Service-Oriented Architecture

Concepts, Technology, and Design

Thomas Erl



PRENTICE HALL PROFESSIONAL TECHNICAL REFERENCE UPPER SADDLE RIVER, NJ • BOSTON • INDIANAPOLIS • SAN FRANCISCO NEW YORK • TORONTO • MONTREAL • LONDON • MUNICH • PARIS • MADRID CAPETOWN • 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.

The publisher offers excellent discounts on this book when ordered in quantity for bulk purchases or special sales, which may include electronic versions and/or custom covers and content particular to your business, training goals, marketing focus, and branding interests. For more information, please contact:

U. S. Corporate and Government Sales (800) 382-3419 corpsales@pearsontechgroup.com

For sales outside the U.S., please contact:

International Sales international@pearsoned.com

Visit us on the Web: www.phptr.com

Library of Congress Number: 2005925019

Copyright © 2005 Pearson Education, Inc. Portions of this work are copyright SOA Systems Inc., and reprinted with permission from SOA Systems Inc. © 2005. Front cover and all photographs by Thomas Erl. Permission to use photographs granted by SOA Systems Inc.

All rights reserved. Printed in the United States of America. This publication is protected by copyright, and permission must be obtained from the copyright holder 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. For information regarding permissions, write to:

Pearson Education, Inc. Rights and Contracts Department One Lake Street Upper Saddle River, NJ 07458

ISBN 0-13-185858-0 Text printed in the United States on recycled paper at R.R. Donnelley in Crawfordsville, Indiana. First printing, July 2005

Contents

xxvii

Preface

Introduction 1		
1.1 Why	y this book is important	. 2
1.1.1	The false SOA	2
1.1.2	The ideal SOA	3
1.1.3	The real SOA	4
1.2 Obj	ectives of this book	. 4
1.2.1	Understanding SOA, service-orientation, and Web services	. 5
1.2.2	Learning how to build SOA with Web services	. 5
1.3 Who	o this book is for	. 6
1.4 Wha	at this book does not cover	. 6
1.5 Hov	v this book is organized	. 7
1.5.1	Part I: SOA and Web Services Fundamentals	8
1.5.2	Part II: SOA and WS-* Extensions	. 10
1.5.3	Part III: SOA and Service-Orientation	13
1.5.4	Part IV: Building SOA (Planning and Analysis)	. 14
1.5.5	Part V: Building SOA (Technology and Design)	16
1.5.6	Conventions	. 19
1.6 Add	litional information	19
1.6.1	The XML & Web Services Integration Framework (XWIF)	. 19
1.6.2	www.serviceoriented.ws	20
1.6.3	Contact the Author	. 20

For more information visit www.serviceoriented.ws.

Х

Contents

29

Chapter 2 **Case Studies** 212.1 2.1.12.1.2 2.1.3 2.2.1 2.2.2 2.2.3 2.2.4 2.3.1 2.3.2 2.3.3 2.3.4

Part I

SOA and Web Services Fundamentals

00

Chapter 3 . .

. .

Introduci	ng SUA 31
3.1 Fun	damental SOA
3.1.1	A service-oriented analogy
3.1.2	How services encapsulate logic
3.1.3	How services relate
3.1.4	How services communicate
3.1.5	How services are designed
3.1.6	How services are built
3.1.7	Primitive SOA
3.2 Cor	nmon characteristics of contemporary SOA 40
3.2.1	Contemporary SOA is at the core of the service-oriented
	computing platform41
3.2.2	Contemporary SOA increases quality of service

Contents

xi

3.	2.3	Contemporary SOA is fundamentally autonomous $\ldots \ldots 42$
3.	2.4	Contemporary SOA is based on open standards $\ldots \ldots 43$
3.	2.5	Contemporary SOA supports vendor diversity $\ldots \ldots 43$
3.	2.6	Contemporary SOA promotes discovery $\ldots \ldots \ldots 44$
3.	2.7	Contemporary SOA fosters intrinsic interoperability $\ldots \ldots 45$
3.	2.8	Contemporary SOA promotes federation $\ldots \ldots 45$
3.	2.9	Contemporary SOA promotes architectural composability $\ldots \ldots 46$
3.2	.10	Contemporary SOA fosters inherent reusability $\ldots \ldots \ldots 47$
3.2	.11	Contemporary SOA emphasizes extensibility $\ldots \ldots 48$
3.2	.12	Contemporary SOA supports a service-oriented
		business modeling paradigm
3.2	.13	Contemporary SOA implements layers of abstraction
3.2	.14	Contemporary SOA promotes loose coupling throughout the
		enterprise
3.2	.15	Contemporary SOA promotes organizational agility
3.2	.16	
3.2	.17	
3.2	.18	Contemporary SOA is still maturing
3.2	.19	
3.2	.20	Defining SOA
3.2	2.21	Separating concrete characteristics
3.3	Com	mon misperceptions about SOA 56
3.	3.1	"An application that uses Web services is service-oriented." 56
3.	3.2	"SOA is just a marketing term used to re-brand Web services." 57
3.	3.3	"SOA is just a marketing term used to re-brand distributed
•	0.4	Computing with web services
3.	3.4	"SOA simplifies distributed computing."
3.	3.5	An application with web services that uses ws-
3	36	"If you understand Web services you won't have a
0.	0.0	problem building SOA."
3.	3.7	"Once you go SOA, everything becomes interoperable."
34	Com	mon tangible benefits of SOA 59
3	4 1	Improved integration (and intrinsic interoperability) 60
3	42	Inherent reuse 60
3	4.3	Streamlined architectures and solutions
3	4.4	Leveraging the legacy investment 61
3.	4.5	Establishing standardized XML data representation
5.	-	J

xii

Contents

3.4.6	Focused investment on communications infrastructure $\ldots \ldots .63$
3.4.7	"Best-of-breed" alternatives
3.4.8	Organizational agility63
3.5 Com	mon pitfalls of adopting SOA64
3.5.1	Building service-oriented architectures like traditional
	distributed architectures $\dots \dots \dots$
3.5.2	Not standardizing SOA65
3.5.3	Not creating a transition plan
3.5.4	Not starting with an XML foundation architecture $\ldots \ldots \ldots 67$
3.5.5	Not understanding SOA performance requirements
3.5.6	Not understanding Web services security
3.5.7	Not keeping in touch with product platforms and standards
	development

The Evo	lution of SOA	71
4.1 An	SOA timeline (from XML to Web services to SOA)	72
4.1.1	XML: a brief history	72
4.1.2	Web services: a brief history	73
4.1.3	SOA: a brief history	74
4.1.4	How SOA is re-shaping XML and Web services	76
4.2 The	e continuing evolution of SOA (standards organizations	
and	d contributing vendors)	78
4.2.1	"Standards" vs. "Specifications" vs. "Extensions"	78
4.2.2	Standards organizations that contribute to SOA	79
4.2.3	Major vendors that contribute to SOA	82
4.3 The	e roots of SOA (comparing SOA to past architectures)	86
4.3.1	What is architecture?	86
4.3.2	SOA vs. client-server architecture	88
4.3.3	SOA vs. distributed Internet architecture	95
4.3.4	SOA vs. hybrid Web service architecture	104
4.3.5	Service-orientation and object-orientation (Part I)	107

Contents

xiii

Web) Serv	vices and Primitive SOA	109
5.1	The	Web services framework 1	111
5.2	Ser	vices (as Web services)	112
5	5.2.1	Service roles	114
5	5.2.2	Service models	126
5.3	Ser	vice descriptions (with WSDL)	131
5	5.3.1	Service endpoints and service descriptions	133
5	5.3.2	Abstract description	134
5	5.3.3	Concrete description	135
5	5.3.4	Metadata and service contracts	136
5	5.3.5	Semantic descriptions	137
5	5.3.6	Service description advertisement and discovery	138
5.4	Mes	ssaging (with SOAP)	142
5	5.4.1	Messages	143
5	5.4.2	Nodes	149
5	5.4.3	Message paths	152

Part II

SOA and WS-* Extensions	155
What is "WS-*"?	. 157

Web Services and Contemporary SOA			
(Part I: Activity Management and Composition) 15			
6.1 Mes	sage exchange patterns	. 162	
6.1.1	Primitive MEPs	. 163	
6.1.2	MEPs and SOAP	. 169	
6.1.3	MEPs and WSDL	. 169	
6.1.4	MEPs and SOA	. 171	
6.2 Serv	vice activity	. 172	
6.2.1	Primitive and complex service activities	. 174	
6.2.2	Service activities and SOA	. 175	

xiv

Contents

6.3 Co	ordination
6.3.1	Coordinator composition
6.3.2	Coordination types and coordination protocols
6.3.3	Coordination contexts and coordination participants
6.3.5	The activation and registration process
6.3.5	The completion process
6.3.6	Coordination and SOA 183
6.4 Ato	mic transactions
6.4.1	ACID transactions
6.4.2	Atomic transaction protocols
6.4.3	The atomic transaction coordinator
6.4.4	The atomic transaction process
6.4.5	Atomic transactions and SOA
6.5 Bus	siness activities
6.5.1	Business activity protocols
6.5.2	The business activity coordinator
6.5.3	Business activity states
6.5.4	Business activities and atomic transactions
6.5.5	Business activities and SOA
6.6 Orc	hestration
6.6.1	Business protocols and process definition
6.6.2	Process services and partner services
6.6.3	Basic activities and structured activities
6.6.4	Sequences, flows, and links
6.6.5	Orchestrations and activities
6.6.6	Orchestration and coordination
6.6.7	Orchestration and SOA
6.7 Cho	preography
6.7.1	Collaboration
6.7.2	Roles and participants
6.7.3	Relationships and channels
6.7.4	Interactions and work units
6.7.5	Reusability, composability, and modularity
6.7.6	Orchestrations and choreographies
6.7.7	Choreography and SOA

Contents

XV

Chapter 7		
Web Ser	vices and Contemporary SOA	
(Dart II: A	Advanced Messaging, Metadata, and Security)	017
(Fart II. <i>F</i>	avanceu messaying, melauala, and Security)	217
7.1 Add	dressing	220
7.1.1	Endpoint references	222
7.1.2	Message information headers	223
7.1.3	Addressing and transport protocol independence	225
7.1.4	Addressing and SOA	225
7.2 Reli	iable messaging	228
7.2.1	RM Source, RM Destination, Application Source,	
	and Application Destination	230
7.2.2	Sequences	230
7.2.3	Acknowledgements	231
7.2.4		233
7.2.5	Reliable messaging and addressing	235
7.2.6	Reliable messaging and SOA	235
7.3 Cor	relation	238
7.3.1	Correlation in abstract	239
7.3.2	Correlation in MEPs and activities	239
7.3.3	Correlation in coordination	240
7.3.4	Correlation in orchestration	240
7.3.5	Correlation in addressing	240
7.3.6	Correlation in reliable messaging	240
7.3.7	Correlation and SOA	241
7.4 Poli	cies	242
7.4.1	The WS-Policy framework	243
7.4.2	Policy assertions and policy alternatives	244
7.4.3	Policy assertion types and policy vocabularies	245
7.4.4	Policy subjects and policy scopes	245
7.4.5	Policy expressions and policy attachments	245
7.4.6	What you really need to know	245
7.4.7	Policies in coordination	246
7.4.8	Policies in orchestration and choreography	246
7.4.9	Policies in reliable messaging	246
7.4.10	Policies and SOA	246

xvi

Contents

277

7.5	Meta	adata exchange
7	.5.1	The WS-MetadataExchange specification
7	.5.2	Get Metadata request and response messages $\ldots \ldots 250$
7	.5.3	Get request and response messages
7	.5.4	Selective retrieval of metadata
7	.5.5	Metadata exchange and service description discovery $\ldots \ldots 252$
7	.5.6	Metadata exchange and version control
7	.5.7	Metadata exchange and SOA $\ldots \ldots \ldots 254$
7.6	Secu	urity
7	.6.1	Identification, authentication, and authorization
7	.6.2	Single sign-on
7	.6.3	Confidentiality and integrity
7	.6.4	Transport-level security and message-level security
7	.6.5	Encryption and digital signatures
7	.6.6	Security and SOA
7.7	Noti	fication and eventing
7	.7.1	Publish-and-subscribe in abstract
7	.7.2	One concept, two specifications
7	.7.3	The WS-Notification Framework
7	.7.4	The WS-Eventing specification
7	.7.5	WS-Notification and WS-Eventing
7	.7.6	Notification, eventing, and SOA

Part III

SOA and Service-Orientation

Princi	ples of Service-Orientation 279
8.1 \$	Service-orientation and the enterprise
8.2	Anatomy of a service-oriented architecture
8.2	2.1 Logical components of the Web services framework
8.2	2.2 Logical components of automation logic
8.2	2.3 Components of an SOA 288
8.2	P.4 How components in an SOA inter-relate

Contents

V\/I	
A V	

8.3	Com	mon principles of service-orientation
8	3.3.1	Services are reusable
8	3.3.2	Services share a formal contract
8	3.3.3	Services are loosely coupled
8	3.3.4	Services abstract underlying logic
8	3.3.5	Services are composable
8	3.3.6	Services are autonomous
8	3.3.7	Services are stateless
8	3.3.8	Services are discoverable
8.4	How	service-orientation principles inter-relate
8	3.4.1	Service reusability
8	3.4.2	Service contract
8	3.4.3	Service loose coupling
8	3.4.4	Service abstraction
8	3.4.5	Service composability
8	3.4.6	Service autonomy
8	3.4.7	Service statelessness
8	3.4.8	Service discoverability
8.5	Serv	rice-orientation and object-orientation (Part II)
8.6	Nativ	ve Web service support for service-orientation principles 324

Ser	vice L	_ayers	327
9.1	Ser	vice-orientation and contemporary SOA	328
ę	9.1.1	Mapping the origins and supporting sources of concrete SOA characteristics	329
ę	9.1.2	Unsupported SOA characteristics	332
9.2	Ser	vice layer abstraction	333
ę	9.2.1	Problems solved by layering services	334
9.3	Арр	Dication service layer	337
9.4	Bus	siness service layer	341
9.5	Orc	hestration service layer	344
9.6	Agr	nostic services	346
9.7	Ser	vice layer configuration scenarios	347
ę	9.7.1	Scenario #1: Hybrid application services only	348
ę	9.7.2	Scenario #2: Hybrid and utility application services	349

xviii

Contents

355

9.7.3	Scenario #3: Task-centric business services and utility application services
9.7.4	Scenario #4: Task-centric business services, entity-centric business services, and utility application services
9.7.5	Scenario #5: Process services, hybrid application services, and utility application services
9.7.6	Scenario #6: Process services, task-centric business services, and utility application services
9.7.7	Scenario #7: Process services, task-centric business services, entity-centric business services, and utility application services
9.7.8	Scenario #8: Process services, entity-centric business services, and utility application services

Part IV

Building SOA (Planning and Analysis)

Chapter 10

- - - -

SOA Deli	ivery Strategies	357
10.1 SO/	A delivery lifecycle phases	. 358
10.1.1	Basic phases of the SOA delivery lifecycle	. 358
10.1.2	Service-oriented analysis	. 359
10.1.3	Service-oriented design	. 359
10.1.4	Service development	. 360
10.1.5	Service testing	. 360
10.1.6	Service deployment	. 361
10.1.7	Service administration	. 361
10.1.8	SOA delivery strategies	. 362
10.2 The	top-down strategy	. 363
10.2.1	Process	. 363
10.2.2	Pros and cons	. 365
10.3 The	bottom-up strategy	. 366
10.3.1	Process	. 367
10.3.2	Pros and cons	. 368
10.4 The	agile strategy	. 370
10.4.1	Process	. 370
10.4.2	Pros and cons	. 373

Contents

xix

Chapter 11		
Service-(Driented Analysis (Part I: Introduction) 37	5
11.1 Intro	oduction to service-oriented analysis	7
11.1.1	Objectives of service-oriented analysis	7
11.1.2	The service-oriented analysis process	7
11.2 Ber	efits of a business-centric SOA	2
11.2.1	Business services build agility into business models	3
11.2.2	Business services prepare a process for orchestration	4
11.2.3	Business services enable reuse	4
11.2.4	Only business services can realize the	
	service-oriented enterprise	5
11.3 Der	iving business services	6
11.3.1	Sources from which business services can be derived	7
11.3.2	Types of derived business services	2
11.3.3	Business services and orchestration	5
Chapter 12		
Service-0	Driented Analysis (Part II: Service Modeling) 39	7
12.1 Ser	vice modeling (a step-by-step process)	8
12.1.1	"Services" versus "Service Candidates"	8
12.1.2	Process description	9
12.2 Ser	vice modeling guidelines	6
12.2.1	Take into account potential cross-process reusability	
	of logic being encapsulated (task-centric business	
	service candidates) 41	6
12.2.2	Consider potential intra-process reusability of logic being encapsulated	

	Service candidates/
12.2.2	Consider potential intra-process reusability of logic being encapsulated
	(task-centric business service candidates)
12.2.3	Factor in process-related dependencies (task-centric
	business service candidates)
12.2.4	Model for cross-application reuse (application
	service candidates)
12.2.5	Speculate on further decomposition requirements $\ldots \ldots 418$
12.2.6	Identify logical units of work with explicit boundaries
12.2.7	Prevent logic boundary creep
12.2.8	Emulate process services when not using orchestration (task-centric business service candidates)

ХХ

Contents

445

12.2.9	Target a balanced model
12.2.10	Classify service modeling logic $\dots \dots \dots$
12.2.11	Allocate appropriate modeling resources $\ldots \ldots \ldots 422$
12.2.12	Create and publish business service modeling standards $\hdots422$
12.3 Clas	ssifying service model logic
12.3.1	The SOE model
12.3.2	The enterprise business model $\ldots \ldots \ldots 426$
12.3.3	"Building Blocks" versus "Service Models" $\dots \dots \dots \dots 426$
12.3.4	Basic modeling building blocks
12.4 Con	trasting service modeling approaches (an example) 430

Part V

Building SOA (Technology and Design)	
--------------------------------------	--

Service-(Oriented Design (Part I: Introduction)	447
13.1 Intro	oduction to service-oriented design	448
13.1.1	Objectives of service-oriented design	. 448
13.1.2	"Design standards" versus "Industry standards"	. 449
13.1.3	The service-oriented design process	449
13.1.4	Prerequisites	451
13.2 WS	DL-related XML Schema language basics	453
13.2.1	The schema element	454
13.2.2	The element element	455
13.2.3	The complexType and simpleType elements	455
13.2.4	The import and include elements	456
13.2.5	Other important elements	456
13.3 WS	DL language basics	457
13.3.1	The definitions element	458
13.3.2	The types element	459
13.3.3	The message and part elements	461
13.3.4	The portType, interface, and operation elements	462
13.3.5	The input and output elements (when used	
	with operation)	462

Contents

13.3.6	The binding element
13.3.7	The input and output elements (when used with binding) . 464
13.3.8	The service, port, and endpoint elements
13.3.9	The import element
13.3.10	The documentation element466
13.4 SOA	AP language basics466
13.4.1	The Envelope element
13.4.2	The Header element
13.4.3	The Body element
13.4.4	The Fault element
13.5 Ser	vice interface design tools
13.5.1	Auto-generation
13.5.2	Design tools
13.5.3	Hand coding

Service-(Oriented Design (Part II: SOA Composition Guidelines)	475
14.1 Ste	ps to composing SOA	476
14.1.1	Step 1: Choose service layers	478
14.1.2	Step 2: Position core standards	478
14.1.3	Step 3: Choose SOA extensions	478
14.2 Cor	nsiderations for choosing service layers	478
14.3 Cor	nsiderations for positioning core SOA standards	481
14.3.1	Industry standards and SOA	481
14.3.2	XML and SOA	482
14.3.3	The WS-I Basic Profile	483
14.3.4	WSDL and SOA	485
14.3.5	XML Schema and SOA	485
14.3.6	SOAP and SOA	486
14.3.7	Namespaces and SOA	487
14.3.8	UDDI and SOA	488
14.4 Cor	nsiderations for choosing SOA extensions	490
14.4.1	Choosing SOA characteristics	490
14.4.2	Choosing WS-* specifications	491
14.4.3	WS-BPEL and SOA	492

xxii

Contents

Chapter	15		
Serv	vice-C	Driented Design (Part III: Service Design)	495
15.1	Serv	vice design overview	497
15	.1.1	Design standards	498
15	.1.2	About the process descriptions	498
15	.1.3	Prerequisites	499
15.2	Enti	ty-centric business service design (a step-by-step	
	proc	cess)	501
15	.2.1	Process description	502
15.3	Арр	lication service design (a step-by-step process)	522
15	.3.1	Process description	523
15.4	Task	k-centric business service design (a step-by-step	
	proc	cess)	540
15	.4.1	Process description	540
15.5	Serv	vice design guidelines	555
15	.5.1	Apply naming standards	555
15	.5.2	Apply a suitable level of interface granularity	556
15	.5.3	Design service operations to be inherently extensible	558
15	.5.4	Identify known and potential service requestors	559
15	.5.5	Consider using modular WSDL documents	559
15	.5.6	Use namespaces carefully	560
15	.5.7	Use the SOAP document and literal attribute values	561
15	.5.8	Use WS-I Profiles even if WS-I compliance isn't required	563
15	.5.9	Document services with metadata	563

Service-C	Driented Design (Part IV: Business Process Design)	565
16.1 WS-	-BPEL language basics	566
16.1.1	A brief history of BPEL4WS and WS-BPEL	567
16.1.2	Prerequisites	568
16.1.3	The process element	568
16.1.4	The partnerLinks and partnerLink elements	569
16.1.5	The partnerLinkType element	570
16.1.6	The variables element	571
16.1.7	The getVariableProperty and getVariableData functions	572

Contents

xxiii

613

16.1.8	The sequence element
16.1.9	The invoke element
16.1.10	The receive element
16.1.11	The <code>reply</code> element
16.1.12	The switch, case, and otherwise elements
16.1.13	The assign, copy, from, and to elements $\ldots\ldots\ldots577$
16.1.14	faultHandlers, catch, and catchAll elements $\ldots\ldots 578$
16.1.15	Other WS-BPEL elements
16.2 WS-	Coordination overview
16.2.1	The CoordinationContext element $\ldots \ldots \ldots 582$
16.2.2	The Identifier and Expires elements
16.2.3	The CoordinationType element
16.2.4	The RegistrationService element
16.2.5	Designating the WS-BusinessActivity coordination type $\ldots \ldots 584$
16.2.6	Designating the WS-AtomicTransaction coordination type $\ \ldots \ 584$
16.3 Serv	vice-oriented business process design (a step-by-step
16.3.1	Process description

Chapter 17

Fundamental WS-* Extensions

You mustU	nderstand this	614
17.1 WS	-Addressing language basics	615
17.1.1	The EndpointReference element	616
17.1.2	Message information header elements	617
17.1.3	WS-Addressing reusability	620
17.2 WS	-ReliableMessaging language basics	622
17.2.1	The Sequence, MessageNumber, and	
	LastMessage elements	623
17.2.2	The SequenceAcknowledgement and	
	AcknowledgementRange elements	625
17.2.3	The Nack element	626
17.2.4	The AckRequested element	627
17.2.5	Other WS-ReliableMessaging elements	628
17.3 WS	-Policy language basics	629
17.3.1	The Policy element and common policy assertions	630
17.3.2	The ExactlyOne element	631

xxiv

Contents

17.3.3	The All element $\ldots \ldots \ldots 632$
17.3.4	The Usage attribute
17.3.5	The Preference attribute
17.3.6	The PolicyReference element
17.3.7	The PolicyURIs attribute $\ldots \ldots 634$
17.3.8	The PolicyAttachment element
17.3.9	Additional types of policy assertions $\ldots \ldots 35$
17.4 WS-	MetadataExchange language basics
17.4.1	The GetMetadata element637
17.4.2	The Dialect element638
17.4.3	The Identifier element639
17.4.4	The Metadata, MetadataSection, and
	MetadataReference elements640
17.4.5	The Get message
17.5 WS-	Security language basics
17.5.1	The $\texttt{Security}$ element (WS-Security) $\ldots \ldots \ldots 644$
17.5.2	The UsernameToken, Username, and Password elements (WS-Security) 644
17.5.3	The BinarySecurityToken element (WS-Security)
17.5.4	The SecurityTokenReference element (WS-Security) 644
17.5.5	Composing Security element contents (WS-Security) 645
17.5.6	The EncryptedData element (XML-Encryption)646
17.5.7	The CipherData, CipherValue, and CipherReference elements (XML-Encryption)
17.5.8	XML-Signature elements

SOA P	atforms 651
18.1 S	OA platform basics652
18.1.1	Basic platform building blocks653
18.1.2	2 Common SOA platform layers654
18.1.3	8 Relationship between SOA layers and technologies
18.1.4	Fundamental service technology architecture
18.1.5	5 Vendor platforms
18.2 S	OA support in J2EE
18.2.1	Platform overview668
18.2.2	2 Primitive SOA support681

Contents	,
----------	---

XXV

18.2.3	Support for service-orientation principles
18.2.4	Contemporary SOA support
18.3 SO/	A support in .NET
18.3.1	Platform overview
18.3.2	Primitive SOA support
18.3.3	Support for service-orientation principles
18.3.4	Contemporary SOA support
18.4 Inte	gration considerations
Appendix A	

Case Studies	s: Conclusion	707
A.1 RailCo I	Ltd	708
A.2 Transit L	Line Systems Inc.	711
A.3 The Oas	isis Car Wash	715

Appendix B

Service Models Reference	717
About the Author	721
About SOA Systems	723
About the Photographs	725
Index	727

Part III



SOA and Service-Orientation

Chapter 8 Principles of Service-Orientation Chapter 9 Service Layers

So far the focus has been on the aspects of SOA as a whole. We have discussed the numerous extensions available to SOA as well as its fundamental concepts. We now turn our attention to the underlying paradigm primarily responsible for defining SOA and distinguishing it as an architectural model.

The principles and concepts covered in the next two chapters discuss the spectrum of service-orientation in detail. This establishes necessary theory that applies to the rudimentary components of primitive SOA, but also provides concepts that can be propagated and leveraged throughout service-oriented environments. For example, topics covered in these chapters form the basis for the service modeling and design processes provided in the subsequent *Building SOA* parts of this book.

Chapter 8



Principles of Service-Orientation

- 8.1 Service-orientation and the enterprise
- 8.2 Anatomy of a service-oriented architecture
- 8.3 Common principles of service-orientation
- 8.4 How service-orientation principles inter-relate
- 8.5 Service-orientation and object-orientation
- 8.6 Native Web service support for service-orientation principles

B efore we can begin building a service-oriented solution, we need to understand what makes a service suitable for SOA. In other words, how can we build Web services that are truly service-oriented?

The answer lies in service-orientation. This approach to modeling business automation logic has resulted in a set of commonly accepted principles applied to each unit of logic that constitutes a service within an SOA. It is through the application of these principles that the primitive components of an SOA (services, descriptions, messages) are shaped in support of service-orientation.

This chapter begins with a look at how service-orientation applies to the enterprise as a whole and then discusses individual principles in-depth.

In Plain English sections

A knowledge of the principles of service-orientation is perhaps even more important than concepts covered in past chapters. They are core to the design of services regardless of what underlying technology is used to implement them. Therefore, our *In Plain English* sections return to supplement the descriptions of individual principles.

How case studies are used: As you might recall from the case study background information provided in Chapter 2, one of RailCo's business goals was to improve their existing automation processes by moving toward SOA.

In this chapter we examine the services built so far as part of RailCo's technical environment and discuss how they comply to or diverge from individual principles of service-orientation. Existing TLS services that already possess serviceorientation characteristics are used for comparison purposes.

8.1 Service-orientation and the enterprise

The collective logic that defines and drives an enterprise is an ever-evolving entity constantly changing in response to external and internal influences. From an IT perspective, this *enterprise logic* can be divided into two important halves: business logic and application logic (Figure 8.1).

Service-orientation and the enterprise



Figure 8.1

The business and application logic domains.

Each exists in a world of its own, and each represents a necessary part of contemporary organization structure. Business logic is a documented implementation of the business requirements that originate from an enterprise's business areas. Business logic is generally structured into processes that express these requirements, along with any associated constraints, dependencies, and outside influences.

Application logic is an automated implementation of business logic organized into various technology solutions. Application logic expresses business process workflows through purchased or custom-developed systems within the confines of an organization's IT infrastructure, security constraints, technical capabilities, and vendor dependencies.

Service-orientation applies to enterprise logic. It introduces new concepts that augment the manner in which this logic is represented, viewed, modeled, and shared. While the principles behind service-orientation exist in a vacuous realm of abstraction and theory, they are a necessary counterpart to the real world environments that require their guidance and structure.

The concepts introduced by service-orientation are realized through the introduction of services. Let's have a look at where services are located within the overall structure of an automated organization. As Figure 8.2 illustrates, services establish a high form of abstraction wedged between traditional business and application layers. When positioned here, services can encapsulate physical application logic as well as business process logic.

Chapter 8: Principles of Service-Orientation

Services modularize the enterprise, forming standalone units of logic that exist within a common connectivity layer. Services can be layered so that parent services can encapsulate child services. This allows for the service layer to consist of multiple layers of abstraction (as explained later in Chapter 9).



Figure 8.2

The service interface layer positioned between enterprise layers that promote application and business logic.

In Figure 8.2 we display a fragmented application layer, where individual applications are confined to the boundaries that represent their respective proprietary platform environments. Though we show services as existing in a single, continuous layer, this only illustrates the open connectivity that exists among service interfaces. Freed from proprietary ties, services can communicate via open protocols.

On a physical level, services are developed and deployed in proprietary environments, wherein they are individually responsible for the encapsulation of specific application

Service-orientation and the enterprise

logic. Figure 8.3 shows how individual services, represented as service interfaces within the service interface layer, represent application logic originating from different platforms.



Figure 8.3

The service interface layer abstracts connectivity from service deployment environments.

SUMMARY OF KEY POINTS

- Enterprise logic can be divided into two domains: business logic and application logic. Service-oriented principles can be applied to both.
- The service interface layer positions services to represent business logic and abstract application logic.

284

Sample Chapter 8 from "Service-Oriented Architecture: Concepts, Technology, and Design" by Thomas Erl For more information visit www.serviceoriented.ws.

Chapter 8: Principles of Service-Orientation

8.2 Anatomy of a service-oriented architecture

Chapter 5 established the components of the basic (first-generation) Web services framework. This framework can be applied to implement services in just about any environment. For example, services can be appended to traditional distributed applications or used as wrappers to expose legacy system logic. However, neither of these environments resembles a "real" service-oriented architecture.

To best understand what constitutes a true SOA, we need to abstract the key components of the Web services framework and study their relationships more closely. To accomplish this, we begin by revisiting these familiar components and altering our perspective of them. First, we re-label them to reflect terminology more associated with service-orientation. Then we position them into a logical view wherein we subsequently re-examine our components within the context of SOA.

8.2.1 Logical components of the Web services framework

The communications framework established by Web services brings us the foundation technology for what we've classified as contemporary SOA. Because we covered this framework in Chapter 5, we will use it as a reference point for our discussion of service-orientation.

Let's first recap some Web services fundamentals within a logical modeling context. As shown in Figure 8.4, each Web service contains one or more operations. Note that this diagram introduces a new symbol to represent operations separately from the service.



Figure 8.4 A Web service sporting two operations.

Anatomy of a service-oriented architecture

Each operation governs the processing of a specific function the Web service is capable of performing. The processing consists of sending and receiving SOAP messages, as shown in Figure 8.5.



Figure 8.5 An operation processing outgoing and incoming SOAP messages.

By composing these parts, Web services form an activity through which they can collectively automate a task (Figure 8.6).



Figure 8.6 A basic communications scenario between Web services.

8.2.2 Logical components of automation logic

The Web services framework provides us not only with a technology base for enabling connectivity, it also establishes a modularized perspective of how automation logic, as a whole, can be comprised of independent units. To illustrate the inherent modularity of Web services, let's abstract the following fundamental parts of the framework:

- SOAP messages
- Web service operations
- Web services
- activities

286

Chapter 8: Principles of Service-Orientation

The latter three items represent units of logic that perform work and communicate using SOAP messages. To better illustrate this in a service-oriented perspective, let's replace these terms with new ones, as follows:

- messages
- operations
- services
- processes (and process instances)

You'll notice that these are quite similar to the terms we used before. The one exception is the use of "process" instead of "activity." In later chapters we actually use the word "activity" in different contexts when modeling service-oriented business processes.

For now, the one discrepancy to be aware of is that while a Web service activity is typically used to represent the temporary interaction of a group of Web services, a process is a static definition of interaction logic. An activity is best compared to an instance of a process wherein a group of services follow a particular path through the process logic to complete a task.

Regardless, for the purposes of our discussion of service-orientation, we'll continue with our look at how automation logic is comprised of the four identified parts. We can further qualify these parts by relating each to different sized units of logic, as follows:

- messages = units of communication
- operations = units of work
- services = units of processing logic (collections of units of work)
- processes = units of automation logic (coordinated aggregation of units of work)

Figure 8.7 provides us with a primitive view of how operations and services represent units of logic that can be assembled to comprise a unit of automation logic.

Next, in Figure 8.8, we establish that messages are a suitable means by which all units of processing logic (services) communicate. This illustrates that regardless of the scope of logic a service represents, no actual processing of that logic can be performed without issuing units of communication (in this case, messages).

Anatomy of a service-oriented architecture



Figure 8.7

A primitive view of how SOA modularizes automation logic into units.



Figure 8.8

A primitive view of how units of communication enable interaction between units of logic.

The purpose of these views is simply to express that processes, services, and operations, on the most fundamental level, provide a flexible means of partitioning and modularizing logic. Regardless of the technology platform used, this remains the most basic concept that underlies service-orientation. In being able to derive this view from the Web services framework, we also have demonstrated the suitability of the Web services platform as a means of implementation for SOA.

288

Chapter 8: Principles of Service-Orientation

8.2.3 Components of an SOA

We'll continue to work with our components of automation logic, but we now broaden our discussion to how the characteristics and behaviors of these components are formed within service-oriented architecture.

Each of the previously defined components establishes a level of enterprise logic abstraction, as follows:

- A message represents the data required to complete some or all parts of a unit of work.
- An operation represents the logic required to process messages in order to complete a unit of work (Figure 8.9).



Figure 8.9

The scope of an operation within a process.

- A service represents a logically grouped set of operations capable of performing related units of work.
- A process contains the business rules that determine which service operations are used to complete a unit of automation. In other words, a process represents a large piece of work that requires the completion of smaller units of work (Figure 8.10).



Figure 8.10

Operations belonging to different services representing various parts of process logic.

Anatomy of a service-oriented architecture

8.2.4 How components in an SOA inter-relate

Having established the core characteristics of our SOA components, let's now look at how these components are required to relate to each other:

- An operation sends and receives messages to perform work.
- An operation is therefore mostly defined by the messages it processes.
- A service groups a collection of related operations.
- A service is therefore mostly defined by the operations that comprise it.
- A process instance can compose services.
- A process instance is not necessarily defined by its services because it may only require a subset of the functionality offered by the services.
- A process instance invokes a unique series of operations to complete its automation.
- Every process instance is therefore partially defined by the service operations it uses.

Figures 8.11 and 8.12 further illustrate these relationships.

A service-oriented architecture is an environment standardized according to the principles of service-orientation in which a process that uses services (a service-oriented process) can execute. Next, we'll take a closer look at what exactly the principles of service-orientation consist of.





290

Chapter 8: Principles of Service-Orientation



Figure 8.12

How the components of a service-oriented architecture define each other.

SUMMARY OF KEY POINTS

- The logical parts of an SOA can be mapped to corresponding components in the basic Web services framework.
- By viewing a service-oriented solution as a unit of automation logic, we establish that SOA consists of a sophisticated environment that supports a highly modularized separation of logic into differently scoped units.
- SOA further establishes specific characteristics, behaviors, and relationships among these components that provide a predictable environment in support of serviceorientation.

8.3 Common principles of service-orientation

In Chapter 3 we established that there is no single definition of SOA. There is also no single governing standards body that defines the principles behind service-orientation. Instead, there are many opinions, originating from public IT organizations to vendors and consulting firms, about what constitutes service-orientation.

Service-orientation is said to have its roots in a software engineering theory known as "separation of concerns." This theory is based on the notion that it is beneficial to break down a large problem into a series of individual concerns. This allows the logic required to solve the problem to be decomposed into a collection of smaller, related pieces. Each piece of logic addresses a specific concern.

Common principles of service-orientation

This theory has been implemented in different ways with different development platforms. Object-oriented programming and component-based programming approaches, for example, achieve a separation of concerns through the use of objects, classes, and components.

Service-orientation can be viewed as a distinct manner in which to realize a separation of concerns. The principles of service-orientation provide a means of supporting this theory while achieving a foundation paradigm upon which many contemporary SOA characteristics can be built. In fact, if you study these characteristics again, you will notice that several are (directly or indirectly) linked to the separation of concerns theory.

As previously mentioned, there is no official set of service-orientation principles. There are, however, a common set of principles most associated with service-orientation. These are listed below and described further in this section.

- Services are reusable—Regardless of whether immediate reuse opportunities exist, services are designed to support potential reuse.
- Services share a formal contract—For services to interact, they need not share anything but a formal contract that describes each service and defines the terms of information exchange.
- Services are loosely coupled—Services must be designed to interact without the need for tight, cross-service dependencies.
- Services abstract underlying logic—The only part of a service that is visible to the outside world is what is exposed via the service contract. Underlying logic, beyond what is expressed in the descriptions that comprise the contract, is invisible and irrelevant to service requestors.
- Services are composable—Services may compose other services. This allows logic to be represented at different levels of granularity and promotes reusability and the creation of abstraction layers.
- Services are autonomous—The logic governed by a service resides within an explicit boundary. The service has control within this boundary and is not dependent on other services for it to execute its governance.
- Services are stateless—Services should not be required to manage state information, as that can impede their ability to remain loosely coupled. Services should be designed to maximize statelessness even if that means deferring state management elsewhere.

Chapter 8: Principles of Service-Orientation

• Services are discoverable—Services should allow their descriptions to be discovered and understood by humans and service requestors that may be able to make use of their logic.

Of these eight, autonomy, loose coupling, abstraction, and the need for a formal contract can be considered the core principles that form the baseline foundation for SOA. As explained in the *How service-orientation principles inter-relate* section later in this chapter, these four principles directly support the realization of other principles (as well as each other).

There are other qualities commonly associated with services and service-orientation. Examples include self-descriptive and coarse-grained interface design characteristics. We classify these more as service design guidelines, and they are therefore discussed as part of the design guidelines provided in Chapter 15.

NOTE

You may have noticed that the reusability and autonomy principles also were mentioned as part of the contemporary SOA characteristics described in Chapter 3. This overlap is intentional, as we simply are identifying qualities commonly associated with SOA as a whole as well as services designed for use in SOA. We further clarify the relationship between contemporary SOA characteristics and service-orientation principles in Chapter 9.

To fully understand how service-orientation principles shape service-oriented architecture, we need to explore the implications their application will have on all of the primary parts that comprise SOA. Let's take a closer look at each of the principles.

8.3.1 Services are reusable

Service-orientation encourages reuse in all services, regardless if immediate requirements for reuse exist. By applying design standards that make each service potentially reusable, the chances of being able to accommodate future requirements with less development effort are increased. Inherently reusable services also reduce the need for creating wrapper services that expose a generic interface over top of less reusable services.

This principle facilitates all forms of reuse, including inter-application interoperability, composition, and the creation of cross-cutting or utility services. As we established earlier in this chapter, a service is simply a collection of related operations. It is therefore the logic encapsulated by the individual operations that must be deemed reusable to warrant representation as a reusable service (Figure 8.13).

Common principles of service-orientation





Messaging also indirectly supports service reusability through the use of SOAP headers. These allow for messages to become increasingly self-reliant by grouping metadata details with message content into a single package (the SOAP envelope). Messages can be equipped with processing instructions and business rules that allow them to dictate to recipient services how they should be processed.

The processing-specific logic embedded in a message alleviates the need for a service to contain this logic. More importantly, it imposes a requirement that service operations become less activity-specific—in other words, more generic. The more generic a service's operations are, the more reusable the service.

CASE STUDY

RailCo delivered the Invoice Submission Service for the sole purpose of being able to connect to TLS's new B2B system. This Web service's primary function therefore is to send electronic invoice documents to the TLS Accounts Payable Service. The service contains the following two operations: SubmitInvoice and GetTLSMetadata (Figure 8.14).

Chapter 8: Principles of Service-Orientation

The SubmitInvoice operation simply initiates the transmission of the invoice document. You might recall in the *Metadata exchange* section of Chapter 7 that an operation was added to periodically check the TLS Accounts Payable Service for changes to its service description. This new operation is GetTLSMetadata.





Because they were built to meet immediate and specific business requirements, these operations have no real reuse potential. The SubmitInvoice operation is designed to forward SOAP messages containing specific headers required by TLS and containing an invoice XML document structured according to a schema also defined by TLS. By its very name, the GetTLSMetadata operation identifies itself as existing for one reason: to query a specific endpoint for new metadata information.

The TLS Accounts Payable Service, on the other hand, provides a series of generic operations related to the processing of accounts payable transactions. This service is therefore used by different TLS systems, one of which is the aforementioned B2B solution.

In Chapters 11 and 12 we will submit the RailCo Invoice Submission Service to a modeling exercise in an attempt to reshape it into a service that implements actual service-orientation principles, including reusability.

Common principles of service-orientation

295

IN PLAIN ENGLISH

One day, a government inspector stops by our car washing operation. Not knowing who he is, I ask if he would like his car washed and waxed or just washed. He responds by asking a question of his own. "Do you have a business license for this operation?"

A subsequent conversation between the inspector and our team results in the revelation that we have indeed been operating without a business license. We are therefore ordered to cease all work until we obtain one. We scramble to find out what needs to be done. This leads us to visit the local Business License Office to start the process of acquiring a license.

The Business License Office provides a distinct service: issuing and renewing business licenses. It is not there to service just our car washing company; it is there to provide this service to anyone requesting it. Because its service is designed to facilitate multiple service requestors, the logic that enables the service can be classified as being reusable.

8.3.2 Services share a formal contract

Service contracts provide a formal definition of:

- the service endpoint
- each service operation
- every input and output message supported by each operation
- rules and characteristics of the service and its operations

Service contracts therefore define almost all of the primary parts of an SOA (Figure 8.15). Good service contracts also may provide semantic information that explains how a service may go about accomplishing a particular task. Either way, this information establishes the agreement made by a service provider and its service requestors.

Because this contract is shared among services, its design is extremely important. Service requestors that agree to this contract can become dependent on its definition. Therefore, contracts need to be carefully maintained and versioned after their initial release.

Chapter 8: Principles of Service-Orientation



Figure 8.15 Service contracts formally define the service, operation, and message components of a service-oriented architecture.

As explained in Chapter 5, service description documents, such as the WSDL definition, XSD schemas, and policies, can be viewed collectively as a communications contract that expresses exactly how a service can be programmatically accessed.

CASE STUDY

From the onset, RailCo and TLS agreed to each other's service contracts, which enabled these two companies to interact via the TLS B2B system. The rules of the contract and the definition of associated service description documents all are provided by TLS to ensure a standardized level of conformance that applies to each of its online vendors.

One day, RailCo is informed that TLS has revised the policy published with the Accounts Payable Service. A new rule has been added where TLS is offering better payment terms to vendors in exchange for larger discounts. RailCo has the choice to continue pricing their products at the regular amounts and face a payment term of 60 days for their invoices or reduce their prices to get a payment term of 30 days.

Both of these options are acceptable contract conditions published by TLS. After some evaluation, RailCo decides not to take advantage of the reduced payment terms and therefore does not adjust its product prices.

Common principles of service-orientation

297

IN PLAIN ENGLISH

For us to get a business license, we must fill out an application form. This process essentially formalizes our request in a format required and expected by the Business License Office.

The completed application form is much like a contract between the service provider and the requestor of the service. Upon accepting the form, the service provider agrees to act on the request.

8.3.3 Services are loosely coupled

No one can predict how an IT environment will evolve. How automation solutions grow, integrate, or are replaced over time can never be accurately planned out because the requirements that drive these changes are almost always external to the IT environment. Being able to ultimately respond to unforeseen changes in an efficient manner is a key goal of applying service-orientation. Realizing this form of agility is directly supported by establishing a loosely coupled relationship between services (Figure 8.16).

Loose coupling is a condition wherein a service acquires knowledge of another service while still remaining independent of that service. Loose coupling is achieved through the use of service contracts that allow services to interact within predefined parameters.

It is interesting to note that within a loosely coupled architecture, service contracts actually tightly couple operations to services. When a service is formally described as being the location of an operation, other services will depend on that operation-to-service association.



Figure 8.16

Services limit dependencies to the service contract, allowing underlying provider and requestor logic to remain loosely coupled.

298

Chapter 8: Principles of Service-Orientation

CASE STUDY

Through the use of service contracts, RailCo and TLS services are naturally loosely coupled. However, one could say that the extent of loose coupling between the two service provider entities is significantly different.

TLS services are designed to facilitate multiple B2B partners, as well as internal reuse and composition requirements. This makes TLS services *very* loosely coupled from any of its service requestors.

RailCo's services, on the other hand, are designed specifically to interact with designated TLS services that are part of the overall TLS B2B solution. No attempt was made to make these services useful for any other service requestors. RailCo services are therefore considered *less* loosely coupled than TLS services.

IN PLAIN ENGLISH

After we have submitted our form, we are not required to remain at the Business License Office, nor do we need to stay in touch with them. We only need to wait until the application is processed and a license is (hopefully) issued.

This is much like an asynchronous message exchange, but it is also a demonstration of a loosely coupled relationship between services or between service provider and requestor. All we need to interact with the Business License Office is an application form that defines the information the office requires to process our request. Prior to and subsequent to the submission of that request, our car washing team (service requestor) and the Business License Office (service provider) remain independent of each other.

8.3.4 Services abstract underlying logic

Also referred to as *service interface-level abstraction*, it is this principle that allows services to act as black boxes, hiding their details from the outside world. The scope of logic represented by a service significantly influences the design of its operations and its position within a process.

There is no limit to the amount of logic a service can represent. A service may be designed to perform a simple task, or it may be positioned as a gateway to an entire automation solution. There is also no restriction as to the source of application logic a service can draw upon. For example, a single service can, technically, expose application logic from two different systems (Figure 8.17).







Operation granularity is therefore a primary design consideration that is directly related to the range and nature of functionality being exposed by the service. Again, it is the individual operations that collectively abstract the underlying logic. Services simply act as containers for these operations.

Service interface-level abstraction is one of the inherent qualities provided by Web services. The loosely coupled communications structure requires that the only piece of knowledge services need to interact is each others' service descriptions.

CASE STUDY

Because both RailCo and TLS employ Web services to communicate, each environment successfully implements service interface-level abstraction. On RailCo's end, this abstraction hides the legacy systems involved with generating electronic invoice documents and processing incoming purchase orders. On the TLS side, services hide service compositions wherein processing duties are delegated to specialized services as part of single activities (Figure 8.18).

300

Chapter 8: Principles of Service-Orientation



Neither of RailCo's or TLS's service requestors require any knowledge of what lies behind the other's service providers.

IN PLAIN ENGLISH

The tasks required for the Business License Office to process our request include:

- A name check to ensure that the name of our company "Oasis Car Wash" isn't already taken.
- A background check of the company principals to ensure that none of us have had past bankruptcies.
- A verification of our sub-lease agreement to ensure that we are, in fact, allowed to operate at the gas station we have been using.

Common principles of service-orientation

These and other tasks are performed completely unbeknownst to us. We don't know or necessarily care what the Business License Office needs to do to process our application. We are just interested in the expected outcome: the issuance of our license.

8.3.5 Services are composable

A service can represent any range of logic from any types of sources, including other services. The main reason to implement this principle is to ensure that services are designed so that they can participate as effective members of other service compositions if ever required. This requirement is irrespective of whether the service itself composes others to accomplish its work (Figure 8.19).



The UpdateEverything operation encapsulating a service composition.

A common SOA extension that underlines composability is the concept of orchestration. Here, a service-oriented process (which essentially can be classified as a service composition) is controlled by a parent process service that composes process participants.

Chapter 8: Principles of Service-Orientation

The requirement for any service to be composable also places an emphasis on the design of service operations. Composability is simply another form of reuse, and therefore operations need to be designed in a standardized manner and with an appropriate level of granularity to maximize composition opportunities.

CASE STUDY

As with RailCo's Invoice Submission Service, its Order Fulfillment Service was created to meet a specific requirement in support of communication with TLS's B2B solution.

The Order Fulfillment Service contains just one public operation called ProcessTLSPO (Figure 8.20). This operation is designed in compliance with TLS vendor service specifications so that it is fully capable of receiving POs submitted by the TLS Purchase Order Service. Part of this compliance requires the operation to be able to process custom SOAP headers containing proprietary security tokens.



Though the Order Fulfillment Service is capable of acting as a composition member, its potential for being useful to any future compositions is limited. Composition support is similar to reusability in that generic functionality exposed by operations make a service more composable. This RailCo service provides one operation that performs a very specialized function, customized to processing a specific document from a specific source. It will likely not be a suitable composition member, but it can act as a controller service, composing other services to complete its PO processing tasks.

The TLS Accounts Payable Service already establishes a well-defined composition, wherein it acts as a controller service that composes the Vendor Profile and Ledger Services (Figure 8.21). Because they each expose a complete set of generic

Common principles of service-orientation



IN PLAIN ENGLISH

Given that the services provided by the Business License Office are distinct and reusable, it can be asked to assist other government offices to participate in the completion of other services. For example, the Business Relocation Office manages all administrative paperwork for businesses that need to be moved when their location is scheduled for demolition.

As part of its many tasks, this office takes care of revising the business license information for the affected company. It does so by enlisting the Business License Office and requesting that they issue a new business license for a particular organization.

By reusing the services offered by the Business License Office, the Business Relocation Office has effectively composed services, much like a controller service reuses and composes other service providers.

8.3.6 Services are autonomous

Autonomy requires that the range of logic exposed by a service exist within an explicit boundary. This allows the service to execute self-governance of all its processing. It also eliminates dependencies on other services, which frees a service from ties that could

Chapter 8: Principles of Service-Orientation

inhibit its deployment and evolution (Figure 8.22). Service autonomy is a primary consideration when deciding how application logic should be divided up into services and which operations should be grouped together within a service context.



Figure 8.22

Autonomous services have control over underlying resources.

Deferring the location of business rules is one way to strengthen autonomy and keep services more generic. Processes generally assume this role by owning the business rules that determine how the process is structured and, subsequently, how services are composed to automate the process logic. This is another aspect of orchestration explored in the *Orchestration service layer* section in Chapter 9.

Note that autonomy does not necessarily grant a service exclusive ownership of the logic it encapsulates. It only guarantees that at the time of execution, the service has control over whatever logic it represents. We therefore can make a distinction between two types of autonomy.

Common principles of service-orientation

- *Service-level autonomy*—Service boundaries are distinct from each other, but the service may share underlying resources. For example, a wrapper service that encapsulates a legacy environment that also is used independently from the service has service-level autonomy. It governs the legacy system but also shares resources with other legacy clients.
- *Pure autonomy*—The underlying logic is under complete control and ownership of the service. This is typically the case when the underlying logic is built from the ground up in support of the service.

CASE STUDY

Given the distinct tasks they perform, the following three RailCo services all are autonomous:

- Invoice Submission Service
- Order Fulfillment Service
- TLS Subscription Service

Each represents a specific boundary of application logic that does not overlap with the boundary of any other services.

Autonomy in RailCo's services was achieved inadvertently. No conscious effort was made to avoid application overlap, as the services were delivered to simply meet specific connectivity requirements.

As shown in Figure 8.23, the Invoice Processing and Order Fulfillment Services encapsulate legacy logic. The legacy accounting system also is used by clients independently from the services, which makes this service-level autonomy. The TLS Notification Service achieves pure autonomy, as it represents a set of custom components created only in support of this service.

In environments where a larger number of services exist and new services are built on a regular basis, it is more common to introduce dedicated modeling processes so pure service autonomy is preserved among individual services. At TLS, for example, services undergo a service-oriented analysis to guarantee autonomy and avoid encapsulation overlap. (Service-oriented analysis is explained in Chapters 11 and 12.)

Chapter 8: Principles of Service-Orientation



RailCo's services luckily encapsulate explicit portions of legacy and newly added application logic.

Common principles of service-orientation

307

IN PLAIN ENGLISH

Let's revisit the three tasks performed by the Business License Office when processing an application for a new business license:

- name check
- background check
- location verification

The Business License Office owns the corporate name database required to perform a name check. Also the office has personnel dedicated to visiting and verifying business site locations. When completing these two tasks, the Business License Office therefore has complete self-governance. However, when having to perform a background check, the office must share a database system with the Revenue Office. When it gets access, it can retrieve an abbreviated credit history for each of the company principals listed on the application.

The Business License Office's reliance on the shared database reduces its independence somewhat. However, its overall ability to perform the tasks within its own boundary give it a degree of autonomy.

8.3.7 Services are stateless

Services should minimize the amount of state information they manage and the duration for which they hold it. State information is data-specific to a current activity. While a service is processing a message, for example, it is temporarily stateful (Figure 8.24). If a service is responsible for retaining state for longer periods of time, its ability to remain available to other requestors will be impeded.

Statelessness is a preferred condition for services and one that promotes reusability and scalability. For a service to retain as little state as possible, its individual operations need to be designed with stateless processing considerations.

A primary quality of SOA that supports statelessness is the use of document-style messages. The more intelligence added to a message, the more independent and self-sufficient it remains. Chapters 6 and 7 explore various WS-* extensions that rely on the use of SOAP headers to carry different types of state data.

308

Chapter 8: Principles of Service-Orientation



Figure 8.24

Stateless and stateful stages a service passes through while processing a message.

CASE STUDY

As with loose coupling, statelessness is a quality that can be measured in degrees. The RailCo Order Fulfillment Service is required to perform extra runtime parsing and processing of various standard SOAP header blocks to successfully receive a purchase order document submitted by the TLS Purchase Order Service. This processing ties up the Order Fulfillment Service longer than, say, the Invoice Submission Service, which simply forwards a predefined SOAP message to the TLS Accounting Service.

Common principles of service-orientation

309

IN PLAIN ENGLISH

During the initial review of the application, our company was briefly discussed by personnel at the Business License Office. But after the application was fully processed, no one really retained any memory of our request.

Though the details of our application have been logged and recorded in various repositories, there is no further need for anyone involved in the processing of our request to remember further information about it once the application processing task was completed. To this extent, the Business License Office simulates a degree of statelessness. It processes many requests every day, and there is no benefit to retaining information about a completed request.

8.3.8 Services are discoverable

Discovery helps avoid the accidental creation of redundant services or services that implement redundant logic. Because each operation provides a potentially reusable piece of processing logic, metadata attached to a service needs to sufficiently describe not only the service's overall purpose, but also the functionality offered by its operations.

Note that this service-orientation principle is related to but distinct from the contemporary SOA characteristic of discoverability. On an SOA level, discoverability refers to the architecture's ability to provide a discovery mechanism, such as a service registry or directory. This effectively becomes part of the IT infrastructure and can support numerous implementations of SOA. On a service level, the principle of discoverability refers to the design of an individual service so that it can be as discoverable as possible.

CASE STUDY

RailCo provides no means of discovery for its services, either internally or to the outside world. Though outfitted with its own WSDL definition and fully capable of acting as a service provider, the Invoice Submission Service is primarily utilized as a service requestor and currently expects no communication outside of the TLS Accounts Payable Service

Similarly, the RailCo Order Fulfillment Service was registered manually with the TLS B2B solution so that it would be placed on the list of vendors that receive purchase orders. This service provides no reusable functionality and is therefore considered to have no immediate requirement for discovery.

Chapter 8: Principles of Service-Orientation

Due to the reusable nature of TLS services and because of the volume of services that are expected to exist in TLS technical environments, an internal service registry was established (as shown in Figure 8.25 and originally explained in Chapter 5). This piece of TLS infrastructure promotes discoverability and prevents accidental redundancy. It further leverages the existing design standards used by TLS that promote the creation of descriptive metadata documents in support of service discoverability.



Figure 8.25

RailCo's services are not discoverable, but TLS's inventory of services are stored in an internal registry.

TLS is not interested in making its services publicly discoverable, which is why it does not register them with a public service registry. Vendors that participate in the TLS B2B system only are allowed to do so after a separate negotiation, review, and registration process.

How service-orientation principles inter-relate

311

IN PLAIN ENGLISH

After some time, our business license is finally issued. Upon receiving the certificate in the mail, we are back in business. Looking back at how this whole process began, though, there is one step we did not discuss. When we first learned that we were required to get a business license, we had to find out where the Business License Office was located. This required us to search through the phone book and locate a listing with contact information.

A service registry provides a discovery mechanism very much like a phone book, allowing potential requestors to query and check candidate service providers. In the same manner in which a registry points to service descriptions, the phone book listing led us to the location at which we were able to obtain the original business license application form.

More relevant to the principle of service discoverability is the fact that steps were taken to make the Business License Office itself discoverable. Examples include signs in the lobby of the high-rise in which the office is located, a sign on the office entrance door, brochures located at other offices, and so on.

SUMMARY OF KEY POINTS

- Different organizations have published their own versions of service-oriented principles. As a result, many variations exist.
- The most common principles relate to loose coupling, autonomy, discoverability, composability, reuse, service contracts, abstraction, and statelessness.

8.4 How service-orientation principles inter-relate

When reading through the previous descriptions, a number of questions might come to mind, such as:

- What's the difference between reusability and composability? (Aren't you reusing a service when you compose it?)
- What's the difference between autonomy and statelessness? (Aren't both a representation of service independence?)
- What's the difference between loose coupling and the use of a service contract? (Doesn't a service contract automatically implement loose coupling?)

312

Chapter 8: Principles of Service-Orientation

To answer these and other questions, this section revisits our service-orientation principles to explore how each relates to, supports, or is affected by others. To accomplish this, we abbreviate the original names we assigned each principle, as follows:

- Services are reusable = service reusability
- Services share a formal contract = service contract
- Services are loosely coupled = service loose coupling
- Services abstract underlying logic = service abstraction
- Services are composable = service composability
- Services are autonomous = service autonomy
- Services are stateless = service statelessness
- Services are discoverable = service discoverability

We intentionally prefix each principle with the word "service" to emphasize that the principle applies to the design of a service only, as opposed to our SOA characteristics, which apply to the design of SOA as a whole.

NOTE

Each relationship is essentially described twice within these sections. This repetitiveness is intentional, as this part of the chapter is provided more for reference purposes. Feel free to skip ahead if you are not interested in learning about each individual principle-to-principle relationship at this point.

8.4.1 Service reusability

When a service encapsulates logic that is useful to more than one service requestor, it can be considered reusable. The concept of reuse is supported by a number of complementary service principles, as follows.

- Service autonomy establishes an execution environment that facilitates reuse because the service has independence and self-governance. The less dependencies a service has, the broader the applicability of its reusable functionality.
- Service statelessness supports reuse because it maximizes the availability of a service and typically promotes a generic service design that defers activity-specific processing outside of service logic boundaries.

How service-orientation principles inter-relate

- Service abstraction fosters reuse because it establishes the black box concept, where processing details are completely hidden from requestors. This allows a service to simply express a generic public interface.
- Service discoverability promotes reuse, as it allows requestors (and those that build requestors) to search for and discover reusable services.
- Service loose coupling establishes an inherent independence that frees a service from immediate ties to others. This makes it a great deal easier to realize reuse.

Additionally, the principle of service reuse itself enables the following related principle:

• Service composability is primarily possible because of reuse. The ability for one service to compose an activity around the utilization of a collection of services is feasible when those services being composed are built for reuse. (It is technically possible to build a service so that its sole purpose is to be composed by another, but reuse is generally emphasized.)

Figure 8.26 provides a diagram showing how the principle of service reusability relates to others.



Figure 8.26

Service reusability and its relationship with other service-orientation principles.

8.4.2 Service contract

A service contract is a representation of a service's collective metadata. It standardizes the expression of rules and conditions that need to be fulfilled by any requestor wanting to interact with the service.

Sample Chapter 8 from "Service-Oriented Architecture: Concepts, Technology, and Design" by Thomas Erl

Chapter 8: Principles of Service-Orientation

Service contracts represent a cornerstone principle in service-orientation and therefore support other principles in various ways, as follows:

- Service abstraction is realized through a service contract, as it is the metadata expressed in the contract that defines the only information made available to service requestors. All additional design, processing, and implementation details are hidden behind this contract.
- Service loose coupling is made possible through the use of service contracts. Processing logic from different services do not need to form tight dependencies; they simply need an awareness of each other's communication requirements, as expressed by the service description documents that comprise the service contract.
- Service composability is indirectly enabled through the use of service contracts. It is via the contract that a controller service enlists and uses services that act as composition members.
- Service discoverability is based on the use of service contracts. While some registries provide information supplemental to that expressed through the contract, it is the service description documents that are primarily searched for in the service discovery process.

The diagram in Figure 8.27 illustrates how the principle of service contract usage relates to others.





How service-orientation principles inter-relate

315

8.4.3 Service loose coupling

Loose coupling is a state that supports a level of independence between services (or between service providers and requestors). As you may have already noticed, independence or non-dependency is a fundamental aspect of services and SOA as a whole. Therefore, the principle of persisting loose coupling across services supports the following other service-orientation principles:

- Service reusability is supported through loose coupling because services are freed from tight dependencies on others. This increases their availability for reuse opportunities.
- Service composability is fostered by the loose coupling of services, especially when services are dynamically composed.
- Service statelessness is directly supported through the loosely coupled communications framework established by this principle.
- Service autonomy is made possible through this principle, as it is the nature of loose coupling that minimizes cross-service dependencies.

Additionally, service loose coupling is directly implemented through the application of a related service-orientation principle:

• Service contracts are what enable loose coupling between services, as the contract is the only piece of information required for services to interact.

Figure 8.28 demonstrates these relationships.



Figure 8.28

Service loose coupling and its relationship with other service-orientation principles.

8.4.4 Service abstraction

Part of building solutions with independent services is allowing those services to encapsulate potentially complex processing logic and exposing that logic through a generic and descriptive interface. This is the primary benefit of service abstraction, a principle that is related to others, as explained here:

- Service contracts, in a manner, implement service abstraction by providing the official description information that is made public to external service requestors.
- Service reusability is supported by abstraction, as long as what is being abstracted is actually reusable.

These relationships are shown in Figure 8.29.



Chapter 8: Principles of Service-Orientation

How service-orientation principles inter-relate



Figure 8.29

Service abstraction and its relationship with other serviceorientation principles.

8.4.5 Service composability

Designing services so that they support composition by others is fundamental to building service-oriented solutions. Service composability therefore is tied to service-orientation principles that support the concept of service composition, as follows:

- Service reusability is what enables one service to be composed by numerous others. It is expected that reusable services can be incorporated within different compositions or reused independently by other service requestors.
- Service loose coupling establishes a communications framework that supports the concept of dynamic service composition. Because services are freed from many dependencies, they are more available to be reused via composition.
- Service statelessness supports service composability, especially in larger compositions. A service composition is reliant on the design quality and commonality of its collective parts. If all services are stateless (by, for example, deferring activityspecific logic to messages), the overall composition executes more harmoniously.
- Service autonomy held by composition members strengthens the overall composition, but the autonomy of the controller service itself actually is decreased due to the dependencies on its composition members.
- Service contracts enable service composition by formalizing the runtime agreement between composition members.

Figure 8.30 further illustrates these relationships.

318

Chapter 8: Principles of Service-Orientation



Figure 8.30

Service composability and its relationship with other service-orientation principles.

8.4.6 Service autonomy

This principle applies to a service's underlying logic. By providing an execution environment over which a service has complete control, service autonomy relates to several other principles, as explained here:

- Service reusability is more easily achieved when the service offering reusable logic has self-governance over its own logic. Service Level Agreement (SLA) type requirements that come to the forefront for utility services with multiple requestors, such as availability and scalability, are fulfilled more easily by an autonomous service.
- Service composability is also supported by service autonomy—for much of the same reasons autonomy supports service reusability. A service composition consisting of autonomous services is much more robust and collectively independent.
- Service statelessness is best implemented by a service that can execute independently. Autonomy indirectly supports service statelessness. (However, it is very easy to create a stateful service that is also fully autonomous.)
- Service autonomy is a quality that is realized by leveraging the loosely coupled relationship between services. Therefore service loose coupling is a primary enabler of this principle.

The diagram in Figure 8.31 shows how service autonomy relates to these other principles.





Service autonomy and its relationship with other service-orientation principles.

8.4.7 Service statelessness

To successfully design services not to manage state requires the availability of resources surrounding the service to which state management responsibilities can be delegated. However, the principle of statelessness is also indirectly supported by the following service-orientation principles:

- Service autonomy provides the ability for a service to control its own execution environment. By removing or reducing dependencies it becomes easier to build statelessness into services, primarily because the service logic can be fully customized to defer state management outside of the service logic boundary.
- Service loose coupling and the overall concept of loose coupling establishes a communication paradigm that is fully realized through messaging. This, in turn, supports service statelessness, as state information can be carried and persisted by the messages that pass through the services.

Service statelessness further supports the following principles:

• Service composability benefits from stateless composition members, as they reduce dependencies and minimize the overhead of the composition as a whole.

320

Chapter 8: Principles of Service-Orientation

• Service reuse becomes more of a reality for stateless services, as availability of the service to multiple requestors is increased and the absence of activity-specific logic promotes a generic service design.

Figure 8.32 illustrates how service statelessness relates to the other service-orientation principles.





8.4.8 Service discoverability

Designing services so that they are naturally discoverable enables an environment whereby service logic becomes accessible to new potential service requestors. This is why service discoverability is tied closely to the following service-orientation principles:

- Service contracts are what service requestors (or those that create them) actually discover and subsequently assess for suitability. Therefore, the extent of a service's discoverability can typically be associated with the quality or descriptiveness of its service contract.
- Service reusability is what requestors are looking for when searching for services and it is what makes a service potentially useful once it has been discovered. A service that isn't reusable would likely never need to be discovered because it would probably have been built for a specific service requestor in the first place.

Service-orientation and object-orientation (Part II) 321

The diagram in Figure 8.33 shows how service discoverability fits in with the other service-orientation principles.



Service discoverability and its relationship with other serviceorientation principles.

SUMMARY OF KEY POINTS

- Service-orientation principles are not realized in isolation; principles relate to and support other principles in different ways.
- Principles, such as service reusability and service composability, benefit from the support of other implemented principles.
- Principles, such as service loose coupling, service contract, and service autonomy, provide significant support for the realization of other principles.

8.5 Service-orientation and object-orientation (Part II)

Having now covered the fundamentals behind service-orientation principles, we can continue the discussion we began in the *Service-orientation and object-orientation (Part I)* section from Chapter 4.

Those of you familiar with object-oriented analysis and design probably will have recognized a similarity between a number of the service-orientation principles discussed and the more established principles of object-orientation.

Indeed, service-orientation owes much of its existence to object-oriented concepts and theory. Table 8.1 provides a look at which common object-orientation principles are related to the service-orientation principles we've been discussing.

322

Chapter 8: Principles of Service-Orientation

 Table 8.1
 An overview of how service-orientation principles relate to object-orientation principles.

Service-Orientation Principle	Related Object-Orientation Principles
service reusability	Much of object-orientation is geared toward the creation of reusable classes. The object-orientation principle of modularity standardized decomposition as a means of application design. Related principles, such as abstraction and encapsula-
	tion, further support reuse by requiring a distinct separa- tion of interface and implementation logic. Service reusability is therefore a continuation of this goal.
service contract	The requirement for a service contract is very compara- ble to the use of interfaces when building object-oriented applications. Much like WSDL definitions, interfaces pro- vide a means of abstracting the description of a class. And, much like the "WSDL first" approach encouraged within SOA, the "interface first" approach also is consid- ered an object-orientation best practice.
service loose coupling	Although the creation of interfaces somewhat decouples a class from its consumers, coupling in general is one of the primary qualities of service-orientation that deviates from object-orientation.
	The use of inheritance and other object-orientation prin- ciples encourages a much more tightly coupled relation- ship between units of processing logic when compared to the service-oriented design approach.
service abstraction	The object-orientation principle of abstraction requires that a class provide an interface to the external world and that it be accessible via this interface. Encapsulation sup- ports this by establishing the concept of information hid- ing, where any logic within the class outside of what is exposed via the interface is not accessible to the external world.
	Service abstraction accomplishes much of the same as object abstraction and encapsulation. Its purpose is to hide the underlying details of the service so that only the service contract is available and of concern to service requestors.

Service-orientation and object-orientation (Part II)

Service-Orientation	Related
Principle	Object-Orientation Principles
service composability	Object-orientation supports association concepts, such as aggregation and composition. These, within a loosely coupled context, also are supported by service- orientation. For example, the same way a hierarchy of objects can be composed, a hierarchy of services can be assembled through service composability.
service autonomy	The quality of autonomy is more emphasized in service- oriented design than it has been with object-oriented approaches. Achieving a level of independence between units of processing logic is possible through service- orientation, by leveraging the loosely coupled relation- ship between services. Cross-object references and inheritance-related depend- encies within object-oriented designs support a lower degree of object-level autonomy.
service statelessness	Objects consist of a combination of class and data and are naturally stateful. Promoting statelessness within services therefore tends to deviate from typical object- oriented design. Although it is possible to create stateful services and stateless objects, the principle of statelessness is generally more emphasized with service-orientation.
service discoverability	Designing class interfaces to be consistent and self- descriptive is another object-orientation best practice that improves a means of identifying and distinguishing units of processing logic. These qualities also support reuse by allowing classes to be more easily discovered. Discoverability is another principle more emphasized by the service-orientation paradigm. It is encouraged that service contracts be as communicative as possible to sup- port discoverability at design time and runtime.

As it stands today, object-orientation and service-orientation are not necessarily competitive paradigms. Service-orientation clearly has several roots in object-orientation, and typical contemporary service-oriented solutions will consist of a mixture of services

Chapter 8: Principles of Service-Orientation

(that adhere to service-orientation principles) and object-oriented components. With a balanced design, each set of principles can be properly positioned and leveraged to complement and support each other.

SUMMARY OF KEY POINTS

- Several principles of service-orientation are related to and derived from objectorientation principles.
- Some object-orientation principles, such as inheritance, do not fit into the serviceoriented world.
- Some service-orientation principles, such as loose coupling and autonomy, are not directly promoted by object-orientation.

8.6 Native Web service support for service-orientation principles

Having now worked through the individual principles of service-orientation in some detail, it becomes evident that Web services provide inherent support for some of these principles. It is important to recognize specifically which principles are built into the structure of common Web services because this allows us to place an emphasis on those that require a conscious effort to realize.

Table 8.2 recaps the principles of service-orientation and explains to what extent they are natively supported by Web services.

Service-Orientation Principle	Web Service Support
service reusability	Web services are not automatically reusable. This quality is related to the nature of the logic encapsulated and exposed via the Web service.
service contract	Web services require the use of service descriptions, mak- ing service contracts a fundamental part of Web services communication.
service loose coupling	Web services are naturally loosely coupled through the use of service descriptions.

Table 8.2	A look at which service-orientation principles are automatically supported by Web
	services.

Native Web service support for service-orientation principles

Service-Orientation Web Service Support **Principle** service abstraction Web services automatically emulate the black box model within the Web services communications framework, hiding all of the details of their underlying logic. service composability Web services are naturally composable. The extent to which a service can be composed, though, generally is determined by the service design and the reusability of represented logic. service autonomy To ensure an autonomous processing environment requires design effort. Autonomy is therefore not automatically provided by a Web service. service statelessness Statelessness is a preferred condition for Web services, strongly supported by many WS-* specifications and the document-style SOAP messaging model. service discoverability Discoverability must be implemented by the architecture and even can be considered an extension to IT infrastructure. It is therefore not natively supported by Web services.

It turns out that half of the principles of service-orientation are natural characteristics of common Web services. This underlines the synergy of the marriage between SOA and the Web services technology platform and gives us a good indication as to why Web services have been so successful in realizing SOA.

It also highlights the principles that require special attention when building serviceoriented solutions. The four principles identified as *not* being automatically provided by Web services are:

- service reusability
- service autonomy
- service statelessness
- service discoverability

Chapter 8: Principles of Service-Orientation

Chapters 11 through 15 discuss service modeling and design in detail and provide guidelines to ensure that these important principles are taken into consideration when building services for use within SOA.

These processes further emphasize the other four principles as well—though they may be automatically implemented through the use of Web services, that does not mean they will necessarily be properly realized. For example, the fact that Web services require the use of service contracts has no bearing on how well individual service descriptions are designed.

SUMMARY OF KEY POINTS

- Service abstraction, composability, loose coupling, and the need for service contracts are native characteristics of Web services that are in full alignment with the corresponding principles of service-orientation.
- Service reusability, autonomy, statelessness, and discoverability are not automatically provided by Web services. Realizing these qualities requires a conscious modeling and design effort.



About the Author

Thomas Erl is an independent consultant with XMLTC Consulting in Vancouver, Canada. His previous book, *Service-Oriented Architecture: A Field Guide to Integrating XML and Web Services*, became the top-selling book of 2004 in both Web Services and SOA categories. This guide addresses numerous integration issues and provides strategies and best practices for transitioning toward SOA.

Thomas is a member of OASIS and is active in related research efforts, such as the XML & Web Services Integration Framework (XWIF). He is a speaker and instructor for private and public events and conferences, and has published numerous papers, including articles for the *Web Services Journal*, *WLDJ*, and *Application Development Trends*.

For more information, visit http://www.thomaserl.com/technology/.

About SOA Systems

SOA Systems Inc. is a consulting firm actively involved in the research and development of service-oriented architecture, service-orientation, XML, and Web services standards and technology. Through its research and enterprise solution projects SOA Systems has developed a recognized methodology for integrating and realizing service-oriented concepts, technology, and architecture.

For more information, visit www.soasystems.com.

One of the consulting services provided by SOA Systems is comprehensive SOA transition planning and the objective assessment of vendor technology products.

For more information, visit www.soaplanning.com.

The content in this book is the basis for a series of SOA seminars and workshops developed and offered by SOA Systems.

For more information, visit www.soatraining.com.