

# SOA

## Principles of Service Design

Thomas Erl



PRENTICE HALL  
UPPER SADDLE RIVER, NJ • BOSTON • INDIANAPOLIS • SAN FRANCISCO  
NEW YORK • TORONTO • MONTREAL • LONDON • MUNICH • PARIS • MADRID  
CAPETOWN • SYDNEY • TOKYO • SINGAPORE • MEXICO CITY

# Contents

## **Preface . . . . . xxv**

## **Chapter 1: Introduction . . . . . 1**

1.1 Objectives of this Book . . . . .	3
1.2 Who this Book Is For . . . . .	3
1.3 What this Book Does Not Cover . . . . .	4
Topics Covered by Other Books . . . . .	4
SOA Standardization Efforts . . . . .	5
1.4 How this Book Is Organized . . . . .	6
Part I: Fundamentals . . . . .	7
Part II: Design Principles . . . . .	9
Part III: Supplemental . . . . .	12
Appendices . . . . .	12
1.5 Symbols, Figures, and Style Conventions . . . . .	13
Symbol Legend . . . . .	13
How Color Is Used . . . . .	13
The Service Symbol . . . . .	13
1.6 Additional Information . . . . .	16
Updates, Errata, and Resources (www.soabooks.com) . . . . .	16
Master Glossary (www.soaglossary.com) . . . . .	16
Referenced Specifications (www.soaspecs.com) . . . . .	16
Service-Oriented Computing Poster (www.soaposters.com) . . . . .	16

The SOA Magazine (www.soamag.com) . . . . .	17
Notification Service . . . . .	17
Contact the Author . . . . .	17

## **Chapter 2: Case Study . . . . . 19**

2.1 Case Study Background: Cutit Saws Ltd. . . . .	20
History . . . . .	20
Technical Infrastructure and Automation Environment . . . . .	21
Business Goals and Obstacles . . . . .	21

## **PART I: FUNDAMENTALS**

## **Chapter 3: Service-Oriented Computing and SOA. . . . . 25**

3.1 Design Fundamentals . . . . .	26
Design Characteristic . . . . .	27
Design Principle . . . . .	28
Design Paradigm . . . . .	29
Design Pattern. . . . .	30
Design Pattern Language . . . . .	31
Design Standard . . . . .	32
Best Practice . . . . .	34
A Fundamental Design Framework. . . . .	35
3.2 Introduction to Service-Oriented Computing . . . . .	37
Service-Oriented Architecture. . . . .	38
Service-Orientation, Services, and Service-Oriented Solution Logic . . . . .	39
Service Compositions . . . . .	39
Service Inventory. . . . .	40
Understanding Service-Oriented Computing Elements . . . . .	40
Service Models . . . . .	43
SOA and Web Services. . . . .	46
Service Inventory Blueprints . . . . .	51
Service-Oriented Analysis and Service Modeling. . . . .	52

## Contents

## xv

Service-Oriented Design . . . . .	53
Service-Oriented Architecture: Concepts, Technology, and Design . . . . .	54
3.3 Goals and Benefits of Service-Oriented Computing . . . . .	55
Increased Intrinsic Interoperability . . . . .	56
Increased Federation . . . . .	58
Increased Vendor Diversification Options . . . . .	59
Increased Business and Technology Domain Alignment . . . . .	60
Increased ROI . . . . .	61
Increased Organizational Agility . . . . .	63
Reduced IT Burden . . . . .	64
3.4 Case Study Background . . . . .	66

**Chapter 4: Service-Orientation . . . . . 67**

4.1 Introduction to Service-Orientation . . . . .	68
Services in Business Automation . . . . .	69
Services Are Collections of Capabilities . . . . .	69
Service-Orientation as a Design Paradigm . . . . .	70
Service-Orientation and Interoperability . . . . .	74
4.2 Problems Solved by Service-Orientation . . . . .	75
Life Before Service-Orientation . . . . .	76
The Need for Service-Orientation . . . . .	81
4.3 Challenges Introduced by Service-Orientation . . . . .	85
Design Complexity . . . . .	85
The Need for Design Standards . . . . .	86
Top-Down Requirements . . . . .	86
Counter-Agile Service Delivery in Support of Agile Solution Delivery . . . . .	87
Governance Demands . . . . .	88
4.4 Additional Considerations . . . . .	89
It Is Not a Revolutionary Paradigm . . . . .	89
Enterprise-wide Standardization Is Not Required . . . . .	89
Reuse Is Not an Absolute Requirement . . . . .	90

4.5 Effects of Service-Orientation on the Enterprise . . . . .	91
Service-Orientation and the Concept of “Application” . . . . .	91
Service-Orientation and the Concept of “Integration” . . . . .	92
The Service Composition . . . . .	94
Application, Integration, and Enterprise Architectures . . . . .	95
4.6 Origins and Influences of Service-Orientation . . . . .	96
Object-Orientation . . . . .	97
Web Services . . . . .	98
Business Process Management (BPM) . . . . .	98
Enterprise Application Integration (EAI) . . . . .	98
Aspect-Oriented Programming (AOP) . . . . .	99
4.7 Case Study Background . . . . .	100

## **Chapter 5: Understanding Design Principles . . . . . 103**

5.1 Using Design Principles . . . . .	104
Incorporate Principles within Service-Oriented Analysis . . . . .	105
Incorporate Principles within Formal Design Processes . . . . .	106
Establish Supporting Design Standards . . . . .	107
Apply Principles to a Feasible Extent . . . . .	108
5.2 Principle Profiles . . . . .	109
5.3 Design Pattern References . . . . .	111
5.4 Principles that Implement vs. Principles that Regulate . . . . .	111
5.5 Principles and Service Implementation Mediums . . . . .	114
“Capability” vs. “Operation” vs. “Method” . . . . .	115
5.6 Principles and Design Granularity . . . . .	115
Service Granularity . . . . .	116
Capability Granularity . . . . .	116
Data Granularity . . . . .	116
Constraint Granularity . . . . .	117
Sections on Granularity Levels . . . . .	118
5.7 Case Study Background . . . . .	119
The Lab Project Business Process . . . . .	119

**PART II: DESIGN PRINCIPLES**

<b>Chapter 6: Service Contracts (Standardization and Design) . . . . .</b>	<b>125</b>
6.1 Contracts Explained . . . . .	126
Technical Contracts in Abstract . . . . .	126
Origins of Service Contracts . . . . .	127
6.2 Profiling this Principle . . . . .	130
6.3 Types of Service Contract Standardization . . . . .	132
Standardization of Functional Service Expression . . . . .	133
Standardization of Service Data Representation . . . . .	134
Standardization of Service Policies . . . . .	137
6.4 Contracts and Service Design . . . . .	140
Data Representation Standardization and Transformation Avoidance. . . . .	140
Standardization and Granularity . . . . .	142
Standardized Service Contracts and Service Models . . . . .	144
How Standardized Service Contract Design Affects Other Principles . . . . .	144
6.5 Risks Associated with Service Contract Design . . . . .	149
Versioning . . . . .	149
Technology Dependencies . . . . .	150
Development Tool Deficiencies. . . . .	151
6.6 More About Service Contracts . . . . .	152
Non-Technical Service Contract Documents . . . . .	152
“Web Service Contract Design for SOA”. . . . .	153
6.7 Case Study Example. . . . .	154
Planned Services . . . . .	154
Design Standards . . . . .	155
Standardized WSDL Definition Profiles . . . . .	155
Standardized XML Schema Definitions. . . . .	157
Standardized Service and Data Representation Layers . . . . .	157
Service Descriptions . . . . .	158
Conclusion . . . . .	160

**Chapter 7: Service Coupling (Intra-Service and Consumer Dependencies) . . . . . 163**

7.1 Coupling Explained . . . . .	164
Coupling in Abstract . . . . .	165
Origins of Software Coupling . . . . .	165
7.2 Profiling this Principle . . . . .	167
7.3 Service Contract Coupling Types . . . . .	169
Logic-to-Contract Coupling (the coupling of service logic to the service contract) . . . . .	173
Contract-to-Logic Coupling (the coupling of the service contract to its logic) . . . . .	174
Contract-to-Technology Coupling (the coupling of the service contract to its underlying technology) . . . . .	176
Contract-to-Implementation Coupling (the coupling of the service contract to its implementation environment) . . . . .	177
Contract-to-Functional Coupling (the coupling of the service contract to external logic) . . . . .	180
7.4 Service Consumer Coupling Types . . . . .	181
Consumer-to-Implementation Coupling . . . . .	182
Standardized Service Coupling and Contract Centralization . . . . .	185
Consumer-to-Contract Coupling . . . . .	185
Measuring Consumer Coupling . . . . .	191
7.5 Service Loose Coupling and Service Design . . . . .	193
Coupling and Service-Orientedness . . . . .	193
Service Loose Coupling and Granularity . . . . .	195
Coupling and Service Models . . . . .	196
How Service Loose Coupling Affects Other Principles . . . . .	197
7.6 Risks Associated with Service Loose Coupling . . . . .	200
Limitations of Logic-to-Contract Coupling . . . . .	200
Problems when Schema Coupling Is “too loose” . . . . .	201
7.7 Case Study Example . . . . .	202
Coupling Levels of Existing Services . . . . .	202
Introducing the InvLegacyAPI Service . . . . .	203
Service Design Options . . . . .	205

**Chapter 8: Service Abstraction (Information Hiding  
and Meta Abstraction Types) . . . . . 211**

8.1 Abstraction Explained . . . . .	212
Origins of Information Hiding . . . . .	213
8.2 Profiling this Principle . . . . .	214
Why Service Abstraction Is Needed . . . . .	214
8.3 Types of Meta Abstraction . . . . .	218
Technology Information Abstraction . . . . .	219
Functional Abstraction . . . . .	221
Programmatic Logic Abstraction . . . . .	222
Quality of Service Abstraction . . . . .	224
Meta Abstraction Types and the Web Service Regions of Influence . . . . .	225
Meta Abstraction Types in the Real World . . . . .	227
8.4 Measuring Service Abstraction . . . . .	231
Contract Content Abstraction Levels . . . . .	231
Access Control Levels . . . . .	232
Abstraction Levels and Quality of Service Meta Information . . . . .	234
8.5 Service Abstraction and Service Design . . . . .	235
Service Abstraction vs. Service Encapsulation . . . . .	235
How Encapsulation Can Affect Abstraction . . . . .	235
Service Abstraction and Non-Technical Contract Documents . . . . .	237
Service Abstraction and Granularity . . . . .	238
Service Abstraction and Service Models . . . . .	239
How Service Abstraction Affects Other Principles . . . . .	239
8.6 Risks Associated with Service Abstraction . . . . .	242
Multi-Consumer Coupling Requirements . . . . .	242
Misjudgment by Humans . . . . .	242
Security and Privacy Concerns . . . . .	243
8.7 Case Study Example . . . . .	244
Service Abstraction Levels . . . . .	244
Operation-Level Abstraction Examples . . . . .	247



xx

Contents

**Chapter 9: Service Reusability (Commercial and  
Agnostic Design) . . . . . 253**

9.1 Reuse Explained . . . . .	254
Reuse in Abstract . . . . .	254
Origins of Reuse . . . . .	257
9.2 Profiling this Principle . . . . .	259
9.3 Measuring Service Reusability and Applying Commercial Design. . . . .	262
Commercial Design Considerations . . . . .	262
Measures of Planned Reuse . . . . .	265
Measuring Actual Reuse . . . . .	267
Commercial Design Versus Gold-Plating . . . . .	267
9.4 Service Reuse in SOA . . . . .	268
Reuse and the Agnostic Service . . . . .	268
The Service Inventory Blueprint . . . . .	269
9.5 Standardized Service Reuse and Logic Centralization . . . . .	270
Understanding Logic Centralization . . . . .	271
Logic Centralization as an Enterprise Standard . . . . .	272
Logic Centralization and Contract Centralization . . . . .	272
Centralization and Web Services . . . . .	274
Challenges to Achieving Logic Centralization . . . . .	274
9.6 Service Reusability and Service Design . . . . .	276
Service Reusability and Service Modeling . . . . .	276
Service Reusability and Granularity . . . . .	277
Service Reusability and Service Models. . . . .	278
How Service Reusability Affects Other Principles . . . . .	278
9.7 Risks Associated with Service Reusability and Commercial Design. . . . .	281
Cultural Concerns . . . . .	281
Governance Concerns . . . . .	283
Reliability Concerns . . . . .	286
Security Concerns. . . . .	286
Commercial Design Requirement Concerns. . . . .	286
Agile Delivery Concerns . . . . .	287

## Contents

## xxi

9.8 Case Study Example . . . . .	288
The Inventory Service Profile . . . . .	288
Assessing Current Capabilities . . . . .	289
Modeling for a Targeted Measure of Reusability . . . . .	289
The New EditItemRecord Operation . . . . .	290
The New ReportStockLevels Operation . . . . .	290
The New AdjustItemsQuantity Operation . . . . .	291
Revised Inventory Service Profile . . . . .	292

## **Chapter 10: Service Autonomy (Processing Boundaries and Control) . . . . . 293**

10.1 Autonomy Explained . . . . .	294
Autonomy in Abstract . . . . .	294
Origins of Autonomy . . . . .	295
10.2 Profiling this Principle . . . . .	296
10.3 Types of Service Autonomy . . . . .	297
Runtime Autonomy (execution) . . . . .	298
Design-Time Autonomy (governance) . . . . .	298
10.4 Measuring Service Autonomy . . . . .	300
Service Contract Autonomy (services with normalized contracts) . . . . .	301
Shared Autonomy . . . . .	305
Service Logic Autonomy (partially isolated services) . . . . .	306
Pure Autonomy (isolated services) . . . . .	308
Services with Mixed Autonomy . . . . .	310
10.5 Autonomy and Service Design . . . . .	311
Service Autonomy and Service Modeling . . . . .	311
Service Autonomy and Granularity . . . . .	311
Service Autonomy and Service Models . . . . .	312
How Service Autonomy Affects Other Principles . . . . .	314
10.6 Risks Associated with Service Autonomy . . . . .	317
Misjudging the Service Scope . . . . .	317
Wrapper Services and Legacy Logic Encapsulation . . . . .	318
Overestimating Service Demand . . . . .	318

10.7 Case Study Example . . . . .	319
Existing Implementation Autonomy of the GetItem Operation . .	319
New Operation-Level Architecture with Increased Autonomy . .	320
Effect on the Run Lab Project Composition . . . . .	322

## **Chapter 11: Service Statelessness (State Management Deferral and Stateless Design) . . . . . 325**

11.1 State Management Explained . . . . .	327
State Management in Abstract . . . . .	327
Origins of State Management . . . . .	328
Deferral vs. Delegation . . . . .	331
11.2 Profiling this Principle . . . . .	331
11.3 Types of State . . . . .	335
Active and Passive . . . . .	335
Stateless and Stateful . . . . .	336
Session and Context Data. . . . .	336
11.4 Measuring Service Statelessness . . . . .	339
Non-Deferred State Management (low-to-no statelessness) . .	340
Partially Deferred Memory (reduced statefulness) . . . . .	340
Partial Architectural State Management Deferral (moderate statelessness) . . . . .	341
Full Architectural State Management Deferral (high statelessness) . . . . .	342
Internally Deferred State Management (high statelessness) . .	342
11.5 Statelessness and Service Design . . . . .	343
Messaging as a State Deferral Option . . . . .	343
Service Statelessness and Service Instances . . . . .	344
Service Statelessness and Granularity . . . . .	346
Service Statelessness and Service Models . . . . .	346
How Service Statelessness Affects Other Principles . . . . .	347
11.6 Risks Associated with Service Statelessness . . . . .	349
Dependency on the Architecture . . . . .	349
Increased Runtime Performance Demands . . . . .	350
Underestimating Delivery Effort . . . . .	350

## Contents

xxiii

11.7 Case Study Example . . . . .	351
Solution Architecture with State Management Deferral . . . . .	352
Step 1 . . . . .	353
Step 2 . . . . .	354
Step 3 . . . . .	355
Step 4 . . . . .	356
Step 5 . . . . .	357
Step 6 . . . . .	358
Step 7 . . . . .	359

## **Chapter 12: Service Discoverability (Interpretability and Communication) . . . . . 361**

12.1 Discoverability Explained . . . . .	362
Discovery and Interpretation, Discoverability and Interpretability in Abstract . . . . .	364
Origins of Discovery . . . . .	367
12.2 Profiling this Principle . . . . .	368
12.3 Types of Discovery and Discoverability	
Meta Information . . . . .	371
Design-Time and Runtime Discovery . . . . .	371
Discoverability Meta Information . . . . .	373
Functional Meta Data . . . . .	374
Quality of Service Meta Data . . . . .	374
12.4 Measuring Service Discoverability . . . . .	375
Fundamental Levels . . . . .	375
Custom Rating System . . . . .	376
12.5 Discoverability and Service Design . . . . .	376
Service Discoverability and Service Modeling . . . . .	377
Service Discoverability and Granularity . . . . .	378
Service Discoverability and Policy Assertions . . . . .	378
Service Discoverability and Service Models . . . . .	378
How Service Discoverability Affects Other Principles . . . . .	378

## xxiv

## Contents

12.6 Risks Associated with Service Discoverability . . . . .	381
Post-Implementation Application of Discoverability . . . . .	381
Application of this Principle by Non-Communicative Resources	381
12.7 Case Study Example . . . . .	382
Service Profiles (Functional Meta Information) . . . . .	382
Related Quality of Service Meta Information . . . . .	386

### **Chapter 13: Service Composability (Composition Member Design and Complex Compositions) . . . . . 387**

13.1 Composition Explained . . . . .	388
Composition in Abstract . . . . .	388
Origins of Composition . . . . .	390
13.2 Profiling this Principle . . . . .	392
13.3 Composition Concepts and Terminology . . . . .	396
Compositions and Composition Instances . . . . .	397
Composition Members and Controllers . . . . .	398
Service Compositions and Web Services . . . . .	401
Service Activities . . . . .	402
Composition Initiators . . . . .	403
Point-to-Point Data Exchanges and Compositions . . . . .	405
Types of Compositions . . . . .	406
13.4 The Complex Service Composition . . . . .	407
Stages in the Evolution of a Service Inventory . . . . .	407
Defining the Complex Service Composition . . . . .	410
Preparing for the Complex Service Composition . . . . .	411
13.5 Measuring Service Composability and Composition	
Effectiveness Potential . . . . .	412
Evolutionary Cycle States of a Composition . . . . .	412
Composition Design Assessment . . . . .	413
Composition Runtime Assessment . . . . .	415
Composition Governance Assessment . . . . .	417
Measuring Composability . . . . .	419

## Contents

xxv

13.6 Composition and Service Design . . . . .	427
Service Composability and Granularity . . . . .	427
Service Composability and Service Models . . . . .	428
Service Composability and Composition Autonomy . . . . .	430
Service Composability and Orchestration . . . . .	430
How Service Composability Affects Other Principles . . . . .	432
13.7 Risks Associated with Service Composition . . . . .	437
Composition Members as Single Points of Failure . . . . .	437
Composition Members as Performance Bottlenecks . . . . .	437
Governance Rigidity of “Over-Reuse” in Compositions . . . . .	438
13.8 Case Study Example . . . . .	439

**PART III: SUPPLEMENTAL****Chapter 14: Service-Oriented and Object-  
Orientation: A Comparison of Principles  
and Concepts . . . . . 445**

14.1 A Tale of Two Design Paradigms . . . . .	446
14.2 A Comparison of Goals . . . . .	449
Increased Business Requirements Fulfillment . . . . .	450
Increased Robustness . . . . .	451
Increased Extensibility . . . . .	451
Increased Flexibility . . . . .	452
Increased Reusability and Productivity . . . . .	452
14.3 A Comparison of Fundamental Concepts . . . . .	453
Classes and Objects . . . . .	453
Methods and Attributes . . . . .	454
Messages . . . . .	454
Interfaces . . . . .	456
14.4 A Comparison of Design Principles . . . . .	457
Encapsulation . . . . .	458
Inheritance . . . . .	459

**xxvi****Contents**

Generalization and Specialization . . . . .	461
Abstraction . . . . .	463
Polymorphism . . . . .	463
Open-Closed Principle (OCP) . . . . .	465
Don't Repeat Yourself (DRY) . . . . .	465
Single Responsibility Principle (SRP) . . . . .	466
Delegation . . . . .	468
Association . . . . .	469
Composition . . . . .	470
Aggregation . . . . .	471
14.5 Guidelines for Designing Service-Oriented Classes . . . . .	472
Implement Class Interfaces . . . . .	473
Limit Class Access to Interfaces . . . . .	473
Do Not Define Public Attributes in Interfaces . . . . .	473
Use Inheritance with Care . . . . .	473
Avoid Cross-Service "has-a" Relationships . . . . .	474
Use Abstract Classes for Modeling, Not Design . . . . .	474
Use Façade Classes . . . . .	474

**Chapter 15: Supporting Practices . . . . . 477**

15.1 Service Profiles . . . . .	478
Service-Level Profile Structure . . . . .	478
Capability Profile Structure . . . . .	480
Additional Considerations . . . . .	482
15.2 Vocabularies . . . . .	483
Service-Oriented Computing Terms . . . . .	484
Service Classification Terms . . . . .	484
Types and Associated Terms . . . . .	485
Design Principle Application Levels . . . . .	487
15.3 Organizational Roles . . . . .	488
Service Analyst . . . . .	490
Service Architect . . . . .	490
Service Custodian . . . . .	491
Schema Custodian . . . . .	491
Policy Custodian . . . . .	492

## Contents

xxvii

Service Registry Custodian . . . . .	492
Technical Communications Specialist . . . . .	493
Enterprise Architect . . . . .	493
Enterprise Design Standards Custodian (and Auditor) . . . . .	494

## **Chapter 16: Mapping Service-Orientation Principles to Strategic Goals . . . . . 497**

16.1 Principles that Increase Intrinsic Interoperability . . . . .	498
16.2 Principles that Increase Federation . . . . .	501
16.3 Principles that Increase Vendor Diversification Options . . . . .	501
16.4 Principles that Increase Business and Technology Domain Alignment . . . . .	502
16.5 Principles that Increase ROI . . . . .	504
16.6 Principles that Increase Organizational Agility . . . . .	505
16.7 Principles that Reduce the Overall Burden of IT . . . . .	507

## **PART IV: APPENDICES**

### **Appendix A: Case Study Conclusion . . . . . 513**

### **Appendix B: Process Descriptions . . . . . 517**

B.1 Delivery Processes . . . . .	518
Bottom-Up vs. Top-Down . . . . .	518
The Inventory Analysis Cycle . . . . .	520
Inventory Analysis and Service-Oriented Design . . . . .	521
Choosing a Delivery Strategy . . . . .	521
B.2 Service-Oriented Analysis Process . . . . .	522
Define Analysis Scope . . . . .	522
Identify Affected Systems . . . . .	523
Perform Service Modeling . . . . .	523



B.3 Service Modeling Process . . . . .	523
B.4 Service-Oriented Design Processes. . . . .	525
Design Processes and Service Models . . . . .	526
Service Design Processes and Service-Orientation . . . . .	527
 <b>Appendix C: Principles and Patterns</b>	
<b>Cross-Reference . . . . .</b>	<b>529</b>
 <b>Additional Resources . . . . .</b>	<b>533</b>
 <b>About the Author . . . . .</b>	<b>535</b>
 <b>About the Photos . . . . .</b>	<b>537</b>
 <b>Index . . . . .</b>	<b>539</b>



# Chapter 4

## Service-Orientation

- 4.1 Introduction to Service-Orientation
- 4.2 Problems Solved by Service-Orientation
- 4.3 Challenges Introduced by Service-Orientation
- 4.4 Additional Considerations
- 4.5 Effects of Service-Orientation on the Enterprise
- 4.6 Origins and Influences of Service-Orientation
- 4.7 Case Study Background

Having covered some of the basic elements of service-oriented computing, we now narrow our focus on service-orientation. The next set of sections establish the paradigm of service-orientation and explain how it is changing the face of distributed computing.

#### 4.1 Introduction to Service-Orientation

In the every day world around us, services are and have been commonplace for as long as civilized history has existed. Any person carrying out a distinct task in support of others is providing a service (Figure 4.1). Any group of individuals collectively performing a task is also demonstrating the delivery of a service.

**Figure 4.1**

Three individuals, each capable of providing a distinct service.



Similarly, an organization that carries out tasks associated with its purpose or business is also providing a service. As long as the task or function being provided is well-defined and can be relatively isolated from other associated tasks, it can be distinctly classified as a service (Figure 4.2).


Certain baseline requirements exist to enable a group of individual service providers to collaborate in order to collectively provide a larger service. Figure 4.2, for example, displays a group of employees that each provide a service for ABC Delivery. Even though each individual contributes a distinct service, for the company to function effectively, its staff also needs to have fundamental, common characteristics, such as availability, reliability, and the ability to communicate using the same language. With all of this in place, these individuals can be composed into a productive working team. Establishing these types of baseline requirements is a key goal of service-orientation.

#### 4.1 Introduction to Service-Orientation

69

**Figure 4.2**

A company that employs these three people can compose their capabilities to carry out its business.



#### Services in Business Automation

In the world of SOA and service-orientation, the term “service” is not generic. It has specific connotations that relate to a unique combination of design characteristics. When solution logic is consistently built as services and when services are consistently designed with these common characteristics, service-orientation is successfully realized throughout an environment.

For example, one of the primary service design characteristics explored as part of this study of service-orientation is reusability. A strong emphasis on producing solution logic in the format of services that are positioned as highly generic and reusable enterprise resources gradually transitions an organization to a state where more and more of its solution logic becomes less dependent on and more agnostic to any one purpose or business process. Repeatedly fostering this characteristic within services eventually results in wide-spread reuse potential.

Consistently realizing specific design characteristics requires a set of guiding principles. This is what the service-orientation design paradigm is all about.


#### Services Are Collections of Capabilities

When discussing services, it is important to remember that a single service can provide a collection of capabilities. They are grouped together because they relate to a functional

context established by the service. The functional context of the service illustrated in Figure 4.3, for example, is that of “shipment.” Therefore, this particular service provides a set of capabilities associated with the processing of shipments.

**Figure 4.3**

Much like a human, an automated service can provide multiple capabilities.



A service can essentially act as a container of related capabilities. It is comprised of a body of logic designed to carry out these capabilities and a service contract that expresses which of its capabilities are made available for public invocation.

References to service capabilities in this book are specifically focused on those that are defined in the service contract. For a discussion of how service capabilities are distinguished from Web service operations and component methods, see the *Principles and Service Implementation Mediums* section in Chapter 5.

### Service-Orientation as a Design Paradigm

As established in Chapter 3, a design paradigm is an approach to designing solution logic. When building distributed solution logic, design approaches revolve around a software engineering theory known as the *separation of concerns*. In a nutshell, this theory states that a larger problem is more effectively solved when decomposed into a set of smaller problems or *concerns*. This gives us the option of partitioning solution logic into capabilities, each designed to solve an individual concern. Related capabilities can be grouped into units of solution logic.

The fundamental benefit to solving problems this way is that a number of the solution logic units can be designed to solve immediate concerns while still remaining agnostic to the greater problem. This provides the constant opportunity for us to reutilize the capabilities within those units to solve other problems as well.

Different design paradigms exist for distributed solution logic. What distinguishes service-orientation is the manner in which it carries out the separation of concerns and how it shapes the individual units of solution logic. Applying service-orientation to a meaningful extent results in solution logic that can be safely classified as “service-oriented”

#### 4.1 Introduction to Service-Orientation

**71**

and units that qualify as “services.” To understand exactly what that means requires an appreciation of the strategic goals covered in Chapter 3 combined with knowledge of the associated design principles documented in Part II.

For now, let’s briefly introduce each of these principles:

##### *Standardized Service Contract*

Services express their purpose and capabilities via a service contract. The Standardized Service Contract design principle is perhaps the most fundamental part of service-orientation in that it essentially requires that specific considerations be taken into account when designing a service’s public technical interface and assessing the nature and quantity of content that will be published as part of a service’s official contract.

A great deal of emphasis is placed on specific aspects of contract design, including the manner in which services express functionality, how data types and data models are defined, and how policies are asserted and attached. There is a constant focus on ensuring that service contracts are both optimized, appropriately granular, and standardized to ensure that the endpoints established by services are consistent, reliable, and governable.

Chapter 6 is dedicated to exploring this design principle in detail.

##### *Service Loose Coupling*

Coupling refers to a connection or relationship between two things. A measure of coupling is comparable to a level of dependency. This principle advocates the creation of a specific type of relationship within and outside of service boundaries, with a constant emphasis on reducing (“loosening”) dependencies between the service contract, its implementation, and its service consumers.

The principle of Service Loose Coupling promotes the independent design and evolution of a service’s logic and implementation while still guaranteeing baseline interoperability with consumers that have come to rely on the service’s capabilities. There are numerous types of coupling involved in the design of a service, each of which can impact the content and granularity of its contract. Achieving the appropriate level of coupling requires that practical considerations be balanced against various service design preferences.

Chapter 7 provides an in-depth exploration of this principle and introduces related patterns and concepts.

**72**

## Chapter 4: Service-Orientation

*Service Abstraction*

Abstraction ties into many aspects of service-orientation. On a fundamental level, this principle emphasizes the need to hide as much of the underlying details of a service as possible. Doing so directly enables and preserves the previously described loosely coupled relationship. Service Abstraction also plays a significant role in the positioning and design of service compositions.

Various forms of meta data come into the picture when assessing appropriate abstraction levels. The extent of abstraction applied can affect service contract granularity and can further influence the ultimate cost and effort of governing the service.

Chapter 8 covers several aspects of applying abstraction to different types of service meta data, along with processes and approaches associated with information hiding.

*Service Reusability*

Reuse is strongly advocated within service-orientation; so much so, that it becomes a core part of typical service analysis and design processes, and also forms the basis for key service models. The advent of mature, non-proprietary service technology has provided the opportunity to maximize the reuse potential of multi-purpose logic on an unprecedented level.

The principle of Service Reusability emphasizes the positioning of services as enterprise resources with agnostic functional contexts. Numerous design considerations are raised to ensure that individual service capabilities are appropriately defined in relation to an agnostic service context, and to guarantee that they can facilitate the necessary reuse requirements.

Variations and levels of reuse and associated agnostic service models are covered in Chapter 9, along with a study of how commercial product design approaches have influenced this principle.

*Service Autonomy*

For services to carry out their capabilities consistently and reliably, their underlying solution logic needs to have a significant degree of control over its environment and resources. The principle of Service Autonomy supports the extent to which other design principles can be effectively realized in real world production environments by fostering design characteristics that increase a service's reliability and behavioral predictability.

#### 4.1 Introduction to Service-Orientation

**73**

This principle raises various issues that pertain to the design of service logic as well as the service's actual implementation environment. Isolation levels and service normalization considerations are taken into account to achieve a suitable measure of autonomy, especially for reusable services that are frequently shared.

Chapter 10 documents the design issues and challenges related to attaining higher levels of service autonomy, and further classifies different forms of autonomy and highlights associated risks.

##### *Service Statelessness*

The management of excessive state information can compromise the availability of a service and undermine its scalability potential. Services are therefore ideally designed to remain stateful only when required. Applying the principle of Service Statelessness requires that measures of realistically attainable statelessness be assessed, based on the adequacy of the surrounding technology architecture to provide state management delegation and deferral options.

Chapter 11 explores the options and impacts of incorporating stateless design characteristics into service architectures.

##### *Service Discoverability*

For services to be positioned as IT assets with repeatable ROI they need to be easily identified and understood when opportunities for reuse present themselves. The service design therefore needs to take the “communications quality” of the service and its individual capabilities into account, regardless of whether a discovery mechanism (such as a service registry) is an immediate part of the environment.

The application of this principle, as well as an explanation of how discoverability relates to interpretability and the overall service discovery process, are covered in Chapter 12.

##### *Service Composability*

As the sophistication of service-oriented solutions continues to grow, so does the complexity of underlying service composition configurations. The ability to effectively compose services is a critical requirement for achieving some of the most fundamental goals of service-oriented computing.



Complex service compositions place demands on service design that need to be anticipated to avoid massive retro-fitting efforts. Services are expected to be capable of participating as effective composition members, regardless of whether they need to be immediately enlisted in a composition. The principle of Service Composability addresses this requirement by ensuring that a variety of considerations are taken into account.

How the application of this design principle helps prepare services for the world of complex compositions is described in Chapter 13.

### Service-Orientation and Interoperability

One item that may appear to be absent from the preceding list is a principle along the lines of “*Services are Interoperable*.” The reason this does not exist as a separate principle is because interoperability is fundamental to every one of the principles just described. Therefore, in relation to service-oriented computing, stating that services must be interoperable is just about as basic as stating that services must exist. Each of the eight principles supports or contributes to interoperability in some manner.

Here are just a few examples:

- Service contracts are standardized to guarantee a baseline measure of interoperability associated with the harmonization of data models.
- Reducing the degree of service coupling fosters interoperability by making individual services less dependent on others and therefore more open for invocation by different service consumers.
- Abstracting details about the service limits all interoperation to the service contract, increasing the long-term consistency of interoperability by allowing underlying service logic to evolve more independently.
- Designing services for reuse implies a high-level of required interoperability between the service and numerous potential service consumers.
- By raising a service’s individual autonomy, its behavior becomes more consistently predictable, increasing its reuse potential and thereby its attainable level of interoperability.
- Through an emphasis on stateless design, the availability and scalability of services increase, allowing them to interoperate more frequently and reliably.

## 4.2 Problems Solved by Service-Orientation

75

- Service Discoverability simply allows services to be more easily located by those who want to potentially interoperate with them.
- Finally, for services to be effectively composable they must be interoperable. The success of fulfilling composability requirements is often tied directly to the extent to which services are standardized and cross-service data exchange is optimized.

A fundamental goal of applying service-orientation is for interoperability to become a natural by-product, ideally to the extent that a level of intrinsic interoperability is established as a common and expected service design characteristic. Depending on the architectural strategy being employed, this extent may or may not be limited to a specific service inventory.

Of course, as with any other design characteristic, there are levels of interoperability a service can attain. The ultimate measure is generally determined by the extent to which service-orientation principles have been consistently and successfully realized (plus, of course, environmental factors such as the compatibility of wire protocols, the maturity level of the underlying technology platform, and adherence to technology standards).

### NOTE

Increased intrinsic interoperability is one of the key strategic goals associated with service-oriented computing (as originally established in Chapter 3). For more detailed information about how service-orientation principles directly support this and other strategic goals, see Chapter 16.

### SUMMARY OF KEY POINTS

- The service-orientation paradigm consists of eight distinct design principles, each of which fosters fundamental design characteristics, such as interoperability. These principles are explored individually in subsequent chapters.
- Interoperability is a natural by-product of applying service-orientation design principles.

## 4.2 Problems Solved by Service-Orientation

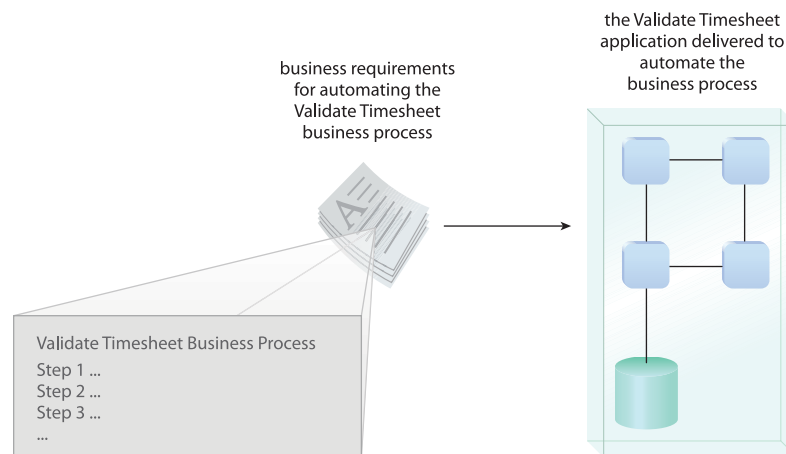
To best appreciate why service-orientation has emerged and how it is intended to improve the design of automation systems, we need to compare before and after perspectives. By studying some of the common issues that have historically plagued IT, we can begin to understand the solutions proposed by this design paradigm.

**NOTE**

This book fully acknowledges that past design paradigms have advocated similar principles and strategic goals as service-orientation. Several of these design approaches, in fact, directly inspired or influenced service-orientation (as explained further in the *Origins and Influences of Service-Oriented Architecture* section of this chapter). The following section is focused specifically on a comparison with the silo-based design approach because it has persisted as the most common means by which applications are delivered.

**Life Before Service-Oriented Architecture**

In the world of business it makes a great deal of sense to deliver solutions capable of automating the execution of business tasks. Over the course of IT's history, the majority of such solutions have been created with a common approach of identifying the business tasks to be automated, defining their business requirements, and then building the corresponding solution logic (Figure 4.4).

**Figure 4.4**

A ratio of one application for each new set of automation requirements has been common.

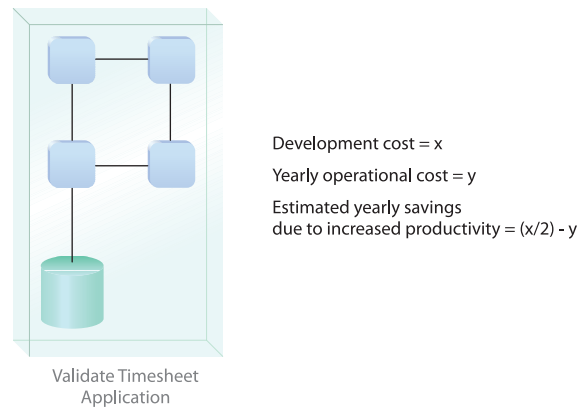
This has been an accepted and proven approach to achieving tangible business benefits through the use of technology and has been successful at providing a relatively predictable return on investment (Figure 4.5).

## 4.2 Problems Solved by Service-Orientation

77

**Figure 4.5**

A sample formula for calculating ROI is based on a predetermined investment with a predictable return.



The ability to gain any further value from these applications is usually inhibited because their capabilities are tied to specific business requirements and processes (some of which will even have a limited lifespan). When new requirements and processes come our way, we are forced to either make significant changes to what we already have, or we may need to build a new application altogether.

In the latter case, although repeatedly building “disposable applications” is not the perfect approach, it has proven itself as a legitimate means of automating business. Let’s explore some of the lessons learned by first focusing on the positive.

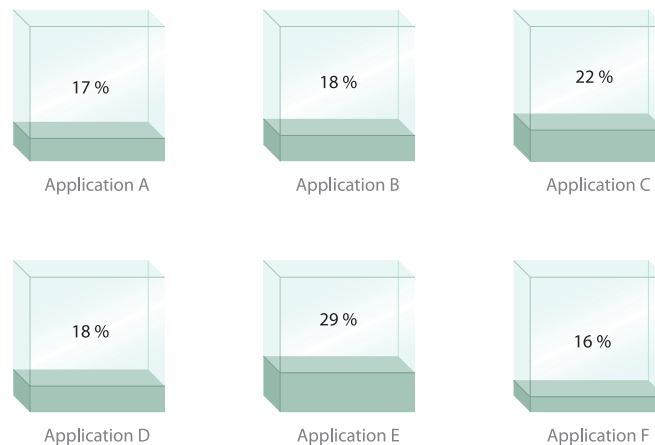
- Solutions can be built efficiently because they only need to be concerned with the fulfillment of a narrow set of requirements associated with a limited set of business processes.
- The business analysis effort involved with defining the process to be automated is straight forward. Analysts are focused only on one process at a time and therefore only concern themselves with the business entities and domains associated with that one process.
- Solution designs are tactically focused. Although complex and sophisticated automation solutions are sometimes required, the sole purpose of each is to automate just one or a specific set of business processes. This predefined functional scope simplifies the overall solution design as well as the underlying application architecture.

- The project delivery lifecycle for each solution is streamlined and relatively predictable. Although IT projects are notorious for being complex endeavors, riddled with unforeseen challenges, when the delivery scope is well-defined (and doesn't change), the process and execution of the delivery phases have a good chance of being carried out as expected.
- Building new systems from the ground up allows organizations to take advantage of the latest technology advancements. The IT marketplace progresses every year to the extent that we fully expect technology we use to build solution logic today to be different and better tomorrow. As a result, organizations that repeatedly build disposable applications can leverage the latest technology innovations with each new project.

These and other common characteristics of traditional solution delivery provide a good indication as to why this approach has been so popular. Despite its acceptance, though, it has become evident that there is still lots of room for improvement.

#### *It Can Be Highly Wasteful*

The creation of new solution logic in a given enterprise commonly results in a significant amount of redundant functionality (Figure 4.6). The effort and expense required to construct this logic is therefore also redundant.



**Figure 4.6**


Different applications developed independently can result in significant amounts of redundant functionality. The applications displayed were delivered with various levels of solution logic that, in some form, already existed.

## 4.2 Problems Solved by Service-Oriented

79

*It's Not as Efficient as it Appears*


Because of the tactical focus on delivering solutions for specific process requirements, the scope of development projects is highly targeted. Therefore, there is the constant perception that business requirements will be fulfilled at the earliest possible time. However, by continually building and rebuilding logic that already exists elsewhere, the process is not as efficient as it could be if the creation of redundant logic could be avoided (Figure 4.7).

**Figure 4.7**

Application A was delivered for a specific set of business requirements. Because a subset of these business requirements had already been fulfilled elsewhere, Application A's delivery scope is larger than it has to be.

*It Bloats an Enterprise*


Each new or extended application adds to the bulk of an IT environment's system inventory (Figure 4.8). The ever-expanding hosting, maintenance, and administration demands can inflate an IT department in budget, resources, and size to the extent that IT becomes a significant drain on the overall organization.

**Figure 4.8**

This simple diagram portrays an enterprise environment containing applications with redundant functionality. The net effect is a larger enterprise.

*It Can Result in Complex Infrastructures and Convolved Enterprise Architectures*

Having to host numerous applications built from different generations of technologies and perhaps even different technology platforms often requires that each will impose unique architectural requirements. The disparity across these “siloeed” applications can lead to a counter-federated environment (Figure 4.9), making it challenging to plan the evolution of an enterprise and scale its infrastructure in response to that evolution.

**Figure 4.9**


Different application environments within the same enterprise can introduce incompatible runtime platforms as indicated by the shaded zones.

*Integration Becomes a Constant Challenge*

Applications built only with the automation of specific business processes in mind are generally not designed to accommodate other interoperability requirements. Making these types of applications share data at some later point results in a jungle of convoluted integration architectures held together mostly through point-to-point patchwork (Figure 4.10) or requiring the introduction of large middleware layers.

## 4.2 Problems Solved by Service-Oriented

81



**Figure 4.10**

A vendor-diverse enterprise can introduce a variety of integration challenges, as expressed by the little lightning bolts that highlight points of concern when trying to bridge proprietary environments.

### The Need for Service-Oriented

After repeated generations of traditional distributed solutions, the severity of the previously described problems has been amplified. This is why service-orientation was conceived. It very much represents an evolutionary state in the history of IT in that it combines successful design elements of past approaches with new design elements that leverage conceptual and technology innovation.

The consistent application of the eight design principles listed earlier results in the widespread proliferation of the corresponding design characteristics:

- increased consistency in how functionality and data is represented
- reduced dependencies between units of solution logic
- reduced awareness of underlying solution logic design and implementation details
- increased opportunities to use a piece of solution logic for multiple purposes
- increased opportunities to combine units of solution logic into different configurations

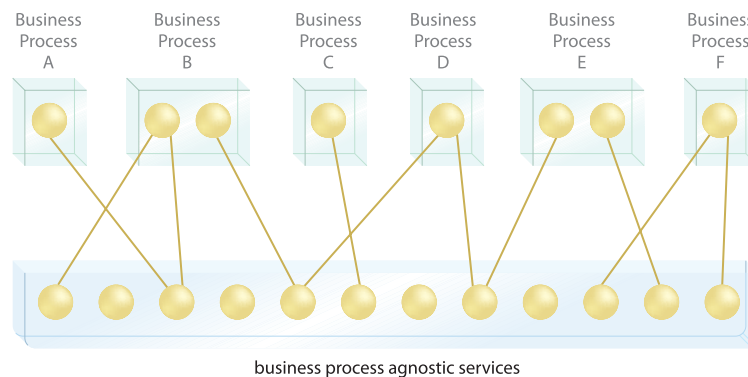


- increased behavioral predictability
- increased availability and scalability
- increased awareness of available solution logic

When these characteristics exist as real parts of implemented services, they establish a common synergy. As a result, the complexion of an enterprise changes as the following distinct qualities are consistently promoted:

#### *Increased Amounts of Agnostic Solution Logic*

Within a service-oriented solution, units of logic (services) encapsulate functionality not specific to any one application or business process (Figure 4.11). These services are therefore classified as agnostic and reusable IT assets.



**Figure 4.11**

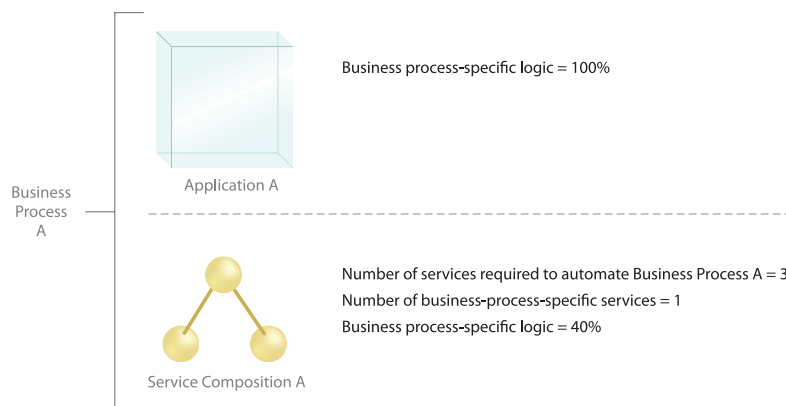
Business processes are automated by a series of business process-specific services (top layer) that share a pool of business process-agnostic services (bottom layer). These layers correspond to the task, entity, and utility service models described in Chapter 3.

#### *Reduced Amounts of Application-Specific Logic*

Increasing the amount of solution logic not specific to any one application or business process decreases the amount of required application-specific logic (Figure 4.12). This blurs the lines between standalone application environments by reducing the overall quantity of standalone applications. (See also the *Service-Oriented and the Concept of “Application”* section later in this chapter.)

## 4.2 Problems Solved by Service-Oriented

83

**Figure 4.12**

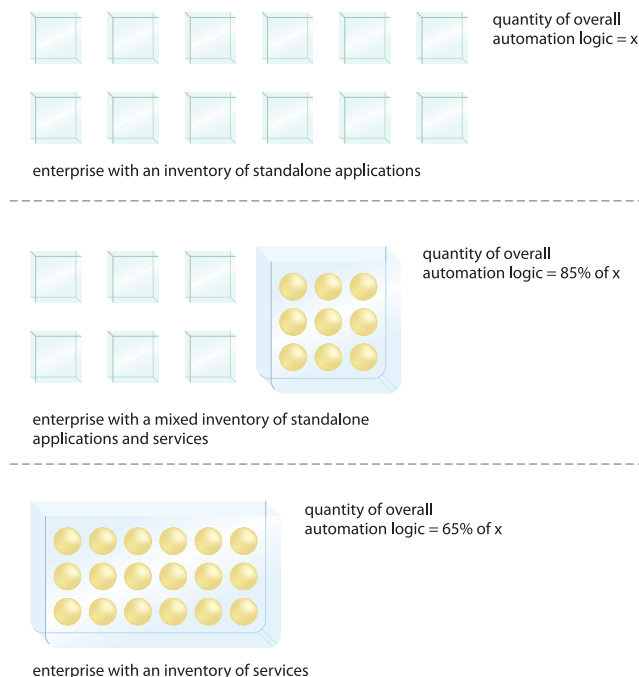
Business Process A can be automated by either Application A or Service Composition A. The delivery of Application A can result in a body of solution logic that is specific to and tailored for the business process. Service Composition A would be designed to automate the process with a combination of agnostic services and 40% of additional logic specific to the business process.

*Reduced Volume of Logic Overall*

The overall quantity of solution logic is reduced because the same solution logic is shared and reused to automate multiple business processes, as shown in Figure 4.13.


**Figure 4.13**

The quantity of solution logic shrinks as an enterprise transitions toward a standardized service inventory comprised of "normalized" services.



*Inherent Interoperability*

Common design characteristics consistently implemented result in solution logic that is naturally aligned. When this carries over to the standardization of service contracts and their underlying data models, a base level of automatic interoperability is achieved across services, as illustrated in Figure 4.14. (See also the *Service-Oriented and the Concept of “Integration”* section later in this chapter.)

**Figure 4.14**

Services from different parts of a service inventory can be combined into new compositions. If these services are designed to be intrinsically interoperable, the effort to assemble them into new composition configurations is significantly reduced.

**SUMMARY OF KEY POINTS**

- The traditional silo-based approach to building applications has been successful at providing tangible benefits and measurable returns on investment.
- This approach has also caused its share of problems, most notably an increase in integration complexity and an increase in the size and administrative burden of IT enterprises.
- Service-orientation establishes a design paradigm that leverages and builds upon previous approaches and proposes a means of avoiding problems associated with silo-based application delivery.

#### 4.3 Challenges Introduced by Service-Orientation

85

### 4.3 Challenges Introduced by Service-Orientation

As much as service-orientation can solve some of the more significant historical problems in IT, its application in the real world can make some serious impositions. It is necessary to be aware of these challenges ahead of time because being prepared is key to overcoming them.

#### Design Complexity

With a constant emphasis on reuse, a significant percentage of a service inventory can ultimately be comprised of agnostic services capable of fulfilling requirements for multiple potential service consumer programs.

Although this can establish a highly normalized and streamlined architecture, it can also introduce an increased level of complexity for both the architecture as well as individual service designs.

Examples include:

- increased performance requirements resulting from the increased reuse of agnostic services
- reliability issues of services at peak concurrent usage times and availability issues of services during off-hours
- single point of failure issues introduced by excessive reuse of agnostic services (and that may require the need for redundant deployments to mitigate risks)
- increased demands on service hosting environments to accommodate autonomy-related preferences
- service contract versioning issues and the impact of potentially redundant service contracts

Design issues such as these can be addressed by a combination of sound technology architecture design, modern vendor runtime platform technology, and the consistent application of service-orientation design principles. Solving service reliability and performance issues in particular are primary goals of those design principles more focused on the underlying service logic, such as Service Autonomy, Service Statelessness, and Service Composability.

### The Need for Design Standards

Design standards can be healthy for an enterprise in that they “pre-solve” problems by making several decisions for architects and developers ahead of time, thereby increasing the consistency and compatibility of solution designs. Their use is required in order to realize the successful propagation of service-orientation.

Although it can be a straight-forward process to create these standards, incorporating them into a (non-standardized) IT culture already set in its ways can be demanding to say the least. The usage of design standards can introduce the need to enforce their compliance, a policing role that can meet with resistance. Additionally, architects and developers sometimes feel that design standards inhibit their creativity and ability to innovate.

A circumstance that tends to aid the large-scale realization of standardization is when the SOA initiative is championed by an executive manager, such as a CIO. When an individual or a governing body has the authority to essentially “lay down the law,” many of these cultural issues resolve themselves more quickly. However, within organizations based on peer-level departmental structures (which are more common in the public sector), the acceptance of design standards may require negotiation and compromise.

The best weapon for overcoming cultural resistance to design standards is communication and education. Those resisting standardization efforts are more likely to become supporters after gaining an appreciation of the strategic significance and ultimate benefits of adopting and respecting the need for design standards.

### Top-Down Requirements

A preferred strategy to delivering services is to first conceptualize a service inventory by defining a blueprint of all planned services, their relationships, boundaries, and individual service models. This approach is very much associated with a top-down delivery strategy in that it can impose a significant amount of up-front analysis effort involving many members of business analysis and technology architecture groups.

Though preferred, achieving a comprehensive blueprint prior to building services is often not feasible. It is common for organizations to face budget and time constraints and tactical priorities that simply won’t permit it. As a result, there are phased and iterative delivery approaches that allow for services to be produced earlier on. These, however, often come with trade-offs in that they can require the service designs to be revisited and revised at a later point. While this can introduce risks associated with

#### 4.3 Challenges Introduced by Service-Orientation

87

the implementation of premature service designs, it is often considered an acceptable compromise.

The principles of service-orientation can be applied to services on an individual basis, allowing a reasonable degree of service-orientation to be achieved regardless of the approach. However, the actual quality of the resulting service designs is typically tied to how much of the top-down analysis work was completed prior to their delivery.

##### BEST PRACTICE

It is recommended that, at minimum, a high-level service inventory blueprint always be defined prior to creating physical service contracts. This establishes an important “broader” perspective in support of service-oriented analysis and service modeling processes and, ultimately, results in stronger and more durable service designs.

#### Counter-Agile Service Delivery in Support of Agile Solution Delivery

Irrespective of the potential top-down efforts needed for some SOA projects, the additional design considerations required to implement a meaningful measure of each of the eight design principles increases both the overall time and cost to deliver service logic.

This may appear contrary to the attention SOA has received for its ability to increase agility. To achieve the state of organizational agility described in Chapter 3 requires that service-orientation already be successfully implemented. This is what establishes an environment in which the delivery of solutions is much more agile.

However, given that it takes more initial effort to design and build services than it does to build a corresponding amount of logic that is not service-oriented, the process of delivering services in support of SOA can actually be *counter-agile*. This can cause issues for an organization that has tactical requirements or needs to be responsive while building a service inventory.

##### BEST PRACTICE

An effective approach, when sufficient resources are available, is to allow SOA initiatives to be delivered alongside existing legacy development and maintenance projects. This way, tactical requirements can continue to be fulfilled by traditional applications while the enterprise works toward a phased transition toward service-oriented computing.

Appendix B provides additional coverage of SOA delivery strategies that address tactical versus strategic service delivery requirements.

### Governance Demands

The eventual existence of one or more service inventories represents the ultimate deliverable of the typical large-scale SOA initiative. A service inventory establishes a powerful reserve of standardized solution logic, a high percentage of which will ideally be classified as agnostic or reusable. Subsequent to their implementation, though, the management and evolution of these agnostic services can be responsible for some of the most profound changes imposed by service-orientation.

In the past, a standalone application was typically developed by a single project team. Members of this team often ended up remaining “attached” to the application for subsequent upgrades, maintenance, and extensions. This ownership model worked because the application’s overall purpose and scope remained focused on the business tasks it was originally built to automate.

The body of solution logic represented by agnostic services, however, is intentionally positioned to *not* belong to any one business process. Although these services may have been delivered by a project team, that same team may not continue to own the service logic as it gets repeatedly utilized by other solutions, processes, and compositions.

Therefore, a special governance structure is required. This can introduce new resources, roles, processes, and even new groups or departments. Ultimately, when these issues are under control and the IT environment itself has successfully adapted to the required changes, the many benefits associated with this new computing platform are there for the taking. However, the process of moving to this new governance model can challenge traditional approaches and demand time, expense, and a great deal of patience.

### SUMMARY OF KEY POINTS

- Applying service-orientation on a broad scale can introduce increased design complexity and the need for a consistent level of standardization.
- The construction of services can be expensive and time-consuming, introducing a more burdensome project delivery lifecycle, further compounded by some of the common top-down analysis requirements that may need to be in place before services can be built.
- Service inventory governance requirements can impose significant changes that can shake up the organizational structure of an IT department.

## 4.4 Additional Considerations

To supplement the benefits and challenges just covered, this section discusses some further aspects of service-orientation.

### It Is Not a Revolutionary Paradigm

Service-orientation is not a brand new paradigm that aims to replace all that preceded it. It, in fact, incorporates and builds upon proven and successful elements from past paradigms and combines these with design approaches shaped to leverage recent technology innovations.

This is why we do not refer to SOA as a revolutionary model in the history of IT. It is simply the next stage in an evolutionary cycle that began with the application of modularity on a small scale (by organizing simple programming routines into shared modules for example) and has now spread to the potential modularization of the enterprise.

### Enterprise-wide Standardization Is Not Required

There is a common misperception that unless design standardization is achieved globally throughout the entire enterprise, SOA will not succeed. Although design standardization is a critical success factor for SOA projects that is *ideally* achieved across an enterprise, it only needs to be realized to a meaningful extent for service-orientation to result in strategic benefit.

For example, service-orientation emphasizes the need for standardizing service data models to avoid unnecessary data transformation and other problematic issues that can compromise interoperability. The extent to which data model standardization is achieved determines the extent to which these problems will be avoided.

The goal is not always to eliminate problems entirely because that can be an unrealistic objective, especially in larger enterprises. Therefore, the goal is sometimes to just minimize problems by taking special considerations into account during service design.

In support of this approach, design patterns exist for organizing the division of an enterprise into more manageable domains. Data standardization is generally more easily attained within each domain, and transformation is then only required when exchanging data across these domains. Even though this does not achieve a global data model, it can still help establish a very meaningful level of interoperability.



### Reuse Is Not an Absolute Requirement

Increasing reusability of solution logic is a fundamental goal of service-orientation, and reuse is clearly one of the most associated benefits of SOA. As a result, organizations that have had limited success with past reuse initiatives, or with concerns that significant amounts of reuse cannot be achieved within their enterprise, are often hesitant about SOA in general.

While reuse, especially over time, can be one of the most rewarding parts of investing in SOA, it is not the sole primary benefit. Perhaps even more fundamental to service-orientation than promoting reuse is fostering interoperability. Enabling an enterprise to connect previously disparate systems or to make interconnectivity an intrinsic quality of new solution logic is extremely powerful.

You could ignore the principle of Service Reusability in service designs and still achieve significant returns on investment based solely on raising the level of enterprise-wide interoperability.

#### NOTE

One could argue that reuse and interoperability are very closely related in that if two services are interoperable, there is always the opportunity for reuse. However, traditional perspectives of reusable solution logic focus on the nature of the logic itself. A service that is designed to be specifically agnostic to business processes and cross-cutting to address multiple concerns will have a particular functional context associated with it. Therefore, reuse can be seen as a separate design characteristic that relies and builds upon interoperability. See Chapter 9 for more details.

#### SUMMARY OF KEY POINTS

- Service-orientation has deep roots in several past computing platforms and design approaches, and is therefore not considered a revolutionary design paradigm.
- Global standardization within an enterprise is not a requirement for creating service-oriented enterprises because individual service inventories can be established (and separately standardized) within different enterprise domains.
- Although fundamental to much of service-orientation, if reusability were to be omitted as a design characteristic, significant interoperability-related benefit would still be attainable.

## 4.5 Effects of Service-Orientation on the Enterprise

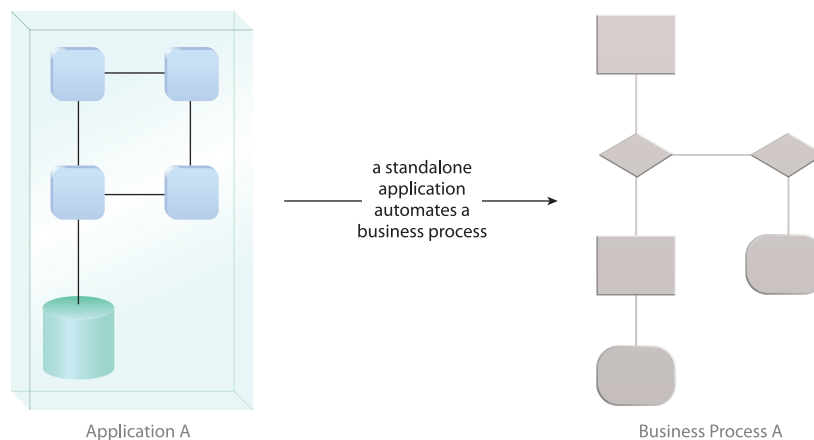
There are good reasons to have high expectations from the service-orientation paradigm. But, at the same time, there is much to learn and understand before it can be successfully applied. The following sections explore some of the more common examples.

### Service-Orientation and the Concept of “Application”

Having just stated that reuse is not an absolute requirement, it is important to acknowledge the fact that service-orientation does place an unprecedented emphasis on reuse. By establishing a service inventory with a high percentage of reusable and agnostic services, we are now positioning those services as the primary (or only) means by which the solution logic they represent can and should be accessed.

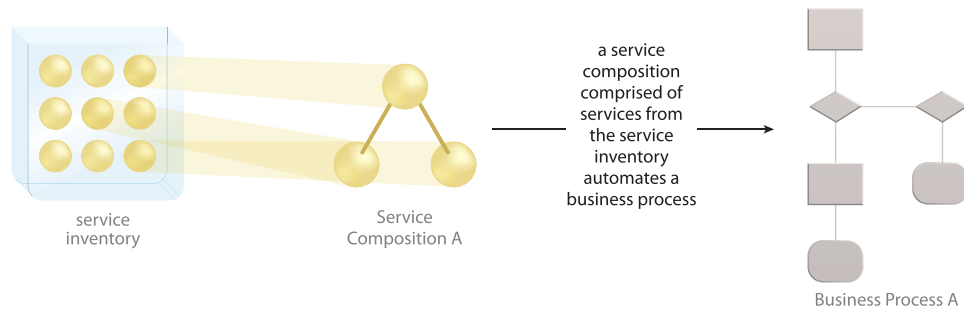
As a result, we make a very deliberate move away from the silos in which applications previously existed. Because we want to share reusable logic whenever possible, we automate existing, new, and augmented business processes through service composition. This results in a shift where more and more business requirements are fulfilled not by building or extending applications, but by simply composing existing services into new composition configurations.

When compositions become more common, the traditional concept of an application, a system, or a solution actually begins to fade, along with the silos that contain them. Applications no longer consist of self-contained bodies of programming logic responsible for automating a specific set of tasks (Figure 4.15). What was an application is now just another service composition. And it's a composition made up of services that very likely participate in other compositions (Figure 4.16).



**Figure 4.15**

The traditional application, delivered to automate specific business process logic.

**Figure 4.16**

The service composition, intended to fulfill the role of the traditional application by leveraging agnostic and non-agnostic services from a service inventory. This essentially establishes a “composite application.”

An application in this environment loses its individuality. One could argue that a service-oriented application actually does not exist because it is, in fact, just one of many service compositions. However, upon closer reflection, we can see that some of the services are actually not business process-agnostic. The task service, for example, intentionally represents logic that is dedicated to the automation of just one business task and therefore is not necessarily reusable.

What this indicates is that non-agnostic services can still be associated with the notion of an application. However, within service-oriented computing, the meaning of this term can change to reflect the fact that a potentially large portion of the application logic is no longer exclusive to the application.

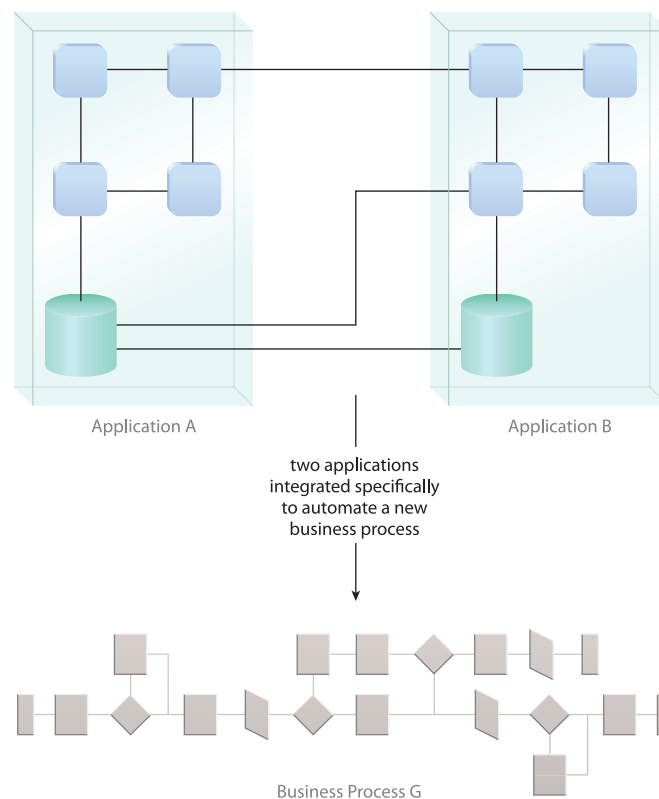
### Service-Orientation and the Concept of “Integration”

When we revisit the idea of a service inventory consisting of services that have, as per our service-orientation principles, been shaped into standardized and (for the most part) reusable units of solution logic, we can see that this can challenge the traditional perception of “integration.”

In the past, integrating something implied connecting two or more applications or programs that may or may not have been compatible (Figure 4.17). Perhaps they were based on different technology platforms or maybe they were never designed to connect with anything outside of their own internal boundary. The increasing need to hook up disparate pieces of software to establish a reliable level of data exchange is what turned integration into an important, high profile part of the IT industry.

## 4.5 Effects of Service-Oriented Architecture on the Enterprise

93

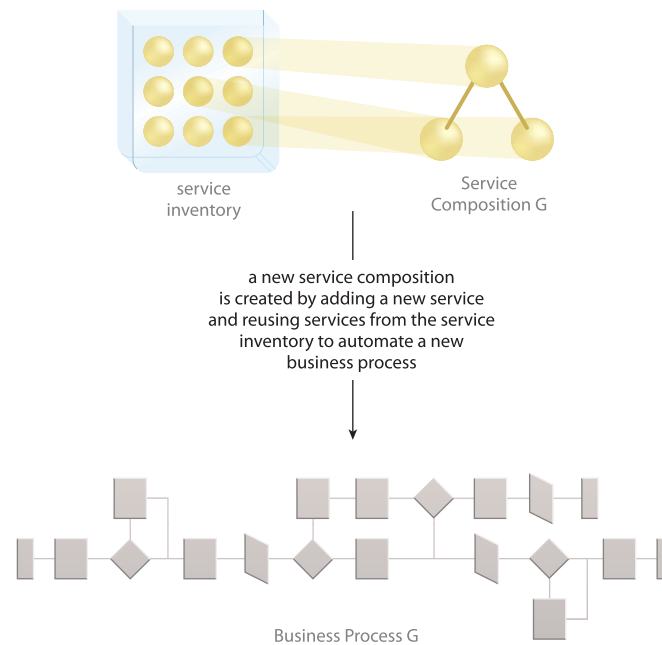
**Figure 4.17**

The traditional integration architecture, comprised of two or more applications connected in different ways to fulfill a new set of automation requirements (as dictated by the new Business Process G).

Services designed to be “intrinsically interoperable” are built with the full awareness that they will need to interact with a potentially large range of service consumers, most of which will be unknown at the time of their initial delivery. If a significant part of our enterprise solution logic is represented by an inventory of intrinsically interoperable services, it empowers us with the freedom to mix and match these services into infinite composition configurations to fulfill whatever automation requirements come our way.

As a result, the concept of integration begins to fade. Exchanging data between different units of solution logic becomes a natural and secondary design characteristic (Figure 4.18). Again, though, this is something that can only transpire when a substantial percentage of an organization’s solution logic is represented by a quality service inventory.

While working toward achieving this environment, there will likely be many requirements for traditional integration between existing legacy systems and also between legacy systems and these services.



**Figure 4.18**

A new combination of services is composed together to fulfill the role of traditional integrated applications.

### The Service Composition

Applications, integrated applications, solutions, systems, all of these terms and what they have traditionally represented can be directly associated with the service composition (Figure 4.19). However, given the fact that many SOA implementations consist of a mixture of legacy environments and services, these terms are sure to survive for quite some time.


In fact, as SOA transition initiatives continue to progress within an enterprise, it can be helpful to make a clear distinction between a traditional application (one which may reside alongside an SOA implementation or which may be actually encapsulated by a service) and the service compositions that eventually become more commonplace.

## 4.5 Effects of Service-Orientation on the Enterprise

95

**Figure 4.19**

A service-oriented solution, application, or system is the equivalent of a service composition. If we were to build an enterprise-wide SOA from the ground up, it would likely be comprised of numerous service compositions capable of fulfilling the traditional roles associated with these terms.

**Application, Integration, and Enterprise Architectures**

Because applications have existed for as long as IT, when technology architecture as a profession and perspective within the enterprise came about, it made perfect sense to have separate architectural views dedicated to individual applications, integrated applications, and the enterprise as a whole.

When standardizing on service-orientation, the manner in which we document technology architecture is also in for a change. The enterprise-level perspective becomes predominant as it represents a master view of the service inventory. It can still encompass the traditional parts of a formal architecture, including conceptual views, physical views, and supporting technologies and governance platforms—but all these views are likely to now become associated with the service inventory.

A new type of technical specification that gains prominence in service-oriented enterprise initiatives is the *service composition architecture*. Even though we talk about the simplicity of combining services into new composition configurations on demand, it is by no means an easy process. It is a design exercise that requires the detailed documentation of the planned composition architecture.

For example, each service needs to be assessed as to its competency to fulfill its role as a composition member, and foreseeable service activity scenarios need to be mapped out.

Message designs, messaging routes, exception handling, cross-service transactions, policies, and many more considerations go into making a composition capable of automating its designated business process.

#### BEST PRACTICE

Although the structure and content of traditional application architecture specifications are augmented when documenting composition architectures, there can still be a natural tendency to refer to these documents as architecture specifications for applications.

While an organization is undergoing a transition toward SOA, it can be helpful to make a clear distinction between an application consisting of a service composition and traditional, standalone or legacy applications.

One approach is to consistently qualify the term “application.” For example, it can be prefixed with “service-oriented,” “composite,” “standalone,” or “legacy.” Another option is to simply limit the use of the term “application” to refer to non-service-composed solutions only.

Furthermore, a composed service encapsulating a legacy application can be documented in separate specifications: a composition architecture specification that identifies the service and points to an application architecture specification that defines the corresponding application.

#### SUMMARY OF KEY POINTS

- The traditional concept of an application can change as more agnostic services become established parts of the enterprise.
- The traditional concept of integration can change as the proliferation of standardized, intrinsic interoperable services increases.
- Architectural views of the enterprise shift in response to the adoption of service-orientation. Principally, the enterprise perspective becomes increasingly prominent.


## 4.6 Origins and Influences of Service-Orientation

It is often said that the best way to understand something is to gain knowledge of its history. Service-orientation, by no means, is a design paradigm that just came out of nowhere. It is very much a representation of the evolution of IT and therefore has many

#### 4.6 Origins and Influences of Service-Oriented

97

roots in past paradigms and technologies (Figure 4.20). At the same time, it is still in a state of evolution itself and therefore remains subject to influences from on-going trends and movements.



**Figure 4.20**

The primary influences of service-orientation also highlight its many origins.

The sections that follow describe some of the more prominent origins and thereby help clarify how service-orientation can relate to and even help further some of the goals from past paradigms.

##### Object-Oriented

In the 1990s the IT community embraced a design philosophy that would lead the way in defining how distributed solutions were to be built. This paradigm was object-orientation, and it came with its own set of principles, the application of which helped ensure consistency across numerous environments. These principles defined a specific type of relationship between units of solution logic classified as objects, which resulted in a predictable set of dynamics that ran through entire solutions.

Service-orientation is frequently compared to object-orientation, and rightly so. The principles and patterns behind object-oriented analysis and design represent one of the most significant sources of inspiration for this paradigm.

In fact, a subset of service-orientation principles (Service Reusability, Service Abstraction, and Service Composability, for example) can be traced back to object-oriented counterparts. What distinguishes service-orientation, though, are the parts of the object-oriented school of thought that were left out and the other principles that were added. See Chapter 14 for a comparative analysis of principles and concepts associated with these two design approaches.



### Web Services

Even though service-orientation as a paradigm and SOA as a technology architecture are each implementation-neutral, their association with Web services has become commonplace—so much so that the primary SOA vendors have shaped their respective platforms around the utilization of Web services technology.

Although service-orientation remains a fully abstract paradigm, it is one that has historically been influenced by the SOA platforms and roadmaps produced by these vendors. As a result, the Web services framework has influenced and promoted several service-orientation principles, including Service Abstraction, Service Loose Coupling, and Service Composability.

### Business Process Management (BPM)

BPM places a significant emphasis on business processes within the enterprise both in terms of streamlining process logic to improve efficiency and also to establish processes that are adaptable and extensible so that they can be augmented in response to business change.

The business process layer represents a core part of any service-oriented architecture. From a composition perspective, it usually assumes the role of the parent service composition controller. The advent of orchestration technology reaffirmed this role from an implementation perspective.

A primary goal of service-orientation is to establish a highly agile automation environment fully capable of adapting to change. This goal can be realized by abstracting business process logic into its own layer, thereby alleviating other services from having to repeatedly embed process logic.

While service-orientation itself is not as concerned with business process reengineering, it fully supports process optimization as a primary source of change for which services can be recomposed.

### Enterprise Application Integration (EAI)

Integration became a primary focal point in the late 90's, and many organizations were ill prepared for it. Numerous systems were built with little thought given to how data could be shared outside of the system boundary. As a result, point-to-point integration

#### 4.6 Origins and Influences of Service-Oriented

99

channels were often created when data sharing requirements emerged. This led to well known problems associated with a lack of stability, extensibility, and inadequate interoperability frameworks.

EAI platforms introduced middleware that allowed for the abstraction of proprietary applications through the use of adapters, brokers, and orchestration engines. The resulting integration architectures were, in fact, more robust and extensible. However, they also became notorious for being overwhelmingly complex and expensive, as well as requiring long-term commitments to the middleware vendor's platform and roadmap.

The advent of the open Web services framework and its ability to fully abstract proprietary technology changed the face of integration middleware. Vendor ties could be broken by investing in mobile services as opposed to proprietary platforms, and organizations gained more control over the evolution of their integration architectures.

Several innovations that became popularized during the EAI era were recognized as being useful to the overall goals associated with building SOA using Web services. One example is the broker component, which allows for services using different schemas representing the same type of data to still communicate through runtime transformation. The other is the orchestration engine, which can actually be positioned to represent an entire service layer within larger SOA implementations. These parts of the EAI platform support several service-orientation principles, including Service Abstraction, Service Statelessness, Service Loose Coupling, and Service Composability.

#### **Aspect-Oriented Programming (AOP)**

A primary goal of AOP is to approach the separation of concerns with the intent of identifying specific concerns that are common to multiple applications or automation scenarios. These concerns are then classified as "cross-cutting," and the corresponding solution logic developed for cross-cutting concerns becomes naturally reusable.

Aspect-orientation emerged from object-orientation by building on the original goals of establishing reusable objects. Although not a primary influential factor of service-orientation, AOP does demonstrate a common goal in emphasizing the importance of investing in units of solution logic that are agnostic to business processes and applications and therefore highly reusable. It further promotes role-based development, allowing developers with different areas of expertise to collaborate.

**NOTE**

The actual events and timeline associated with the emergence of SOA are documented in Chapter 4 of the book *Service-Oriented Architecture: Concepts, Technology, and Design*.

**SUMMARY OF KEY POINTS**

- Service-orientation represents a design paradigm that has its roots in several origins. It emphasizes successful and proven approaches and supplements them with new principles that leverage recent conceptual and technology innovation.
- Service-orientation, as a design paradigm, is comparable with object-orientation. In fact, several key object-oriented principles have persisted in service-orientation.
- The Web services technology platform is primarily responsible for the popularity of SOA and is therefore also a significant influence in service-orientation. Conversely, the rise of service-oriented computing has repositioned and formalized the Web services technology set from its original incarnation.

**4.7 CASE STUDY BACKGROUND**

Cutit's immediate priority is to streamline their internal supply chain process. The order process in particular needs to be supported by the planned services so that orders and back-orders can be fulfilled as soon as possible.

Below are brief descriptions of the service candidates shown in Figure 4.21 in relation to how they inter-relate based on their entity-centric functional contexts:

- Everything originates with the manufacturing of chain blades in the Cutit lab, which requires the use of specific *materials* that are applied as per predefined *formulas*.
- The assembly of *chains* results in products being added to their overall *inventory*.
- *Saws* and *kits* are items Cutit purchases from different manufacturers to complement their chain models.
- *Notifications* need to be issued when stock levels fall below certain levels or if other urgent conditions occur.

## 4.7 Case Study Background

101

- Finally, a periodic *patent sweep* is conducted to search for recently issued patents with similarities to Cutit's planned chain designs.

Note that all services shown are entity services, with the exception of Patent Sweep and Notifications, which are based on the utility service model. A task service is added in Part II.

**Figure 4.21**

The initial set of services planned to support the following types of processes: keeping track of orders and back-orders, chain manufacturing, tracking required manufacturing materials, and inventory management of manufactured and purchased products. All of the displayed services are based on the entity service model, except for the bottom two, which are utility services.

# About the Author

Thomas Erl is the world's top-selling SOA author, the Series Editor of the *Prentice Hall Service-Oriented Computing Series from Thomas Erl*, and Editor of *The SOA Magazine*.

With over 65,000 copies in print, his first two books, *Service-Oriented Architecture: A Field Guide to Integrating XML and Web Services* and *Service-Oriented Architecture: Concepts, Technology, and Design* have become international bestsellers and have been translated into several languages. Books by Thomas Erl have been formally reviewed and endorsed by senior members of major software organizations, including IBM, Sun, Microsoft, Oracle, BEA, HP, SAP, Google, and Intel.

Thomas is also the founder of SOA Systems Inc. ([www.soasystems.com](http://www.soasystems.com)), a company specializing in SOA training and strategic consulting services with a vendor-agnostic focus. Through his work with standards organizations and independent research efforts, Thomas has made significant contributions to the SOA industry, most notably in the areas of service-orientation and SOA methodology.

Thomas is a speaker and instructor for private and public events, and has delivered many workshops and keynote speeches. For a current list of his workshops, seminars, and courses, see [www.soatraining.com](http://www.soatraining.com).

Papers and articles written by Thomas have been published in numerous industry trade magazines and Web sites, and he has delivered Webcasts and interviews for many publications, including the *Wall Street Journal*.

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

## THE PRENTICE HALL SERVICE-ORIENTED COMPUTING SERIES FROM THOMAS ERL


**Service-Oriented Architecture:  
A Field Guide to Integrating XML and Web Services**


ISBN 0131428985

This top-selling field guide offers expert advice for incorporating XML and Web services technologies within service-oriented integration architectures.


**Service-Oriented Architecture:  
Concepts, Technology, and Design**

ISBN 0131858580

Widely regarded as the definitive “how-to” guide for SOA, this best-selling book presents a comprehensive end-to-end tutorial that provides step-by-step instructions for modeling and designing service-oriented solutions from the ground up.


**SOA: Principles of Service Design**

ISBN 0132344823

Published with over 240 color illustrations, this hands-on guide contains practical, comprehensive, and in-depth coverage of service engineering techniques and the service-orientation design paradigm. Proven design principles are documented to help maximize the strategic benefit potential of SOA.


**SOA: Design Patterns**

ISBN 0136135161

Software design patterns have emerged as a powerful means of avoiding and overcoming common design problems and challenges. This new book presents a formal catalog of design patterns specifically for SOA and service-orientation. All patterns are documented using full-color illustrations and further supplemented with case study examples.

Several additional series titles are currently in development and will be released soon.  
For more information about any of the books in this series, visit [www.soabooks.com](http://www.soabooks.com).

