom-up, or meet-in-the-middle.
  - For REST services, a resource implementation artifact is used to model a Web resource that responds to HTTP operations. Apart from XML, support for resource representations can encompass a wide range of other media types. REST services can also use MIME features, such as multipart messages, to deal with binary content. JAX-RS implementations provide content handlers which can be customized for dealing with different representations.
  - Standards-based service contracts can map the relevant XML schema constructs to and from Java. For both SOAP and REST services, the JAXB 2.0 standard offers support for the entire set of XML schema features. Industry standards are widely used in service contracts and do not introduce particular challenges when using Java.
  - For SOAP-based Web services, special considerations apply when leveraging specific WSDL features, such as header fields, attachments with binary data, or the use of wrapper elements.
- 

## 7.3 Service Loose Coupling

In a service-oriented environment, components can be coupled at a number of different levels. Coupling can occur at the level of the service contract, the service implementation logic, or the underlying technology platform the service is running on if a service

consumer is coupled with a particular service. In general, SOA promotes the notion of decoupled systems by which the parts of a service-oriented solution are as decoupled as possible. Reducing the dependency of one part on another allows one area, such as a service or an aggregation, to be changed without requiring changes to other areas.

The following design characteristics are helpful in the creation of decoupled systems:

- separation of contract and implementation
- functional service context with no dependency on outside logic
- minimal requirements for the service consumer

When service logic is implemented in Java and executed in a JRE, a coupling is created between the logic and the underlying technology platform, Java. Additional dependencies arise when Java EE is used as the hosting platform. For an in-depth explanation of the variations of positive and negative coupling, see Chapter 7 in *SOA Principles of Service Design*.

### Separation of Contract and Implementation

The separation of a service contract from its implementation is a key characteristic of service-oriented design that was first established as part of the definition of remote procedure calls and then in object-oriented programming. The interface of a component is exposed to a point that allows other logic to invoke this component, but no details about the implementation are added. Changes to the implementation can then be made without affecting the client. SOA took this concept further by establishing the notion of a standards-based contract. For example, the service interface can be expressed in a programming language-neutral way to allow for the integration of logic written in different languages and running on different platforms.

The choice of programming language can be exposed in the service contract. In the context of a SOAP Web service from the *Top-Down vs. Bottom-Up* section, a common approach is to generate the service contract from existing Java logic. In JAX-WS, the `wsgen` tool can be used to generate a complete WSDL document, including the XML schema definition from an existing JavaBean or Java interface. However, the existing Java code may not be easily mapped into XML schema.

For example, assume that the `public java.util.Hashtable<String, String> getTable(int i);` method is part of an interface to be turned into a service. Using the `wsgen` tool to generate both the WSDL and XML schema definition from this interface creates the type definitions presented in Example 7.33.

```

<xs:element name="getTableResponse" type="ns2:getTableResponse"
  xmlns:ns2="http://the.package/" />
<xs:complexType name="getTableResponse">
  <xs:sequence>
    <xs:element name="return" type="ns3:hashtable" minOccurs="0"
      xmlns:ns3="http://the.package/" />
  </xs:sequence>
</xs:complexType>
<xs:complexType name="hashtable">
  <xs:complexContent>
    <xs:extension base="ns4:dictionary"
      xmlns:ns4="http://the.package/" />
    <xs:sequence />
  </xs:extension>
</xs:complexContent>
</xs:complexType>
<xs:complexType name="dictionary" abstract="true" />

```

---

### Example 7.33

The result is a valid XML schema definition, although a Java `Hashtable` is exposed in the service interface which renders the schema not useful or usable outside of Java. The same is true for many other classes and interfaces in Java, which, while useful as part of the implementation, do not map well into XML schema and non-Java environments. Many of the more technically-oriented classes, including the members of the `java.io` package, fall into this category.

The use of internal, platform-specific information carried in a generic type is another aspect of decoupling the contract and implementation. For example, the name of a file usually stored as a string can be mapped into XML schema and added to a service contract. However, exposing file names on a service contract is considered poor practice, and revealing details of the service implementation to the service consumer should be avoided. The same holds true for names of databases, database tables, machine names, and addresses.

Using the reverse approach is another way of creating this coupling, particularly when generating a service logic skeleton directly from a service contract. In JAX-WS, the `wsimport` tool is used to create Java code skeletons representing a given WSDL definition. Despite being the recommended approach, be aware that the generated Java code is now tightly coupled to its contract, which prevents the code from being easily reused to serve other types of service contracts or updated versions of the current contract. In most cases, the tight coupling is acceptable because the logic is created for a particular contract and nothing else.

However, instances occur where service logic must be created for reuse despite changes to the underlying contract changes or the need to concurrently serve multiple versions of the same contract. The JAX-WS `Provider` API is equipped for such instances. In this model, the service logic parses the incoming message at runtime for dynamic processing with no direct dependency on the types and elements defined in the WSDL contract. Use of the JAX-WS `Provider` API is detailed in Chapter 8.

For a REST service, coupling the code to the service contract is inconsequential because generating implementation artifacts from a machine-readable contract, unless WADL is being used, is uncommon. Implementing a Web resource for a platform, such as JAX-RS, out of the box supports only a finite subset of Java classes that can be automatically mapped to appropriate content-types. For custom types not mapped automatically by the built-in content handlers, developers must provide implementations of `javax.ws.rs.core.MessageBodyReader/MessageBodyWriter` interfaces to map such types to a known set of resource representations or media types.

### Independent Functional Contexts

Besides the direct compile-time of coupling services, consider the coupling of a service to its outside functional context. Service invocations happen as part of a business transaction or process to establish a context that effectively binds the services, which are invoked downstream, into an aggregated set. Having a set of services as part of the same process establishes a type of coupling between the services, a coupling that should always be top-down. Particularly when leveraging other services to fulfill functionality, a service implementation can be coupled with or have dependency on those services. However, the service should not have a dependency on any services at a higher level of the process hierarchy, or be coupled with the invoking service or with peer services.

For example, assume that a service offers Credit Check functionality. The service is implemented as a business process that invokes a number of other finer-grained services, such as Credit Lookup, Credit Eligibility, Update Profile, and Notification. All four downstream services are peers within the Credit Check business process. Service peers should have no dependency on each other, or be aware of or coupled with their upstream Credit Check service.

Services representing business processes, such as Credit Check and Maintain Customer Information, are decomposed into a set of fine-grained services to create a downstream dependency. A service should not introduce a dependency on another service that is higher level. For example, the implementation of the Update Profile entity service should not introduce any dependency on the Maintain Customer Information task service.

In Java, the same principles of downstream dependency hold true. Unwanted dependencies can be detected by examining the classes that are used by a piece of Java logic. Organizing classes into packages directly identifying the service and/or affiliated business process and ensuring that logically decoupled functions are not packaged together is recommended. A package name should be selected with consideration for possible reuse opportunities. For Web services, the same requirements for namespaces are used in the service contract.

### Service Consumer Coupling

For SOAP-based Web services, a service consumer will often be tightly coupled with the service contract and not the implementation of the service being invoked. However, a looser coupling lessens the impact when the service contract changes.

Using the JAX-WS Dispatch API, service consumer logic can dynamically assemble a request message at runtime and manage any response messages. Additional effort is required to build service consumer logic that can build request messages that the intended service can process.

#### CASE STUDY EXAMPLE

After using the Account service in production for some time, a new operation is added to enhance account services. NovoBank wants to allow all required information about new accounts to be sent to the service as XML on top of the binary image that the initial version supported. Different branch offices use different systems to capture the data required to open a new account, such as traditional “fat client” or browser-based solutions. Additionally, the details of the information stored with new accounts change consistently.

The development team will design the new operation to process different types of input XML formats, and deliver a generic piece of service consumer code that shows how to invoke the new operation from within a JAX-WS supported client. Example 7.34 illustrates the updated WSDL definition for the enhanced Account service to accept the input of XML formats.

```

<definitions targetNamespace="http://personalbanking.services.
  novobank.com/" ...>
  <types>
    <xsd:schema>
      <xsd:element name="openAccount" type="ns1:openAccount"
        xmlns:ns1="http://personalbanking.services.novobank.com/" />
      <xsd:complexType name="openAccount">
        <xsd:sequence>
          <xsd:element name="arg0" type="xs:base64Binary"
            xmime:expectedContentTypes="image/jpeg" minOccurs="0" />
        </xsd:sequence>
      </xsd:complexType>
      <xsd:element name="openAccountXML" type="ns1:openAccountXML"
        xmlns:ns1="http://personalbanking.services.novobank.com/" />
      <xsd:complexType name="openAccountXML">
        <xsd:sequence>
          <xsd:any />
        </xsd:sequence>
      </xsd:complexType>
      <xsd:element name="openAccountResponse"
        type="ns2:openAccountResponse"
        xmlns:ns2="http://personalbanking. services.novobank.com/" />
      <xsd:complexType name="openAccountResponse">
        <xsd:sequence>
          <xsd:element name="return" type="xsd:string" minOccurs="0" />
        </xsd:sequence>
      </xsd:complexType>
    </xsd:schema>
  </types>
  <message name="openAccount">
    <part name="parameters" element="tns:openAccount" />
  </message>
  <message name="openAccountResponse">
    <part name="parameters" element="tns:openAccountResponse" />
  </message>
  <message name="openAccountXML">
    <part name="parameters" element="tns:openAccountXML" />
  </message>
  <portType name="Account">
    <operation name="openAccount">
      <input message="tns:openAccount" />
      <output message="tns:openAccountResponse" />
    </operation>
    <operation name="openAccountXML">
      <input message="tns:openAccountXML" />
    </operation>
  </portType>

```

```

    </operation>
  </portType>
  ...
</definitions>

```

### Example 7.34

The updated WSDL definition for the enhanced Account service accommodates the new operation, `openAccountXML`.

Note how the element named `openAccountXML`, which acts as the wrapper element for the `openAccountXML` operation, contains only one `<xsd:any/>` element. The contract indicates that any kind of XML content can be sent to the service without providing any further details, which allows for decoupling of the service logic from the contract.

Use of the `<xsd:any/>` element minimizes the requirements for service consumers of this service. Any XML document can be passed to the service, allowing the development of generic service consumer logic. As an example of completely decoupling the service consumer from the service, the NovoBank development team delivers the following piece of client code to the users of the Account service in Example 7.35.

```

package com.novobank.services.personalbanking.client;

import java.io.StringReader;
import java.net.URL;
import javax.xml.namespace.QName;
import javax.xml.transform.Source;
import javax.xml.transform.stream.StreamSource;
import javax.xml.ws.Dispatch;
import javax.xml.ws.Service;

public class Account2Client {
    public static String testMessage =
        "<ns1:openAccountXML xmlns:ns1=\"http://personalbanking.
        services.novobank.com/\"><someDocument><personalData>Here goes
        the information</personalData></someDocument> </
        ns1:openAccountXML>";

    public static void main(String[] args) throws Exception {
        QName serviceName =
            new QName("http://personalbanking.services.novobank.com/",
                "AccountService");
    }
}

```



```
QName portName = new QName("http://personalbanking.services.
    novobank.com/", "AccountPort");

Service service = Service.create(
    new URL("http://localhost:8080/account2/account2?wsdl"),
    serviceName);
Dispatch<Source> dispatch = service.createDispatch(
    portName, Source.class, Service.Mode.PAYLOAD);
    dispatch.invoke(new StreamSource(new StringReader(testMessage)));
}
}
```

**Example 7.35**

Java client code decouples the service consumer from the service for the Account service.

First, an instance of the `Service` class providing the location of the appropriate WSDL file and the name of the targeted service is created. A `Dispatch` object is then created from the `Service` class, and the `Service.Mode.PAYLOAD` is defined to indicate that only the content within the SOAP `<body>` element and not the entire message will be passed. Finally, the service can be invoked dynamically via the `Dispatch` object. The service consumer has no compile-time dependency on the service or its contract.

The payload XML document has one root element called `openAccountXML`. Given that the service definition used the wrapped document/literal style described in the *Mapping Between Java and WSDL* section, this `openAccountXML` element indicates which operation of the service is being invoked. Similarly, the service code can be developed in a flexible way with regard to any kind of input message being sent, of which a detailed case study can be found in Chapter 8. Even though any XML content can be sent to the service, the implementation logic will always have constraints on what can be processed. The message must contain data meaningful in the context of the invoked operation.

Another way of further decoupling a service consumer from the service is to insert an intermediary between the two. The intermediary, generally deployed as part of an ESB, can mediate the differences in message format and network protocol to further decouple the service consumer and service. Chapter 12 explores this further as part of its coverage of ESBs.

### SUMMARY OF KEY POINTS

---

- For SOAP-based Web services, coupling between service contract and implementation can also occur. Generating service contracts directly from existing Java logic often exposes language-specific details and creates an unwanted tight coupling between the logic and the contract. However, service consumers can be developed in a generic way independent of a particular service contract using the JAX-WS `Dispatch` API.
  - With REST services, the service contract is tightly constrained by a known set of media types and a handful of HTTP operations.
  - Namespaces and Java package names can help structure code to control and minimize the dependencies among service implementation pieces.
- 

## 7.4 Service Abstraction

The appropriate level of abstraction at which services are described achieves additional agility and alignment between business and IT areas of an enterprise. Abstracting information means taking technical details out of a problem to be solved on a higher level. Since the beginning of computer technology, information has been abstracted into higher levels in a number of ways, for example:

- Assembler constructs, which are instructions for a processor, are translated into a series of 1 and 0.
- Operating systems offer access to system resources and APIs that encapsulate lower-level constructs.
- Programming languages, such as Java, introduce a more abstract way of describing logic, which is then compiled into a format that the operating system can understand.

Service-orientation continues the evolution of higher-level abstraction use to make creating and changing solutions easier. For example, a business process defined with WS-BPEL describes a sequence of service invocations and the data flows between them by expressing this sequence in XML form without writing any actual code. A side effect of this increased abstraction is the ability to utilize visual programming tools that support the creation of process definitions via drag-and-drop-style interfaces. The Service Abstraction principle advocates that the technical details of the technology platform

underlying a service contract are hidden from the service consumer. It also promotes hiding non-essential details about the service itself.

### **Abstracting Technology Details**

The service contract represents an abstraction of the functionality implemented in the service logic. Included in this notion is the abstraction of technical resources utilized to fulfill a service's functionality.

Filenames, database tables, machine names, and network addresses should be omitted from the service contract and completely hidden from service consumers. Given that a service contract should always be designed with a particular business purpose in mind, this should never be a problem.

Concerns arise when services are generated straight out of existing code, because technology details which should have otherwise been abstracted away will often be exposed. For example, whether or not a service is implemented in Java or running in a Java environment such as Java EE should be completely irrelevant to the service consumer.

### **Hiding Service Details**

Maximum flexibility is achieved when the technology used to implement a service and additional details about that service are hidden, which can be divided into information about the input or output message format and contextual information.

Hiding the information about the input or output message format may seem counterintuitive. If a service's input and output messages are hidden, what is left to put into the service contract? Message formats can be abstracted to a generic level without surrendering the message definition altogether.

For example, assume a Credit Check service receives customer information as input. The customer information can be defined and represented by a Customer complex type in the XML schema definition to a detailed level, adding constraint information to each attribute and element of that schema. The length of the `lastName` character field is limited to 35 characters in Example 7.36.

```
<complexType name="Customer">
  <sequence>
    <element name="firstName" type="string"/>
    <element name="lastName">
      <simpleType>
        <restriction base="string">
          <length value="35"/>
        </restriction>
      </simpleType>
    </element>
    <element name="customerNumber" type="string"/>
  </sequence>
</complexType>
```

---

**Example 7.36**

XML schema definition with added constraint inheritance can limit the `lastName` field to a set number of characters.

The XML schema definition in Example 7.36 maps to a Java class, `Customer.java`, as seen in Example 7.37. (The generated Javadoc comments are omitted.)

```
public class Customer {
    @XmlElement(required = true)
    protected String firstName;
    @XmlElement(required = true)
    protected String lastName;
    @XmlElement(required = true)
    protected String customerNumber;

    public String getFirstName() {
        return firstName;
    }
    public void setFirstName(String value) {
        this.firstName = value;
    }
    public String getLastName() {
        return lastName;
    }
    public void setLastName(String value) {
        this.lastName = value;
    }
    public String getCustomerNumber() {
        return customerNumber;
    }
    public void setCustomerNumber(String value) {
```

```

        this.customerNumber = value;
    }
}

```

**Example 7.37**

The generated Java type does not include the schema type restriction.

**NOTE**

The limit on the length of the `lastName` element was not carried over into Java because Java has no concept of a fixed-length string.

On the opposite end of the abstraction spectrum would be a message definition stating that the incoming message is an XML document with a root `Customer` element. No information is given about individual attributes or elements contained in the document, as seen in Example 7.38.

```

<complexType name="Customer">
  <sequence>
    <any/>
  </sequence>
</complexType>

```

**Example 7.38**

An XML schema definition using the `<any/>` element

The generic type definition leads to a generic Java class, as shown in Example 7.39.

```

public class Customer {
    @XmlAnyElement(lax = true)
    protected Object any;
    public Object getAny() {
        return any;
    }
    public void setAny(Object value) {
        this.any = value;
    }
}

```

**Example 7.39**

The `<any/>` element is mapped to `java.lang.Object`.

Hiding service details and increasing what is abstracted about a service may always appear to be a prudent step. However, the service consumer is sometimes not provided with all of the necessary information on what exactly the service expects and what will be returned in response. Details of an interaction on both sides are left to be resolved at design-time or runtime (outside of the service contract).

In Example 7.38, the generic version states that a `Customer` object contains a `java.lang.Object`, which must be defined at runtime to allow for processing. Generally, more Java code must be written for errors that can occur at runtime. For XML payloads, this consideration is equally valid for SOAP and REST services, as the mechanics perform the same role for both in mapping XML to Java via a binding tool such as JAXB. An abstract service contract can be appropriate in some instances, such as utility services.

Contextual data can also be hidden, as this type of data is commonly about an interaction which does not contain any business-relevant payload information. When using SOAP-based Web services, SOAP header fields store contextual data, such as unique message identifiers, timestamps, and service consumer identity. For greater abstraction, detailed information about contextual header fields can be left out of the service contract altogether. This contextual information can be added or removed depending on the environment in which a service runs, is not relevant for its business purpose, and can often be left out of the contract.

For REST services, in the absence of any kind of a payload envelope, contextual information must be part of a resource representation. Such resource metadata can still be packaged inside specially designated header elements. The technical details of a service that are not part of the service interface, such as a WSDL or service-level information about response times and availability, can be abstracted. The technical details can be important to know, but often change and depend on a particular runtime environment and deployment of a service. For REST services, such service-level agreement characteristics can be described in a separate document.

### Document Constraints

Non-technical information about a service cannot be articulated in a standard format. A generic example of this is a service-level agreement, but may also include other constraints about the usage of a service, valid ranges of input data beyond what can be expressed in WSDL and XML schema, and any additional applicable documentation.

A service can be implemented and deployed in different ways throughout an enterprise. As such, this documentation should not be directly linked with a service. Java, for

example, is well suited for deployment on multiple platforms and operating systems. A Unix-based environment has different performance, scalability, and availability characteristics than a Windows-based system. Additionally, a Java EE application server can be leveraged to host the Java logic. A service instance can run on just one server instance on a small machine. As reuse of the service increases, the service instance can be moved to a clustered environment with greater computing power.

As per the Dual Protocols pattern, a Web service offered over HTTP can be later exposed via JMS for additional reliability requirements by particular service consumers. Abstract service contracts provide the freedom to make changes throughout the lifetime of the service, without breaking or violating previous versions. REST service implementations are synonymous with HTTP, making transport mechanism abstraction a non-issue. As the information about a service grows in abstraction, the service implementation and service consumer logic must become more flexible to anticipate future changes.

---

#### SUMMARY OF KEY POINTS

---

- Details about a specific technology used to implement and/or host a service should be abstracted out of a service contract, which is achieved by REST services over HTTP by default.
  - Services can be built in a more abstract fashion by using abstract and generic message specifications and leaving contextual information out of the service contract altogether.
  - Non-technical information about a service often assumes a separate life-cycle from the service contract and its implementation.
- 

## 7.5 Service Composability

The composability of a service is an implicit byproduct of the extent to which the other service-orientation design principles have been successfully applied. The ability to compose services into new, higher-level services is an inherent, logical consequence if the other design principles outlined in this chapter are followed. For example, a service contract that is standardized allows interaction and composition between services implemented in different languages using different runtime environments. Decoupling a service implementation from its contract allows the same logic to be reused in other compositions.

With regards to the implications that service composability has for the service contract, this section highlights some of the issues and requirements for the runtime environment in which the services run. See Chapter 11 for further exploration of service composition with Java.

### Runtime Environment Efficiency

A highly efficient and robust runtime environment is required for service compositions. The ability for services to participate in multiple compositions places severe challenges on runtime environments and must be taken into account when designing a service inventory.

If a service is used by multiple compositions, applying different non-functional characteristics can be necessary. One composition rarely invokes a service with no particular requirement for fast response times, whereas another composition using the same service can require support for high transaction loads with short response times. Similarly, one composition can require a reliable connection between the composition controller and its members, whereas another can be tolerant of lost messages. Applying different QoS definitions to the same service, depending on the composition, must be possible without requiring code changes in the service logic itself.

Java, as a programming language, has no built-in features that help or hinder the runtime environment. In most cases, Java-based services are hosted in a Java EE-compliant application server or an alleged Web server environment, which is Java SE-compliant. In either case, the runtime environments provide advantages over other traditional server platforms in terms of composability.

#### NOTE

Depending on how a given service or service consumer participates in a service composition at runtime, it may assume one or more roles during the service composition's lifespan. For example, when a service is being composed, it acts as a composition member. When a service composes other services, it acts as a composition controller. To learn more about these roles, visit [www.serviceorientation.com](http://www.serviceorientation.com) or read Chapter 13 of *SOA Principles of Service Design*.



Java EE defines the management of solution logic according to the roles people perform, such as component provider, deployer, and assembler. As a result, much of the information on the execution of solution logic at runtime is not hardcoded into the logic itself but is instead stored in declarative configuration files, called deployment descriptors. The same piece of Java code can be used differently with different deployment descriptors and changed at runtime without requiring recompilation of the code. Despite being undefined by the Java EE standard, most Java EE application servers support defining individual server instances with specific runtime definitions. One instance usually runs in one process with its own configuration for elements like thread pools or memory heap sizes. In many cases, instances can be clustered to provide one logical application server across multiple physical processes or machines.

Separating runtime information about a component from its actual logic is possible because the Java EE application server runs as a container. This means that all incoming and outgoing data is intercepted and processed by the application server based on the current configuration. For example, if a certain piece of Java logic can only run within the scope of a transaction, the container can be configured to ensure a transaction is present whenever this piece of logic is invoked.

The same approaches to the runtime environment apply to Web services on two levels:

1. The runtime environment that processes an incoming message, such as a SOAP or REST request message, can perform several tasks before forwarding the request to the actual service. These tasks include the ability to decrypt data sent in an encrypted form, establish handshaking between service consumer and service by returning acknowledgement that the message has been received, interpret information in the request message header indicating that the invocation of the service must be part of a distributed transaction, and convert incoming XML data into Java objects.
2. For SOAP-based Web services, JAX-RPC and JAX-WS provide a mechanism and an API that allow the insertion of custom logic to parse, interpret, or change incoming and outgoing messages. The custom logic runs inside a handler. For JAX-RS-based implementations of REST services, such interception logic can be implemented by developers in special entity handlers known as entity providers. The JAX-RS 2.0 release provides added filters and interceptors otherwise not included in previous versions.

Both approaches are similar regardless of the Web services used. One is controlled by the user of the system, whereas the other is implicitly included with the runtime. Reading

and manipulating incoming and outgoing messages separate from the service logic is crucial to supporting the Service Composability principle. Composing hosted services, including both composition members and composition controllers, is supported by the concept of containers and deployment descriptors and enhanced by SOAP-based Web services, such as handlers.

Ultimately, implementing a highly efficient runtime allows the developer to focus on the business logic of the actual service implementation, leaving everything else to the underlying Java platform. More advanced technologies, such as the SCA, expand on this concept by separating core business logic implementation further away from aspects of a service component, such as the protocol bindings used to interact with service consumers and other services used to fulfill functionality.

### Service Contract Flexibility

Service contracts can be designed to increase the ability of a service for multiple compositions. Generally, multiple compositions are only applicable to the composition members for increasing the reusability of a service in different business contexts or business tasks. A service contract can be rewritten to enable reuse of a service without changing the core functionality of the service.

To write a flexible service contract that is reusable, recall the following approaches:

- Use generic data types or supertypes instead of concrete subtypes. If an enterprise deals with both commercial customers (`CommercialCustomer`) and personal customers (`PersonalCustomer`), evaluate whether a common supertype can be established for both (`Customer`) to be used in the service contract. JAXB supports polymorphism to ensure that the appropriate Java object is created when a message containing a subtype is received.
- Decouple the service contract from its underlying runtime platform by hiding details about QoS characteristics, which can change over time and will vary depending on service consumer requirements and how a service is deployed.
- Decouple the service implementation from its contract by utilizing generic APIs, such as the JAX-WS Provider API for SOAP-based Web services. For REST services, deal with generic Java types in resource methods, such as `String`, `byte[]`, and `InputStream`. Note that generated generic service consumer or service logic results in additional code that must be developed and tested.

### Standards-Based Runtime

Composition members and controllers benefit from a runtime environment that supports a wide range of accepted standards. Composing services means the services interact and interoperate. Interoperability of services is supported by a runtime environment upheld by relevant standards, such as the WS-I. Java's APIs for SOAP-based Web services, JAX-RPC and JAX-WS, require support for the WS-I Basic Profile, which allow services to be designed with a high degree of interoperability. REST services achieve full interoperability inherently through HTTP.

Advanced standards relevant in a composition of services include the WS-Security standards for which a WS-I profile also exists, the WS-Transaction and related standards, WS-Addressing, and WS-ReliableMessaging, which supports the reliable exchange of messages between services. These advanced standards are combined in another WS-I profile known as the Reliable Secure profile.

---

#### SUMMARY OF KEY POINTS

---

- Composability is supported in utilizing a runtime environment that allows hosting both composition members and composition controllers efficiently and flexibly, such as Java EE-compliant application servers.
  - Creating flexible service contracts can facilitate the use of services as composition members.
  - Platforms that support accepted and established standards can be utilized to improve service interoperability and composability.
-

## 7.6 Service Autonomy

Service-orientation revolves around building flexible systems. Flexibility is, to a large degree, achieved through making services decoupled and autonomous to enable them to be composed, aggregated, and changed without affecting other parts of the system. For a service to be autonomous, the service must be as independent as possible from other services with which it interacts, both functionally and from a runtime environment perspective. Java and Java EE provide a highly efficient runtime environment that supports service composition. For example, a Java EE application server supports concurrent access to its hosted components, making each access to such a component autonomous from the others.

### Well-Defined Functional Boundary

Occasionally, the functional boundary is defined by a certain business domain that a service lives within, as is the case if a service implements a particular business process or task within that domain. Alternatively, the functional boundary of a service can be described by the type of data the service operates on, such as entity services.

Translating this requirement into Java and XML schema utilizes namespaces and Java packages as structuring elements. Checking the list of imported classes and packages for a particular service implementation will help identify dependencies throughout the system and provide an indication of whether the functional boundary of the service is maintained in its implementation.

For example, an entity service called `Customer` delivers customer data retrieved from a variety of data sources for reuse across many business processes. The service is defined in the `http://entity.services.acme.com/Customer` namespace and uses `com.acme.services.entity` as the Java package for its implementation. The `Customer` service imports a package called `com.acme.services.accounting`, which immediately identifies that the Java service implementation contains a potentially undesired dependency on a business domain-specific piece of code. This warrants further investigation of the underlying logic and removal of the dependency.

The `Customer` service has a well-defined functional boundary in delivering relevant customer data to its service consumer. However, a dependency on logic specific to the `Accounting` service business domain naturally reduces the autonomy of the `Customer` service.

### Runtime Environment Control

The underlying runtime influences the degree of autonomy a service can achieve. For each service to have control over its runtime environment, the environment must be partitioned to allocate dedicated resources accordingly. The JVM offers all internal code a degree of autonomy by isolating the executed code from the operating system and providing controlled access to physical resources, such as files or communication ports. Java EE application servers leverage the concept of a container in which components run.

For SOAP-based Web services, runtime control in Java can be achieved (while maintaining a high degree of autonomy) by exposing plain JavaBeans as Web services or utilizing Stateless Session EJBs. The service implementation and all other relevant artifacts, such as WSDL files, are packaged in a module, such as a JAR or WAR file, which can then be installed on an application server independent of other code running on that server.

The same is true for non-Web services or services implemented as regular EJBs. The components related to the service have individual private deployment descriptors that can be configured on the application server as independent entities.

Java EE allows for the packaging of multiple modules and multiple services in one Enterprise ARchive file. This EAR file is deployed and installed on the application server as an independent enterprise application. To increase a service's autonomy, use of only one service packaged per EAR file is recommended. This allows each service to be treated as a completely independent unit to configure, start, stop, or replace without affecting any other services running on the same system.

To decrease the number of moving parts in the environment, however, multiple services can be packaged into one enterprise application. Co-locating services is suitable when the services interact frequently with each other in a performance-critical manner.

In JAX-RS, POJOs are more commonly used to model Web resources, although stateless and single session beans can be designated as root resource classes. The JAX-RS runtime packages the resource artifacts into a WAR or EAR file which can be deployed as a standalone module in an application server. Some JAX-RS implementations support embeddable containers, in which a JAX-RS runtime is bootstrapped from inside a driver program for testing purposes.

## High Concurrency

Developing services for reuse across multiple service consumers is a benefit of implementing service-oriented design. The considerations presented in the *Agnostic Functional Contexts* and *Concurrent Access to Service Logic* sections help illustrate the benefits of providing each service consumer exclusive access to a certain instance of the service implementation, which is true regardless of whether the service implementation is located in a `JavaBean` or is a stateless session `EJB`.

For REST services, the default lifecycle of root resource classes is per-request. A new instance of a root resource class is created every time the request URL path matches the `@Path` annotation of the root resource. With this model, resource class fields can be utilized without concern for multiple concurrent requests to the same resource. However, a resource class can also be annotated with the `javax.inject.Singleton` annotation, creating only one instance per Web application. Using the default lifecycle model for resources is recommended, unless compelling reasons arise to do otherwise.

In general, each Java component that implements a service is accessible concurrently by definition via the application server. However, installing the same service separately on different machines is also acceptable. The ultimate autonomy of a service is achieved if one instance of a service, running in its own process and controlling all of its underlying resources, serves only one specific service consumer. While inefficient from a resource utilization and maintenance perspective, requirements can dictate a service as part of a mission-critical business process which requires high performance and high transaction loads that force a dedicated instance of the service to be deployed for that specific purpose.

### NOTE

Establishing environments that support high levels of concurrent access introduces scalability considerations that some IT enterprises may not be equipped to handle, relational directly to the amount of service compositions a given service participates in and the amount of access within each composition the service is subjected to.

Cloud computing platforms provide infrastructure with IT resources that can dramatically improve the extent to which the Service Autonomy principle can be applied to a service implementation, by reducing or eliminating the need for the service implementation to share or compete for resources within the enterprise boundary. For more information about scalability and elasticity as part of cloud computing environments, see the series title *Cloud Computing: Concepts, Technology & Architecture*.

---

### SUMMARY OF KEY POINTS

---

- A well-defined functional boundary of a service, reflected in its contract, ensures a high degree of autonomy. However, this boundary must also be maintained in the service implementation by avoiding any dependencies on other services outside of that functional boundary.
  - A service can increase autonomy by having complete control over its runtime. Java and Java EE support control over the runtime with the concept of containers, which depict virtual runtime boundaries that can be controlled individually and independently.
  - Despite being accessed concurrently by multiple service consumers, services running in a Java application server can run autonomously on behalf of each service consumer. The JAX-WS and JAX-RS programming models automatically ensure that a new thread is started for each new service request.
  - Highly critical services can have total autonomy by being deployed on a server for exclusive access by one particular service consumer.
- 

## 7.7 Service Statelessness

Each invocation of a service operation is completely independent from any other invocation, whether by the same service consumer or any other service consumer. The Service Statelessness principle offers various benefits centered around improved scalability by which additional stateless service instances can be easily provisioned on available environments. With the advent of cloud computing, on-demand scaling out of services is considered as a natural evolutionary step for stateless services.

Many real-life business scenarios can be expressed as business processes that include automated steps, which can require manual intervention. Designing and implementing such a process requires some state to be maintained for the duration of the process. Executing an instance of the process definition forms the notion of a session or transaction across multiple service invocations. Therefore, a service implementing the execution of a business process cannot be stateless and may need to even maintain context information over extended periods.

## Orchestration Infrastructure

An orchestration infrastructure, such as a WS-BPEL execution environment, will support state maintenance either by storing state in the local stack of the thread executing the process or in a permanent relational data store. A permanent relational data store ensures that a process instance can be continued after a system crash or other interruption. This style of statefulness is built into WS-BPEL and its execution environment, so there is little to be aware of when designing such a service. Designing a service that aggregates other service invocations in its implementation without utilizing WS-BPEL will require the developer to decide whether to store temporary data in a relational data store for later recovery in case of failure.

## Session State

Another aspect of achieving service statelessness occurs when a service must establish some form of session with a service consumer, such as when the state is not kept in the calling logic but in the called service itself. Compare this to the shopping cart scenario in which a service consumer uses a service repeatedly to collect a set of data, which is then committed all in one final step. Java Servlets, which at the core also offer a stateless programming model, leverage the `HTTPSession` information and the concept of cookies to enable data to be stored on behalf of a certain client.

For REST services, using cookies as handles to store client state violates the stateless constraint of a REST-style architecture. Any maintenance of server-side session state is not recommended, and REST limits application state to be held on the client and not on the server.

Using cookies breaks the statelessness model as the server is expected to maintain a reference to a session and use the information in the incoming cookie for discovery. Each invocation is expected to carry the complete contextual information without any references to any previous interaction, promoting scalability where any available service instance can process this request. For idempotent requests, the failed request can be redirected to another functioning stateless service. Note that the convenience of server-side session maintenance and the associated personalization benefits are sacrificed for superior scalability and optimum resource usage.

REST services can read or write state information for various domain entities from or to databases or other persistent stores, but storing client state violates the statelessness constraint with implications for scalability and resilience. SOAP-based Web services



have no such constraints and the servlet model is leveraged by JAX-WS to handle stateful interactions.

### Storing State

Two methods are available for clients to store state as part of a repeated interaction with a servlet. The first allows the data to be sent back and forth with each request, and the amount of data grows with each invocation. The second method enables the servlet to maintain the state, and the servlet sends back a key to this data (the cookie). The client sends the same key along for a subsequent request, which allows the servlet to retrieve the appropriate state for this particular client. The interface used to store the data is called `javax.servlet.http.HttpSession`.

The behavior described in maintaining a key is a well-defined part of the standard servlet programming model, and APIs are available to enable its use. The `javax.servlet.http.HttpSession` interface offers methods for storing simple key-value pairs. JAX-WS describes a mechanism to pass a reference of the `HttpSession` instance into the service implementation class.

A Web service implementation can be made stateful by leveraging HTTP sessions and the JAX-WS support for the sessions. This requires that the service consumer receives and reuses the cookie sent back with the response to the original HTTP POST request.

#### CASE STUDY EXAMPLE

Several of NovoBank's business processes require that certain forms and brochures be mailed to customers. The need for mailing each form is established in a separate part of the process, but in each case, the same back-end system is called via a Web service interface offering an `addOrderForm` operation. To reduce mailing costs, the design team makes the service establish a session with each process instance so that all orders can be bundled together. This requires the service to keep track of all forms requested. All orders are confirmed at once with the invocation of the `confirm` operation.

Example 7.40 shows an excerpt of the WSDL definition for the Order Form service.

```

<definitions targetNamespace="http://utility.services.novobank.com/"
  name="OrderFormService"
  xmlns:tns= "http://utility.services.novobank.com/"
  xmlns:xsd= "http://www.w3.org/2001/XMLSchema"
  xmlns:soap= "http://schemas.xmlsoap.org/wsdl/soap/"
  xmlns= "http://schemas.xmlsoap.org/wsdl/"
  ...
  <xs:element name="addOrderForm" type="ns1:addOrderForm"
    xmlns:ns1="http://utility.services.novobank.com/" />
  <xs:complexType name="addOrderForm">
    <xs:sequence>
      <xs:element name="arg0" type="xs:string" minOccurs="0" />
    </xs:sequence>
  </xs:complexType>
  <xs:element name="addOrderFormResponse" type="ns2:
    addOrderFormResponse" xmlns:ns2="http://utility.services.
    novobank.com/" />

  <xs:complexType name="addOrderFormResponse" />
  <xs:element name="confirm" type="ns3:confirm" xmlns:ns3= "http://
    utility.services.novobank.com/" />

  <xs:complexType name="confirm">
    <xs:sequence>
      <xs:element name="arg0" type="xs:string" minOccurs="0" />
    </xs:sequence>
  </xs:complexType>
  ...
  <xs:element name="confirmResponse" type="ns4:confirmResponse"
    xmlns:ns4="http://utility.services.novobank.com/" />
  <xs:complexType name="confirmResponse" />
  <portType name="OrderForm">
    <operation name="addOrderForm">
      <input message="tns:addOrderForm" />
      <output message="tns:addOrderFormResponse" />
    </operation>
    <operation name="confirm">
      <input message="tns:confirm" />
      <output message="tns:confirmResponse" />
    </operation>
  </portType>
</definitions>

```

**Example 7.40**

The WSDL definition for the Order Form service

Nothing in the contract indicates that the `addOrderForm()` operation keeps state from previous invocations by the same client in its `HTTPSession`, meaning this information must be documented elsewhere.

Example 7.41 identifies the implementation class for the Order Form service with the `addOrderForm()` operation in Java.

```
package com.novobank.services.utility;

import java.util.Vector;
import java.util.Iterator;
import javax.annotation.Resource;
import javax.jws.WebService;
import javax.servlet.http.HttpServletRequest;
import javax.servlet.http.HttpSession;
import javax.xml.ws.WebServiceContext;
import javax.xml.ws.handler.MessageContext;

@WebService
public class OrderForm {

    @Resource
    private WebServiceContext webServiceContext;

    public void addOrderForm(String formNumber) {
        System.out.println("Form with number "+formNumber+" was
            ordered.");
        HttpSession session = retrieveSession();
        Vector<String> formList = (Vector<String>)session.
            getAttribute("formList");
        if (formList==null) {
            formList = new Vector<String>();
        }
        formList.add(formNumber);
        session.setAttribute("formList", formList);
    }

    private HttpSession retrieveSession() {
        MessageContext messageContext = webServiceContext.
            getMessageContext();
        HttpServletRequest servletRequest =
            (HttpServletRequest)messageContext.get(MessageContext.
                SERVLET_REQUEST);
        return servletRequest.getSession();
    }
}
```

```

    }
    ...
}

```

### Example 7.41

The implementation class for the Order Form service

Each invocation of the `addOrderForm()` method retrieves the `HTTPSession` instance by using the `WebServiceContext` attribute injected via the `@Resource` annotation. This attribute provides access to the `MessageContext` for the request, which in turn stores a pointer to the servlet request object. Finally, the servlet request allows the session to be retrieved. These steps are all encapsulated in the private `retrieveSession()` method.

The list of ordered form numbers is stored in a `Vector`, which is kept in the `HTTPSession` under the name `formList`. Each time the `addOrderForm()` method is called, the new form number is added to that `Vector`.

What is missing from Example 7.41 is the implementation for the `confirm()` operation. This is where the accumulated information from previous invocations is used, as seen in Example 7.42.

```

public void confirm(String customerNumber) {
    HttpSession session = retrieveSession();
    List<String> formList = (Vector<String>)session.
        getAttribute("formList");
    if (formList==null) {
        System.out.println("No orders found.");
    } else {
        System.out.println("Confirming "+formList.size()+" orders.");
        for (String s : formList) {
            System.out.println("Order for Form " + s + ".");
        }
        session.removeAttribute("formList");
    }
}
}

```

### Example 7.42

Access to the `HTTPSession` is provided using the same private method called `retrieveSession()` which provides the complete list of ordered forms. Note how the list is reset by setting the `formList` attribute in the `Vector` to `null` after processing.

The service consumer of this Web service must store the cookie returned from the first invocation to send along with any subsequent invocations, which JAX-WS is equipped to manage. The NovoBank team decides to provide a sample client along with their service for service consumer reuse, as shown in Example 7.43.

```
package com.novobank.services.utility.client;

import java.util.Map;
import javax.xml.ws.BindingProvider;

public class OrderFormClient {

    public static void main(String[] args) {
        OrderForm orderForm =
            new OrderFormService().getOrderFormPort();

        Map<String, Object> requestContext =
            ((BindingProvider)orderForm).getRequestContext();

        requestContext.put(BindingProvider.SESSION_MAINTAIN_PROPERTY,
            true);

        orderForm.addOrderForm("123");
        orderForm.addOrderForm("456");
        orderForm.confirm("any customer");
    }
}
```

**Example 7.43**

Sample client code for the Order Form service

Note that the local service proxy is cast to the `javax.xml.ws.BindingProvider` interface so that the request context can be retrieved and the `SESSION_MAINTAIN_PROPERTY` value can be set to `true`.

HTTP sessions and JAX-WS are only applicable for services with SOAP/HTTP bindings, although similar behavior in a service not accessed over HTTP can be created using the same principles. At invocation by a specific service consumer, the service can return a unique identifier to the service consumer, such as in the SOAP response header, and expect that identifier to be returned with every subsequent request. The service uses the identifier as a key into a table where the data is stored.

How the data is physically stored depends on the developer's requirements. The data stored on behalf of specific service consumers must be recoverable across a server restart persistently in a relational database, which can be accessed using JDBC or a similar mechanism.

Note that state as discussed in this section is usually transient and only exists for the duration of the interaction with the relevant service consumer. Business-relevant data, which needs to be persisted permanently, would not be stored using the mechanisms described. Business data should be stored using data access APIs, such as JDBC, or persistence frameworks, such as JPA or Hibernate.

### SUMMARY OF KEY POINTS

---

- Many scenarios require data to be stored beyond a single invocation of a service in the calling service, such as from either within a WS-BPEL process or in the called service. Data is often stored for long durations.
  - JAX-WS provides a mechanism to utilize the `HTTPSession` object to store temporary state on behalf of a specific service consumer.
  - Business-relevant data that must be stored permanently should be handled through entity services that explicitly wrap the handling of such data.
- 

## 7.8 Service Discoverability

The two primary aspects of the Service Discoverability principle are discovery at design-time (which promotes service reuse in a newly developed solution), and discovery at runtime (which involves resolving the appropriate endpoint address for a given service or retrieving other metadata). Even though the information can be physically stored in one place, the way in which the information is accessed in each scenario varies.

### Design-Time Discoverability

At design-time, it is important for project teams to be able to effectively identify the existence of services that contain logic relevant to the solution they plan to build. This way they can either discover services that they can reuse or confirm that new service logic they plan to build does not already exist. For example, a service is designed to address

an Order Management business process for which customer information is required. The service designer must investigate whether a `Customer` data type already exists, and if so, determine whether it meets the requirements for the Order Management Process service. If the data sent into the newly designed service is missing information, the designer can also check whether an entity service encapsulating all data access to this type of data exists. Additional customer information required can be built or retrieved via a `Customer` entity service.

During the design of a new service, several types of artifacts must be evaluated for reuse directly on the new service's interface and within the service via some kind of aggregation. This includes non-functional aspects, which may not be expressed in machine-readable form. Meta information, such as performance and reliability of existing services, can influence whether a service is reusable in a new context.

Code can exist for existing data type definitions. JAXB, for example, defines a mapping between XML schema and Java. Java code can be directly generated from a schema definition and reused wherever that particular data type is used.

All of the relevant information must be available to the designer during design-time, ideally via the development environment directly. This relevant information includes the artifacts themselves, such as the WSDL, XML schema, and Java source code as well as relationships and dependencies between them. These dependencies must be documented thoroughly, as part of the service profile so that a complete metamodel of the service exists once the design is complete.

Since the design of a service is manually performed by a service designer, this information must be accessible and searchable by humans. How this feature is enabled depends on the registry type used and the access mechanisms offered, without depending on the programming language used to implement services or on the runtime platform that the service will run on. The underlying mechanism should offer a way to control access to the information and support versioning of artifacts.

### **Runtime Discoverability**

Runtime service discovery refers to the ability of software programs to programmatically search for services using APIs exposed by the service registry. Doing so allows for the retrieval of the physical location or address of services on the network. Because services may need to be moved from one machine to another or perhaps redundantly deployed on multiple machines, it may be advisable for service addresses not to be hardcoded into the service consumer logic.

For SOAP services using JAX-WS, the location of the target service is included in the `<port>` element of the WSDL definition by default. In turn, the location of the WSDL file is automatically added to the generated `...Service` class and passed to the tooling, or `wsimport`, that generates the appropriate service consumer artifacts. Pointing to the local WSDL file in the `filesystem` will result in a service consumer that cannot be moved to a different environment, because functionality depends on that particular location for the WSDL file and the endpoint address contained by the WSDL file.

Using the URL and appending `"?wsdl"` to the service is a more flexible approach. For example, a service can be located at `myhost.acme.com:8080/account/account`, in which case the WSDL definition can be retrieved via the URL `http://myhost.acme.com:8080/account/account?wsdl`. During the installation and deployment of a service, the endpoint information in this WSDL file is updated to reflect the real address of the hosting system.

The address points to a fixed location for the WSDL file despite now residing on the network, meaning the address of any target services should always be resolved separately from the generated service consumer code. JAX-WS provides a way to set a target endpoint address on a proxy instance at runtime that utilizes the `javax.xml.ws.BindingProvider` interface. Each proxy implements this interface and allows properties to be dynamically set on an exchange between a service consumer and service, with the endpoint address being one of the pre-defined properties. The code in Example 7.44 illustrates how a service consumer sets a new endpoint address on a service proxy before invocation.

```
String endpointAddress = ...; //retrieve address somehow
OrderForm orderForm =
    new OrderFormService().getOrderFormPort();
java.util.Map<String, Object> requestContext = (javax.xml.
    ws.BindingProvider)orderForm).getRequestContext();
requestContext.put (BindingProvider.ENDPOINT_ADDRESS_PROPERTY,
    endpointAddress);
orderForm.addOrderForm ("123");
```

---

#### Example 7.44

The WSDL location can alternatively be set for a service at runtime instead of the endpoint address, to maintain the address of the service together with the rest of the service contract in its WSDL file and not in multiple locations, as shown in Example 7.45.



```
URL wsdlLocation = ...; //retrieve WSDL location
OrderFormService orderFormService =
    new OrderFormService(wsdlLocation, new QName("http://utility.
        services.novobank.com/", "OrderFormService"));
OrderForm orderForm = orderFormService.getOrderFormPort();
orderForm.addOrderForm("123");
```

---

**Example 7.45**

In addition to the lookup of a service endpoint address, other artifacts can be used by service consumers at runtime to identify an appropriate service and build the applicable request message. The JAX-WS `Dispatch` API allows a service request to be completely built at runtime. Theoretically, a service consumer could look up a WSDL definition at runtime, read its `portType` and XML schema definitions, and build the right messages from scratch. However, building a service request completely at runtime is impractical in a real-life scenario, as the service will likely perform poorly and require plenty of tedious coding.

Non-functional characteristics of a service are other examples of information that a service consumer can discover at runtime. For example, assume that two physical instances of a service exist on both a slower machine and a faster machine. The service consumer can retrieve this information and select the appropriate service to be used, depending on the business context of the call. For instance, a silver customer is routed to the slower machine, whereas a gold customer is routed to the faster machine. The service is still implemented in both invocations, as the differentiation in routing is a non-functional characteristic.

API endpoints are unnecessary in a REST-based system because no RPC-style invocation is involved. Recall that URI-based addressability, like statelessness, is a formal REST constraint. Resources that offer services are the fundamental entities and must be discoverable through URIs before clients can invoke operations on them. The hypermedia constraint, if followed accurately, requires only the URI of the root resource to be provided to the service consumer.

Various operations on the root resource will publish URIs of related resources, which can then be used by the service consumers to reduce the amount of foreknowledge required. However, service consumers would still require knowledge of the semantics associated with the URIs to be able to make meaningful use of them, which makes this approach impractical in real-life application.

## Service Registries

Information can be retrieved by a service consumer at runtime with the service registry. As the following methods are inapplicable to REST services, the remainder of this discussion will only apply to SOAP-based Web services. Storing information about WSDL locations and endpoint addresses in a file accessible through a proprietary API at runtime is a retrieval mechanism appropriate for small environments. However, the format in which the information is stored, such as XML or a character-delimited format, must be defined, and the location of this file must always be maintained throughout the enterprise for easy accessibility.

Alternatively, storing the information in a relational database allows for remote access and query using a standard language, such as SQL, although a proprietary relational model for this information must still be invented. Other options include leveraging the JNDI or LDAP to serve the same purpose.

In the early days of Web services, a platform-independent standard describing how to uniformly store, find, and retrieve business and technical information about services was developed as the UDDI registry, which offers a query and publish interface. As a Web service, the UDDI registry allows the API to be described by the WSDL so that any Web service-capable service consumer can access a UDDI registry by generating a proxy from the standard WSDL. For accessing registries at runtime, many IDEs such as Eclipse's Web Tools Platform have built-in support for existing UDDI registries.

### NOTE

Discoverability processing can be delegated into the ESB, where tasks such as service lookup can be handled centrally and uniformly. The service lookup logic does not then clutter the service consumer, which can instead focus on the business purpose.

### SUMMARY OF KEY POINTS

---

- For Web services, JAX-WS provides runtime APIs to set the endpoint address of a target service. UDDI defines a standardized way for storing service meta information, which includes access via a Web service-based API.
  - REST services use embedding-related resource links in resource representations to leverage the hypermedia constraint and lead the client through a discovery of networked resources.
-

*This page intentionally left blank*

# Index

## A

abstraction. *See* Service Abstraction  
design principle

ACC (Applet and Application Client  
Container), 42

ACCEPT header, HTTP, 109-110

Agnostic Capability design pattern, 427

Agnostic Context design pattern, 428

agnostic functional contexts, 140

agnostic logic, defined, 31

Agnostic Sub-Controller design pattern,  
429

annotated SOA Manifesto, 28, 38,  
520-532

annotations

defined, 54

in JAX-RS standard, 125-129

in JAX-WS standard, 113-115

## APIs

Java EE (Java Platform, Enterprise  
Edition), 49-59

*JPA (Java Persistence API),*  
53-54

*JTA (Java Transactions API),*  
50-51

Java SE (Java Platform, Standard  
Edition), 46-49

*JDBC (Java Database  
Connectivity), 49*

*JNDI (Java Naming and  
Directory Interface), 48*

*RMI (Remote Method  
Invocation), 46-47*

*RMI over IIOP, 47-48*

**Applet and Application Client Container  
(ACC), 42**

application programming interfaces. *See*

APIs

application servers

IBM WebSphere, 75-77

architecture

of entity services, 263-264

Java EE (Java Platform, Enterprise Edition), 41-44

tiers, 45-46

Java SE (Java Platform, Standard Edition), 41

messaging

*with ESBs*, 392-396

*without ESBs*, 389-391

of utility services, 212-220

asynchronous communications, 135-136

asynchronous invocation, 338-339

Asynchronous Queuing design pattern, 384, 430

Atomic Service Transaction design pattern, 431

availability, 369

## B

Big Data Science Certified Professional (BDSCP), 9

binary data

in REST services, 170-172

in Web services, 165-166

bindings, 153

Brokered Authentication design pattern, 432

## C

cacheability, 131

caching, 371-374

Canonical Expression design pattern, 433

Canonical Protocol design pattern, 434

Canonical Resources design pattern, 435

Canonical Schema Bus compound pattern, 437

Canonical Schema design pattern, 436

Canonical Versioning design pattern, 438

Capability Composition design pattern, 439

capability granularity, defined, 34

Capability Recomposition design pattern, 440

capacity, 369

case study examples

NovoBank

*background*, 13-16

*conclusion*, 406-407

SmartCredit Co.

*background*, 16-18

*conclusion*, 407-408

CCP (Cloud Certified Professional), 9  
 change notifications, 268-269  
 Cloud Certified Professional (CCP), 9  
 cloud computing, defined, 29-30  
*Cloud Computing: Concepts, Technology  
 & Architecture*, 196  
 Common Object Request Broker  
 Architecture (CORBA), 47-48  
 Commons Logging, 231  
 Compatible Change design pattern, 441  
 Compensating Service Transaction  
 design pattern, 442  
 components  
   symbols for, 29  
 composability. *See* Service  
   Composability design principle  
 Composition Autonomy design pattern,  
 443  
 composition controllers, 336  
 composition initiators, 337  
 composition sub-controllers, 337  
 compound patterns  
   Canonical Schema Bus, 437  
   Enterprise Service Bus, 95-96, 461  
   Federated Endpoint Layer, 466  
   Official Endpoint, 480  
   Orchestration, 481  
   Service Broker, 392, 495  
   Three-Layer Inventory, 512  
   Uniform Contract, 515

Concurrent Contracts design pattern,  
 444  
 CONNECT method, HTTP, 107  
 Connection API, 279  
 constraint granularity, defined, 34  
 Content Negotiation design pattern, 445  
 contexts and dependency injection,  
 49-50  
 Contract Centralization design pattern,  
 446  
 Contract Denormalization design  
 pattern, 447  
 cookies, 198  
 coupling. *See* Service Loose Coupling  
   design principle  
 CORBA (Common Object Request  
 Broker Architecture), 47-48  
 Cross-Domain Utility Layer design  
 pattern, 448

## D

DAOs (data access objects), 230  
 data access objects (DAOs), 230  
 data aggregation, 266-267  
 Data Confidentiality design pattern, 449  
 Data Format Transformation design  
 pattern, 392, 450  
 data granularity, defined, 34  
 data grids, 371-372

**Data Model Transformation design pattern, 392, 451**

**Data Origin Authentication design pattern, 452**

**data transfer objects (DTOs), 280**

**Decomposed Capability design pattern, 453**

**Decoupled Contract design pattern, 454**

**DELETE method, HTTP, 107**

**dependency injection, 49-50**

**deployment**

of composed services, 362-368

of Java EE applications, 44-45

**deployment descriptor, 44-45**

**design**

of entity services, 270-278

of micro-utility services, 253

of omni utility services, 241

of resource utility services, 249-251

of utility services, 221-225

of wrapper utility services, 258-259

**design patterns**

Agnostic Capability, 427

Agnostic Context, 428

Agnostic Sub-Controller, 429

Asynchronous Queuing, 384, 430

Atomic Service Transaction, 431

Brokered Authentication, 432

Canonical Expression, 433

Canonical Protocol, 434

Canonical Resources, 435

Canonical Schema, 436

Canonical Versioning, 438

Capability Composition, 439

Capability Recomposition, 440

Compatible Change, 441

Compensating Service Transaction, 442

Composition Autonomy, 443

Concurrent Contracts, 444

Content Negotiation, 445

Contract Centralization, 446

Contract Denormalization, 447

Cross-Domain Utility Layer, 448

Data Confidentiality, 449

Data Format Transformation, 392, 450

Data Model Transformation, 392, 451

Data Origin Authentication, 452

Decomposed Capability, 453

Decoupled Contract, 454

defined, 36-37

Direct Authentication, 455

Distributed Capability, 456

Domain Inventory, 457

Dual Protocols, 458

Endpoint Redirection, 459



- Enterprise Inventory, 460
  - Entity Abstraction, 462
  - Entity Linking, 463
  - Event-Driven Messaging, 464
  - Exception Shielding, 465
  - File Gateway, 467
  - Functional Decomposition, 468
  - Idempotent Capability, 469
  - Intermediate Routing, 392, 470
  - Inventory Endpoint, 471
  - Legacy Wrapper, 472
  - Lightweight Endpoint, 473
  - Logic Centralization, 474
  - Message Screening, 475
  - Messaging Metadata, 476
  - Metadata Centralization, 477
  - Multi-Channel Endpoint, 478
  - Non-Agnostic Context, 479
  - Partial State Deferral, 482
  - Partial Validation, 483
  - Policy Centralization, 484
  - Process Abstraction, 485
  - Process Centralization, 486
  - Protocol Bridging, 392, 487
  - Proxy Capability, 488
  - Redundant Implementation, 489
  - Reliable Messaging, 490
  - Reusable Contract, 491
  - Rules Centralization, 492
  - Schema Centralization, 493
  - Service Agent, 494
  - Service Callback, 496
  - Service Data Replication, 497
  - Service Decomposition, 498
  - Service Encapsulation, 499
  - Service Façade, 500
  - Service Grid, 501
  - Service Instance Routing, 502
  - Service Layers, 503
  - Service Messaging, 504
  - Service Normalization, 505
  - Service Perimeter Guard, 506
  - Service Refactoring, 507
  - State Messaging, 508
  - State Repository, 374, 509
  - Stateful Services, 375, 510
  - Termination Notification, 511
  - Trusted Subsystem, 513
  - UI Mediator, 514
  - Utility Abstraction, 516
  - Validation Abstraction, 517
  - Version Identification, 518
- design principles**
- Service Abstraction, 184-185, 336
    - profile*, 414
  - Service Autonomy, 194
    - profile*, 417
  - Service Composability, 189-190, 336
    - profile*, 422-423

- Service Discoverability, 204
  - profile, 420-421*
- Service Loose Coupling, 176-177
  - profile, 413*
- Service Reusability
  - profile, 415-416*
- Service Statelessness, 197
  - profile, 418-419*
- Standardized Service Contract, 34, 151
  - profile, 411-412*
- Direct Authentication design pattern, 455**
- disconnected approach (data access), 268**
- discoverability. See Service Discoverability design principle**
- Distributed Capability design pattern, 456**
- Document Object Model (DOM), 63**
- document-style SOAP messaging, 87**
- DOM (Document Object Model), 63**
- domain entities**
  - designing, 271-272
  - message entities versus, 265-266
- Domain Inventory design pattern, 457**
- DTD (XML), 62**
- DTOs (data transfer objects), 280**
- Dual Protocols design pattern, 458**

**E**

- EAR (enterprise archive), 44, 239**
- ebXML (Electronic Business using XML), 120-121**
- EJBs (Enterprise JavaBeans), 43, 51-52**
  - entity, 281-282
  - session, 52-53
- Electronic Business using XML (ebXML), 120-121**
- encryption, 134**
- Endpoint Redirection design pattern, 459**
- endpoints, composition member, 344**
- enterprise archive (EAR), 44, 239**
- Enterprise Inventory design pattern, 460**
- Enterprise JavaBeans (EJBs), 43, 51-52**
  - entity, 281
  - session, 52-53
- Enterprise Service Bus (ESB), 95-96, 461**
  - service message processing with, 392-396
  - service message processing without, 389-391
- Entity Abstraction design pattern, 462**
- Entity Linking design pattern, 463**
- entity services**
  - architecture, 263-264
  - as Web-based services, 282-283
  - defined, 31

- design, 270-278
- implementation, 278-282
- REST, 291-299
- SOAP-based Web, 283-285
- testing, 302-303

**EntityManager class, 54**

**error handling, 345**

**ESB (Enterprise Service Bus), 95-96, 461**

- service message processing with, 392-396
- service message processing without, 389-391

**Event-Driven Messaging design pattern, 464**

**eventing, defined, 240**

**Exception Shielding design pattern, 465**

**explicit headers, 154-157**

**Extensible Markup Language (XML), 59-61**

- schema, 61-62

**Extensible Stylesheet Language Transformations (XSLT), 63**

## **F**

**failure handling, 375-376**

**Federated Endpoint Layer compound pattern, 466**

**figures, 7**

**File Gateway design pattern, 467**

**functional boundaries, 194**

**functional context dependencies, 140-141, 179-180**

**Functional Decomposition design pattern, 468**

## **G**

**generic entity services, 273-275**

**generic service contracts, 144-145**

**generic service logic, 141-144**

**GET method, HTTP, 107**

**GlassFish Enterprise Server, 74-75**

**granularity, defined, 34-35**

## **H**

**handlers, 118-119**

**HEAD method, HTTP, 107**

**Hibernate, 231, 302**

## **HTTP**

- ACCEPT header, 109-110

- characteristics, 67

- CONNECT method, 107

- DELETE method, 107

- GET method, 107

- HEAD method, 107

- methods, 106-108

- OPTIONS method, 107

POST method, 107  
 PUT method, 107  
 response codes, 102-103  
 TRACE method, 107

**I**

**IBM WebSphere Application Server,**  
 75-77  
**IBM WebSphere Application Server**  
 Community Edition, 78  
**Idempotent Capability design pattern,**  
 469  
**implicit headers,** 157-159  
**in-memory data grids,** 371-372  
**Intermediate Routing design pattern,**  
 392, 470  
**Inventory Endpoint design pattern,** 471  
**IT resources, defined,** 30

**J**

**JAAS (Java Authentication and**  
**Authorization Service),** 58-59  
**JAF (JavaBeans Activation Framework),**  
 165  
**JAR (Java archive),** 239  
**Java**  
 mapping with WSDL, 152-153  
 platforms, 40-41  
 WS-\* standards implementation,  
 122-124  
**Java API for RESTful Web Services. See**  
**JAX-RS standard**  
**Java API for XML Binding (JAXB),** 64  
**Java API for XML Processing (JAXP),**  
 63-64  
**Java API for XML Registries (JAXR),**  
 120-121  
**Java API for XML Web Services. See**  
**JAX-WS standard**  
**Java API for XML-based RPC (JAX-**  
**RPC) standard,** 112, 152-153  
**Java Archive (JAR),** 239  
**Java Authentication and Authorization**  
**Service (JAAS),** 58-59  
**Java Database Connectivity (JDBC)**  
 API, 49  
 standard, 278-279  
**Java Development Kit (JDK),** 41  
**Java EE Connector Architecture (JCA),**  
 51, 231  
 in Spring framework, 231  
**Java EE (Java Platform, Enterprise**  
**Edition)**  
 API, 49-59  
*contexts and dependency*  
*injection,* 49-50

- EJBs (Enterprise JavaBeans)*, 51-52
- JCA (Java EE Connector Architecture)*, 51
- JMS (Java Message Service)*, 56-58
- JPA (Java Persistence API)*, 53-54
- JTA (Java Transactions API)*, 50-51
- application, 44-45
- application server
  - IBM WebSphere*, 75-77
- architecture, 41-44
  - tiers*, 45-46
- defined, 40-41
- Java ME (Java Platform, Micro Edition)**, defined, 40
- Java Message Service (JMS)**, 56-58
  - characteristics, 67
- Java Naming and Directory Interface (JNDI)**, 48
- Java Persistence API (JPA)**, 53-54
- Java runtime environment (JRE)**, 41
- Java SE (Java Platform, Standard Edition)**
  - API, 46-49
    - JDBC (Java Database Connectivity)*, 49
    - JNDI (Java Naming and Directory Interface)*, 48
    - RMI (Remote Method Invocation)*, 46-47
    - RMI over IIOP*, 47-48
  - architecture, 41
  - defined, 40
  - utility services, 226
- Java Transaction Processing**, 51
- Java Transactions API (JTA)**, 50-51
- Java Virtual Machine (JVM)**, 41
- JavaBeans Activation Framework (JAF)**, 165
- JavaScript Object Notation (JSON)**, 108
  - conversion with POJOs, 161-162
- JAXB (Java API for XML Binding)**, 64
- JAXP (Java API for XML Processing)**, 63-64
- JAXR (Java API for XML Registries)**, 120-121
- JAX-RPC (Java API for XML-based RPC) standard**, 112, 152-153
- JAX-RS (Java API for RESTful Web Services) standard**, 124-136
  - annotations, 125-129
  - goals, 125
- JAX-WS (Java API for XML Web Services) standard**, 112-121, 153
  - handlers, 118-119

JCA (Java EE Connector Architecture), 51, 231

JDBC (Java Database Connectivity), 49  
API, 49  
standard, 278-279

JDK (Java Development Kit), 41

JMS (Java Message Service), 56-58  
characteristics, 67

JMX, 230

JNDI (Java Naming and Directory Interface), 48

JPA (Java Persistence API), 53-54

JRE (Java runtime environment), 41

JSON (JavaScript Object Notation), 108  
conversion with POJOs, 161-162

JTA (Java Transactions API), 50-51

JVM (Java Virtual Machine), 41

**L**

layered SOA model, 27

Legacy Wrapper design pattern, 472

Lightweight Endpoint design pattern, 473

link relation, 106

literal, defined, 153

literal SOAP messaging, 87

load testing, 370

locking, 280-281  
optimistic, 281  
pessimistic, 281

Log4J, 231

Logic Centralization design pattern, 474

logical handlers, 118

loose coupling. *See* Service Loose  
Coupling design principle

**M**

Management of Web Services (MOWS)  
specification, 95

Management Using Web Services  
(MUWS) specification, 95

MDBs (message-driven beans), 58

MEPs (message exchange patterns), 337-338

message consumers, 385-386

message entities  
designing, 271-272  
domain entities versus, 265-266

message exchange patterns (MEPs), 337-338

message producers, 385-386

Message Screening design pattern, 475

message-driven beans (MDBs), 58

messages  
in service contracts, 69-70  
parsing and marshaling, 376-378  
processing  
*with* ESBs, 392-396  
*without* ESBs, 389-391

SOAP, 86-87

**messaging**

- in resource utility services, 248
- RPC (remote procedure calls)
  - versus, 381-386
- traditional frameworks, 380-386

**Messaging Metadata design pattern, 476****Metadata Centralization design pattern, 477****methods, HTTP, 106-108****micro-utility services, 253-256**

- defined, 220

**middleware**

- ESB (Enterprise Service Bus), 95-96
- orchestration, 97-98

**MOWS (Management of Web Services)**

- specification, 95

**MTOM, 122, 166****Multi-Channel Endpoint design pattern,**

- 478

**MUWS (Management Using Web Services) specification, 95****N****NFR (non-functional requirements),**

- defined, 213

**Non-Agnostic Context design pattern,**

- 479

**non-agnostic logic, defined, 31****non-functional requirements (NFR), defined, 213****NovoBank. See case study examples, NovoBank****O****object-relational mapping, 230****Official Endpoint compound pattern, 480****omni utility services, 240-247**

- defined, 220

**open-source frameworks**

- REST entity service, 302
- utility services and, 229-231

**optimistic approach (data access), 268****optimistic locking, 280-281****OPTIONS method, HTTP, 107****Oracle WebLogic Server, 79-80****orchestration, 97-98**

- coding versus, 340-341

**Orchestration compound pattern, 481****orchestration infrastructure, 198****P****packaging**

- composed services, 362-368
- entity services, 303

- task services, 334
- utility services, 239
- parsing messages, 376-378
- Partial State Deferral design pattern, 482
- Partial Validation design pattern, 483
- performance
  - measuring, 368-369
  - testing, 370
  - of task services, 315-316
- pessimistic locking, 281
- Plain Old Java Objects (POJOs), 43
  - conversion with JSON, 161-162
- platforms, 40-41
- policies, defined, 94
- Policy Centralization design pattern, 484
- POST method, HTTP, 107
- Prentice Hall Service Technology Series*
  - from Thomas Erl, 7-8
- Process Abstraction design pattern, 485
- Process Centralization design pattern, 486
- Protocol Bridging design pattern, 392, 487
- protocol handlers, 118
- Proxy Capability design pattern, 488
- publish/subscribe model, 56-57
- PUT method, HTTP, 107

## R

- recovery from failure, 375-376
- Redundant Implementation design pattern, 489
- registries, 99-100
- Reliable Messaging design pattern, 490
- Remote Method Invocation (RMI), 46-47
- remote procedure calls (RPC), 87
  - messaging versus, 381-386
- repositories, 99-100
- resource representations, 108-109
- resource utility services, 247-252
  - defined, 220
- resource utilization, 369
- response codes (HTTP), 102-103
- response time, 368-369
- REST, 101-110
  - caching, 372-374
  - data mapping with, 159-175
  - entity service, 291-299
  - error handling, 345
  - HTTP methods, 106-108
  - HTTP response codes, 102-103
  - resource location, 103-106
- Reusable Contract design pattern, 491
- RMI (Remote Method Invocation), 46-47



**RMI over IIOP, 47-48**

- characteristics, 67

**routing messages**

- with ESBs, 392

- without ESBs, 390-391

**RPC (remote procedure calls), 87**

- messaging versus, 381-386

**Rules Centralization design pattern, 492****runtime discoverability, 205-207****runtime environments**

- control, 195

- efficiency, 190-192

- standards-based, 193

**S****SAAJ (SOAP with Attachments API for Java), 115-118****SAX (Simple API for XML), 63****SCA (Service Component Architecture), 123-124****scalability**

- of REST services, 130-132

- of stateful services, 374-375

**Schema Centralization design pattern, 493****schemas (XML), 61-62**

- type reuse, 353-354

**SDO (Service Data Objects)**

- specification, 123-124

**security**

- in Java EE, 58-59

- JAX-RS standard, 132-134

**SEI (service endpoint interface), 113****services**

- defined, 29

- Java components as, 64-73

- local, invocation, 360

- message processing

- with ESBs, 392-396*

- without ESBs, 389-391*

- REST, 101-110

- SOAP-based Web, 82-101

- middleware, 95-101*

**Service Abstraction design principle, 25, 184-185, 336**

- profile, 414

**Service Agent design pattern, 494****Service Autonomy design principle, 25, 194**

- profile, 417

**Service Broker compound pattern, 392, 495****Service Callback design pattern, 496****service candidates, defined, 33****service catalogs, 32-33****Service Component Architecture (SCA), 123-124**

- Service Composability design principle,**
  - 25, 189-190, 336
  - profile, 422-423
- service composition**
  - with ESBs, 401-402
  - initiators, 337
  - members versus controllers, 190, 336
  - and MEPs (message exchange patterns), 337-338
  - sub-controllers, 337
- service consumer coupling, 180-183**
- service contracts, 68-71**
  - defined, 33-34
  - flexibility, 192
  - generic, 144-145
  - separation from implementation, 177-179
  - standardized, 71
- Service Data Objects (SDO)**
  - specification, 123-124
- Service Data Replication design pattern,**
  - 497
- Service Decomposition design pattern,**
  - 498
- Service Discoverability design principle,**
  - 25, 204
  - profile, 420-421
- Service Encapsulation design pattern,**
  - 499
- service endpoint interface (SEI), 113**
- service endpoint selection with ESBs,**
  - 397
- Service Façade design pattern, 500**
- service granularity, defined, 34**
- Service Grid design pattern, 501**
- Service Instance Routing design pattern,**
  - 502
- service interface, 222**
- service inventory, defined, 32**
- service inventory blueprints, 32**
- Service Layers design pattern, 503**
- service level agreements (SLAs), 339-340**
- Service Loose Coupling design principle,**
  - 25, 176-177
  - profile, 413
- Service Messaging design pattern, 504**
- service models, defined, 30-31**
- Service Normalization design pattern,**
  - 505
- Service Perimeter Guard design pattern,**
  - 506
- service portfolio, defined, 32-33**
- service portfolio management, defined,**
  - 33
- service processing with ESBs, 399-400**
- service profiles, defined, 35-36**
- Service Refactoring design pattern, 507**
- service registries, 99-100, 208**

- service repositories, 99-100
- Service Reusability design principle, 25
  - profile, 415-416
- Service Statelessness design principle, 25, 197
  - profile, 418-419
- service-orientation
  - defined, 24-26, 28
  - EJBs (Enterprise JavaBeans) and, 55-56
- Service-Oriented Architecture: Concepts, Technology, and Design*, 4, 38
- service-oriented architecture (SOA), defined, 27-28
- service-oriented computing, defined, 22-24, 28
- session EJBs (Enterprise JavaBeans), 52-53
- session state, 198-199
  - storing, 199-204
- Simple API for XML (SAX), 63
- Simple Object Access Protocol. *See* SOAP
- SLAs (service level agreements), 339-340
- SmartCredit Co. *See* case study
  - examples, SmartCredit Co.
- SOA (service-oriented architecture), defined, 27-28
- SOA Certified Professional (SOACP), 9
  - SOA Design Patterns*, 4, 38, 426
  - SOA Governance: Governing Shared Services On-Premise and in the Cloud*, 38
  - SOA Manifesto, 28, 520-532
  - SOA Principles of Service Design*, 3, 38, 140, 410
  - SOA with REST: Principles, Patterns & Constraints for Building Enterprise Solutions with REST*, 4, 82, 101, 343, 426
  - SOACP (SOA Certified Professional), 9
  - soak testing, 370
  - SOAP (Simple Object Access Protocol)
    - defined, 82
    - explicit headers, 154-157
    - implicit headers, 157-159
    - no headers, 159
  - SOAP with Attachments API for Java (SAAJ), 115-118
  - SOAP with Attachments (SwA), 90-91
  - SOAP-based Web services, 82-101
    - binary data in, 165-166
    - entity services as, 283-285
    - error handling, 345
    - Java components versus, 359-360
    - packaging, 363
    - task services as, 320
    - utility services as, 232-236

- spatial coupling, 382
- Spring framework, 229-231
- Spring MVC, 231
- Spring-WS framework, 124
- Standardized Service Contract design
  - principle, 25, 34, 151
  - profile, 411-412
- standards-based runtime environment, 193
- State Messaging design pattern, 508
- State Repository design pattern, 374, 509
- stateful services, scalability, 374-375
- Stateful Services design pattern, 375, 510
- stateless entity services, 272
- statelessness, 131
- storing state, 199-204
- stress testing, 370
- supertypes, 143
- SwA (SOAP with Attachments), 90-91
- symbols
  - for components, 29
  - legend, 7
- synchronous invocation, 338-339
- T**
- task services
  - as REST services, 324
  - as SOAP-based Web services, 320
  - defined, 31
  - design, 308-309
  - implementation, 316-320
  - performance, 315-316
  - testing, 332-334
- taxonomy, utility service, 220-221
- technology coupling, 382
- technology detail abstraction, 185
- temporal coupling, 382-383
- Termination Notification design pattern, 511
- testing
  - composed services, 362-368
  - entity services, 302-303
  - task services, 332-334
  - utility services, 238
- Three-Layer Inventory compound pattern, 512
- throughput, 369
- toolkits, 119-120
- TRACE method, HTTP, 107
- transaction
  - management, 229
  - recovery mechanism, 375-376
  - resources, 249
  - support, 135
- transactionality, 122
- transformation
  - with ESBs, 392-396
  - without ESBs, 391
- Trusted Subsystem design pattern, 513

**U**

**UDDI (Universal Description, Discovery, and Integration), 83**  
**UI Mediator design pattern, 514**  
**Uniform Contract compound pattern, 515**  
**uniform contracts, 131**  
**Uniform Resource Identifiers (URIs), 61**  
**Uniform Resource Locations (URLs), 103-106**  
**Universal Description, Discovery, and Integration (UDDI), 83**  
**URIs (Uniform Resource Identifiers) for namespaces, 61**  
**URLs (Uniform Resource Locations), REST services, 103-106**  
**UTF-8 encoding, 60**  
**Utility Abstraction design pattern, 516**  
**utility services**  
     and open-source frameworks, 229-231  
     architecture, 212-220  
     as a Web-based service defined, 31  
     design, 221-225  
     in Java EE, 227-229  
     in Java SE, 226  
     JAX-WS, 234-236  
     packaging, 239

REST, 236-238  
 taxonomy, 220-221  
 testing, 238  
 types of, 240-260

**V**

**Validation Abstraction design pattern, 517**  
**validation constraint granularity, 34**  
**Version Identification design pattern, 518**

**W**

**WAR (Web archive), 239**  
**WAS (WebSphere Application Server), 75-77**  
**WAS-CE (WebSphere Application Server Community Edition), 78**  
**Web archive (WAR), 239**  
**Web container (Java EE architecture), 42-43**  
**Web Profile (Java EE), 43**  
*Web Service Contract Design & Versioning for SOA, 4, 82*  
**Web services. See SOAP-based Web services**  
**Web Services Description Language (WSDL), 83-88**

**Web Services Distributed Management (WSDM) specification, 95**

**Web services standards. *See* WS-\* standards**

**Web sites**

- [www.arcitura.com/notation](http://www.arcitura.com/notation), 7
- [www.bigdatascienceschool.com](http://www.bigdatascienceschool.com), 9
- [www.cloudpatterns.org](http://www.cloudpatterns.org), 8
- [www.cloudschool.com](http://www.cloudschool.com), 9
- [www.serviceorientation.com](http://www.serviceorientation.com), 3, 8, 38, 140, 190, 410
- [www.servicetechbooks.com](http://www.servicetechbooks.com), 4, 8, 9, 38
- [www.servicetechmag.com](http://www.servicetechmag.com), 8
- [www.servicetechspecs.com](http://www.servicetechspecs.com), 3, 8, 100, 110
- [www.soaglossary.com](http://www.soaglossary.com), 3, 38
- [www.soa-manifesto.com](http://www.soa-manifesto.com), 3, 28, 38, 520
- [www.soa-manifesto.org](http://www.soa-manifesto.org), 28, 520
- [www.soapatterns.org](http://www.soapatterns.org), 3, 8, 426
- [www.soaprinciples.com](http://www.soaprinciples.com), 410
- [www.soaschool.com](http://www.soaschool.com), 9
- [www.whatiscloud.com](http://www.whatiscloud.com), 8
- [www.whatisrest.com](http://www.whatisrest.com), 8, 101

**WebLogic Server, 79-80**

**WebSphere Application Server Community Edition (WAS-CE), 78**

**WebSphere Application Server (WAS), 75-77**

**workload characterization, 370**

**wrapped document/literal contracts, 153-154**

**wrapper utility services, 256-260**  
defined, 221

**WS-\* standards**

- Java implementation, 122-124
- SOAP-based Web services and, 88-95

**WS-Addressing standard, 89-90, 122**

**WS-AtomicTransaction standard, 93, 122**

**WS-BPEL standard, 97-98, 341**

**WS-BusinessActivity standard, 93**

**WS-Coordination standard, 93, 122**

**WSDL (Web Services Description Language), 83-88**

- mapping with Java, 152-153

**WSDM (Web Services Distributed Management) specification, 95**

**WS-Policy standard, 94-95**

**WS-ReliableMessaging standard, 91-92, 122, 339**

**WS-Security standard, 93-94, 122, 134**

**WS-Transaction standard, 92-93**

**X****XML (Extensible Markup Language),****59-61**

schema, 61-62

*type reuse*, 353-354

sending in utility services, 232-234

sending with JAX-WS standard,

234-236

**XSLT (Extensible Stylesheet Language  
Transformations), 63**