




SEI SERIES IN SOFTWARE ENGINEERING

Documenting Software Architectures



Views
and
Beyond

SECOND EDITION

Paul Clements • Felix Bachmann • Len Bass
David Garlan • James Ivers • Reed Little
Paulo Merson • Robert Nord • Judith Stafford



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 Addison-Wesley was aware of a trademark claim, the designations have been printed with initial capital letters or in all capitals.

CMM, CMMI, Capability Maturity Model, Capability Maturity Modeling, Carnegie Mellon, CERT, and CERT Coordination Center are registered in the U.S. Patent and Trademark Office by Carnegie Mellon University.

ATAM; Architecture Tradeoff Analysis Method; CMM Integration; COTS Usage-Risk Evaluation; CURE; EPIC; Evolutionary Process for Integrating COTS Based Systems; Framework for Software Product Line Practice; IDEAL; Interim Profile; OAR; OCTAVE; Operationally Critical Threat, Asset, and Vulnerability Evaluation; Options Analysis for Reengineering; Personal Software Process; PLTP; Product Line Technical Probe; PSP; SCAMPI; SCAMPI Lead Appraiser; SCAMPI Lead Assessor; SCE; SEI; SEPG; Team Software Process; and TSP are service marks of Carnegie Mellon University.

IEEE Std 1471 is a trademark of the Institute of Electrical and Electronics Engineers, Inc.

Special permission to reproduce portions of the following is granted by the Software Engineering Institute:

- Robert L. Nord, Paul C. Clements, David Emery, and Rich Hilliard, "A Structured Approach for Reviewing Architecture Documentation" (CMU/SEI-2009-TN-030). Copyright © 2009 by Carnegie Mellon University.
- Felix Bachmann, Len Bass, Paul Clements, David Garlan, James Ivers, Reed Little, Robert Nord, and Judith Stafford, "Documenting Software Architecture: Documenting Behavior" (CMU/SEI-2002-TN-001). Copyright © 2002 by Carnegie Mellon University.
- Felix Bachmann, Len Bass, Paul Clements, David Garlan, James Ivers, Reed Little, Robert Nord, and Judy Stafford, "Documenting Software Architectures: Organization of Documentation Package" (CMU/SEI-2001-TN-010). Copyright © 2001 by Carnegie Mellon University.
- Felix Bachmann, Len Bass, Jeromy Carriere, Paul Clements, David Garlan, James Ivers, Robert Nord, and Reed Little, "Software Architecture Documentation in Practice: Documenting Architectural Layers" (CMU/SEI-2000-SR-004). Copyright © 2000 by Carnegie Mellon University.
- Felix Bachmann, Len Bass, Paul Clements, David Garlan, James Ivers, Robert Nord, Reed Little, and Judith Stafford, "Software Architecture Documentation in Practice: Documenting Software Interfaces" (CMU/SEI-2002-TN-015). Copyright © 2002 by Carnegie Mellon University.

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@pearsoned.com

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

International Sales
international@pearsoned.com

Visit us on the Web: informit.com/aw

Library of Congress Cataloging-in-Publication Data

Documenting software architectures : views and beyond / Paul Clements
... [et al.]. — 2nd ed.
p. cm.

Includes bibliographical references and index.
ISBN 978-0-321-55268-6 (hardcover : alk. paper)

1. Computer architecture. 2. Software documentation. I. Clements, Paul, 1955– II. Title.
QA76.9.A73D63 2010
005.1'5—dc22

2010024318

Copyright © 2011 Pearson Education, Inc.

All rights reserved. Printed in the United States of America. This publication is protected by copyright, and permission must be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise. For information regarding permissions, write to:

Pearson Education, Inc.
Rights and Contracts Department
501 Boylston Street, Suite 900
Boston, MA 02116
Fax: (617) 671-3447

ISBN-13: 978-0-321-55268-6

ISBN-10: 0-321-55268-7

Text printed in the United States on Recycled paper at Courier in Westford, Massachusetts.
First printing, October 2010

Contents

<i>About the Cover</i>	xxi
<i>Foreword to the Second Edition</i>	xxiii
<i>Foreword to the First Edition</i>	xxv
<i>Preface</i>	xxix
<i>Acknowledgments</i>	xxxiii
<i>Reader's Guide</i>	xxxv

Prologue: Software Architectures and Documentation **1**

The prologue establishes the necessary concepts and vocabulary for the remainder of the book. It discusses how software architecture documentation is used and why it is important. It defines the concepts that provide the foundation of the book's approach to documentation. It also contains seven basic rules for sound documentation.

P.1	A Short Overview of Software Architecture	1
P.1.1	Overview	1
P.1.2	Architecture and Quality Attributes	2
	Coming to Terms: What Is Software Architecture?	3
	Perspectives: What's the Difference Between Architecture and Design?	6
P.2	A Short Overview of Architecture Documentation	9
P.2.1	Why Document Software Architecture?	9
	Coming to Terms: Specification, Representation, Description, Documentation	10
P.2.2	Uses and Audiences for Architecture Documentation	12
P.2.3	Architecture Documentation and Quality Attributes	17
P.2.4	Economics of Architecture Documentation	18
P.2.5	The Views and Beyond "Method"	19
P.2.6	Views and Beyond in an Agile Environment	20
P.2.7	Architectures That Change Faster Than You Can Document Them	20

P.3	Architecture Views	22
	Coming to Terms: A Short History of Architecture Views	23
P.4	Architecture Styles	25
	P.4.1 Three Categories of Styles	29
	Coming to Terms: Module, Component	29
	Coming to Terms: “Architecture Style” and “Architecture Pattern”	32
P.5	Seven Rules for Sound Documentation	36
	Perspectives: Beware Notations Everyone “Just Knows”	38
	Perspectives: Quivering at Arrows	41
P.6	Summary Checklist	45
P.7	Discussion Questions	46
P.8	For Further Reading	47

Part I A Collection of Software Architecture Styles 49

Part I introduces the basic tools for software architecture documentation: architecture styles. A style is a specialization of element and relationship types, together with constraints on how they may be used. By identifying element and relationship types, styles identify the architecture structures that architects design to achieve the system’s quality and behavioral goals. There are three fundamental kinds of structures: module structures, component-and-connector structures, and allocation structures. Within each category reside a number of architecture styles. The introduction to Part I includes a brief catalog of the styles that are described in Chapters 1–5.

I.1	Three Categories of Styles	49
I.2	Style Guides: A Standard Organization for Explaining a Style	50
I.3	Choosing Which Element and Relation Properties to Document	52
I.4	Notations for Architecture Views	53
I.5	Examples	54

Chapter 1 Module Views 55

A module is an implementation unit of software that provides a coherent unit of functionality. Modules form the basis of many standard architecture views. This chapter defines modules and outlines the information required for documenting module views.

1.1	Overview	55
1.2	Elements, Relations, and Properties of Module Views	56
	1.2.1 Elements	56
	1.2.2 Relations	57
	1.2.3 Properties	57
1.3	What Module Views Are For	59

1.4	Notations for Module Views	60
1.4.1	Informal Notations	60
1.4.2	Unified Modeling Language	61
1.4.3	Dependency Structure Matrix	62
1.4.4	Entity-Relationship Diagram	62
1.5	Relation to Other Views	63
1.6	Summary Checklist	63
1.7	Discussion Questions	64
1.8	For Further Reading	64

Chapter 2 A Tour of Some Module Styles 65

This chapter introduces some common and important styles in the module category. Each style is presented in terms of how it specializes the overall elements and relations found in module styles.

2.1	Decomposition Style	65
2.1.1	Overview	65
2.1.2	Elements, Relations, and Properties	66
2.1.3	What the Decomposition Style Is For	67
2.1.4	Notations for the Decomposition Style	67
2.1.5	Relation to Other Styles	68
2.1.6	Examples Using the Decomposition Style	69
	Coming to Terms: Subsystem	73
2.2	Uses Style	74
2.2.1	Overview	74
2.2.2	Elements, Relations, and Properties	75
2.2.3	What the Uses Style Is For	75
2.2.4	Notations for the Uses Style	76
2.2.5	Relation to Other Styles	79
2.2.6	Examples Showing the Uses Style	79
	Coming to Terms: Uses	81
2.3	Generalization Style	82
2.3.1	Overview	82
2.3.2	Elements, Relations, and Properties	83
2.3.3	What the Generalization Style Is For	84
2.3.4	Notations for the Generalization Style	84
2.3.5	Relation to Other Styles	84
2.3.6	Examples Using the Generalization Style	85
2.4	Layered Style	87
2.4.1	Overview	87
2.4.2	Elements, Relations, and Properties	89
2.4.3	What the Layered Style Is For	90
2.4.4	Notations for the Layered Style	92
2.4.5	Relation to Other Styles	96

2.4.6	Examples Using the Layered Style	97
	Coming to Terms: Virtual Machines	99
	Perspectives: Calling Higher Layers	100
	Perspectives: Using a DSM to Maintain a Layered Architecture	101
2.5	Aspects Style	104
2.5.1	Overview	104
2.5.2	Elements, Relations, and Properties	104
2.5.3	What the Aspects Style Is For	105
2.5.4	Notations for the Aspects Style	105
2.5.5	Relation to Other Styles	106
2.5.6	Examples Using the Aspects Style	106
	Coming to Terms: Aspect-Oriented Programming	107
2.6	Data Model	109
2.6.1	Overview	109
2.6.2	Elements, Relations, and Properties	111
2.6.3	What the Data Model Is For	114
2.6.4	Notations for the Data Model Style	116
2.6.5	Relations to Other Styles	117
2.6.6	Examples	118
	Coming to Terms: Entity	118
2.7	Summary Checklist	120
2.8	Discussion Questions	120
2.9	For Further Reading	121

Chapter 3 Component-and-Connector Views 123

Component-and-connector views represent units of execution plus the pathways and protocols of their interaction. This chapter defines components and connectors and describes the rules for documenting them.

3.1	Overview	123
3.2	Elements, Relations, and Properties of C&C Views	126
3.2.1	Elements	127
3.2.2	Component-and-Connector Types and Instances	129
3.2.3	Relations	131
3.2.4	Properties	133
	Perspectives: Are Complex Connectors Necessary?	135
3.3	What C&C Views Are For	136
	Perspectives: Choosing Connector Abstractions	137
3.4	Notations for C&C Views	139
3.4.1	Informal Notations	139
3.4.2	Formal Notations	139
3.4.3	Semiformal Notations: UML	139
	Perspectives: Data Flow and Control Flow Models	146

3.5	Relation to Other Kinds of Views	148
3.6	Summary Checklist	150
3.7	Discussion Questions	151
3.8	For Further Reading	152

Chapter 4 A Tour of Some Component-and-Connector Styles 155

This chapter introduces some important component-and-connector (C&C) styles. The chapter describes how each style is a specialization of the generic elements and relations of C&C styles, discusses what makes each style useful, and explains how each style is documented.

4.1	An Introduction to C&C Styles	155
4.2	Data Flow Styles	157
4.2.1	Pipe-and-Filter Style	158
4.3	Call-Return Styles	161
4.3.1	Client-Server Style	162
4.3.2	Peer-to-Peer Style	166
4.3.3	Service-Oriented Architecture Style	169
4.4	Event-Based Styles	172
4.4.1	Publish-Subscribe Style	174
4.5	Repository Styles	178
4.5.1	Shared-Data Style	178
4.6	Crosscutting Issues for C&C Styles	182
4.6.1	Communicating Processes	182
4.6.2	Tiers	183
4.6.3	Dynamic Creation and Destruction	184
4.7	Summary Checklist	185
4.8	Discussion Questions	186
4.9	For Further Reading	187

Chapter 5 Allocation Views and a Tour of Some Allocation Styles 189

Software architects are often obliged to document nonsoftware structures and show how the software designs are mapped to the structures: the computing environment in which their software will run, the organizational environment in which it will be developed, and so on. This chapter introduces the allocation view category, which is used to express the allocation of software elements to nonsoftware structures, and three major allocation styles.

5.1	Overview	189
5.2	Deployment Style	191
5.2.1	Overview	191
5.2.2	Elements, Relations, and Properties	192
5.2.3	What the Deployment Style Is For	194

5.2.4	Notation for the Deployment Style	196
5.2.5	Relation to Other Styles	197
5.3	Install Style	198
5.3.1	Overview	198
5.3.2	Elements, Relations, and Properties	199
5.3.3	What the Install Style Is For	200
5.3.4	Notations for the Install Style	200
5.3.5	Relation to Other Styles	202
5.4	Work Assignment Style	202
5.4.1	Overview	202
5.4.2	Elements, Relations, and Properties	202
5.4.3	What a Work Assignment Style Is For	203
5.4.4	Notations for the Work Assignment Style	203
5.4.5	Relation to Other Styles	204
	Perspectives: Why Is a Work Assignment View Architectural?	205
5.5	Other Allocation Styles	206
	Perspectives: Coordination Views	209
5.6	Summary Checklist	213
5.7	Discussion Questions	213
5.8	For Further Reading	214

Part II Beyond Structure: Completing the Documentation 215

Part II concentrates on the rest of the information an architect should include in architecture documentation, such as context diagrams, variation points, interfaces, and software behavior.

Chapter 6 Beyond the Basics 217

This chapter introduces documentation approaches to handle some special architecture issues and situations, such as breaking a view into chunks, documenting context and variation points, and combining views.

6.1	Refinement	218
6.1.1	Decomposition Refinement	218
6.1.2	Implementation Refinement	219
6.1.3	Spectrum of Design	220
6.1.4	Style Specialization	221
6.2	Descriptive Completeness	222
6.3	Documenting Context Diagrams	225
6.3.1	Create Context Diagrams Using the Vocabulary of the View	226
6.3.2	Content of a Context Diagram	228
6.3.3	Context Diagrams and Other Supporting Documentation	229
6.3.4	Notations for Context Diagrams	229

6.4	Documenting Variation Points	231
6.4.1	What Are Variation Points?	231
6.4.2	Variation Mechanisms	232
	Coming to Terms: Product-Line Architectures	234
6.4.3	Dynamism and Dynamic Architectures	234
6.4.4	Documenting Variation Points	235
6.5	Documenting Architectural Decisions	239
6.5.1	Why Document Architectural Decisions?	239
6.5.2	A Template for Documenting Architectural Decisions	239
6.5.3	Documenting Alternatives	242
6.5.4	Which Decisions to Document	242
	Perspectives: "It may sound like a lot of effort to do this, but here's how we do it in the trenches."	244
6.5.5	The Payback for Documenting Architectural Decisions	245
	Perspectives: From Documenting Architectures to Architecting As Decision Making	246
	Perspectives: An Ontology of Architecture Decisions	247
6.6	Combining Views	250
6.6.1	Types of Associations Between Views	251
6.6.2	Combined Views	252
6.6.3	When to Combine Views	254
6.6.4	Examples of Combined Views	255
6.7	Summary Checklist	258
6.8	Discussion Questions	259
6.9	For Further Reading	260

Chapter 7 Documenting Software Interfaces 261

The interfaces of the elements are critical parts of any architecture, and documenting them is an important responsibility for the architect. This chapter tells you how to specify an interface.

7.1	Overview	261
	Coming to Terms: Provided vs. Required Interfaces	264
7.2	Interface Documentation	265
7.2.1	Showing the Existence of Interfaces in Diagrams	268
7.3	A Standard Organization for Interface Documentation	271
	Coming to Terms: Error Handling	277
7.4	Stakeholders of Interface Documentation	278
7.5	Conveying Syntactic Information	279
7.6	Conveying Semantic Information	280
	Coming to Terms: Signature, Interface, API	280

7.7	Examples of Interface Documentation	281
7.7.1	Zip Component API	281
7.7.2	Interface to a SOAP Web Service	285
7.8	Summary Checklist	285
7.9	Discussion Questions	286
7.10	For Further Reading	286
Chapter 8 Documenting Behavior		289
Documenting behavior is an essential counterpoint to documenting structure. This chapter covers the techniques and notations available for expressing the behavior of elements, groups of elements, and the system as a whole.		
8.1	Beyond Structure	289
8.2	How to Document Behavior	290
8.2.1	Step 1: Decide What Kinds of Questions You Need to Answer	290
8.2.2	Step 2: Determine What Types of Information Are Available or Can Be Constrained	291
8.2.3	Step 3: Choose a Notation	293
8.3	Notations for Documenting Behavior	295
8.3.1	Notations for Capturing Traces	295
8.3.2	Notations for Capturing Comprehensive Models	303
8.4	Where to Document Behavior	306
8.5	Why to Document Behavior	306
8.5.1	Driving Development Activities	306
8.5.2	Analysis	307
8.6	Summary Checklist	308
8.7	Discussion Questions	309
8.8	For Further Reading	311
Part III Building the Architecture Documentation		313
Part III covers what you have to do to create and maintain the documentation artifacts: choosing views to include, laying out and packaging the information, and reviewing the document.		
Chapter 9 Choosing the Views		315
This chapter provides guidance for selecting views, given the intended use of an architecture: analysis, reconstruction, achieving common understanding, the basis for deriving code, and so on.		
9.1	Stakeholders and Their Documentation Needs	316

9.2	A Method for Choosing the Views	326
	Perspectives: Listening to the Stakeholders	327
9.3	Example	329
	Perspectives: How Not to Introduce an Architecture	333
9.4	Summary Checklist	335
9.5	Discussion Questions	335
9.6	For Further Reading	335

Chapter 10 Building the Documentation Package 337

This chapter explains how the documentation is organized to serve its stakeholders. The chapter shows how the elements discussed in the prior chapters fit together to produce usable documentation. The chapter includes templates for architecture documentation.

10.1	Documenting a View	337
10.1.1	A Standard Organization for Documenting a View	337
	Perspectives: From Context Diagrams to a Context View	341
10.1.2	Useful Variations in the Standard Organization for a View	344
10.1.3	Avoiding Unnecessary Repetition Across Views or View Packets	349
10.2	Documentation Beyond Views	350
10.2.1	A Standard Organization for Documenting Information Beyond Views	351
10.2.2	Useful Variations in the Standard Organization for Documentation Beyond Views	356
10.3	Documenting a Mapping to Requirements	357
	Perspectives: A Mapping to Requirements: You Might Already Have It	362
10.4	Packaging the Architecture Documentation	362
10.4.1	Packaging Schemes	362
10.4.2	Online Documentation, Hypertext, and Wikis	365
	Coming to Terms: Wiki	366
10.4.3	Configuration Management	368
10.4.4	Follow a Release Strategy	368
	Perspectives: Presentation Is Also Important	369
	Perspectives: Tooling Requirements	370
10.5	Summary Checklist	372
10.6	For Further Reading	373

Chapter 11 Reviewing an Architecture Document 375

This chapter describes a step-by-step approach for conducting a structured review of an architecture document, and it includes a large selection of review questions.

11.1	Steps of the Procedure	376
	Coming to Terms: Active Design Reviews	380

11.2	Sample Question Sets for Reviewing the Architecture Document	382
11.2.1	Example Question Set for Capturing the Right Stakeholders and Concerns	385
11.2.2	Example Question Set for Supporting Evaluation	386
11.2.3	Example Question Set for Supporting Development	389
11.2.4	Example Question Set for Reviewing for Conformance to ISO/IEC 42010	391
11.3	An Example of Constructing and Conducting a Review	393
11.4	Summary Checklist	395
11.5	Discussion Questions	396
11.6	For Further Reading	396

Epilogue: Using Views and Beyond with Other Approaches **399**

The epilogue compares the “Views and Beyond” approach to other documentation approaches. It ties related work to the prescriptions given in this book.

E.1	ISO/IEC 42010, <i>née</i> ANSI/IEEE Std 1471-2000	400
E.1.1	Overview	400
E.1.2	42010 and Views and Beyond	404
E.2	Rational Unified Process/Kruchten 4+1	406
E.2.1	RUP/4+1 and Views and Beyond	406
E.3	Using the Rozanski and Woods Viewpoint Set	408
	Coming to Terms: Architecture Perspectives	410
E.3.1	Rozanski and Woods Viewpoints and Views and Beyond	411
E.4	Documenting Architecture in an Agile Development Project	414
E.4.1	Overview	414
E.4.2	Agile Development and Views and Beyond	415
E.5	U.S. Department of Defense Architecture Framework	419
E.5.1	Overview of DoDAF	419
E.5.2	DoDAF and Software Architecture	421
E.5.3	DoDAF and Views and Beyond	421
E.5.4	A Strategy to Use DoDAF to Document Software Architecture	426
E.6	Where Architecture Documentation Ends	428
E.7	A Final Word	429
E.8	For Further Reading	429

Appendix A UML—Unified Modeling Language 431

This appendix gives an overview of the Unified Modeling Language and tells how it should be used to document the architecture constructs described in this book.

A.1	Introduction	431
A.2	Documenting a Module View	433
A.2.1	Decomposition Style	433
A.2.2	Uses Style	433
A.2.3	Generalization Style	434
A.2.4	Layered Style	434
A.2.5	Aspects Style	434
A.2.6	Data Model Style	435
	Perspectives: UML Class Diagrams: Too Much, Too Little	436
A.3	Documenting a Component-and-Connector View	438
A.4	Documenting an Allocation View	443
A.4.1	Deployment Style	443
A.4.2	Install and Implementation Styles	445
A.4.3	Work Assignment Style	446
A.5	Documenting Behavior	449
A.5.1	Activity Diagram	450
A.5.2	Sequence Diagram	450
A.5.3	Communication Diagram	453
A.5.4	Timing Diagram	454
A.5.5	Interaction Overview Diagram	455
A.5.6	State Machine Diagram	457
A.5.7	Use Case Diagram	458
A.6	Documenting Interfaces	460
	Perspectives: UML Tools	461

Appendix B SysML—Systems Modeling Language 465

The Systems Modeling Language (SysML) is a general-purpose systems modeling language intended to support a broad range of analysis and design activities for systems engineering applications. This appendix gives a short overview of using SysML to represent an architecture.

B.1	Architecture Documentation	466
B.2	Requirements	466
B.3	Documenting a Module View	468
B.4	Documenting a Component-and-Connector View	469
B.5	Documenting an Allocation View	470
B.6	Documenting Behavior	471
B.7	Documenting Interfaces	472
B.8	Summary	472

Appendix C AADL—The SAE Architecture Analysis and Design Language **473**

The Architecture Analysis and Design Language (AADL) provides a textual and graphical language to represent the runtime architecture of software systems as a component-based model in terms of tasks and their interactions, the hardware platform on which the system executes, and the physical environment with which it interfaces. This appendix summarizes AADL and briefly describes how it can be used to document architectures.

C.1	Introduction	473
C.2	Documenting a Module Style	475
C.3	Documenting a Component-and-Connector View	478
C.4	Documenting a Deployment View	481
C.5	Documenting Behavior	482
C.6	Documenting Interfaces	484
C.7	Summary	484
<i>Acronyms</i>		487
<i>Glossary</i>		491
<i>References</i>		497
<i>About the Authors</i>		509
<i>About the Contributors</i>		513
<i>Index</i>		517

Foreword to the Second Edition

A colleague of mine, in the market for a home, fell in love with an older property that had been designed by a student of Frank Lloyd Wright himself. Curious about its history, its structure, its evolution, he contacted the local planning office, which happily and quickly provided him with a copy of the original blueprints.

Why, my friend asked me, can we get the drawings for a house that's several decades old, but we are unable to see the architecture of software written last year?

In this book, the authors offer some pragmatic wisdom that helps attend to my friend's lament.

The theory and the practice of the architecture of software-intensive systems are in a very vibrant phase. The early work of Mary Shaw and David Garlan in particular gave rise to software architecture as an identifiable domain of study, and in the years since, we've seen the emergence of architecture-as-an-artifact as a mainstream concern for the development and evolution of systems. This has manifest itself in notations such as the Unified Modeling Language (which was explicitly influenced by Philippe Kruchten's 4+1 model view of software architecture) as well as a panoply of architectural frameworks, such as The Open Group Architecture Framework and the Department of Defense Architecture Framework. Add to these methods such as IBM's Unified Process and, at another extreme, the Federal Segment Architecture Methodology, and it is clear that architecture-as-an-artifact has found an important role in the reasoning about and governing of software-intensive systems.

There are some things we can say with confidence. Every system has an architecture. All complex systems are hierarchical in nature, but also exhibit other patterns of regularity. There's

an intimate dance that occurs between the processes of architecting and of implementation. And, to understand and reason about the architecture of a software-intensive system, one has to consider multiple views from the perspectives of specific concerns from multiple classes of stakeholders.

The most commonly used notation and tool for describing a system's architecture is a boxes-and-lines sketch created on a whiteboard. Such documentation is both expeditious and useful, but it is neither enduring nor rigorous nor complete. In this book the authors offer the definitive reference on the documenting of the architecture of software-intensive systems, in ways that are enduring and rigorous and complete. And useful, by the way!

I remember reading the first edition of this book, and e-mailing my compliments to the authors for producing such a comprehensive reference. Well, they've outdone themselves. This new edition is brighter, shinier, more complete, more pragmatic, more focused than the previous one, and I wouldn't have thought it possible to improve on the original. As the field of software architecture has grown over these past decades, there is much more to be said, much more that we know, and much more that we can reflect upon of what's worked and what hasn't—and the authors here do all that, and more.

So, my hope for you, dear reader, is this: May the software you write today have an architecture that your children's children may discern and celebrate.

—Grady Booch
IBM Fellow

Foreword to the First Edition

Ten years ago, I was brought in to lead the architecture team of a new and rather ambitious command-and-control system. After some rocky beginnings, the architectural design work started to proceed full speed, and the architects were finally forging ahead, inventing and resolving and designing and trying, almost in a euphoric state. We had many brainstorming sessions, filling whiteboards with design fragments and notebooks with scribbles; various prototypes validated or invalidated our reasoning. As the development team grew in size, the architects had to explain the principles of the nascent architecture to a wider and wider audience, consisting of not only new developers but also many parties external to the development group. Some were intrigued by this new concept of a software architecture. Some wanted to know how this architecture would impact them: for planning, for organizing the teams and the contractors, for delivery of the system, for acquisition of some of the system parts. Some parties wanted to influence the design of this architecture. Further removed from development, customers and prospects wanted a peek, too. So the architects had to spend hours and days describing the architecture in various forms and levels and tones to varied audiences, so that each party could better understand it.

Becoming this center of communication slowly stretched our capacity. On the one hand, we were busy designing the architecture and validating it; on the other hand, and at the same time, we were communicating to a large audience what it was and why it was that way and why we did not choose some other solution. A few months into the project, overwhelmed, we even began having a difficult time agreeing among ourselves about what it was we had actually decided.

This led me to the conclusion that “if it is not written down, it does not exist.” This became sort of a leitmotiv in the architecture team for the following two years. As the ancient Chinese poet Lao-Tsu says in the *Tao Te Ching*:

Let your workings remain a mystery.
Just show people the results.
(Tablet #36)

The architecture *could be* whatever we had talked about, argued, imagined, or even drafted on a board, and so on. But the architecture of this system *was* only what was described in one major document: the *Software Architecture Document (SAD)*. Architectural elements and architectural decisions not captured in this document simply did not exist. This one rule—“*If it is not in the SAD, it does not exist.*”—became our incentive to evolve and to keep the document up-to-date, almost to the week; there was also an incentive to not include anything and everything and untried ideas, as this was the project’s definite arbiter.

The SAD rapidly became a central element in the life of the project. It became our best display window for showing off our stuff, our comfort when we were down, and our shield when attacked.

The key problem we faced at the time was: What do we document for a software architecture? How do we document it? What outline do we use? What notation? How much or how little? There were few exemplars of architectural description for systems as ambitious as ours. Driven by necessity, we improvised. We made some mistakes and corrected some. We discovered rapidly that architecture is not flat but rather a multidimensional reality, with several intertwined facets, and some facets—or views—of interest to only a few parties. We found out that many readers would not even open a document that weighed more than a pound, and we would have a difficult time updating it anyhow. We realized that without capturing the reasons for our choices, we were doomed to reconstruct them again and again, every time a new stakeholder with a sharp mind came around. We picked a visual notation, not too vague and fuzzy but not too esoteric and convoluted, either, in order to not discourage most parties.

Today, software architects have a great starting point for deciding how to document their software architectures. You have it in your hands. The authors went through many experiences similar to mine and extracted the important lessons learned. They read many software architecture documents.

They reviewed the academic literature, studied all the published books, checked the standards, and synthesized all this wisdom in this handbook: the essential things you need to know to define your own software architecture document. You will find guidance for the scope of software architecture; its organization; the techniques, tools, and notation to use or not to use; and comparisons, advice, and rules of thumb. In here, you'll find the templates to get you started and the continuing guidance for when you get lost or despairing on the way.

This book is of immense value. The description and communication of software architecture is quite crucial to its many stakeholders, and this handbook should save you months of trials and errors, lots of undeserved hassle, and many costly mistakes that could potentially jeopardize the whole endeavor. It will become an important reference on the shelf of the software architect.

—Philippe Kruchten
Director of Process Development
Rational Software Canada, Vancouver

This page intentionally left blank

Preface

The purpose of this book is to answer the following question:

How do you document an architecture so that others can successfully use it, maintain it, and build a system from it?

The audience for this book includes all the people involved in the production and consumption of architecture documentation. The goal of this book is to help you decide what information about an architecture is important to capture and to provide guidelines, notations, and examples for capturing it. We intend this book to be a practitioner-oriented guide to the various kinds of information that constitute an architecture. We give practical guidance for choosing what information should be documented and show—with examples in various notations, including but not limited to the Unified Modeling Language (UML)—how to describe that information in writing so that others can use it to carry out their architecture-based work: implementation, analysis, and recovery. We also show how to create a comprehensive software architecture document that others can use.

Although piles of books exist about how to use a particular notation (UML comes to mind), we believe what an architect really needs is guidance in which architecture and its stakeholders are the first-class citizens, and language is relegated more appropriately to a supporting role. That's what we've tried to provide with this book.

Languages and Tools for Architecture

Commercial languages and tool suites are available for capturing design information, especially in the realm of object-oriented

systems. Some of these tools are bound up with associated design methods, notations, and commercial products. Some tools are aimed at points in the design space other than architecture. If you have decided to adopt one of these tools and/or notations, will this book relate to you?

Very few things become obsolete faster than references to specific tools, so we've avoided those. Instead, we have concentrated on the information you should capture about an architecture. We believe that is the approach you should take, too: Concentrate on the information you need to capture, and then figure out how to capture it using an available tool. Almost all tools provide ways to add free-form annotations to the building blocks they provide; if all else fails, these annotations will let you capture and record information in ways you see fit. Remember that not all the people for whom architecture documentation is prepared will be able to use the tool environment you've chosen or understand the commercial notation you've adopted.

Having said that, however, we acknowledge that a few standard languages and notations have come to dominate, chief among them UML. And so this book provides a plethora of examples showing UML 2 representing the architecture views we cover, as well as other concepts such as refinement and behavior. If you have chosen UML as your modeling language, you'll feel at home.

Appendix A contains a summary of UML's visual notation and its applicability to document the concepts in this book. Appendices B and C summarize the Systems Modeling Language (SysML) and the Architecture Analysis and Design Language (AADL), respectively. Our purpose is not to teach these languages, but to offer a quick refresher for those familiar with them and a flavor-providing introduction for everyone else.

What's New in the Second Edition

- A number of new architecture styles have entered the mainstream, and this edition talks about documenting those. These include service-oriented architectures, multi-tier architectures, and architectures for aspect-oriented systems. We also treat the architecture-level documentation of a software system's data model, as well as its installation and production environment, as first-class styles.

- This edition is much more Agile-friendly, orienting its advice to be consistent with the Agile Manifesto's entreaty to value working software over comprehensive documentation.
- We treat the systematic documentation of rationale with much greater depth, reflecting best industrial practices. We've added a new chapter about reviewing an architecture document to make sure it's serving its stakeholders as intended.
- The suggested templates for architecture documentation have several improvements, reflecting years of use and feedback. They are also more flexible, and we lay out different options for arranging your documentation.
- We have replaced the comprehensive example of a documented software architecture with a new one. The architecture is for a Web-based service-oriented system, more in today's industrial mainstream. To make the book smaller and allow us to maintain the example over time, we put the example online. And many of our in-line examples have been replaced or updated.
- Since the first edition was published, the Unified Modeling Language has graduated to version 2.0 and beyond. That opened up new possibilities for more straightforwardly documenting various architecture constructs, especially components and connectors. Where necessary, our figures are updated to reflect the new constructs.
- This edition has concise appendices summarizing three important languages and notations useful for documenting architectures: UML, AADL, and SysML. Each appendix constitutes a mini-reference guide on the language.
- Finally, this edition reflects the experience we've gained with Views and Beyond in the intervening years since the first edition was published. This experience has come from creating documented architectures for very challenging systems, and helping other people do so. It also comes from using architecture documentation in practice, such as when we evaluate other organizations' software architectures. Finally, it has come from interacting with more than a thousand participants in our two-day industrial course based on the book. These interactions with practicing software architects have let us make our advice more prescriptive and crisp and reflect the problems and situations that architects face daily.

Complete Example of a Software Architecture Document Online

You can see a fully worked-out example of a software architecture document using the approaches and templates described in this book at wiki.sei.cmu.edu/sad.

—P.C.

Austin, Texas

—F.B., L.B., D.G., J.I., R.L., R.N.

Pittsburgh, Pennsylvania


—P.M.

Brasilia, Brazil

—J.S.

Boston, Massachusetts

Prologue: Software Architectures and Documentation



The prologue establishes a small but fundamental set of concepts that will be used throughout the book. We begin with short overviews of software architecture (Section P.1) and architecture documentation (Section P.2), and then we go on to discuss the following topics:

- Section P.3: Architecture views
- Section P.4: Architecture styles (and their relation to architecture patterns) and the classification of styles into three categories: module styles, component-and-connector styles, and allocation styles
- Section P.5: Rules for sound documentation

P.1 A Short Overview of Software Architecture

P.1.1 Overview

Software architecture has emerged as an important subdiscipline of software engineering. Architecture is roughly the prudent partitioning of a whole into parts, with specific relations among the parts. This partitioning is what allows groups of people—often separated by organizational, geographical, and even time-zone boundaries—to work cooperatively and productively together to solve a much larger problem than any of them could solve individually. Each group writes software that interacts with the other groups' software through carefully crafted interfaces that reveal the minimal and most stable information necessary for interaction. From that interaction emerges the functionality and quality attributes—security, modifiability, performance, and so forth—that the system's stakeholders demand. The larger and more complex the sys-



The **software architecture** of a computing system is the set of structures needed to reason about the system, which comprise software elements, relations among them, and properties of both.



Many projects make the mistake of trying to impose a single partition in multiple component domains, such as equating threads with objects, which are equated with modules, which in turn are equated with files. Such an approach never succeeds fully, and adjustments eventually must be made, but the damage of the initial intent is often hard to repair. This invariably leads to problems in development and occasionally in final products.

—Jazayeri, Ran, and van der Linden (2000, pp. 16–17)

tem, the more critical is this partitioning—and hence, architecture. And as we will see, the more demanding those quality attributes are, the more critical the architecture is.

A single system is almost inevitably partitioned simultaneously in a number of different ways. Each partitioning results in the creation of an architectural structure: different sets of parts and different relations among the parts. Each is the result of careful design, carried out to satisfy the driving quality attribute requirements and the most important business goals behind the system.

Architecture is what makes the sets of parts work together as a coherent and successful whole. Architecture documentation help architects make the right decisions; it tells developers how to carry them out; and it records those decisions to give a system’s future caretakers insight into the architect’s solution.

P.1.2 Architecture and Quality Attributes

For nearly all systems, quality attributes such as performance, reliability, security, and modifiability are every bit as important as making sure that the software computes the correct answer. A software system’s ability to produce correct results isn’t helpful if it takes too long doing it, or the system doesn’t stay up long enough to deliver it, or the system reveals the results to your competition or your enemy. Architecture is where these concerns are addressed. For example:

- If you require high performance, you need to
 - Exploit potential parallelism by decomposing the work into cooperating or synchronizing processes.
 - Manage the interprocess and network communication volume and data access frequencies.
 - Be able to estimate expected latencies and throughputs.
 - Identify potential performance bottlenecks.
- If your system needs high accuracy, you must pay attention to how the data elements are defined and used and how their values flow throughout the system.
- If security is important, you need to
 - Legislate usage relationships and communication restrictions among the parts.
 - Identify parts of the system where an unauthorized intrusion will do the most damage.
 - Possibly introduce special elements that have earned a high degree of trust.

- If you need to support modifiability and portability, you must carefully separate concerns among the parts of the system, so that when a change affects one element, that change does not ripple across the system.
- If you want to deploy the system incrementally, by releasing successively larger subsets, you have to keep the dependency relationships among the pieces untangled, to avoid the “nothing works until everything works” syndrome.

The solutions to these concerns are purely architectural in nature. It is up to architects to find those solutions and communicate them effectively to those who will carry them out. Architecture documentation has three obligations related to quality attributes. First, it should indicate which quality attribute requirements drove the design. Second, it should capture the solutions chosen to satisfy the quality attribute requirements. Finally, it should capture a convincing argument why the solutions provide the necessary quality attributes. The goal is to capture enough information so that the architecture can be analyzed to see if, in fact, the system(s) derived from it will possess the necessary quality attributes.

COMING TO TERMS

What Is Software Architecture?

If we are to agree on what it means to document a software architecture, we should establish a common basis for what it is we’re documenting. No universal definition of software architecture exists. The Software Engineering Institute’s Web site collects definitions from the literature and from practitioners around the world; so far, more than 150 definitions have been collected.

It seems that new fields try to nail down standard definitions or their key terms as soon as they can. As the field matures, basic concepts become more important than ironclad definitions, and this urge seems to fade. When object-oriented development was in its infancy, you could bring any OO meeting to a screeching halt by putting on your best innocent face and asking, “What exactly is an object?” This largely ended when people realized that the scatter plot of definitions had an apparent (if unarticulated) centroid, from which very useful progress could be made. Sometimes “close enough” is, well, close enough.



Chapter 10 will show where in the documentation to record the driving quality attribute requirements, the solutions chosen, and the rationale for those solutions.



Software architecture is the set of design decisions which, if made incorrectly, may cause your project to be cancelled.

—Eoin Woods (SEI 2010)



You can read the SEI collection of definitions, or contribute your own, at www.sei.cmu.edu/architecture.

This seems to be the case with software architecture. Looking at the major attempts to nail down its definition gives us a good glimpse at our own centroid. With that in mind, here are a few influential definitions:

By analogy to building architecture, we propose the following model of software architecture: Software Architecture = {Elements, Form, Rationale}. That is, a software architecture is a set of architectural (or, if you will, design) elements that have a particular form. We distinguish three different classes of architectural elements: processing elements; data elements; and connecting elements. The processing elements are those components that supply the transformation on the data elements; the data elements are those that contain the information that is used and transformed; the connecting elements (which at times may be either processing or data elements, or both) are the glue that holds the different pieces of the architecture together. (Perry and Wolf 1992, p. 44)

. . . beyond the algorithms and data structures of the computation; designing and specifying the overall system structure emerges as a new kind of problem. Structural issues include gross organization and global control structure; protocols for communication, synchronization, and data access; assignment of functionality to design elements; physical distribution; composition of design elements; scaling and performance; and selection among design alternatives. (Garlan and Shaw 1993, p. 1)

The structure of the components of a program/system, their interrelationships, and principles and guidelines governing their design and evolution over time. (Garlan and Perry 1995, p. 269)

An architecture is the set of significant decisions about the organization of a software system, the selection of the structural elements and their interfaces by which the system is composed, together with their behavior as specified in the collaborations among those elements, the composition of these structural and behavioral elements into progressively larger subsystems, and the architecture style that guides this organization—these elements and their interfaces, their collaborations, and their composition. (Booch, Rumbaugh, and Jacobson 1999, p. 31)

The fundamental organization of a system embodied in its components, their relations to each other, and to the environment, and the principles guiding its design and evolution. (IEEE 1471 2000, p. 9)

The software architecture of a program or computing system is the structure or structures of the system, which

comprise software elements, the externally visible properties of those elements, and the relations among them. By “externally visible properties,” we are referring to those assumptions other components can make of a component, such as its provided services, performance characteristics, fault handling, shared resource usage, and so on. (Bass, Clements, and Kazman 2003, p. 27)

The set of principal design decisions governing a system. (Taylor, Medvidovic, and Dashofy 2009, p. xv)

A few other “mainstream” definitions have emerged since then, but they are largely restatements and recombinations of the ones we just listed. The centroid seems to have stabilized.

That centroid takes a largely structural perspective on software architecture: Software architecture is composed of elements, connections or relations among them, and, usually, some other aspect or aspects, such as (take your pick) configuration; constraints or semantics; analyses or properties; or rationale, requirements, or stakeholders’ needs.

These perspectives do not preclude one another, nor do they represent a fundamental conflict about what software architecture is. Instead, they represent a spectrum in the software architecture community about the emphasis that should be placed on architecture: its constituent parts, the whole entity, the way it behaves once built, or the building of it. Taken together, they form a consensus view of software architecture.

In this book we use a definition similar to the one from Bass, Clements, and Kazman (2003). We chose it because it helps us know what to document about an architecture. The definition emphasizes the plurality of structures present in every software system. These structures, carefully chosen and designed by the architect, are the key to achieving and reasoning about the system’s design goals. And those structures are the key to understanding the architecture. Therefore, they are the focus of our approach to documenting a software architecture. Structures consist of elements, relations among the elements, and the important properties of both. So documenting a structure entails documenting those things.

PERSPECTIVES

What's the Difference Between Architecture and Design?

The question of how architecture is different from design has nipped at the heels of the software development community for years. It is a question I often hear when teaching an introductory course on architecture. It matters here because the question deals with what we should put in an architecture document and what we should put somewhere else.

The first thing we can say is that clearly architecture *is* design, but not all design is architecture. That is, many design decisions are left unbound by the architecture and are happily left to the discretion and good judgment of downstream designers and even implementers. The architecture establishes constraints on downstream activities, and those activities must produce artifacts—finer-grained designs and code—that comply with the architecture.

It's tempting to stop there, but if you're paying attention you've seen that we've just translated the question: Architecture consists of architectural design decisions, and all others are nonarchitectural. So what decisions are nonarchitectural? That is, what design decisions does the architect leave to the discretion of others?

To answer this question, we return to the primary purpose of architecture, which is to assure the satisfaction of the system's quality and behavioral requirements and business goals. The architect does this by making design decisions that manifest themselves in the system's architectural structures.

Thus, architectural decisions are ones that permit a system to meet its quality attribute and behavioral requirements. All other decisions are nonarchitectural.

Clearly any design decisions resulting in element properties that are *not* visible—that is, make no difference outside the element—are nonarchitectural. A typical example is the selection of a data structure, along with the algorithms to manage and access that data structure.

You may have been hoping for a more concrete answer, such as “the first three levels of module decomposition

are architectural, but any subsequent decomposition is not.” Or, “the classes, packages and their relations in a UML class diagram are architectural, but sequence diagrams are not.” Or “defining the services of an SOA system is architectural, but designing the internal structure of each service provider component is not.”

But those don’t work because they draw arbitrary and artificial boundaries. Attempts like that to be practical end up being impractical because true architecture bleeds across those boundaries.

Here are some more sometimes-heard artificial definitions.

First, “architecture is the small set of big design decisions.” Some people define “small set” by insisting that an architecture document should be no more than 50 pages. Or 80. Or 30. Their feeling, apparently, is that architecture is the set of design decisions that you can squeeze into a given page quota, and everything beyond that is not. This is, of course, utter nonsense.

Another oft-heard nonanswer is “architecture is what you get before you start adding detail to the design.” Terminology often directs our thinking, rather than serves it. A pernicious example that puts us in the wrong mind set is “detailed design.” Detailed design is what many people say follows architecture. The term is everywhere, and needs to be stamped out. It implies that the difference between architectural and nonarchitectural design is something called “detail.” Architecture is apparently not allowed to be detailed, because if it is, well, you’re doing detailed design then, aren’t you? Never mind that we have no idea how to measure “detail” nor to set a threshold for when there is too much of it to be architectural. If your design starts to look “detailed” then you aren’t doing architecture and you’ll be reported to the Detailed Design Police for overstepping your authority. More utter nonsense.

It’s true that some architectural design decisions may lack much specificity; that is, they preserve freedom of choice for downstream designers. Some architectural design decisions may not be “decisions” at all, but broad constraints. Plug-ins that populate your Web browser are an example. No architecture nails down the complete set, but the architecture does constrain new ones to meet certain standards and interfaces. Or the architect might describe an element by saying, “The element delivers its computational result through this published interface, is



Don’t use the term “detailed design”! Use “nonarchitectural design” instead.

thread-safe, puts no more than three messages on the network per invocation, and returns its answer in less than 20 ms.” The team implementing that element is free to make whatever design decisions they wish as long as they satisfy the architect’s prescription for it.

On the other hand, some architectural decisions can be quite “detailed,” such as the adoption of specific protocols, an XML schema, or communication or technology standards. Such decisions are usually made for purposes of interoperability or various flavors of modifiability (such as scalability or extensibility).

Even interfaces of elements, which some decry as “obviously” outside the realm of architecture, can be supremely architectural. For instance, in a service-oriented architecture (SOA), components interact through published interfaces. Important design decisions made when defining these interfaces include the granularity of the operations, the data format, and the type of interaction (synchronous or asynchronous) for each operation. Or consider an element that processes data from a real-time sensor. Making this element’s interface process a stream as opposed to individual data elements will make an enormous difference in the ability of the element (and hence the system) to meet real-time performance requirements. This decision cannot be left up to the element’s development team; everything depends on it.

A legitimate question about detail does arise when considering modules and other **hierarchical elements**: When do you stop? When have you designed enough levels in the hierarchy? Are submodules enough, or does the architect need to design sub-sub-sub-submodules? Here’s a good test of our claim for when architecture stops. Module decomposition is about achieving independent development and modifiability. Both are achieved by carefully assigning coherent responsibilities to each module. When the modules you’ve designed are fine-grained enough to satisfy the system’s modifiability and independent development requirements, you’ve discharged your obligation as an architect.

Finally, what is architectural is sensitive to context. Suppose the architect identifies an element but is content to sketch the element’s interface and behavior in broad terms. If the element being prescribed is very large and complex, the team developing it may choose to give it an



A **hierarchical element** is any kind of element that can consist of like-kind elements. A module is a hierarchical element because modules consist of submodules, which are themselves modules. A task or a process is not a hierarchical element.

internal substructure of its own, which for all the world looks like an architecture. And within the context of that element, it is. But in the context of the overall system, the substructure is not architectural but merely an internal design decision made by the development team for that element.

To summarize, architecture is design, but not all design is architectural. The architect draws the boundary between architectural and nonarchitectural design by making those decisions that need to be bound in order for the system to meet its development, behavioral, and quality goals. All other decisions can be left to downstream designers and implementers. Decisions are architectural or not, according to context. If structure is important to achieve your system's goals, that structure is architectural. But designers of elements, or subsystems, that you assign may have to introduce structure of their own to meet their goals, in which case such structures are architectural: *to them* but not to you.

And (repeat after me) we all promise to stop using the phrase “detailed design.” Try “nonarchitectural design” instead.

—P.C.

P.2 A Short Overview of Architecture Documentation

P.2.1 Why Document Software Architecture?

Even the best architecture, most perfectly suited for the job, will be essentially useless if the people who need to use it do not know what it is, cannot understand it well enough to apply it, or (worst of all) misunderstand it and apply it incorrectly. All of the effort, analysis, hard work, and insightful design on the part of the architecture team will have been wasted. They might as well have gone on vacation for all the good their architecture will do.

Creating an architecture isn't enough. It has to be communicated in a way to let its stakeholders use it properly to do their jobs. If you go to the trouble of creating a strong architecture, you *must* go to the trouble of describing it in enough detail, without ambiguity, and organized so that others can quickly find needed information.

Documentation speaks for the architect. It speaks for the architect today, when the architect should be doing other things besides answering a hundred questions about the architecture.



Doing business without advertising [or designing an architecture without documenting it] is like winking at a girl in the dark. You know what you're doing, but nobody else does.

—Stewart Henderson Britt

And it speaks for the architect tomorrow, when he or she has left the project and now someone else is in charge of its evolution and maintenance.

Documentation is often treated as an afterthought, something people do because they have to. Maybe a contract requires it. Maybe a customer demands it. Maybe a company's standard process calls for it. In fact, these may be legitimate reasons. But none of them are compelling enough to produce high-quality documentation. Why should the architect spend valuable time and energy just so a manager can check off a deliverable?

The best architects produce the best documentation not because it's "required," but because they see that it is essential to the matter at hand: producing a high-quality product, predictably and with as little rework as possible. They see their immediate stakeholders as the people most intimately involved in this undertaking: developers, deployers, testers, and analysts.

But the best architects also see documentation as delivering value to themselves. Documentation serves as the receptacle to hold the results of design decisions as they are made. A well-thought-out documentation scheme can make the process of design go much more smoothly and systematically. Documentation helps the architect while the architecting is in progress, whether in a six-month design phase or a six-day Agile sprint.

COMING TO TERMS

Specification, Representation, Description, Documentation

What shall we call the activity of writing down a software architecture for the benefit of others or for our own benefit at a later time? Leading contenders are *documentation*, *representation*, *description*, and *specification*. None of these terms has a standardized meaning in our field: the difference between them is unclear. For the most part, we use *documentation* throughout this book, and we want to explain why.

Specification tends to connote an architecture rendered in a formal language. Now, we are all for formal specs. But formal specs are not always practical, nor are they always necessary. Sometimes, they aren't even useful: How, for example, do you capture in a formal language the rationale behind your architectural decisions, and why would you try?

Representation connotes a model, an abstraction, a rendition of a thing that is separate or different from the thing itself. Is architecture something more than what someone writes down about it? Arguably yes, but it's certainly pretty intangible in any case. We felt that raising the issue of a model versus the thing being modeled would only elicit needlessly diverting questions best left to those whose hobby, or calling, is philosophy: Does an abstraction of a tree falling in a model of a forest make a representation of a sound? This does not seem like the start of a productive conversation.

Description has been staked out by the architecture description language (ADL) community, and more recently by the standards community coming up with mandates for how to write down an architecture. It's curious that the people you'd think would be the most formal snagged the least rigorous sounding term of the bunch. (The next time you board a jet, sit in front of a computer-controlled X-ray machine, or watch the launch of a billion-dollar space vehicle your tax dollars paid for, ask yourself whether you hope the control software has been specified to the implementers, or merely described.) We eschewed *description*, then, because it all at once sounds too formal—we didn't want people to think that writing down an architecture requires an architecture description language—and too informal. Descriptions can be notoriously vague, such as when your friends *describe* the blind date they set you up with. Sometimes we need a little more specificity in our lives, and certainly we need it in our architectures.

That leaves *documentation*. *Documentation* connotes the creation of an artifact: namely, a document, which may of course consist of electronic files, Web pages, a snapshot of a whiteboard, or paper. Thus, documenting a software architecture becomes a concrete task: producing a software architecture document. Viewing the activity as creating a tangible product has advantages. We can describe good architecture documents and bad ones. We can use completeness criteria to judge how much work is left in producing this artifact and determining when the task is done. Planning or tracking a project's progress around the creation of artifacts, or documents, is an excellent way to manage. Making the architecture information available to its consumers and keeping it up to date reduces to a solved problem of configuration



ADLs are discussed in Section 3.4.2 and in the For Further Reading section of Chapter 8. For an overview of ADLs, see the work by Stafford and Wolf (2001).



Section 6.1.3 (“Spectrum of Design”) discusses how architecture documentation captures the very abstract to the very detailed.

control. Documentation can be formal or not, as appropriate, and may contain models or not, as appropriate. Documents may describe, or they may specify. Hence, the term is appropriately general.

No matter what you call it, the essence of the activity is writing down—and keeping current—the results of architectural decisions so that the stakeholders of the architecture—people who need to know what it is to do their job—have the information they need in an accessible, nonambiguous form.



In Chapter 9, the documentation’s expected uses, along with the documentation obligations each use imparts, become the basis for helping an architect plan the documentation package.

P.2.2 Uses and Audiences for Architecture Documentation

Architecture documentation must serve varied purposes. It should be sufficiently abstract to be quickly understood by new employees. It should be sufficiently concrete to serve as a blueprint for construction. It should have enough information to serve as a basis for analysis.

Architecture documentation is both prescriptive and descriptive. For some audiences, it prescribes what *should* be true, placing constraints on decisions yet to be made. For other audiences, it describes what *is* true, recounting decisions already made about a system’s design.

The best architecture documentation for, say, performance analysis may well be different from the best architecture documentation we would wish to hand to an implementer. And both of these will be different from what we put in a new hire’s “welcome aboard” package or a briefing we put together for an executive. The process of documentation planning and review needs to ensure support for all the relevant needs.

We can see that many different kinds of people are going to have a vested interest in an architecture document. They hope and expect that the architecture document will help them do their respective jobs. Understanding their uses of architecture documentation is essential, as those uses determine the important forms.

Fundamentally, architecture documentation has three uses.

1. *Architecture serves as a means of education.* The educational use consists of introducing people to the system. The people may be new members of the team, external analysts, or even a new architect. In many cases, the “new” person is the customer to whom you’re showing your solution for the first time, a presentation you hope will result in funding or go-ahead approval.



Chapter 9 discusses planning the contents of a documentation package. Chapter 11 discusses reviewing documentation.

2. *Architecture serves as a primary vehicle for communication among stakeholders.* An architecture's precise use as a communication vehicle depends on which stakeholders are doing the communicating. Some examples are described in Table P.1.

Perhaps one of the most avid consumers of architecture documentation is none other than the architect in the project's future. The future architect may be the same person as the present one, or he or she may be a replacement, but in either case he or she is guaranteed to have an enormous stake in the documentation. New architects are interested in learning how their predecessors tackled the difficult issues of the system and why particular decisions were made. Even if the future architect is the same person, he or she will use the documentation as a repository of thought, a storehouse of design decisions too numerous and hopelessly intertwined ever to be reproducible from memory alone.

Even in the short term, documenting an architecture helps in the process of designing the architecture. First, the documentation provides dedicated compartments for recording various kinds of design decisions as soon as they are made. Second, the documentation gives you a rough but helpful way to gauge progress and the work remaining: As "TBD"s disappear from the document, completion draws near. Finally, documentation provides a framework for systematic attack on designing the architecture. Key design decisions, usually made early, should be written down so that the shadow they cast on subsequent design decisions is explicit and remembered.



A **stakeholder** of an architecture is someone who has a vested interest in it. (Many of an architecture's stakeholders are listed in Table P.1.)

Chapter 9 is about how stakeholders' needs will help determine the contents of the architecture documentation.



Stakeholders (explicitly or implicitly) drive the whole shape and direction of the architecture, which is developed solely for their benefit and to serve their needs. . . . Without stakeholders, there would be no point in developing the architecture because there would be no need for the system it will turn into, nor would there be anyone to build it, deploy it, run it, or pay for it. . . . Architectures are created solely to meet stakeholder needs.

—Rozanski and Woods (2005, p. 21)

QUOTE

In our organization, a development group writes design documents to communicate with other developers, external test organizations, performance analysts, the technical writers of manuals and product helps, the separate installation package developers, the usability team, and the people who manage translation testing for internationalization. Each of these groups has specific questions in mind that are very different from the ones that other groups ask:

- What test cases will be needed to flush out functional errors?
- Where is this design likely to break down?
- Can the design be made easier to test?

- How will this design affect the response of the system to heavy loads?
 - Are there aspects of this design that will affect its performance or ability to scale to many users?
 - What information will users or administrators need to use this system, and can I imagine writing it from the information in this design?
 - Does this design require users to answer configuration questions that they won't know how to answer?
 - Does it create restrictions that users will find onerous?
 - How much translatable text will this design require?
 - Does the design account for the problems of dealing with double-byte character sets or bi-directional presentation?
- Kathryn Heninger Britton (Hoffman and Weiss 2001, pp. 337–338)



Get the habit of analysis—analysis will in time enable synthesis to become your habit of mind.
—Frank Lloyd Wright

3. *Architecture serves as the basis for system analysis and construction.*
 - Architecture tells implementers what to implement.
 - For those interested in the ability of the design to meet the system's quality objectives, the architecture documentation serves as the fodder for evaluation. The architecture documentation must contain the information necessary to evaluate a variety of attributes, such as security, performance, usability, availability, and modifiability. Analyses of each one of these attributes have their own information needs.
 - For system builders who use automatic code-generation tools, the documentation may incorporate the models used for generation.

Table P.1 Some of the stakeholders of architecture documentation, their roles, and how they might use it

Name	Description	Use for Architecture Documentation
Analyst	Responsible for analyzing the architecture to make sure it meets certain critical quality attribute requirements. Analysts are often specialized; for instance, performance analysts, safety analysts, and security analysts may have well-defined positions in a project.	Analyzing satisfaction of quality attribute requirements of the system based on its architecture.

Table P.1 Some of the stakeholders of architecture documentation, their roles, and how they might use it (*continued*)

Name	Description	Use for Architecture Documentation
Architect	Responsible for the development of the architecture and its documentation. Focus and responsibility is on the system.	Negotiating and making trade-offs among competing requirements and design approaches. A vessel for recording design decisions. Providing evidence that the architecture satisfies its requirements.
Business manager	Responsible for the functioning of the business/organizational entity that owns the system. Includes managerial/executive responsibility, responsibility for defining business processes, and more.	Understanding the ability of the architecture to meet business goals.
Conformance checker	Responsible for assuring conformance to standards and processes to provide confidence in a product's suitability.	Basis for conformance checking, for assurance that implementations have been faithful to the architectural prescriptions.
Customer	Pays for the system and ensures its delivery. The customer often speaks for or represents the end user, especially in a government acquisition context.	Assuring required functionality and quality will be delivered, gauging progress, estimating cost, and setting expectations for what will be delivered, when, and for how much.
Database administrator	Involved in many aspects of the data stores, including database design, data analysis, data modeling and optimization, installation of database software, and monitoring and administration of database security.	Understanding how data is created, used, and updated by other architectural elements, and what properties the data and database must have for the overall system to meet its quality goals.
Deployer	Responsible for accepting the completed system from the development effort and deploying it, making it operational, and fulfilling its allocated business function.	Understanding the architectural elements that are delivered and to be installed at the customer's or end user's site, and their overall responsibility toward system function.
Designer	Responsible for systems and/or software design downstream of the architecture, applying the architecture to meet specific requirements of the parts for which they are responsible.	Resolving resource contention and establishing performance and other kinds of runtime resource consumption budgets. Understanding how their part will communicate and interact with other parts of the system.
Evaluator	Responsible for conducting a formal evaluation of the architecture (and its documentation) against some clearly defined criteria.	Evaluating the architecture's ability to deliver required behavior and quality attributes.
Implementer	Responsible for the development of specific elements according to designs, requirements, and the architecture.	Understanding inviolable constraints and exploitable freedoms on development activities.

continues

Table P.1 Some of the stakeholders of architecture documentation, their roles, and how they might use it (*continued*)

Name	Description	Use for Architecture Documentation
Integrator	Responsible for taking individual components and integrating them, according to the architecture and system designs.	Producing integration plans and procedures, and locating the source of integration failures.
Maintainer	Responsible for fixing bugs and providing enhancements to the system throughout its life (including adaptation of the system for uses not originally envisioned).	Understanding the ramifications of a change.
Network administrator	Responsible for the maintenance and oversight of computer hardware and software in a computer network. This may include the deployment, configuration, maintenance, and monitoring of network components.	Determining network loads during various use profiles and understanding uses of the network.
Product line manager	Responsible for development of an entire family of products, all built using the same core assets (including the architecture).	Determining whether a potential new member of a product family is in or out of scope and, if out, by how much.
Project manager	Responsible for planning, sequencing, scheduling, and allocating resources to develop software components and deliver components to integration and test activities.	Helping to set budget and schedule, gauging progress against established budget and schedule, and identifying and resolving development-time resource contention.
Representative of external systems	Responsible for managing a system with which this one must interoperate, and its interface with our system.	Defining the set of agreement between the systems.
System engineer	Responsible for design and development of systems or system components in which software plays a role.	Assuring that the system environment provided for the software is sufficient.
Tester	Responsible for the (independent) test and verification of the system or its elements against the formal requirements and the architecture.	Creating tests based on the behavior and interaction of the software elements.
User	The actual end users of the system. There may be distinct kinds of users, such as administrators, superusers, and so on.	Users, in the role of reviewers, might rely on architecture documentation to check whether desired functionality is being delivered. Users might also refer to the documentation to understand what the major system elements are, which can aid them in emergency field maintenance.

P2.3 Architecture Documentation and Quality Attributes

If architecture is largely about the achievement of quality attributes, and if one of the main uses of architecture documentation is to serve as a basis for analysis (to make sure the architecture will achieve its required quality attributes), where do quality attributes show up in the documentation? There are five major ways:

1. Any major design approach (such as an architecture pattern or style) chosen by the architect will have quality attribute properties associated with it. Client-server is good for scalability, layering is good for portability, an information-hiding-based decomposition is good for modifiability, services are good for interoperability, and so forth. Explaining the choice of approach is likely to include a discussion about the satisfaction of quality attribute requirements and trade-offs incurred. Look for the place in the documentation where such an explanation occurs. In our approach, we call that *rationale*.
2. Individual architectural elements that provide a service often have quality attribute bounds assigned to them. Consumers of the services need to know how fast, secure, or reliable those services are. These quality attribute bounds are defined in the interface documentation for the elements, sometimes in the form of a Quality of Service contract. Or they may simply be recorded as *properties* that the elements exhibit.
3. Quality attributes often impart a “language” of things that you would look for. Security involves things like security levels, authenticated users, audit trails, firewalls, and the like. Performance brings to mind buffer capacities, deadlines, periods, event rates and distributions, clocks and timers, and so on. Availability conjures up mean time between failure, failover mechanisms, primary and secondary functionality, critical and noncritical processes, and redundant elements. Someone fluent in the “language” of a quality attribute can search for the kinds of architectural elements (and properties of those elements) that were put in place precisely to satisfy that quality attribute requirement.
4. Architecture documentation often contains a *mapping to requirements* that shows how requirements (including quality attribute requirements) are satisfied. If your requirements document establishes a requirement for availability, for instance, then you should be able to look up that requirement by name or reference in your architecture document to see the place(s) where that requirement is satisfied.



For more on styles and patterns, see “Coming to Terms: ‘Architecture Style’ and ‘Architecture Pattern’ ” on page 32, in this chapter.



Documenting rationale is covered in Section 6.5.



Interface documentation is covered in Chapter 7.

Properties are discussed in Section I.3, in the introduction to Part I.



Documenting a mapping to requirements is covered in Section 10.3.

5. Every quality attribute requirement will have a constituency of stakeholders who want to know that that quality attribute requirement is going to be satisfied. For these stakeholders, the architect should provide a special place in the documentation’s introduction that either provides what the stakeholder is looking for or tells the stakeholder where in the document to find it. It would say something like “If you are a performance analyst, you should pay attention to the processes and threads and their properties (defined [here]), and their deployment on the underlying hardware platform (defined [here]).” In our documentation approach, we put this here’s-what-you’re-looking-for information in a section called the documentation roadmap.



The documentation roadmap is described in Section 10.2.



The man who stops advertising to save money is like the man who stops the clock to save time. [The same could be said for the architect who stops documenting.]

—Thomas Jefferson

P.2.4 Economics of Architecture Documentation

We’d all like to make our stakeholders happy, of course. Giddy, in fact. So why is producing high-quality architecture documentation often relegated to the “I’ll do it if I have time” category of an architect’s many tasks? Why do project managers often fail to insist that architecture documentation accompany the other archival artifacts produced during development? The answer, of course, is that an architecture document, let alone one that induces giddiness, costs time and money.

Project managers are, by and large, rational people. (No, seriously, they are.) They are willing to invest resources in activities that yield demonstrable benefit, and not so much otherwise. As architects, we should be able to make a business case for producing and maintaining architecture documentation. And here it is: Activities that the project manager is going to have to fund will be less costly in the presence of high-quality, up-to-date documentation than they would otherwise.

A formula to show the savings looks like this:

$$\sum_{\text{over all activities } A} (\text{Cost of } A \text{ without AD} - \text{Cost of } A \text{ with AD}) > \text{Cost of AD},$$

where “Cost of A without AD” and “Cost of A with AD” are the cost of performing activity A without and with (respectively) an architecture document. “Cost of AD” is the cost of producing and maintaining the architecture documentation. In other words, the payback from good architecture documentation should exceed the effort to create it. Payback is measured in terms of effort saved.

This formula gives us a way to think about documentation, its effort, and its payoff. When deciding whether you should produce a particular piece of documentation, ask yourself how

much effort it will take to do so, and what activities will be cheaper as a result. By choosing even a small number of key activities that will benefit from the presence of documentation, you should be able to make a convincing back-of-the-envelope argument that the effort invested will more than pay for itself.

And if you can't—that is, if the effort doesn't pay for itself—then you shouldn't expend it. Put your resources elsewhere.

The formula is nicely general; it does not require that you actually enumerate all the activities involved. The ones that are not affected by the presence or absence of architecture documentation at all simply wash out of the formula. But other activities such as coding, re-engineering, launching a change effort, and so on should have significant cost savings.

P2.5 The Views and Beyond “Method”

We call our approach to documentation Views and Beyond. This is to emphasize that we use the concept of a view—explained in the next section—as the fundamental organizing principle for architecture documentation, but also because we go beyond views to include additional information that belongs in an architecture document.

Views and Beyond is not actually a method. It does not have a sequence of steps, with entry and exit criteria for each. Rather, it is more a collection of techniques that carry out an underlying philosophy. The philosophy is that an architecture document should be helpful to the people who depend on it to do their work (far from least of which is the architect). The techniques can be bundled into a few categories:

1. Finding out what stakeholders need. If you don't do this, you're going to end up with documentation that may serve no one.
2. Providing the information to satisfy those needs by recording design decisions according to a variety of views, plus the beyond-view information.
3. Checking the resulting documentation to see if it satisfied the needs.
4. Packaging the information in a useful form to its stakeholders.

While items 3 and 4 denote document-centric activities, items 1 and 2 denote activities that should be carried out in conjunction with performing the architecture design. That is, we *don't want* Views and Beyond to be an architecture *documentation* method; rather, we want it to help the architect identify and record the necessary design decisions as they are made. Documentation should be the helpful result of making an



Chapter 9 covers a way to use stakeholder needs to determine the views you include in your architecture document.

Chapter 11 covers reviewing documentation.

Chapter 10 covers packaging and organization of documentation.



Don't consider architecture documentation as a task separate from design; rather, make it an essential part of the architecture design process, serving as a ready vessel for holding the output of architectural decisions as soon as those decisions are made.

architecture decision, not a separate step in the architecture process. The more that documentation is treated like a follow-on to design, with its own separate method, the less likely it is to be done at all.



[W]e have come to value . . . working software over comprehensive documentation.

—The Agile Manifesto (Agile Alliance 2002)



Section E.4 in the epilogue elaborates on architecture documentation in an Agile environment.

P.2.6 Views and Beyond in an Agile Environment

It is an unfortunate myth that Agile development and documentation (particularly architecture documentation) are at odds with each other. They aren't, and there are many examples of Agile leaders saying exactly that. Nevertheless, it is possible to interpret the advice in this book as prescribing a heavyweight and cumbersome approach to documentation. You can imagine an architect lagging hopelessly behind the project, which has gone on to deliver the product while he or she is still struggling to complete a Views-and-Beyond-style documentation package from six iterations ago. Neither the architect (nor this book) would likely be invited back to the next project.

Here is some advice that applies to *all* projects but especially to Agile projects: The Views and Beyond approach provides guidance for documenting many kinds of architecture information: structures, elements, relations, behavior, interfaces, rationale, traces to requirements, style guides, system context, and a whole lot more. But nowhere is it written that you have to do all of that. Decide what is useful (you can use the formula in Section P.2.4 to help you decide). Then, for example, if you decide that documenting the rationale behind a certain design decision is going to pay off in the future, then you can use the available guidance to help you do it. If you decide that documenting certain views is useful, then you can use the available guidance to help you do it. And so forth.

Choose what's useful and cost-effective to document. Document that. Period.

P.2.7 Architectures That Change Faster Than You Can Document Them

When your Web browser encounters a file type it's never seen before, odds are that it will go to the Internet, download the appropriate plug-in to handle the file, install it, and reconfigure itself to use it. Without even needing to shut down, let alone go through the code-integrate-test development cycle, the browser is able to change its own architecture by adding a new component.

Service-oriented systems that utilize dynamic service discovery and binding also exhibit these properties. More challenging systems that are highly dynamic, self-organizing, and reflective

(meaning self-aware) are on the horizon. In these cases, the identities of the components interacting with each other cannot be pinned down, let alone their interactions, in any static architecture document.

Another kind of architectural dynamism, equally challenging from a documentation perspective, is found in systems that are rebuilt and redeployed with great rapidity. Some development shops, such as those responsible for commercial Web sites, build and “go live” with their system many dozens of times every single day.

Whether an architecture changes at runtime, or as a result of a high-frequency release-and-deploy cycle, both share something in common with respect to documentation: They change much faster than the documentation cycle. In either case, nobody is going to hold up things until a new architecture document is produced, reviewed, and released.

But knowing the architecture of these systems is every bit as important, and arguably more so, than for systems in the world of more traditional life cycles. Here’s what you can do if you’re an architect in a highly dynamic environment:

1. **Document what is true about all versions of your system.** Your Web browser doesn’t go out and grab just any piece of software when it needs a new plug-in; a plug-in must have specific properties and a specific interface. And it doesn’t just plug in anywhere, but in a predetermined location in the architecture. Record those invariants as you would for any architecture. This may make your documented architecture more a description of constraints or guidelines that any compliant version of the system must follow. That’s fine.
2. **Document the ways the architecture is allowed to change.** In the previous examples, this will usually mean adding new components and/or replacing components with new implementations. In the Views and Beyond approach, the place to do this is called the variability guide.
3. **Make your system capture its own architecture-of-the-moment automatically.** When your Web browser or SOA system crashes, your recovery team is going to want to know exactly what configuration was running when the problem occurred. This ability can run the spectrum from primitive (write changes in a log file) to sophisticated (drive a real-time display of the components and their interactions, much like what is found in network service centers).



Using a variability guide to document an architecture’s variation points is covered in Section 6.4.

P.3 Architecture Views



A **view** is a representation of a set of system elements and the relationships associated with them.



For more information about the bird wing analogy, see “About the Cover” on page xxi.



Chapter 9 shows how to choose the relevant views. Section 10.1 shows how to document a view, and Section 10.2 shows how to document the information that applies to more than one view.



Layered views are covered in Section 2.4. Deployment views are covered in Section 5.2.

Perhaps the most important concept associated with software architecture documentation is that of the **view**. A software architecture is a complex entity that cannot be described in a simple one-dimensional fashion. Our analogy with the bird wing proves illuminating. If you are interested in any but the most superficial understanding, then no single rendition of a bird wing will do. Instead, you need many: feathers, skeleton, circulation, muscular views, and many others. Which of these views *is* the “architecture” of the wing? None of them. Which views *convey* the architecture? All of them.

In this book, we use the concept of views to give us the most fundamental principle of architecture documentation, illustrated in Figure P.1:

Documenting an architecture is a matter of documenting the relevant views and then adding documentation that applies to more than one view.


What are the relevant views? It depends on your goals. As we saw previously, architecture documentation can serve many purposes: a mission statement for implementers, a basis for analysis, the specification for automatic code generation, the starting point for system understanding and asset recovery, or the blueprint for project planning.

Different views also expose different quality attributes to different degrees. Therefore, the quality attributes that are of most concern to you and the other stakeholders in the system’s development will affect the choice of what views to document. For instance, a *layered view* will tell you about your system’s portability, a *deployment view* will let you reason about your system’s performance and reliability, and so forth.

Different views support different goals and uses. This is fundamentally why we do not advocate a particular view or collection

Figure P.1

A documentation package for a software architecture can be composed of one or more view documents and documentation that explains how the views relate to one another, introduces the package to its readers, and guides them through it.



of views. The views you should document depend on the uses you expect to make of the documentation. Different views will highlight different system elements and/or relations.

It may be disconcerting that no single view can fully represent an architecture. Additionally, it feels somehow inadequate to see the system only through discrete, multiple views that may or may not relate to one another in any straightforward way. The essence of architecture is the suppression of information not necessary to the task at hand, and so it is somehow fitting that the very nature of architecture is such that it never presents its whole self to us but only a facet or two at a time. This is its strength: Each view emphasizes certain aspects of the system while deemphasizing or ignoring other aspects, all in the interest of making the problem at hand tractable. Nevertheless, no one of these individual views adequately documents the software architecture for the system. That is accomplished by the complete set of views along with information that transcends them.

The documentation for a view contains

- A primary presentation, usually graphical, that depicts the primary elements and relations of the view
- An element catalog that explains and defines the elements shown in the view and lists their properties
- A specification of the elements' interfaces and behavior
- A variability guide explaining any built-in mechanisms available for tailoring the architecture
- Rationale and design information

The documentation that applies to all of the views contains

- An introduction to the entire package, including a reader's guide that helps a stakeholder find a desired piece of information quickly
- Information describing how the views relate to one another, and to the system as a whole
- Constraints and rationale for the overall architecture
- Such management information as may be required to effectively maintain the whole package



An object-oriented program's runtime structure often bears little resemblance to its code structure. The code structure is frozen at compile-time; it consists of classes in fixed inheritance relationships. A program's runtime structure consists of rapidly changing networks of communicating objects. In fact, the two structures are largely independent. Trying to understand one from the other is like trying to understand the dynamism of living ecosystems from the static taxonomy of plants and animals, and vice versa.

—Gamma et al. (1995, p. 22)



Section 10.1 substantially elaborates this outline.



Section 10.2 substantially elaborates this outline.

COMING TO TERMS

A Short History of Architecture Views

Nearly all modern approaches to designing and documenting architectures rely on the concept of an architectural view. Where did this concept come from?



More than three decades ago, David Parnas (1974) observed that software consists of many structures, which he defined as partial descriptions showing a system as a collection of parts and showing some relations among the parts. This definition largely survives in architecture papers today. Parnas identified several structures prevalent in software. A few were fairly specific to operating systems, such as the structure that defines what process owns what memory segment, but others are more generic and broadly applicable. These include the *module structure*, in which the units are work assignments and the relation is *is-a-part-of* or *shares-part-of-the-same-secret-as*; the *uses structure*, in which the units are programs, and the relation is *depends on the correctness of*; and the *process structure*, in which the units are processes, and the relation is *gives computational work to*.



Quite a bit later, DeWayne Perry and Alexander Wolf recognized that, similar to building architecture, a variety of views of a system are required. Each view emphasizes certain architectural aspects that are useful to different stakeholders or for different purposes (Perry and Wolf 1992).



Later, Philippe Kruchten (1995) of the Rational Software Corporation wrote an influential paper describing four main views of software architecture (logical, process, development, physical) that can be used to great advantage in system building, along with a distinguished fifth view that ties the other four together by showing how they satisfy key use cases: the “4+1” approach to architecture. The 4+1 approach has since been embraced as a foundation piece of the Rational Unified Process.



At about the same time, Dilip Soni, Robert Nord, and Christine Hofmeister of Siemens Corporate Research made a similar observation about views of architecture they found in use in industrial practice (Soni, Nord, and Hofmeister 1995). They wrote about the conceptual view, module interconnection view, execution view, and code view. These views, which correspond more or less to Kruchten’s four views, have become known as the Siemens Four View model for architecture.



To see how the 4+1 views correspond to views described in this book, see Section E.2 of the epilogue.



The Siemens Four View model is explained in the book by Hofmeister, Nord, and Soni (2000).

Other “view sets” have emerged since these. In their book *Software Systems Architecture*, Rozanski and Woods (2005) advocate using functional, information, concurrency, development, deployment, and operational views. Philips Research, the R&D arm of the giant Dutch electronics company, has created the “CAFCR” model of architecture, which calls for five views: the **c**ustomer, **a**pplication, **f**unctional, **c**onceptual, and **r**ealization views.

In the year 2000, the IEEE adopted a standard (IEEE 1471-2000) for architecture descriptions. Unlike approaches that prescribe a fixed set of views, this standard advocates creating your own views that best serve the stakeholders and their concerns associated with your system. (The Views and Beyond approach also advises flexibility in choosing your view set.)



IEEE 1471-2000 is now known as ISO/IEC 42010:2007. We describe this standard in Section E.1 of the epilogue.

P.4 Architecture Styles

Recurring forms have been widely observed, even if written for completely different systems. These forms occur often enough that they are worth writing and learning about in their own right. We call these forms **architecture styles**. (In this book, we usually just say *styles*.) Styles have implications for architecture documentation and deserve definition and discussion in their own right.

Styles allow one to apply specialized design knowledge to a particular class of systems and to support that class of system design with style-specific tools, analysis, and implementations. The literature is replete with a number of styles, and most architects have a wide selection in their repertoires.

For example, we’ll see that modules can be arranged into a useful configuration by restricting what each one is allowed to use. The result is a layered style that imparts to systems that use it qualities of modifiability and portability. Different systems will have a different number of layers, different contents in each layer, and different rules for what each layer is allowed to use. However, the layered style is abstract with respect to these options and can be studied and analyzed without binding them.

For another example, we’ll see that client-server is a common architecture style. The elements in this style are clients, servers, and the protocol connectors that depict their interaction. When used in a system, the client-server style imparts desirable



An **architecture style** is a specialization of element and relation types, together with a set of constraints on how they can be used.



In all processes of life people imitate, and so must artists. They are influenced by their peers as by their antecedents because this is the way of organic development. Late Beethoven and early Schubert, for instance, are almost indistinguishable; while Brahms took certain themes, note for note, from Beethoven; and Shakespeare stole nearly all of his plots—all the good ones certainly.

—Agnes de Mille, American dancer and choreographer (Atlantic 1956)



The layered style is described in Section 2.4.

The client-server style is described in Section 4.3.1.



A **style guide** is the description of an architecture style that specifies the vocabulary of design (sets of element and relationship types) and the rules (sets of topological and semantic constraints) for how that vocabulary can be used.

The contents of a style guide are given in Section 1.2, in the introduction to Part I. Section 6.1.4 discusses how to create and document a new style.



Combining views is an important concept covered in Section 6.6.



A **bridging element** is an element that is common to two views and is used to provide the continuity of understanding from one view to the other. A bridging element appears in both views and has supporting documentation, usually a mapping between views, that makes the correspondence clear, perhaps by showing the combined picture.

properties to the system, such as the ability to add clients with little effort. Different systems will have different protocols, different numbers of servers, and different numbers of clients each can support. However, the client-server style is abstract with respect to these options and can be studied and analyzed without binding them.

Some styles are applicable in every software system. For example, every system is decomposed into modules to divide the work; hence, the decomposition style applies everywhere. Other examples of “universal styles” are uses, deployment, and work assignment. Some styles occur only in systems in which they were explicitly chosen and designed in by the architect: layered, service oriented, and multi-tier, for example.

Choosing a style, whether it’s one covered in this book or somewhere else, imparts a documentation obligation to record the specializations and constraints that the style imposes and the characteristics that the style imparts to the system. We call this piece of documentation a **style guide**. The obligation to document a style can usually be discharged by citing a description of the style in the literature: this book, for example. If you invent your own style, however, you should write a style guide for it because it will help you and your peers to apply that style in other systems.

No system is built exclusively from a single style. On the contrary, every system can be seen to be an amalgamation of many different styles. Some (such as decomposition and work assignment) occur in every system, but in addition to these, systems can exhibit a combination of one or more “chosen” styles as well.

Even restricting our attention to component-and-connector styles, it’s possible for one system to exhibit several styles in the following ways:

- Different “areas” of the system might exhibit different styles. For example, a system might use a pipe-and-filter style to process input data but route the result to a database that is accessed by many elements. This system would be a blend of pipe-and-filter and shared-data styles. Documentation for this system would include (1) a pipe-and-filter view that showed one part of the system and (2) a shared-data view that showed the other part. In a case like this, one or more elements must occur in both views and have properties of both kinds of elements. (Otherwise, the two parts of the system could not communicate with each other.) These **bridging elements** provide the continuity of understanding from one view to the next. They likely have multiple interfaces, each

providing the mechanisms for letting the element work with other elements in each of the views to which it belongs. The filter/database connector in Figure P.2 is an example.

- An element playing a part in one style may itself be composed of elements arranged in another style. For example, a service provider in an SOA system might, unknown to other service providers or its own service users, be implemented using a multi-tier style. Documentation for this system would include an SOA view showing the overall system, as well as a multi-tier view documenting that server, as illustrated in Figure P.3.
- Finally, the same system might simply be seen in different lights, as though you were looking at it through filtered glasses. For example, a system featuring a database repository, as in Figure P.4, may be seen as embodying either a shared-data style or a client-server style. The glasses you choose will determine the style that you “see.”

In the last case, your choice of style-filtered glasses depends, once again, on the uses to which you and your stakeholders intend to put the documentation. For instance, if the shared-data style is more easily understood by the stakeholders that will consume that view, you might choose it. If you need the perspective afforded by more than one style, however, you have a choice. You can document the corresponding views separately, or you can combine them into a single view that is, roughly speaking, the union of what the separate views would be.



This combined view is called an overlay. Overlays are discussed in Section 6.6.

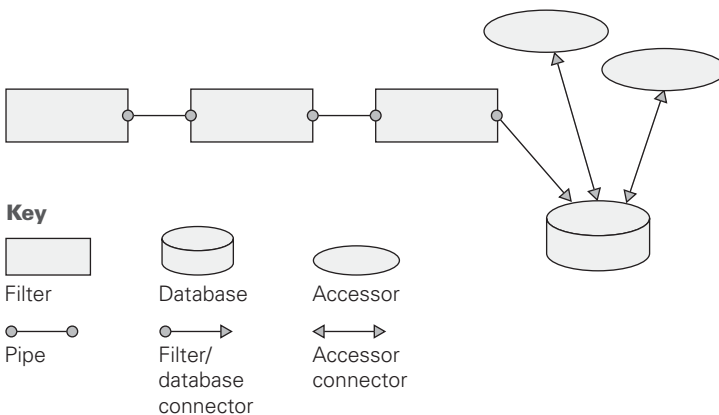


Figure P.2
A system combining a pipe-and-filter style with a shared-data style. The “filter/database connector” is a bridging element.

Figure P.3

A system combining two styles. Here a service provider is composed internally in a multi-tier style.

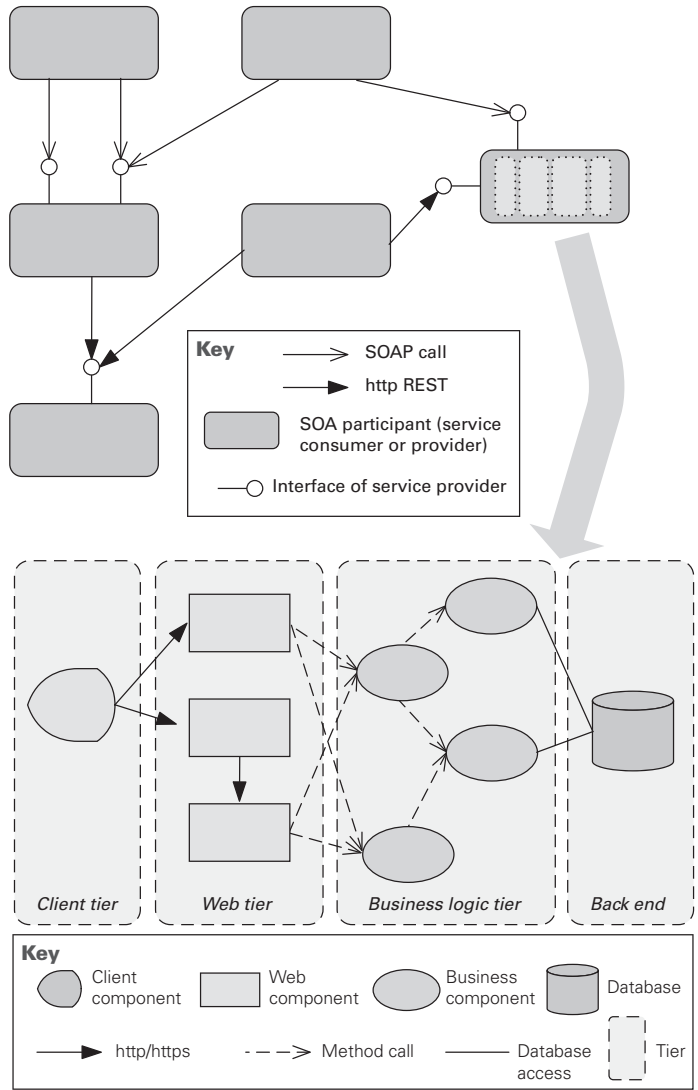



Figure P.4

This system could be in the shared-data style, or the client-server style, depending on your perspective.



P.4.1 Three Categories of Styles

Although no fixed set of views is appropriate for every system, broad guidelines can help us gain a footing. Architects need to think about their software in three ways simultaneously:

1. How it is structured as a set of implementation units
2. How it is structured as a set of elements that have runtime behavior and interactions
3. How it relates to nonsoftware structures in its environment

Each style we present in this book falls into one of these three categories:

1. *Module* styles
2. *Component-and-connector (C&C)* styles
3. *Allocation* styles

When we apply a style to a system, the result is a view. Module views document a system's principal units of implementation. C&C views document the system's units of execution. And allocation views document the relations between a system's software and nonsoftware resources of the development and execution environments.



A selection of module styles is presented in Chapter 2. A selection of C&C styles is presented in Chapter 4. A selection of allocation styles is presented in Chapter 5.

COMING TO TERMS

Module, Component

In this book, we rely on three categories of styles: module, component-and-connector, and allocation. This three-way distinction allows us to structure the information we're presenting in an orderly way and, we hope, allows you to recall it and access it in an orderly way, so that you can write an architecture document that presents *its* information in an orderly way. But for this strategy to succeed, the distinctions have to be meaningful. Two of the categories rely on words for which we give precise meanings, but which are not historically well differentiated: *module* and *component*.

Like many words in computing, these two have meanings outside our field. Furthermore, both terms have come to be associated with movements in software engineering that have overlapping goals.

During the 1960s and 1970s, software systems increased in size and were no longer able to be produced by one



One of the best ways to avoid confusion in your architecture is to be meticulous about making it clear whether each architecture element is a module or a component.

person. It became clear that new techniques were needed to manage software complexity and to partition work among programmers. To address such issues of “programming in the large,” various criteria were introduced to help programmers decide how to partition their software. Encapsulation, information hiding, and abstract data types became the dominant design paradigms of the day. Until this movement, computer programs were largely about calculating the correct answer, but thought leaders were now saying that how you structure your code determines other important properties of the system. *Module* became the carrier of their meaning. The 1970s and 1980s saw the advent of “module interconnection languages” and features of new programming languages such as Modula modules, Smalltalk classes, and Ada packages. Today’s dominant design paradigm—object-oriented programming—has these module concepts at its heart. Components, by contrast, are in the limelight with component-based software engineering and the component-and-connector perspective in the software architecture field.

Both movements aspire to achieve rapid system construction and evolution through the selection, assembly, and wholesale replacement of independent subpieces. Both modules and components are about the decomposition of a whole software system into constituent parts. But beyond that, the two terms take on different shades of meaning.

- A module refers first and foremost to a unit of implementation. Parnas’s foundational work in module design (Parnas 1972) used information hiding as the criterion for allocating responsibility to a module. Information that was likely to change over the lifetime of a system, such as the choice of data structures or algorithms, was assigned to a module, which had an interface through which its facilities were accessed. Modules have long been associated with source code, but information models, XML files, config files, BNF files for parsers, and other implementation artifacts are all perfectly fine modules.
- A component refers to a runtime entity. Szyperski says that a component “can be deployed independently and is subject to composition by third parties” (Szyperski 1998, p. 30). The emphasis is clearly on the finished product and not on the implementation considerations that went into it. Indeed, the operative

model is that a component is delivered in the form of an executable binary only: Nothing upstream from that is available to the system builder.

In short, a module suggests implementation units and artifacts, with less emphasis on the delivery medium and what goes on at runtime. A component is about units of software active at runtime with no visibility into the implementation structure.

Who cares? If every module turned into exactly one component at runtime, it would be easy to sweep the difference under the rug. But this is often far from reality! In many systems, a single module might turn into many components, or it might take many modules to turn into a single component. An easy way to see this is to imagine a trivially simple client-server system. Suppose our system has a single server, which at runtime serves up some interesting piece of data to ten interested clients, all of which do the same thing. This system has *eleven components* but only *two modules*. The server module maps 1:1 onto the server component S1. The client module maps 1:10 to the client components C1–C10. Failing to distinguish between modules and components makes it too easy to blithely assume that every unit of implementation turns into exactly one unit of execution. It isn't so.

Our use of the terms in this book reflects their pedigrees. Module styles described in this book reflect implementation artifact considerations: decompositions that assign parts of the problem to units of design and implementation,

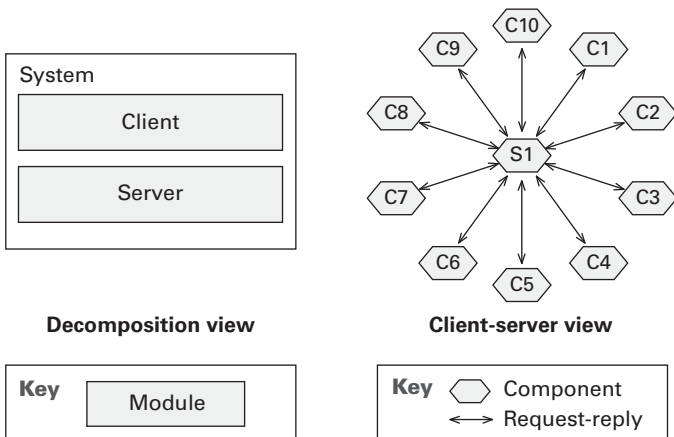


Figure P.5
A client-server system might consist of two modules but eleven components.



Section 10.2 describes how to document the mapping between a system's modules and its components. Sections 1.5 and 3.5 discuss how modules and components relate to each other.

layers that reflect what uses are allowed when software is being written, and classes that factor out commonality from a set of instances. Modules in these styles are often units of source code, but there's also the data model style, where the module is a model of the data that the system manipulates. Of course, all these module styles have runtime implications; that's the end game of software design, after all. C&C styles described in this book focus on how processes interact and data travels around the system during execution.

In many architectures, there is a one-to-one mapping between modules and components. Further, the module and its component counterpart are usually given the same name in this case. This makes it tempting to believe that the modules and components are the same, which in turn makes it tempting to believe there is no difference. Don't be tempted. Although a one-to-one mapping does no harm, the truth is that the module and component are different elements sharing the same name. In such an architecture, the module will show up in a module view, and a component with the same name will show up in one or more component-and-connector views.

Modules and components represent the current bedrock of the software engineering approach to rapidly constructed, easily changeable software systems. As such, modules and components serve as fundamental building blocks for creating and documenting software architectures.

COMING TO TERMS

“Architecture Style” and “Architecture Pattern”

What do the two terms mean?

In this book we use “architecture style” as the term for a package of design decisions that explains a generic design approach for a software system. Another term for a similar concept, used by many architects and authors, is “architecture pattern.” What is the difference between these two concepts and why did we choose style over pattern?

An **architecture style** is a “specialization of element and relation types, together with a set of constraints on how they can be used” (Bass, Clements, and Kazman 2003).

An **architecture pattern** “expresses a fundamental structural organization schema for software systems” (Buschmann et al. 1996, p. 12). It is, above all, a *pattern*, which in the context of architecture “describes a particular recurring design problem that arises in specific design contexts, and presents a well-proven generic scheme for its solution. The solution scheme is specified by describing its constituent components, their responsibilities and relations, and the ways in which they collaborate” (Buschmann et al. 1996, p. 8).

An essential part of an architecture pattern is its focus on the problem and context as well as how to solve the problem in that context. That last part we’ll call the architecture approach. An architecture style focuses on the architecture approach, with more lightweight guidance on when a particular style may or may not be useful. Very informally, we can put it this way (where the arrow means “suggests”):


- Architecture pattern: {problem, context} → architecture approach
- Architecture style: architecture approach

How did these two terms come about?

“Architecture style” as we use it today traces to some early writing from the formative days of software architecture study.



In 1990 and 1991, Mary Shaw was noticing and describing recurring architecture concepts she found in many systems. She called these “elements of a design language for software architecture” or “design idioms” (Shaw 1990, 1991). In 1992 Dewayne Perry and Alexander Wolf wanted to “build an intuition” about the still-new field of software architecture (Perry and Wolf 1992). Looking around at other kinds of architecture—network architecture, computer architecture, and others—they hit upon building architecture as rich in fertile (and borrowable) concepts. One of those concepts was **architecture style**. Like Shaw before them, they were also noticing recurring design forms in software architectures, and they saw that this would be a useful term to appropriate to describe those forms. Styles, then, were observed phenomena, approaches (manifest in the kinds of elements and relations employed) that the authors noticed were being



Thus, we find in building architecture some fundamental insights about software architecture: multiple views are needed to emphasize and to understand different aspects of the architecture; styles are a cogent and important form of codification that can be used both descriptively and prescriptively; and, engineering principles and material properties are of fundamental importance in the development and support of a particular architecture and architectural style.

—Perry and Wolf (1992)

[In building architecture,] architectural styles classify architecture in terms of form, techniques, materials, time period, region, etc. . . . leading to a terminology such as Gothic “style.”

—Wikipedia (2010a)



“An **architectural pattern** expresses a fundamental structural organization schema for software systems. It provides a set of predefined sub-systems, specifies their responsibilities, and includes rules and guidelines for organizing the relationships between them.”
(Buschmann et al. 1996, p. 12)



We must not forget that the wheel is reinvented so often because it is a very good idea; I've learned to worry more about the soundness of ideas that were invented only once.

—D. L. Parnas (1996)

used over and over. The emphasis was on discovery and categorization of utilized forms.



In 1996 Frank Buschmann and his colleagues at Siemens made the inevitable connection between two powerful concepts: software architecture and design patterns (the latter having electrified software engineering the previous year). Their book, *Pattern-Oriented Software Architecture, Volume 1: A System of Patterns* (Buschmann et al. 1996; PoSA, for short), is where the term **architectural pattern** was first used. Followed over the years by (at this writing) four sequels, the PoSA series does for architects what *Design Patterns* (Gamma et al. 1995) did for designers and programmers.



Both design patterns and (software) architecture patterns owe their meaning to the building architect Christopher Alexander, who in the 1970s wrote several books detailing architecture approaches to solve common building design problems. People love to sit next to windows, he wrote, so make every room have a place where they can comfortably do so. People love balconies, he wrote, but observations show they won't spend time on a balcony less than 10 feet wide. So make your balconies at least 10 feet wide. People love outdoor spaces, he wrote, but not if they're in the shadow of a building. So in the northern hemisphere put your courtyards on the south side. He called these design nuggets patterns: “a three-part rule, which expresses a relation between a certain context, a problem, and a solution” (Alexander 1979, p. 247). The patterns community (of whatever flavor) has tried to remain faithful to his meaning.

Why do patterns seem more specific?

It has turned out, not as a matter of the intrinsic nature of these things but rather as a matter of practice, that the published architecture patterns tend to be more constraining—that is, they embed more design decisions—than the published architecture styles. Patterns often look “more detailed” or “less abstract” than styles. Styles tend to tell people what the element and relation types of interest are, and give topological constraints: Put layers on top of layers; pipes connect to filters, not pipes; and so on. Patterns tend to be more specific, showing instances of the element type interacting with each other.

That's because the collectors of styles were motivated to find commonality where none had been observed before. Broad categories are more inclusive. Pattern writers have tended to record very specific and context-dependent problems; hence their solutions are correspondingly specific.

Architects can use this *de facto* distinction to their advantage. For instance, if you're handling a lot of data in your system, you might want to consider a style (the shared-data style is a good candidate) and ask yourself if the element and relation types are what you need: That is, do you *really* need a database? Yes? OK, now go look for a more constrained architecture approach (which might very well be given as a pattern).

Why did we use “architecture style” in this book?

In this book, which is about documenting software architectures and not so much about designing them, we concentrate on presenting a variety of solution approaches—architecture styles—so that we can show how to document systems built using them. In a software architecture document, one doesn't document a pattern, one documents an application of it—that is, the instantiated solution approach.

How do I document the use of a style or pattern in a software architecture document?

Architects can use either patterns or styles as a starting point for their design. They might be published in existing catalogs, stored in an organization's proprietary repository of standard designs, or created specifically for the problem at hand by the architect. In either case, they provide a generic (that is, incomplete) solution approach that the architect will have to refine and instantiate.

First, record the fact that the given style or pattern is being used. Then say why this solution approach was chosen—why it is a good fit to the problem at hand. If the chosen approach comes from a pattern, show that the problem at hand fits the problem and context of the pattern. If the chosen approach comes from a style, explain why the style does the needed job.

Using a pattern or a style means making successive design decisions that eventually result in an architecture. These design decisions manifest themselves as newly instantiated elements and relations among them. The architect can document a snapshot of the architecture at



The shared-data style is described in Section 4.5.1.



The software architecture document templates in Chapter 10 will provide a place for all of this information.



The concept of making successively more constrained design decisions is called a “spectrum of design” and is discussed in Section 6.1.3.



Styles are described using a common set of information; this layout is called a *style guide*. The style guide we use to describe the styles covered in this book is explained in the introduction to Part I.



These are the rules for any technical documentation, including software architecture documentation:

1. Write documentation from the reader's point of view.
2. Avoid unnecessary repetition.
3. Avoid ambiguity.
4. Use a standard organization.
5. Record rationale.
6. Keep documentation current but not too current.
7. Review documentation for fitness of purpose.



The consumer isn't a moron. She is your wife.

—David Ogilvy, writing about advertising

each stage. How many stages there are depends on many things, not the least of which is the ability of readers to follow the design process in case they have to revisit it in the future.

Summary

Architecture styles represent observed architecture approaches. A style description does not generally include detailed problem/context information. Architecture patterns do. An architecture approach might be documented (and several are) as an architecture style *and* an architecture pattern. Both styles and patterns are a set of prepackaged design decisions involving the choice of element types, relation types, properties, and constraints on the topology and interaction among the elements via the relations. Both provide vocabularies that shortcut explanation and allow greatly facilitated communication (“My system is layered.” “Ah, I understand. What are the layers?”), and help chart a course to the satisfaction of specific quality attribute requirements. Both can be used in combination—it is a rare system that uses only one style or one architecture pattern. And both represent essential elements of an architect’s vocabulary.

P.5 Seven Rules for Sound Documentation

Architecture documentation is much like the documentation we write in other facets of our software development projects. As such, it obeys the same fundamental rules for what distinguishes good, usable documentation from poor, ignored documentation. We close the prologue with seven rules for sound software documentation. Use this checklist when you write technical documentation. (You can also use it when you read technical documentation: the rules provide objective criteria for judging a document’s quality, and they let you say something constructive in a critical review.)

Rule 1: Write Documentation from the Reader's Point of View

This rule simply reminds us to keep the end game in mind as we produce our documentation: Make your document serve its stakeholders and their intended uses of it. It is surprisingly easy to forget that rule in the midst of looming deadlines, an overflowing e-mail queue, and a cell phone that won’t shut up.

The great computing scientist Edsger Dijkstra (1930–2002), the inventor of many of the software engineering principles we

now take for granted, once said that he would happily spend two hours pondering how to make a single sentence clearer. He reasoned that if the paper were read by a couple of hundred people—a decidedly modest estimate for someone of Dijkstra’s caliber—and he could save each reader a minute or two of confusion, it was well worth the effort. Professor Dijkstra’s consideration for the reader reflects his classic manners, but it also gives us a new and useful concept of the effort associated with a document. Usually we just count how long it takes to write. Dijkstra taught us to be concerned with how long it takes to *use*. Writing a document that a reader finds easy to use will help tilt the economics of documentation in our favor, as defined in the formula in Section P.2.4.

Writing for the reader is just plain polite, but it has a practical advantage as well. A reader who feels that the document was written with him or her in mind appreciates the effort but, more to the point, will come back to the document again and again in the future. Documents written for the reader will be read; documents written for the convenience of the writer will not. All of us like to shop at stores that seem to want our business, and we avoid stores that do not. This is no different.

Tips on how to write for the reader include:

- Find out who your readers are, what they know, and what they expect of the document. Have an informal chat with some representatives of various kinds of readers and see what their expectations are. Don’t make uninformed assumptions about what your readers know.
- Avoid stream of consciousness writing. If you find yourself writing things down in the order they occur to you, without an overall organizational plan, stop. Work out where specific kinds of information should go and put them where they belong. Make sure that you know what question(s) are being answered by each section of a document.
- Avoid unnecessary insider jargon. The documentation may be read by someone new to the field or from a company that does not share the same jargon. Add a glossary to define specialized terms.
- Avoid overuse of acronyms. Resist using an acronym when the spelled-out phrase is short or it appears only a few times. Always provide a dictionary that decodes whatever acronyms you do use.

Rule 2: Avoid Unnecessary Repetition

Each kind of information should be recorded in exactly one place. This makes documentation easier to use and *much* easier



I have made this letter rather long only because I have not had time to make it shorter.

—Blaise Pascal, French mathematician, physicist, and moralist



The true measure of a man is how he treats someone who can do him absolutely no good.

—Attributed to Samuel Johnson



Rozanski and Woods’s book *Software Systems Architecture* (2005) lists the following properties of an “effective architectural description”: correctness, sufficiency, conciseness, clarity, currency, and precision.

to change as it evolves. It also avoids confusion: information that is repeated is likely to be in a slightly different form, and now the reader must wonder “Was the difference intentional? If so, what is the meaning of the difference? Did the author change one place and forget to update the other?”

It should be a goal that information never be repeated. However, at times the cost to the reader of not repeating information in the other places where it’s needed is high. Readers don’t like to flip pages or click hyperlinks unnecessarily. The information may be repeated in two or more different places for clarity or to make different points. Also, expressing the same idea in different forms is often useful for achieving a thorough understanding. If keeping the information separate comes at too high a cost to the reader, repeat the information.

In a document maintained and viewed online, hyperlinks make this rule easier to follow. For example, each term can be hyperlinked to its definition; a concept can be hyperlinked to an explanation or elaboration.



The data flow diagrams . . . don’t seem to be much use. They’re just vague pictures suggesting what someone thinks might be the shape of a system to solve a problem, and no one’s saying what the problem is. [T]he big picture isn’t much use if it doesn’t say anything you can understand. You’re all just guessing what Fred’s diagram means. It wouldn’t mean anything at all to you if you didn’t already have a pretty good idea of what the problem is and how to solve it.

—A character in a parable about data flow diagrams written by Michael Jackson (1995)

PERSPECTIVES

Beware Notations Everyone “Just Knows”

Rule 3 admonishes us to avoid ambiguity. “A well-defined notation with precise semantics,” we say, “goes a long way toward eliminating whole classes of linguistic ambiguity from a document.” Here we want to emphasize the part about “precise semantics.” Just having a well-defined notation is not enough.

Consider data flow diagrams. Years ago Michael Jackson wrote a wonderful Socratic dialogue that showed how a data flow diagram is largely incapable of conveying useful information about a software design unless you already have a pretty good idea what the design is by the time you start looking at it (Jackson 1995, pp. 42–47; we reprinted the dialogue in Chapter 11 of the first edition of this book [Clements et al. 2003]). Data flow diagrams, for heaven’s sake! They’ve been around for decades. Can it really be that nobody understands what they mean? Jackson was able to show convincingly how easily they can be misinterpreted.

Consider layer diagrams. Layered systems were first described more than four decades ago. We’ve all seen them; we’ve all written them. Yet how many times have

we stopped to ask exactly what they mean? A layer diagram is about the only graphical representation of architecture in which position is significant. Box 1 on top of Box 2 is quite a different system than Box 2 on top of Box 1. What does it mean, exactly, that some rectangles are stacked up on top of each other? “Oh, the programs on top can call programs below” is an answer I often get when I ask this question in class. Well, can programs at the top call *any* programs below, or just the programs in the next lower layer? Ask this question in a room full of professional software engineers, and (if my experience teaching to these groups is any measure) you’ll usually get one-third nods, one-third head shakes, and one-third looking as though you just told them the sun is made of really shiny cheese. Can programs in a layer call other programs in the same layer? Generally the same response. And everyone, absolutely everyone, forgets to tell me that programs below are *not* allowed to call programs above, which is a rather important thing to remember about layers.

So, surprise: Simple layer diagrams are inherently ambiguous. Common variants, such as what I call “layers with a sidecar,” where a vertical box is smooshed up against the stack on one side, are even more ambiguous. (The good news is that they can be easily disambiguated.)

A well-defined notation is one in which you can look at an example and tell whether it’s a legal example of using the notation or not. Layers and data flow diagrams both have this property. But neither, traditionally presented, have precise enough semantics to be unambiguous.

Notations like this, where software engineers “just know” what they mean, are the most dangerous. We all might “know” what a layer diagram means. The problem is that what I “know” it means will be different from what you “know” it means, and different still from what the architect meant. So we’ll all go merrily along with no hint of a problem until late in the project when our errors in understanding may cause us to miss a deadline or suffer an operating failure.

—P.C.



It is far better to be explicit and wrong than to be vague.

—Frederick Brooks, Jr. (1995, p. 259)



Clarity is our only defense against the embarrassment felt on completion of a large project when it is discovered that the wrong problem has been solved.

—C. A. R. Hoare (1985, p. 85)

Rule 3: Avoid Ambiguity

Ambiguity occurs when documentation can be interpreted in more than one way and at least one of those ways is incorrect. The most dangerous kind of ambiguity is undetected ambiguity. Here, each reader will think he or she understands the document, but unwittingly each reader will come to different conclusions about what it is saying.

Following two of the other rules will help you avoid ambiguity:

- By avoiding needless repetition (rule 2), you avoid the “almost but not quite alike” form of ambiguity.
- Reviewing the document with members of its intended audience (rule 7) will help spot and weed out ambiguities.

A well-defined notation with precise semantics goes a long way toward eliminating whole classes of linguistic ambiguity from a document. This is one area where standard languages and notations help a great deal, but using a formal language isn’t always necessary. Simply adopting a set of notational conventions and then using them consistently and rigorously will help eliminate many sources of ambiguity. But if you do adopt a notation, then the following corollary applies:

ADVICE

We have several things to say about box-and-line diagrams masquerading as architecture documentation.

- Don’t be guilty of drawing one and claiming that it’s anything more than a start at an architecture description.
- If you draw one yourself, make sure that you explain precisely what the boxes and lines mean.
- If you see one, ask its author what the boxes mean and what, precisely, the arrows connote. The result is usually illuminating, even if the only thing illuminated is the author’s confusion.

Rule 3a: Explain Your Notation

The ubiquitous box-and-line diagrams that people always draw on whiteboards are one of the greatest sources of ambiguity in architecture documentation. Although not a bad starting point, these diagrams are certainly not good architecture documentation. First, most such diagrams suffer from ambiguity.

Are the boxes supposed to be modules, objects, classes, services, clients, servers, databases, processes, functions, tiers, procedures, processors, or something else? Do the arrows mean calls, uses, data flow, I/O, inheritance, communication, processor migration, or something else?

Make it as easy as possible for your reader to determine the meaning of the notation. The best way to do this is always to include a key in your diagrams. If you're using a standard visual language defined elsewhere, the key can simply name it or refer readers to the source of the language's semantics. Even if the language is standard or widely used, different versions often exist. Let your reader know, by citation, which one you're using. For example, "Key: UML 2.0" is a perfectly fine key, and it puts readers and authors on the same page. For a home-grown informal notation, include a key to the symbology. This is good practice because it compels you to understand what the pieces of your system are and how they relate to one another; it's also courteous to your readers.



Every diagram in the architecture documentation should include a key that explains the meaning of every symbol used. The key should identify the notation. If a predefined notation is being used (such as UML), the key should name it and if necessary cite the document that defines the version being used. Otherwise, the key should define the symbology and the meaning, if any, of colors, shapes, position, and other information-carrying aspects of the diagram. If your diagram uses color but the color has no particular meaning or is only there to enhance readability, say so in the key.


If you define an informal notation for your diagrams, try to use the same notation consistently across diagrams of the same type. Use different symbols for different types of elements and relations. For example, if you used a rounded rectangle for Web components in a diagram, avoid using a different shape for Web components in other diagrams.

PERSPECTIVES

Quivering at Arrows

Many architecture diagrams with an informal notation use arrows to indicate a directional relationship among architecture elements. Although this might seem like a good and innocuous way to indicate that two elements interact, it creates a great source of confusion in many cases. What do the arrows mean?

Consider the following architecture snippet:



What does the arrow mean? Here are some possibilities:

- C1 calls C2.
- Data flows from C1 to C2.
- C1 instantiates C2.
- C1 sends a message to C2.
- C1 is a subtype of C2. (Usually C2 would be positioned above C1, but that is not mandatory.)

- C2 is a data repository and C1 is writing data to C2.
- Conversely, C1 is a repository and C2 is reading data from C1.

Any of these might make sense, and people use arrows to mean all these things and more, often using multiple interpretations in the same diagram.

Suppose we know the arrow indicates that component C1 calls component C2. If your system uses different kinds of calls, it's a good idea to differentiate them in the diagrams. In particular, it is important to distinguish synchronous from asynchronous calls, and local from remote calls. Both aspects may have implications for behavior, performance, modifiability, and reliability of the interaction. It may also be useful to differentiate the technology used to implement the call when the solution will accommodate different ones. For example, a synchronous remote call can be implemented via a Web service such as SOAP, REST, Java RMI, or .NET remoting, among other options. To differentiate the types of interaction in the diagram, use distinct arrowheads (open, closed, solid, hollow) and lines (solid, dotted, dashed, double).

Suppose that we know that C1 calls C2. Sometimes we feel tempted to also show a data flow between the two. We could use the preceding figure and assume the arrow indicates data flow (instead of "calls"), but if C2 returns a value to C1, shouldn't an arrow go both ways? Or should a single arrow have two arrowheads? These two options are not interchangeable. A double-headed arrow typically denotes a symmetric relationship between two elements, whereas two single-headed arrows suggest two asymmetric relationships at work. In either case, the diagram will lose the information that C1 initiated the interaction. Suppose that C2 also invokes C1. Would we need to put *two* double-headed arrows between C1 and C2? When a component C1 calls a component C2, C1 may pass data as arguments to C2 and C2 may return data back to C1. Therefore, it's often a better idea to use the arrow to indicate the call's relation rather than data flow; otherwise the diagram may easily end up full of double-headed arrows that don't tell much.

Although arrows are often used to indicate interactions, often one can avoid confusion by not using them where they are likely to be misinterpreted. For example, one can



SOAP and REST are defined in Section 4.3.3. In previous versions of the SOAP specification, SOAP was an acronym, but this is no longer the case. See www.w3.org/TR/soap12-part1/#intro.

use lines without arrowheads. Sometimes physical placement, rather than lines, can convey the same information. For example, a layer A on top of a layer B indicates that modules in A can use modules in B. Nesting one element inside another often means “is part of.”

Finally, a good key is essential for understanding the meaning of arrows, even ones that represent “simple” interactions such as “calls.” A useful arrow, suitably explained in the key, will leave no doubt as to which is the calling end and which is the called end of a call-return connector, and which way the data flows.

—D.G. and P.M.

Rule 4: Use a Standard Organization

Establish a standard, planned organization scheme, make your documents adhere to it, and ensure that readers know about it. A standard organization, also called a template, offers many benefits.

- It helps the reader navigate the document and find specific information quickly. Thus, this benefit is also related to the write-for-the-reader rule.
- It also helps the document writer plan and organize the contents. The writer doesn’t have to start with a blank page when answering the question “What topics and in what order should I have in this document?” The template already provides an outline of the important topics to cover.
- It allows the writer to record information as soon as it’s known. For example, pieces of section 4 may be written before sections 1–3 are there.
- It reveals what work remains to be done by the number of sections labeled “TBD” (to be determined) or “To Do.”
- It embodies completeness rules for the information; the sections of the document constitute the set of important aspects that need to be conveyed. Hence, the standard organization can form the basis for a first-order validation check of the document at review time.

Corollaries to this rule are these:

1. *Organize documentation for ease of reference.* Software documentation may be read from cover to cover at most once, probably never. But a document is likely to be referenced hundreds or thousands of times. Do what you can to make it easy to find information quickly. Adding a table of contents,



Section I.2, in the introduction to Part I, contains a standard organization for a style guide. Sections 10.1 and 10.2 contain a standard organization that we recommend for documenting views and information beyond views. Chapter 7 contains a standard organization for the documentation of a software interface.



Take any long explanations of figures that are in the main text and move these to the figures’ captions. In-text explanations would serve first-time readers well, but putting explanations in captions will serve second-time readers better: When they see a figure they’re looking for they won’t have to go search the text for its explanation.

—Instructions to the editors of this book, explaining one way in which we tried to organize the book for ease of reference



Don't leave sections blank. Mark them as "not applicable" or "to be determined," as appropriate. Better: "Not applicable because [reason]" and "To be determined by [date or milestone]."

an index, a glossary, and an acronym list are all good ways to help readers look up specific information.

2. *Don't leave any section blank; mark as "TBD" what you don't yet know or "NA" what you know is not applicable.* Many times, we can't fill in a document completely because we don't yet know the information, or because decisions have not been made, or because we didn't yet have time to do it. In that case, mark the document accordingly (for example, "TBD" or "To Do"). Templates are by nature generic and hence comprehensive. If a given section of the template does not apply for the document you're creating, mark it as "NA." If the section is blank, the reader will wonder whether the information is coming later or whether it is indeed supposed to be blank. Thus this advice is related to the rule about avoiding ambiguity.

Rule 5: Record Rationale



"Well, it's an idea, and even a bad idea is better than none," said Master Li. "Error can point the way to truth, while empty-headedness can only lead to more empty-headedness or to a career in politics."

—Barry Hughart, *Bridge of Birds* (1984)

Architecture is the result of making a set of important design decisions, and architecture documentation records the outcomes of those decisions. For the most important decisions, you should record why you made them the way you did. You should also record the important or most likely alternatives you rejected and state why. Later, when those decisions come under scrutiny or pressure to change, you will find yourself revisiting the same arguments and wondering why you didn't take another path. Recording your rationale will save you enormous time in the long run, although it requires discipline to record your rationale in the heat of the moment.



Section 6.5 discusses the documentation of rationale.

Of course, not every single design decision should have the rationale captured in the architecture documentation. If a design decision is key to achieve a quality requirement of the system, its rationale is probably worth capturing. If a design decision required a long meeting with stakeholders, that's a good decision to capture. If you conducted technical experiments and studies or created prototypes to evaluate design alternatives, the conclusions of this effort should be captured as rationale for the chosen alternative. Keep in mind that one week, one month, or one year from now, you may not remember why you did things that way, and other people will not know either.

Rule 6: Keep Documentation Current but Not Too Current

Documentation that is incomplete or out of date does not reflect truth, does not obey its own rules for form and internal consistency, and is not used. Documentation that is kept current and accurate is used. Why? Because questions about the

software can be most easily and most efficiently answered by referring to the appropriate document. Documentation that is somehow inadequate to answer the question needs to be fixed. Updating it and *then* referring the questioner to it will deliver a strong message that the documentation is the final, authoritative source for information.

During the design process, on the other hand, decisions are made and reconsidered with great frequency. Revising documentation to reflect decisions that will not persist is an unnecessary expense.

Your development plan should specify particular points at which the documentation is brought up to date or the process for keeping the documentation current. For example, the end of each iteration or sprint, or each incremental release, could be associated with providing revised documentation. Every design decision should not be recorded and distributed the instant it is made; rather, the document should be subject to version control and have a release strategy, just as every other artifact does.

Rule 7: Review Documentation for Fitness of Purpose

Only the intended users of a document will be able to tell you whether it contains the right information presented in the right way. Enlist their aid. Before a document is released, have it reviewed by representatives of the community or communities for which it was written.

P.6 Summary Checklist

- The goal of documenting an architecture is to write it down so that others can successfully use it, maintain it, and build a system from it.
- Documentation exists to further architecture's uses as a means of education, as a vehicle for communication among stakeholders, and as the basis for analysis.
- Documenting an architecture is a matter of documenting the relevant views and then adding documentation that applies to more than one view.
- Documentation should pay for itself by making development activities less costly.
- Module styles help architects think about their software as a set of implementation units. C&C views help architects think about their software as a set of elements that have runtime behavior and interactions. Allocation views help architects think about how their software relates to the non-software structures in its environment.



Even with the best intentions, sometimes budget and schedule preclude conscientious updating of an architecture document as the system undergoes change. In that case, as happens all too often, the code becomes the final source of authority. Try to use the formula in Section P.2.4 to justify maintaining the document by making a case that doing so is worth the investment. If that fails, then at least mark the sections of the document that are out of date so that readers can still have confidence in the remainder.




Chapter 11 covers reviewing architecture documents.

- An *architecture style* is a specialization of elements and relations, together with a set of constraints on how they can be used. A style defines a family of architectures that satisfy the constraints.
- Some styles are applicable in every software system. Other styles occur only in systems in which they were explicitly chosen and designed in by the architect.
- Follow the seven rules for sound documentation.
 1. Write documentation from the point of view of the reader, not the writer.
 2. Avoid unnecessary repetition.
 3. Avoid ambiguity. Always explain your notation.
 4. Use a standard organization.
 5. Record rationale.
 6. Keep documentation current but not too current.
 7. Review documentation for fitness of purpose.

P.7 Discussion Questions

1. Think of a technical document that you remember as being exceptionally useful. What made it so?
2. Think of a technical document that you remember as being dreadful. What made it so?
3. List several architectural aspects of a system you're familiar with, and state why they are. List several aspects that are not architectural, and state why they are not. List several aspects that are "on the cusp," and make a compelling argument for putting each into "architectural" or "nonarchitectural" categories.
4. If you visit Seoul, Korea, you might see the following sign presiding over one of the busy downtown thoroughfares:



What does it mean? Is the information this sign conveys structural, behavioral, or both? What are the elements in this system? Are they more like modules or like components? What qualities about the notation make this sign understandable or not understandable? Does the sign convey a dynamic architecture, or dynamic behavior within a static architecture? Who are the stakeholders of this sign? What quality attributes is it attempting to achieve? How would you validate it, to assure yourself that it was satisfying its requirements?

5. How much of a project's budget would you devote to software architecture documentation? Why? How would you measure the cost and the benefit?

P.8 For Further Reading

The full treatment of software architecture—how to build one, how to evaluate one to make sure it's a good one, how to recover one from a jumble of legacy code, and how to drive a development effort once you have one—is beyond the scope of this book. However, general books on software architecture are plentiful. Several authors provide good coverage: Bass, Clements, and Kazman (2003); Hofmeister, Nord, and Soni (2000); Shaw and Garlan (1996); Bosch (2000); and Gorton (2006). Also, Jeff Garland and Richard Anthony's *Large-Scale Software Architecture: A Practical Guide Using UML* is a good resource (Garland and Anthony 2003).

The Software Engineering Institute's software architecture Web page—at www.sei.cmu.edu/architecture—provides a wide variety of software architecture resources and links, including a broad collection of definitions of the term (SEI 2010).

One of the goals of documentation is to provide sufficient information so that an architecture can be analyzed for fitness of purpose. For more about analysis and evaluation of software architectures, see the book by Clements, Kazman, and Klein (2002).

The seven rules of sound documentation are adapted from a paper by Parnas and Clements (1986), which also espouses a philosophy directly relevant to this book. That paper holds that although system design is almost always subject to errors, false starts, and resource-constrained compromises, systems should be documented as though they were the product of an idealized, step-by-step, smoothly executed design process. That is the documentation that will be the most helpful in the long run. This book is consistent with that philosophy, in that it lays out what the end state of your documentation should be.

If you want a deeper appreciation of the field of architecture and its roots, then diving into some of the early papers will be worth your time:

David Parnas (1974) first made the observation that software can be described by many structures, not just one. This insight led directly to the concept of views that we use today. Architecture views in general, and “4+1 views” in particular, are a fundamental aspect of the Rational (now IBM Rational) Unified Process for object-oriented software (Kruchten 1995).

An early paper on software architecture that tied us to building architecture and our “architecture styles” to the architecture styles of buildings is by Perry and Wolf (1992).

A tour de force in style comparison is found in the paper by Shaw (1995), in which the author examines 11 different previously published solutions to the automobile cruise-control problem and compares each solution through the lens of architecture style. Chapter 3 of the book by Shaw and Garlan (1996) continues the theme. A number of example problems are presented. For each one, several architecture solutions are presented, each based on the choice of a different style. These side-by-side comparisons not only reveal qualities of the styles themselves, but also richly illustrate the overall concept.

For encyclopedic catalogs of architecture patterns, see the *Pattern-Oriented Software Architecture* series of books by the following authors: Buschmann et al. (1996); Schmidt et al. (2000); Kircher and Jain (2004); and Buschmann, Henney, and Schmidt (2007a and 2007b). Also see Martin Fowler’s book *Patterns of Enterprise Application Architecture* (2002).

Smith and Williams (2002) include three chapters of principles and guidance for architecting systems in which performance is an overriding concern.

Note: *Italicized* page locators indicate figures/tables.

4+1 approach, 404, 406–408, *408*, 429

A

A-7E avionics system, 71–73, 72, 96, *96*, 122

Abstract component in AADL, *474*

Abstract machines, 99

layered style, 90

Abstract module, 83, 84

Abstractions for connectors, 137–138, *138*

Access dependencies in layered style, 95

Accuracy as documentation need, 323

Acme language, 153

Acronyms

in directories, 355

list of, 487–489

overuse, 37

ACSPP (architecture-centered software project planning), 397

Active design reviews, 380–382

Active Reviews for Intermediate Designs (ARID), 396

Activity diagrams, 450, *450*

behavior documentation, 300–302, *301*

SysML, 471, *471*

Actors

element use by, 262–263

use case diagrams, 458

Ada packages, 30

Adventure Builder system

activity diagram, *301*

C&C view, *230*

decomposition style example, 69

interface documentation, 285

multi-tiered system, 184, *184*

SOA view, 172, *173*

state machine diagrams, *305*

uses style example, 79

Advice for question sets, 384

Affected artifacts, 241

Aggregation relations in data model style, 111, 113

Agile Alliance, 429–430

Agile environment

overview, 414–415, *415*

Views and Beyond approach in, 20, 415–418, *418*

Akerman, Art, 239, 246, *247*, 260, 513

Alexander, Christopher, 34, 338, 355

All-view (AV) in DoDAF, 419–421, *420*

Allen, R. J., 153, 287

Allocated-to relations

allocation style, 191

deployment style, 192, 196

install style, 199

SysML module view, 470

work assignment style, 203

Allocation styles, 29, 50

deployment. *See* Deployment style

elements and relations, 191

for end users, 322, *322*

install, 198–202, *201*

miscellaneous, 206–209

overview, 189–191, *190*

work assignment, 202–205, *204*

Allocation views, 29, 50

analysts, 322–324, *324*

application builders, 320–321, *321*

- Allocation views (*continued*)
 - customers, 321–322, 321
 - designers of other systems, 319–320, 319
 - development teams, 317–318, 317
 - DoDAF, 426
 - future architects, 325–326, 325
 - infrastructure support personnel, 324–325, 324
 - maintainers, 320, 320
 - new stakeholders, 325, 325
 - project managers, 316–317, 317
 - SysML, 470–471, 470
 - testers and integrators, 318–319, 318
 - UML, 443–447, 443–447
- Allowed-to-use relations
 - layered style, 87–90, 92–93, 92
 - module views, 57
 - UML module views, 434
 - uses style, 79
- Alternative frames in sequence diagrams, 298
- Alternatives in architectural decisions, 240, 242
- Ambiguity
 - architectural documentation, 40–43
 - UML, 447–449, 448–449
- Ambler, Scott W., 117
- Analysis
 - architecture document reviews, 380, 395
 - module views for, 59–60
- Analysis role, layered style for, 91
- Analysts
 - architecture documentation for, 14
 - documentation needs, 322–324, 324
 - interface documentation for, 279
- Andres, Cynthia, 414, 430
- Angle brackets (<< >>), 433
- Anonymous instances, 438
- ANSI/IEEE Std 1471-2000. *See* ISO/IEC 42010
- Anticrises in architectural decisions, 248
- Application builders' documentation needs, 320–321, 321
- Application programming interfaces (APIs), 280
- Application-specific types in C&C views, 130
- Approved status for architectural decisions, 249
- Araujo, I., 247
- ArchE tool, 85
 - combined views, 255–256
 - generalization style example, 85, 86
 - publish-subscribe style, 177–178, 177
- Architects
 - allocation styles for, 325–326, 325
 - architecture documentation for, 15
- Architectural decisions documentation
 - alternatives, 242
 - choices, 242–244
 - ontology, 247–250, 248–249
 - payback, 245–246
 - purpose, 239
 - templates, 239–242
- Architectural frameworks in ISO/IEC 42010, 404
- Architectural knowledge field, 247
- Architecture Analysis and Design Language (AADL), 153, 197, 305, 312
 - behavior documentation, 482–484, 483
 - C&C view, 478–481, 479, 481
 - deployment view, 481–482, 482
 - interface documentation, 484
 - introduction, 473–474, 475
 - module style, 475–477, 476–478
 - summary, 484–485
- Architecture cartoons, 339
- Architecture-centered software project planning (ACSPP), 397
- Architecture description languages (ADLs)
 - C&C views, 139
 - description, 11
 - notations, 53
- Architecture document reviews, 375
 - active design reviews, 380–382
 - example construction and conducting, 393–395
 - question sets, 382–393, 384
 - steps, 376–380
- Architecture documentation
 - for changing architectures, 20–21
 - as decision making, 246–247, 247
 - economics, 18–19
 - packages, 362–369
 - purpose, 9–10
 - and quality attributes, 17–18
 - seven rules, 36–44
 - SysML, 466
 - terms, 10–12
 - uses and audiences, 12–16
 - beyond views, 351, 351
 - Views and Beyond approach, 19–20
- Architecture Expert (ArchE) tool
 - combined views, 255–256
 - generalization style example, 85, 86
 - publish-subscribe style, 177–178, 177
- Architecture overview presentations, 364–365
- Architecture patterns, 32–36

- Architecture perspectives, 410–411
 - Architecture styles. *See* Styles
 - Architecture Tradeoff Analysis Method (ATAM), 378, 394
 - Architecture views. *See* Views
 - Architectures, product-line, 234
 - Arguments in architectural decision templates, 240, 242
 - Army Training Information Architecture-Migrated (ATIA-M)
 - decomposition style example, 69, 70
 - uses style example, 79, 80
 - Arnold, Ken, 187
 - Arrows
 - C&C diagrams, 440, 447–448
 - confusion from, 41–43
 - UML diagrams, 300–302, 301, 440
 - Artifacts
 - architectural decision templates, 241
 - documents as, 11
 - UML, 444, 444
 - Aspect-oriented programming (AOP), 104, 107–108, 108–109
 - Aspect-oriented software development (AOSD) movement, 108
 - AspectJ language, 107, 122
 - Aspects style
 - elements, relations, and properties, 104–105
 - examples, 106, 107
 - notations, 105–106, 106
 - overview, 104
 - purpose, 105
 - relation to other styles, 106
 - UML module views, 434–435
 - Aspects views, combined, 255
 - Associations
 - UML, 448, 448
 - view-to-view, 251, 353–354, 354
 - Assumptions in architectural decision templates, 240–242
 - Asynchronous calls, 42
 - Asynchronous communication
 - behavior documentation, 292
 - SOA style, 171
 - ATAM (Architecture Tradeoff Analysis Method), 378, 394
 - ATM banking system
 - AOP example, 107–108, 108–109
 - client-server example, 165–166, 165
 - Attachment relations
 - C&C views, 126–127, 131–132
 - client-server style, 163
 - peer-to-peer style, 167
 - pipe-and-filter style, 158
 - publish-subscribe style, 174
 - shared-data style, 179
 - SOA style, 171
 - UML connectors, 142–143, 143
 - Attack resistance in deployment style, 195
 - Attribute-Driven Design (ADD) method, 345–347
 - Audiences for architecture documentation, 12–16
 - Authors in architectural decisions, 248–249
 - Availability
 - deployment style, 195
 - documentation needs, 323
 - Available options for variation points, 237
 - Avritzer, A., 213, 214
- B**
- Bach, Maurice, 98
 - Bachmann, Felix, 373, 509
 - Background threads, 478
 - Backgrounds for presentations, 369–370
 - Backup modes in AADL, 482
 - Balconies, 34
 - Ban decisions in architectural decisions, 248
 - Bandwidth in deployment style, 194
 - Barker, Richard, 116
 - Bass, Len, 5, 32, 47, 71, 122, 510
 - Batch sequential style, 157
 - Beck, Kent, 414, 430
 - Behavior Annex standard in AAL, 483
 - Behavior documentation, 289–290
 - AADL, 482–484, 483
 - communication types, 291–293, 292
 - DoDAF, 426
 - location, 306
 - notation, 293–306, 295–301, 304–305
 - purpose, 306–308
 - questions to answer, 290–291
 - SysML, 471, 471
 - UML, 449–459, 449–459
 - view documentation, 340
 - Behavior Hiding Module in A-7E avionics system, 71, 72
 - Bell LaPadula framework, 482–483
 - Bianco, Phil, 187, 287
 - Big design up front (BDUF), 415
 - Binding time of options for variation points, 237

- Black-box testers, documentation needs, 318–319, 319
 - Blank sections in documentation, 44
 - Bloch, Joshua, 286
 - Boehm, B., 428
 - Booch, Grady, 4, 122, 311
 - Boolean algebra, 280
 - Box-and-line diagrams, 40
 - Bridging, layer, 88
 - Bridging elements, 26, 259
 - Britton, Kathryn Heninger, 14
 - Brooks, Frederick P., Jr., 40, 81, 147
 - Build-versus-buy decisions, 66
 - Buschmann, Frank, 33–34, 48, 187, 356
 - Buses in AADL, 474
 - Business managers, architecture documentation for, 15
 - Business Process Execution Language (BPEL) standard
 - notations, 302–303
 - obtaining, 311
 - orchestration server for, 170
 - Business Process Modeling Notation (BPMN), 303
- C**
- C&C styles. *See* Component-and-connector (C&C) styles
 - C&C views. *See* Component-and-connector (C&C) views
 - CAFRCR model, 25
 - Cai, Y., 213, 214
 - Call-return connectors
 - peer-to-peer style, 167
 - SOA style, 170–171
 - Call-return styles, 156, 157, 161–162
 - client-server, 162–166
 - peer-to-peer, 166–169, 169
 - service-oriented architecture, 169–172, 172
 - Callbacks in layered architecture, 100
 - Calls procedure in C&C view, 150
 - Candidate view lists, 330–331
 - Carroll, Lewis, 352
 - Cataldo, M., 210
 - Categories
 - architectural decisions, 249
 - architecture styles, 29, 49–50
 - Challenged status for architectural decisions, 249
 - Chen, Peter, 116, 118, 122
 - Child modules in generalization style, 82
 - Child view packets in context diagrams, 349
 - Chinese Wall framework, 483
 - Class diagrams in UML, 436–438, 437, 447
 - Class inheritance in generalization style, 84
 - Clements, Paul, 5, 32, 47, 71, 122, 234, 320, 396, 509
 - Client-server connectors in C&C views, 124
 - Client-server style
 - elements, relations, and properties, 162–164
 - examples, 164–166, 165
 - overview, 162
 - purpose, 163–164
 - relation to other styles, 164
 - Cockburn, Alistair, 311, 414, 416, 430
 - Color in layered style, 95
 - Combining primary presentations and context diagrams, 347–348, 348
 - Combining views, 250–251
 - associations between views, 251
 - considerations, 254–255, 328–329
 - examples, 255–257, 256–258
 - process, 252–254, 252–253, 331–332
 - Comment boxes, 225
 - Commercial off-the-shelf (COTS) components, 330
 - Common Language Runtime (CLR), 99
 - Communicating processes, 182–183
 - Communication
 - architectural decision payback, 246
 - behavior documentation, 291–293, 292
 - layered style for, 91
 - module views for, 60
 - SOA style, 171
 - Communication capacity matrices (CCMs), 211–212
 - Communication diagrams
 - behavior documentation, 299–300, 300
 - UML, 453–454, 454
 - Competence-center style, 208
 - Complex connectors in C&C views, 135–136, 135
 - Complexity in module relations, 210–212
 - Component-and-connector (C&C) styles, 29, 50
 - communicating processes, 182–183
 - crosscutting issues, 182–185, 184
 - data flow. *See* Data flow styles
 - dynamic creation and destruction, 184
 - introduction, 155–157, 157
 - repository, 178–182, 181
 - tiers, 183–184, 184

- Component-and-connector (C&C) views
 - AADL, 478–481, 479, 481
 - analysts, 323–324, 324
 - application builders, 321, 321
 - combined, 255
 - connectors, 135–138, 135
 - customers, 321, 321
 - data flow and control flow models, 146–148
 - designers of other systems, 319–320, 319
 - development teams, 318, 318
 - DoDAF, 425–426
 - elements, 126–129
 - end users, 322, 322
 - future architects, 325–326, 325
 - infrastructure support personnel, 324–325, 324
 - maintainers, 320, 320
 - new stakeholders, 325, 325
 - notations, 132, 133, 139–146, 140–145
 - overview, 123–126, 125
 - primary presentations, 347
 - properties, 126, 133–134, 134
 - purpose, 127, 136–137
 - relation to other kinds of views, 148–150, 149
 - relations, 126, 131–132, 133
 - style specialization, 221–222
 - SysML, 469, 469
 - testers and integrators, 318, 319
 - types and instances, 129–131
 - UML, 438–443, 439–442, 460
- Components, 29–32
 - C&C views, 123, 127–128
 - in multiple views, 349
 - replicating, 233, 235–236
 - UML, 139–140
 - UML for C&C views, 438–443, 439–442
- Composite state in state machine diagrams, 304, 457
- Comprehensive models
 - behavior, 294
 - notations, 303–306, 304–305
- Comprises relations in architectural decisions, 250
- Computational models, 155–157
 - client-server style, 163
 - peer-to-peer style, 167
 - pipe-and-filter style, 158
 - publish-subscribe style, 174
 - shared-data style, 179
 - SOA style, 172
- Concept phase in architecture document reviews, 376
- Concepts of operations (CONOPS) in DoDAF, 420
- Conceptual data model, 110, 110
- Concurrency
 - activity diagrams, 301–302
 - C&C views, 134
- Concurrency views for viewpoint sets, 409, 412–413
- Conditional branching in activity diagrams, 302
- Conditions of applicability for variation points, 237
- Configuration management, 368
- Configuration parameters for interfaces, 276
- Conflicts With relations in architectural decisions, 250
- Conformance checkers, architecture documentation for, 15
- Conformance points, 389
- Connecting elements, 4–5
- Connectors
 - abstractions, 137–138, 137
 - C&C views, 123, 128–129, 135–138, 135
 - publish-subscribe style, 174
 - shared-data style, 179
 - SOA style, 170–171
 - UML, 142–143, 438–443, 439–442
- Consistency in presentations, 370
- Constants in interfaces, 274
- Constrains relations in architectural decisions, 250
- Constraints
 - allocation styles, 191
 - aspects style, 105
 - behavior documentation, 293
 - C&C views, 126
 - client-server style, 163
 - data model style, 111
 - decomposition style, 67
 - deployment style, 192
 - generalization style, 83
 - install style, 199
 - layered style, 89
 - model, 172
 - module views, 56
 - peer-to-peer style, 167
 - pipe-and-filter style, 158
 - publish-subscribe style, 175
 - sequence diagrams, 298
 - shared-data style, 179

- Constraints (*continued*)
 - style guides, 51
 - uses style, 75
 - work assignment style, 203
 - Construction, module views for, 59
 - Containment relations in install style, 199
 - Context diagrams, 225–226
 - child view packets, 349
 - combining primary presentations with, 347–348, 348
 - content, 228–229
 - notations, 229–231, 230–231
 - supporting documentation, 229
 - view documentation, 340–343
 - vocabulary of view for, 226–228, 226–228
 - Context views, 341–343
 - Control flow models in C&C views, 146–148
 - Control information in documentation, 351–356, 354–355
 - Conveying risk, architectural decisions for, 246
 - Conway, M. E., 209, 214
 - Coordination views, 209–213
 - Copy-migrates-to relations in deployment style, 192–193
 - Costs
 - architectural decisions, 249
 - architecture documentation, 18–19
 - combined views, 255
 - CPU properties in deployment style, 193
 - Creates, reads, updates, or deletes data (CRUD) matrices, 118
 - Cripps, P., 360
 - Criticality in question sets, 384
 - Crosscuts relations
 - aspects style, 104–105
 - module views, 57, 435, 435
 - Crow’s foot ERD notation, 116, 116
 - Crystal Clear approach, 414
 - CSP language, 304–305
 - Currency in architectural documentation, 44–45
 - Customers, 15, 321–322, 321
- D**
- Dahl, Ole-Johan, 122
 - Dashofy, E. M., 5
 - Data accessor component in shared-data style, 179
 - Data-centric approach in DoDAF, 427
 - Data communication in behavior documentation, 292
 - Data elements, 4
 - Data entities in data model style, 111
 - relationships, 57, 62, 109, 110, 111, 116, 117, 118, 120
 - Data flow diagrams, 38–43
 - Data flow models in C&C views, 146–148
 - Data flow styles, 156–157, 157
 - call-return, 161–162
 - client-server, 162–166
 - event-based, 172–178, 173, 177
 - peer-to-peer, 166–169, 169
 - pipe-and-filter, 158–161, 161
 - service-oriented architecture, 169–172, 172
 - Data in AADL, 474
 - Data integrity, 115
 - Data model style
 - elements, relations, and properties, 111–113, 113–114
 - entities, 118–119
 - examples, 118, 119
 - notations, 116–117, 116–117
 - overview, 109–111, 110–111
 - purpose, 111, 114–115
 - relationships, 57, 62, 109, 110, 111, 116, 117, 118, 120
 - relations to other styles, 117–118
 - UML module views, 435–436, 436
 - Data ports in AADL, 479–480
 - Data reading and writing connectors in shared-data style, 179
 - Data stores style, 207
 - Data transfer objects (DTOs), 98
 - Data Types and Constants section in interface documentation, 283
 - Data types in interfaces, 274, 283
 - Database access connectors in C&C views, 124
 - Database administrators, architecture documentation for, 15
 - Database management systems (DBMSs), 178
 - Date, C. J., 113, 122
 - de Boer, R. C., 260
 - de Mille, Agnes, 25
 - Debugging, uses style for, 82
 - Decided status in architectural decisions, 249
 - Decision making, architecting as, 246–247, 247
 - Decisions in architectural templates, 240
 - Decomposition refinement, 218–219, 218–219
 - Decomposition relation, 66

- Decomposition style, 65
 - elements, relations, and properties, 66–67, 67
 - examples, 69–73, 70
 - notations, 67–68, 68
 - overview, 65–66
 - purpose, 67
 - relation to other styles, 68–69
 - UML module views, 433, 433
- Decomposition views in combined views, 255–256, 256
- Denial of service, 323
- Department of Defense Architecture Framework (DoDAF), 377, 404, 419
 - documentation strategies, 426
 - overview, 419–420, 420
 - and software architecture, 421
 - version 2.0, 427, 427
 - and Views and Beyond, 421–425
- Dependability in client-server style, 164
- Dependencies
 - in layered style, 95
 - UML module views, 433–434, 434
 - variation point options, 237
- Dependency arrows in C&C diagrams, 447–448
- Dependency structure matrices (DSMs)
 - layered style, 101–102, 102–103
 - module views, 62
 - uses style, 77, 78
- Depends-on relations
 - module views, 57, 60, 61
 - modules, 49
 - SysML module view, 468
 - uses, 74–75, 81
- Depends on the correctness of relation in
 - architecture views, 24
- Deployers, architecture documentation for, 15
- Deployment style
 - elements, relations, and properties, 192–194
 - notation, 196–198, 196
 - overview, 191–192
 - purpose, 194–195
 - relation to other styles, 197
 - UML allocation view, 443–445, 443–445
- Deployment views, 22
 - 4+1 approach, 407
 - AADL, 481–482, 482
 - combined views, 255, 257, 257–258
 - Rozanski and Woods viewpoint sets, 409, 412–413
- Deprecated methods, 273
- DeRemer, F., 64
- Descriptions, architecture, 11
- Descriptive architecture documentation, 12
- Descriptive completeness, 222–225, 223–225
- Design approaches, documenting, 356
- Design idioms, 33
- Design issues for interfaces, 277
- Design reviews, active, 380–382
- Design Structure System, 101
- Design vs. architecture, 6–9
- Designers, architecture documentation for, 15
- Designers of other systems, documentation needs, 319–320, 319
- Detailed design, 7–8
- Details in architecture document reviews, 379, 394–395
- Developers, interface documentation for, 278–279
- Development activities in behavior documentation, 306–307
- Development phase in architecture document reviews, 377
- Development teams, documentation needs, 317–318, 318
- Development views in viewpoint sets, 409, 412–413
- Devices in AADL, 474
- Diacrisis in architectural decisions, 248
- Diagrams
 - activity. *See* Activity diagrams
 - class, 436–438, 437, 447
 - communication, 299–300, 300, 453–454, 454
 - context. *See* Context diagrams
 - ERDs, 62, 116, 116
 - interaction overview, 455–456, 456
 - interfaces, 268–270, 269–270
 - limitations, 38–43
 - sequence. *See* Sequence diagrams
 - state machine, 303–304, 304, 457, 458
 - timing, 302, 454, 455
 - use case, 458, 459
- Dickinson, Emily, 375
- Dijkstra, Edsger W., 36–37, 121, 242
- Directional flow in AADL, 480
- Directional relationships, 41–43
- Directories, 355–356
- Disk capacity in deployment style, 194
- Distributed computing applications, 168
- Document control information, 351–356, 354–355

- Documentation beyond views, 350–351
 - behavior documentation in rationale, 306
 - standard organization for, 351–356, 351, 354–355
 - variations, 356–357
- Documentation capacity matrices (DCMs), 211–212
- Documentation overview, 9–12
- Documentation packages. *See* Packages
- Documentation roadmap, 352
- DoDAF. *See* Department of Defense Architecture Framework (DoDAF)
- Downstream filters in pipe-and-filter style, 159
- Drawing tools in UML, 462
- Dynamic architecture, 234–235
- Dynamic creation and destruction in C&C styles, 184
- Dynamism and dynamic architectures, 234–235

- E**
- e-Business Reference Architecture Framework, 260
- E-mail communication in behavior documentation, 292
- Ease of reference, organization for, 43
- Eclipse UI event manager, 177
- ECMA toaster model, 121, 121
- Economics in architecture documentation, 18–19
- Education, architecture as, 12
- Eeles, Peter, 360–361, 361, 513
- Element catalogs
 - behavior documentation, 306
 - view documentation, 340
 - whole architecture, 356–357
- Element developers, interface documentation for, 278
- Element-level behavior documentation in DoDAF, 426
- Element properties in styles, 52–53
- Element substitution, 233, 235
- Elements, 4–5
 - allocation styles, 191
 - aspects style, 104–105
 - bridging, 26, 259
 - C&C views, 126–129
 - client-server style, 162–163
 - data model style, 111–113
 - decomposition style, 66–67, 67
 - deployment style, 192
 - generalization style, 83
 - hierarchical, 8
 - install style, 199–200
 - interfaces, 262–263
 - layered style, 89
 - module views, 56–57
 - peer-to-peer style, 166–167
 - pipe-and-filter style, 158–159
 - publish-subscribe style, 174–175
 - shared-data style, 179
 - SOA style, 169–171
 - style guides, 50
 - tools for, 371
 - uses style, 75
 - in variability, 233
 - variation points, 236
 - work assignment style, 202–203
- Ellipses (...), 224
- Emery, David, 375, 400, 513
- Emphasis in presentations, 369
- Enables relations in architectural decisions, 250
- Encapsulation in module views, 58
- End users, documentation needs, 322, 322
- Enterprise Architect tool, 472
- Enterprise service bus (ESB), 170
- Entities in data model style, 118–119
- Entity-relationship diagrams (ERDs)
 - data model style, 116, 116
 - module views, 62
- Entry actions in state machine diagrams, 457
- Environmental elements
 - allocation styles, 190–191
 - allocation views, 443, 470
 - deployment style, 192, 196, 443
 - install style, 199–200, 445
 - work assignment style, 203, 446
- EOSDIS Core System (ECS)
 - uses style example, 79, 80
 - view selection example, 329–332
 - work assignment views, 204, 204
- Error handling
 - in interface documentation, 283–284
 - interfaces, 273–276, 276
 - terms, 277–278
- Error Modeling Annex in AADL, 473, 483, 484
- Error state machine in AADL, 483
- Evaluators, architecture documentation for, 15
- Event-based styles, 156, 157, 172–178, 173, 177
- Event ports in AADL, 479

- Execution-migrates-to relations in deployment style, 192–193
 - Executive decisions in architectural decisions, 248
 - Existence decisions in architectural decisions, 248
 - Exit actions in state machine diagrams, 457
 - Expected answers in question sets, 383
 - Exploratory robot, 73, 74
 - Extends relation
 - AADL module style, 475
 - use case diagrams, 459, 459
 - Extension, generalization style for, 84
 - External memory from architectural decisions, 245
 - Extreme Programming, 414
- F**
- Failover Request/Reply connectors, 144
 - Fairbanks, George, 289, 513–514
 - Farenhorst, Rik, 246, 260, 514
 - Fault tolerance
 - AADL, 482
 - deployment style, 194
 - Feature-Driven Development, 414
 - Feiler, Peter H., 312, 473–474, 476, 483, 514
 - Felsing, John, 414, 430
 - Filters in pipe-and-filter style, 158–159
 - Final state in state machine diagrams, 457
 - Finkelstein, Clive, 116
 - Fitness of purpose
 - architectural documentation, 44–45
 - document reviews for, 375
 - Flight Operations Segment (FOS) in ECS module, 204
 - Flint, Emily, 375
 - Flow charts in activity diagrams, 450
 - Flurry, G., 260
 - Focus in behavior documentation, 291
 - Forbids relations in architectural decisions, 250
 - Fork nodes in activity diagrams, 301
 - Formal notations
 - architecture views, 53
 - C&C views, 139
 - deployment style, 197
 - Fowler, Martin, 48, 122, 311
 - Frames in UML, 451, 452
 - Frameworks in variability, 233, 237
 - Freeman, Eric, 187
 - FSP language, 304–305
 - Functional views in viewpoint sets, 408, 411–413
 - Functionality in C&C views, 133
 - Future architects, documentation needs, 325–326, 325
- G**
- Gamma, E., 23, 34, 178
 - Garlan, David, 4, 153, 187, 287, 510
 - Gelernter, D., 187
 - Generalization concept in AADL module style, 475
 - Generalization/specialization relations in data model style, 111, 113
 - Generalization style
 - elements, relations, and properties, 83–84
 - examples, 85–87, 86
 - notations, 84, 85
 - overview, 82–83
 - purpose, 83–84
 - relation to other styles, 84
 - UML module views, 434
 - Generalization views, combined, 255–256, 256
 - Generalization vs. inheritance, 120
 - Generators in variability, 233
 - Gives computational work to relation, 24
 - Global policies, 349–350
 - Glossary of terms
 - in documentation beyond views, 355
 - list, 491–496
 - Gluch, David P., 312
 - Gnutella peer-to-peer networks, 168–169, 169
 - Gorton, Ian, 220
 - Graphical user interfaces in publish-subscribe style, 174
 - Groups in architectural decision templates, 240
 - Guard constraints on transitions, 457
 - Guillemets (<< >>), 433
- H**
- Hansson, Jörgen, 483
 - Hardware components in AADL, 474
 - Hardware Hiding Module in A-7E avionics system, 71, 72
 - Harel, David, 303, 311
 - Harvey, Miles, 350
 - Henderson, Stuart, 9
 - Henney, Kevin, 48, 187
 - Herbsleb, James D., 209–213, 514
 - Hierarchical elements, 8

- High Level Architecture (HLA), 187
 - Hilliard, Rich, 375, 400, 514
 - History in architectural decisions, 248–249
 - History state in state machine diagrams, 457
 - Hoare, C. A. R., 40
 - Hofmeister, Christine, 24, 64, 214, 336
 - Hohpe, Gregor, 187
 - Hudak, John, 312
 - Hughart, Barry, 44
 - Humphrey, Watts S., 379
 - Hupfer, Susanne, 187
 - Hybrid styles in combined views, 253
 - Hybrid threads, 478
 - Hyman, Isabelle, 399
 - Hypertext, 365–366
 - Hypertext Markup Language (HTML), 164–165
 - Hypertext Transfer Protocol (HTTP), 165
- I**
- Idea status in architectural decisions, 249
 - IDEFIX notation, 116
 - Identification and overview information in
 - ISO/IEC 42010, 405
 - Identifying relation in data model style, 113
 - Identity in interfaces, 271
 - Implementation constraints in module views, 59
 - Implementation information property, 58–59
 - Implementation refinement, 219, 220
 - Implementation style, 207
 - UML allocation view, 445–446, 446
 - Implementation views
 - 4+1 approach, 406–407
 - combined views, 254
 - use by infrastructure support personnel, 324
 - Implementers, architecture documentation for, 15
 - Implications in architectural decision templates, 240–241
 - Implicit invocation in publish-subscribe style, 175
 - Include relations in use case diagrams, 459, 459
 - Incremental development, uses style for, 81–82
 - Indexes in directories, 355
 - Informal notations
 - architecture views, 53
 - C&C views, 139
 - context diagrams, 229–230, 230
 - decomposition style, 67–68
 - deployment style, 196
 - install style, 200, 201
 - layered style, 92–95, 92–95
 - module views, 60, 61
 - uses style, 76
 - work assignment style, 203
 - Information Engineering approach, 116
 - Information systems, data model style for, 115
 - Information views in viewpoint sets, 409, 412–413
 - Infrastructure support personnel, documentation needs, 324–325, 324
 - Inheritance
 - vs. generalization, 120
 - generalization style, 84
 - as a variation mechanism, 233
 - Initial state in state machine diagrams, 457
 - Install style
 - elements, relations, and properties, 199–200
 - notations, 200, 201
 - overview, 198–199
 - purpose, 200
 - relation to other styles, 202
 - UML allocation view, 445–446, 446
 - Install views in combined views, 255
 - Instances
 - AADL, 480, 481
 - of components and connectors, 129–131
 - UML, 448
 - UML C&C views, 438, 439
 - Integration testing, uses style for, 82
 - Integrators
 - architecture documentation for, 16
 - documentation needs, 318–319, 319
 - interface documentation for, 278–279
 - Interaction overview diagrams in UML, 455–456, 456
 - Interactions of connectors, 129
 - Interface Definition Language (IDL), 279
 - Interface delegation in C&C views, 126, 132
 - Interface documentation, 261
 - AADL, 484
 - behavior documentation in, 306
 - data types and constants, 274
 - diagrams, 268–270, 269–270
 - error handling, 273–276, 276, 283–284
 - examples, 281–285
 - extended by generalization, 263
 - guidelines, 266–267
 - identity, 271, 281
 - multiple, 270
 - overview, 261–263, 265–268

- provided vs. required, 264–265
- quality attributes, 276–277, 284
- resources, 271–274, 279
- stakeholders, 278–279
- standard organization, 271–277, 271, 276
- syntactic information, 279–280
- terms, 280–281
- UML, 460–461, 460–461
- variability, 276
- Interface inheritance in generalization style, 84
- Interface realization in generalization style, 83–84
- Interfaces
 - layered style, 90–91
 - SysML, 472, 472
 - view documentation, 340
- International Council on Systems Engineering (INCOSE), 465
- Invokes-services roles in C&C views, 128
- Is-a relations
 - generalization style, 82
 - module views, 57
 - modules, 49
 - SysML module view, 468
 - UML module views, 434
- Is an Alternative To relations in architectural decisions, 250
- Is Bound To relations in architectural decisions, 250
- Is-part-of relations
 - A-7E avionics system, 71
 - architecture views, 24
 - decomposition style, 65–67
 - module views, 57, 60, 61
 - modules, 49
 - SysML module view, 468
- Is Related To relations in architectural decisions, 250
- ISO/IEC 42010
 - overview, 400–404, 401, 403
 - question sets for, 391–393
 - Views and Beyond approach, 404, 405
- Issues in architectural decision templates, 240
- Iterations in Agile projects, 414
- Ivers, James, 510
- J**
- Jackson, Michael, 38, 260
- Jacobson, Ivar, 4, 122, 187, 311
- Jain, Prashant, 187
- Jargon, 37
- Java EE application, 98–99, 99
- Java modules, ATIA server-side, 69, 70
- Java Virtual Machine (JVM), 99
- JavaScript Object Notation (JSON), 279
- JavaServer Faces (JSF) action classes, 98–99
- Jazayeri, Mehdi, 2
- Jefferson, Thomas, 18
- Johnson, Samuel, 36
- Join nodes, activity diagrams, 301
- K**
- Kazman, R., 5, 32, 47, 71, 122, 396
- Kiczales, Gregor, 122
- Kircher, Michael, 187
- Klein, John, 333–335, 514
- Klein, M., 47, 396
- Knuth, Donald, 241
- Koala language, 306
- Komiya, S., 260
- Kotermanski, Rick, 187
- Kron, H. J., 64
- Kruchten, Philippe, 24, 48, 221, 239, 515
 - 4+1 approach, 360, 404, 406–408, 408, 429
 - architectural decisions, 247–249, 248–249, 260
- Kylmäkoski, Roope, 336
- L**
- Laddad, Ramnivas, 122
- Lago, P., 247
- Language of quality attributes, 17
- Layer bridging, 88
- Layer diagrams, 38–39
- Layered process for combined views, 253
- Layered style
 - calling higher layers, 100–101
 - Dependency Structure Matrix for, 101–102, 102–103
 - elements, relations, and properties, 89–90
 - examples, 97–99, 97, 99
 - notations, 92–95, 92–95
 - overview, 87–89, 87–88
 - purpose, 89–91
 - relation to other styles, 96–97, 96
 - UML module views, 434, 435
- Layered views, 404
 - combined views, 255
 - and portability, 22
- Lewis, Grace, 287, 474

- Life-cycle phases in architecture document reviews, 376–377
 - Linda programming language, 187
 - Links in communication diagrams, 299–300, 300
 - Liskov, Barbara, 122
 - Little, Reed, 510
 - Local calls vs. remote calls, 42
 - Local change and variation, generalization style for, 84
 - Local communication in behavior documentation, 293
 - Logical data model, 110, 110
 - Logical views in 4+1 approach, 406
 - Lollipop/socket notation, 141, 141, 268, 269
 - Loop frames in sequence diagrams, 298
- M**
- MagicDraw tool, 472
 - Mailing lists, publish-subscribe style for, 174
 - Maintainers
 - architecture documentation for, 16
 - documentation needs, 320, 320
 - interface documentation for, 278
 - Major design approaches, documenting, 356
 - Management information in module views, 59
 - Many-to-many associations in views, 251
 - Many-to-many relations in data model style, 111
 - Many-to-one associations, 251
 - Mapping
 - module and C&C views, 149
 - to requirements, 17, 357–362, 358, 361, 371
 - to source code units, 59
 - between views, 353–354, 354
 - Mars Climate Orbiter, 281
 - Mars exploratory robot, 73, 74
 - Martin, James, 116
 - McGregor, John D., 465, 515
 - Medvidovic, N., 5, 153
 - Memory for AADL, 474
 - Memory properties in deployment style, 193
 - Merge nodes in activity diagrams, 301–302
 - Merson, Paulo, 187, 287, 373, 511
 - Message sequence charts, 302
 - Messaging connectors in SOA style, 171
 - Meta-Model in AADL, 473
 - Microsoft .NET Pet Shop application, 118
 - Migrates-to relations in deployment style, 192–193
 - Migration triggers in deployment style, 194
 - Mode transitions in AADL, 482
 - Model-view-controller (MVC) pattern, 174
 - Modeling tools in UML, 462
 - Modifiability
 - documentation needs, 323
 - layered style, 98–99
 - Modula modules, 30
 - Modularization, 209–212
 - Module decomposition and layered style, 96, 96
 - Module interfaces in UML, 460
 - Module structure in architecture views, 24
 - Module styles, 29, 65
 - AADL, 475–477, 476–478
 - aspects, 104–108, 106–109
 - data model. *See* Data model style
 - decomposition. *See* Decomposition style
 - generalization. *See* Generalization style
 - layered. *See* Layered style
 - use. *See* Uses style
 - Module views
 - analysts, 323–324, 324
 - application builders, 321, 321
 - vs. C&C view, 148–149, 149
 - customers, 321, 321
 - designers of other systems, 319–320, 319
 - development teams, 318, 318
 - DoDAF, 426
 - elements, 56–57
 - end users, 322, 322
 - future architects, 325–326, 325
 - infrastructure support personnel, 324, 324
 - maintainers, 320, 320
 - new stakeholders, 325, 325
 - notations, 60–62, 61–62
 - overview, 55–56, 56
 - project managers, 317, 317
 - properties, 57–59, 58
 - purpose, 59–60
 - relation to other views, 63
 - relations, 57
 - subsystems, 73–74, 74
 - SysML, 468, 468
 - testers and integrators, 318, 319
 - UML, 433–436, 433–436
 - Modules, 29–32, 31
 - in multiple views, 349
 - styles, 49
 - Morale, uses relation for, 81
 - Multi-part primary presentations, 349
 - Multiple interfaces, 262–263, 270

N

- Name property in module views, 57
- Named frames in sequence diagrams, 297
- Naming conventions in UML, 140
- Natural language, 280
- Navigable associations in UML, 440
- Neighbors, J. M., 64
- Nesting notation in decomposition style, 68
- Nesting of states in state machine diagrams, 304
- Network administrators, architecture documentation for, 16
- New stakeholders, documentation needs, 325, 325
- Nii, H. P., 187
- Nonarchitectural design, 7
- Nonexistence decisions in architectural decisions, 248
- Nonidentifying relations in data model style, 113
- Nord, Robert, 24, 64, 214, 336, 511
- Normalization in data model style, 113, 113–114
- Northrop, Linda, 234, 320
- Notations
 - architecture views, 53
 - formal, 53
 - informal, 53
 - semiformal, 53
 - aspects style, 105–106, 106
 - behavior documentation, 293–306, 295–301, 304–305
 - C&C views, 132, 133, 139–146, 140–145
 - comprehensive models, 303–306, 304–305
 - context diagrams, 229–231, 230–231
 - data model style, 116–117, 116–117
 - decomposition style, 67–68, 68
 - deployment style, 196–198, 196
 - explaining, 40–41
 - generalization style, 84, 85
 - install style, 200, 201
 - layered style, 92–95, 92–95
 - module views, 60–62, 61–62
 - style guides, 51
 - for traces, 295–303, 295–301
 - uses style, 76–77, 77–78
 - work assignment style, 203–204, 204
- Notes in architectural decision templates, 241
- Nygaard, Kristen, 122

O

- Object Constraint Language (OCL), 460–461
 - Object Management Group (OMG), 465
 - AADL standards, 474
 - IDL standard, 279
 - OCL standard, 460
 - SysML standard, 465
 - UML specification, 122, 311, 433
 - Object-oriented designs, generalization style for, 84
 - Object-relational impedance matching, 117
 - Object-relational mapping (ORM) tools, 117
 - Obsolesced status in architectural decisions, 250
 - Ockerbloom, J., 153, 287
 - O’Connell, Don, 419, 515
 - Ogilvy, David, 36
 - OMG. *See* Object Management Group (OMG)
 - One-to-many associations in views, 251
 - One-to-many relations in data model style, 111
 - One-to-one relations, 111
 - Online documentation for packages, 365–366
 - Ontocrisis in architectural decisions, 248
 - Ontology in architectural decisions, 247–250, 248–249
 - Open questions, 357
 - Open-source style, 208
 - Operational view (OV)
 - DoDAF, 419, 422
 - Rozanski and Woods viewpoint sets, 410, 413
 - Optional inclusion, 233, 236
 - Orchestration servers in SOA style, 170–171
 - Order Processing Center (OPC) component in Adventure Builder, 173, 184, 230, 285
 - Ordering constraints in behavior documentation, 293
 - OSATE tools, 485
 - OSGi framework, 121, 121
 - Overhead in layered style, 91
 - Overlays in combined views, 252–253
 - Overrides relations in architectural decisions, 250
 - Overview presentations, 364–365
- P**
- Packages
 - AADL module style, 475
 - architecture documentation, 349–350, 362–369

- Packages (*continued*)
 - configuration management, 368
 - documentation beyond views, 350–357, 351, 354–355
 - mapping to requirements, 357–362, 358, 361
 - online documentation and hypertext, 365–366
 - release strategy, 368–369
 - schemes, 362–364
 - tooling requirements, 370–372
 - UML, 61, 76, 77–78, 355, 433
 - wikis, 365–368
- Pair programming, 414
- Palmer, Stephen, 414, 430
- Parameterization in variability, 233
- Parent modules in generalization style, 82
- Parnas, David L., 24, 30, 34, 47, 48, 104, 121, 122, 209, 286–287, 396
- Partial results in error handling, 278
- Pascal, Blaise, 37
- Patterns, architectural
 - and styles, 32–36
 - and views, 343
- Paulish, D., 213, 214, 397
- Payback from architectural decisions, 245–246
- Peer-to-peer style
 - elements, relations, and properties, 166–167
 - examples, 168–169, 169
 - overview, 166
 - purpose, 167
 - relation to other styles, 168
- Performance
 - C&C views, 133
 - client-server style, 164
 - deployment style, 194
 - documentation needs, 323
- Pericrisis in architectural decisions, 248
- Perry, Dewayne, 4, 24, 33, 48
- Perspectives, architecture, 410–411
- Pet Shop application, 118, 119
- PetStore application, 85–87, 86
- Physical data model, 110, 111
- Pipe-and-filter style
 - elements, relations, and properties, 158–159
 - example, 160–161, 161
 - overview, 158
 - purpose, 158–160
 - relation to other styles and models, 160
 - UML, 442–443, 442
- Plain old Java objects (POJOs), 98
- Platform style, 208
- Plus One View, 360
- Policies, global, 349–350
- Politi, Michael, 311
- Portability in layered style, 90, 98–99
- Ports
 - AADL, 478–480, 479
 - C&C views, 127–128, 134
 - UML, 140, 141, 460
- Prabhakar, T. V., 515
- Prescriptive architecture documentation, 12
- Presentations
 - C&C views, 347
 - combining with context diagrams, 347–348, 348
 - guidelines, 369–370
 - multi-part, 349
 - overview, 364–365
 - view documentation, 338–339
 - view packets, 344
- Prieto-Diaz, R., 64
- Primary presentations
 - C&C views, 347
 - combining with context diagrams, 347–348, 348
 - multi-part, 349
 - view documentation, 338–339
 - view packets, 344
- Priorities in view selection, 329, 332
- Process-steps style, 208–209
- Process structure in architecture views, 24
- Process views in 4+1 approach, 407
- Processes in AADL, 474
- Processing elements, 4
- Processors in AADL, 474
- Product-line architectures, 234
- Product line implementation in decomposition style, 66
- Product line managers, architecture documentation for, 16
- Programming in the large, 30
- Project managers
 - architecture documentation for, 16
 - documentation needs, 316–317, 317
 - interface documentation for, 279
- Properties
 - AADL, 482–483, 483
 - in architectural decisions, 248
 - in architecture documentation, 17
 - aspects style, 104–105

- C&C views, 126, 133–134, *134*
 - client-server style, 162–164
 - data model style, 111–113
 - decomposition style, 66–67, *67*
 - deployment style, 192–194
 - generalization style, 83
 - install style, 199–200
 - layered style, 89
 - module views, 57–59, *58*
 - peer-to-peer style, 166–167
 - pipe-and-filter style, 158–159
 - publish-subscribe style, 174–176
 - SOA style, 169–170
 - style guides, 50–51
 - tools for, 371
 - uses style, 75
 - view documentation, 340
 - work assignment style, 202–203
 - Provided interfaces, 264–265
 - Provides-services roles in C&C views, 128
 - Pseudo-states in state machine diagrams, 457
 - Publish-subscribe connectors in C&C views, 124
 - Publish-subscribe style
 - elements, relations, and properties, 174–176
 - example, 177–178, *177*
 - overview, 174
 - purpose, 175–176
 - relation to other styles, 176–177
 - Published styles and views, 343
 - Publisher roles in C&C views, 128
- Q**
- Quality attributes
 - and architecture documentation, 17–18
 - decomposition style, 66
 - interface documentation, 284
 - interfaces, 276–277
 - and software architecture, 2–3
 - Questions and question sets
 - architecture document reviews, 379, 382–394, *384*
 - behavior documentation, 290–291
 - open, 357
 - for supporting development, 389–391
 - for supporting evaluation, 386–388
- R**
- Ran, Alexander, 2
 - Rapid7 approach, 335–336
 - Rational Unified Process (RUP)
 - architecture views, 24
 - Views and Beyond, 406–408, *408*
 - Rationale
 - in architectural decisions, 239
 - architecture documentation, 17, 44
 - behavior documentation, 306
 - documentation beyond views, 354–355
 - interface issues, 277
 - tools for, 371
 - view documentation, 340–341
 - Rationale and Design Issues section in interface documentation, 284
 - Reader point of view in architectural documentation, 36–37
 - Reader roles in C&C views, 128
 - Realization concept in AADL module style, 475
 - Reallocating resources in variation points, 236
 - Redundant discussion, architectural decisions for, 246
 - Reentrant calls in UML, 452, *453*
 - Refactoring Agile projects, 414
 - Referenced material in directories, 356
 - References in sequence diagrams, 298
 - Refinement, 218
 - decomposition, 218–219, *218–219*
 - implementation, 219, *220*
 - spectrum of design, 220–221
 - style specialization, 221–222
 - Reid, Thomas, 54
 - Rejected status in architectural decisions, 250
 - Related decisions in architectural decision templates, 241
 - Related requirements in architectural decision templates, 241
 - Relations, 5
 - allocation styles, 191
 - aspects style, 104–105
 - C&C views, 126, 131–132, *133*
 - client-server style, 162–163
 - data model style, 111–113, *113–114*
 - decomposition style, 66–69, *67*
 - deployment style, 192–193
 - directional, 41–43
 - generalization style, 83
 - install style, 199–200
 - layered style, 89
 - module views, 57

Relations (*continued*)

- peer-to-peer style, 166–167
 - pipe-and-filter style, 158–159
 - publish-subscribe style, 174–175
 - selecting, 52–53
 - shared-data style, 179
 - SOA style, 169–172
 - style guides, 50–51
 - uses style, 75
 - view documentation, 340
 - work assignment style, 202–203
- Release-based style, 209
- Release strategy in packages, 368–369
- Relevant views, 22
- Reliability
- C&C views, 133
 - deployment style, 195
- Remote calls vs. local calls, 42
- Remote communication in behavior documentation, 293
- Remote procedure call (RPC), 292
- Repetition
- architectural documentation, 37–38
 - across views and view packets, 349–350
- Repository component in shared-data style, 179
- Repository styles, 156, 157, 178–182, 181
- Representation, defined, 11
- Representatives of external systems, architecture documentation for, 16
- Request/reply connectors in client-server style, 163
- Required interfaces, 264–265
- Requirements
- mapping to, 17, 357–362, 358, 361
 - tooling, 370–372
 - viewpoints, 360–361
- Resources, 262
- C&C views, 133
 - deployment style, 194
 - in interface documentation, 281–283
 - interfaces, 271–274
- Respondents in question sets, 383
- Responsibilities in modules, 56
- Responsibility property in module views, 57–58
- REST connectors in SOA style, 170–171
- Retrying in error handling, 278
- Reuse
- generalization style for, 84
 - interface documentation for, 279

- Rhapsody tool, 472
- Rich C&C connectors, 143, 143
- Rings notation in layered style, 93–94, 94
- Risk in architectural decisions, 246, 249
- Roles in C&C views, 128, 134
- Rozanski, Nick, 13, 25 37, 341–343, 350, 408–413, 409, 413, 430, 515
- RSS feeds, 160, 161
- Rugina, Ana-Elena, 483
- Rumbaugh, James, 4, 122, 311

S

- SaaS (software as a service) model, 287
- Safety critical properties in deployment style, 194
- Safety views, 404
- Saini, Darpan, 516
- SARA report, 397
- Satellite system, 79, 80
- Satisfy relationships in SysML, 466
- Scaffidi, Christopher, 287
- Schemes, packaging, 362–364
- Schmerl, Bradley, 153
- Schmidt, Douglas C., 187
- Schwaber, Ken, 414, 430
- Science Data Processing Segment (SDPS) in ECS module, 204
- SCM (software configuration management), 214
- Scope
- architectural decisions, 248
 - documentation roadmaps, 352
- Scrum approach, 414
- Security
- AADL, 482–483
 - C&C views, 133–134
 - client-server style, 164
 - deployment style, 195
 - documentation needs, 323
- Security views, 404
- Segmented layers in layered style, 92–93, 93
- Self calls in UML, 452, 453
- Semantics for resource, 272–274
- Semiformal notations
- architecture views, 53
 - C&C views, 139–146, 140–145
 - uses style, 76, 77–78
- Send-data-to relations in decomposition refinement, 218
- Separate documents in scheme packaging, 363–364

- Sequence diagrams
 - Agile projects, 417, 418
 - behavior documentation, 297–298, 298–299
 - SysML, 471
 - UML, 450–453, 451–453
- Servers in client-server style, 162
- Service-level agreements (SLAs), 277
- Service-oriented architecture (SOA) style
 - elements, 169, 171
 - example, 172, 173
 - module view, 266
 - overview, 169
 - properties, 169
 - purpose, 172
 - relations, 169, 171–172
- Service-oriented views, 404
- Service providers in SOA style, 171
- Service registry in SOA style, 170
- Services in call-return styles, 161
- Servlets, 98
- Shared-data stores, 179
- Shared-data style, 178
 - elements, relations, and properties, 179
 - example, 181–182
 - overview, 178–179
 - purpose, 179–180
 - relation to other styles, 180–181
- Shares-part-of-the-same-secret-as relations in
 - architecture views, 24
- Sharing and Reusing Architectural Knowledge (SHARK) workshops, 260
- Shaw, Mary, 4, 33, 48, 153, 187, 286–287
- Short iterations approach, 414
- Sidecars, layers with, 94–95, 94
- Siemens Four View model, 24, 214
- Signatures, 280–281
- Simplicity in presentations, 370
- Size in layered style, 95
- Slices, 73
- Smalltalk classes, 30
- Snyder, Alan, 122
- SOA. *See* Service-oriented architecture (SOA) style
- SOAP connectors
 - C&C views, 129
 - SOA style, 170–172
- SOAP Web service
 - interface documentation, 285
 - synchronous remote calls, 42
- Social networks, publish-subscribe style for, 174
- Socializing decisions, architectural decisions
 - for, 245
- Society of Automotive Engineers (SAE), 312
- Software architecture overview, 1–5
 - vs. design, 6–9
 - documentation overview. *See* Architecture documentation overview
 - and quality attributes, 2–3
- Software as a service (SaaS) model, 287
- Software components in AADL, 474
- Software configuration management (SCM), 214
- Software Decision Hiding Module in A-7E avionics system, 71–73, 72
- Software elements
 - allocation styles, 191
 - deployment style, 192
 - install style, 199
 - work assignment style, 203
- Software Engineering Institute, 47
- Software interfaces. *See* Interfaces
- Software product lines, 320
- Soni, Dilip, 24, 64, 214, 336
- Specializations
 - allocations, 208
 - C&C views, 130–131
 - styles, 221–222
- Specification, defined, 10
- Specification and Description Language (SDL), 302, 306
- Spectrum of design, 220–221
- Spivey, J. M., 311
- Spring framework, 222
- Stack notation in layered style, 92, 92
- Stafford, Judith A., 11, 511
- Staging in view selection, 329
- Stakeholder/view tables, 326–327
- Stakeholders, 15–16
 - architecture documentation for, 13
 - documentation roadmaps, 352
 - interface documentation, 278–279
 - ISO/IEC 42010, 405
 - quality attributes for, 18
 - question sets, 383, 385–386
 - view selection for, 316–328, 317–322, 324–326
- Standard organization
 - architectural documentation, 43–44
 - for documentation beyond views, 351–356, 351, 354–355
 - variations, 356–357

- Standard organization (*continued*)
 - for interface documentation, 271–277, 271, 276
 - for style guides, 50–51
 - for views, 337–341
 - variations, 344–349, 348
 - State machine diagrams, 303–304, 304, 457, 458
 - Status in architectural decision templates, 240
 - Status indicators in interfaces, 277
 - Stereotypes in UML, 62, 62, 433, 441
 - Steward, Donald, 101
 - Stimulation communication, 292
 - Storage unit capacity in deployment style, 194
 - Stream of consciousness writing, 37
 - Style guides, 26
 - outline, 50–51
 - presentations, 369
 - Style specialization, 221–222
 - Styles, 25
 - allocation. *See* Allocation styles
 - categories, 29, 49–50
 - C&C. *See* Component-and-connector (C&C) styles
 - element and relation properties, 52–53
 - examples, 54
 - module. *See* Module styles
 - notations for architecture views, 53
 - overview, 25–28, 27–28
 - terms, 29–36
 - Subjects
 - architecture document reviews, 378–379, 394
 - use case diagrams, 459
 - Subprograms in AADL, 474
 - Subscriber role in C&C views, 128
 - Substates in state machine diagrams, 304
 - Subsumes relations in architectural decisions, 250
 - Subsystems, 73–74, 74
 - decomposition style, 66
 - Support phase in architecture document reviews, 377
 - Supporting documentation, 341
 - C&C views, 125
 - context diagrams, 229
 - view packets, 344
 - Swim lanes in UML sequence diagrams 300–301, 450
 - Symbols in use case diagrams, 459, 459
 - Synchronous calls, 42
 - Synchronous communication, 292
 - Syntactic information, interface documentation for, 279–280
 - Syntax, resource, 272
 - SysML. *See* Systems Modeling Language (SysML)
 - System analysis and construction, architecture documentation for, 14
 - System context views, 341–343
 - System engineers, architecture documentation for, 16
 - System-level behavior documentation in DoDAF, 426
 - System overview in documentation beyond views, 353
 - Systems and services view (SV) in DoDAF, 419, 423–425
 - Systems in AADL, 474
 - Systems Modeling Language (SysML), 197, 465–466, 466
 - allocation view, 470–471, 470
 - architecture documentation, 466
 - behavior documentation, 471, 471
 - C&C view, 469, 469
 - interfaces, 472, 472
 - module view, 468, 468
 - requirements, 466, 467
 - summary, 472
 - Systems of interest in ISO/IEC 42010, 401
 - Szyperski, C., 30
- T**
- Tables
 - DoDAF, 419
 - SysML, 197
 - uses style, 76, 80
 - work assignment style, 203, 204
 - Tagged values in UML, 441, 441
 - Taylor, R. N., 5, 153
 - TDDT (Training and Doctrine Development Tool), 69
 - Team allocation in decomposition style, 66
 - Team morale, uses relation for, 81
 - Technical standards view (TV) in DoDAF, 419, 425
 - Telephone communication, 292
 - Templates. *See also* Standard organization.
 - architectural decisions, 239–242
 - architecture perspectives, 411
 - question sets, 384, 384
 - in variability, 233

- Tentative status in architectural decisions, 249
 - Terms
 - in directories, 355
 - glossary, 491–496
 - Test-driven development in Agile projects, 414
 - Test information in module views, 59
 - Testers
 - architecture documentation for, 16
 - documentation needs, 318–319, 318
 - interface documentation for, 278
 - Textual representation of views
 - decomposition view for A-7E example, 72
 - module views, 61
 - visibility in module views, 68
 - work assignment view for ECS, 79–80
 - The Open Group Architecture Framework (TOGAF), 377
 - Threads in AADL, 474, 478–480
 - Three-part rule in architecture patterns, 34
 - Tiered client-server views, 257, 257–258
 - Tiers
 - C&C styles, 183–184, 184
 - C&C views, 134
 - layered style, 96–97
 - Time-based stimulation in behavior documentation, 293
 - Time constraints in sequence diagrams, 298
 - Time stamps in architectural decisions, 248–249
 - Timed threads, 478
 - Timely development, architectural decisions for, 246
 - Timing diagrams in UML, 302, 454, 455
 - Tools
 - for combined views, 255
 - requirements, 370–372
 - UML, 461–463
 - Top-level context diagrams (TLCDs), 225–226
 - Topcased project, 472
 - Traceability of requirements viewpoints, 360
 - Traces
 - behavior documentation, 294–295
 - notations for, 295–303, 295–301
 - for semantic information, 280
 - Trachtenberg, Marvin, 399
 - Training and Doctrine Development Tool (TDDT), 69
 - Transclusion process, 368
 - Transitions in state machine diagrams, 303–304, 304, 457
 - Tuple spaces, 187
 - Types
 - of components and connectors, 129–131
 - UML, 448, 449
 - Tyree, Jeff, 239, 516, 246, 247, 260
- ## U
- UML. *See* Unified Modeling Language
 - Unified Modeling Language (UML), 431
 - activity diagrams, 300–302, 301, 450, 450
 - allocation view, 443–447, 443–447
 - ambiguity traps, 447–449, 448–449
 - aspects style, 105–106, 106
 - assembly connectors, 142, 146, 440, 440
 - associations, 57, 117, 117, 197, 251–255, 251–253, 436–437, 440, 448, 448, 459, 459
 - behavior documentation, 296–305, 296–301, 449–459, 449–459
 - C&C views, 132, 133, 139–146, 140–145, 438–443, 439–442
 - class diagrams, 436–438, 437, 447
 - communication diagrams, 299–300, 300, 453–454, 454
 - components, 74, 133, 135, 137, 139–146, 140–145, 151, 197, 198, 219, 219, 268, 269, 286, 431, 432, 438–445, 439–444, 448, 449, 451, 451, 460, 461
 - context diagrams, 230–231, 231
 - data model style, 117, 117
 - decomposition refinement, 218–219, 219
 - decomposition style, 68, 68
 - delegation connectors, 132, 133, 219, 441, 442
 - deployment style, 197, 198
 - generalization style, 84, 85
 - interaction overview diagrams, 455–456, 456
 - interface documentation, 268, 269, 460–461, 460–461
 - introduction, 431–433, 432
 - layered style, 95, 95
 - module views, 61–62, 61–62, 433–436, 433–436
 - packages, 61, 61, 68, 68, 74–77, 77–78, 95, 231, 268, 269, 355, 431, 432, 433–435, 433, 435, 437, 446, 449
 - port, 127, 140–141, 141, 146, 149, 151, 169, 220, 286, 438, 439, 441, 460
 - provided interface, 61, 141–142, 264–265, 268, 269, 285, 437, 438, 439, 440–441, 440, 460, 460

- Unified Modeling Language (UML) (*continued*)
 - required interface, 141–142, 264–265, 268, 269, 285, 437, 438, 439, 440–441, 440, 460, 460
 - sequence diagrams, 297–298, 298–299, 450–453, 451–453
 - state machine diagrams, 303–304, 304, 457, 458
 - SysML, 465–466, 466
 - timing diagrams, 302, 454, 455
 - tools, 461–463
 - use case diagrams, 458, 459
 - use cases, 296, 296–297
 - uses style, 76, 77–78
- Uncertainty relations between modules (UM), 210–212
- Unit Training Management Configuration (UTMC), 69
- Universal styles, 26
- UNIX System V operating system, 97–98, 97
- Usability in documentation needs, 323
- Usage Guide section in interface documentation, 277, 285
- Use case diagrams, 458, 459
- Use cases
 - 4+1 approach, 407
 - behavior documentation, 296, 296–297
 - documenting, 356
- User stories in Agile projects, 414
- Users, architecture documentation for, 16
- Uses of architecture documentation, 12–16
- Uses relations, 81–82
 - layered style, 97
 - module views, 57
 - UML, 448
- Uses structure in architecture views, 24
- Uses style, 74–75
 - elements, relations, and properties, 75
 - examples, 79–80, 80
 - notations, 76–77, 77–78
 - purpose, 75–76, 76
 - relation to other styles, 79
 - UML module views, 433
 - uses relation, 81–82
- Uses views in combined views, 255–256, 256
- Utilization phase in architecture document reviews, 377
- UTMC (Unit Training Management Configuration), 69
- V**
 - van der Linden, Frank, 2
 - van Gogh, Vincent, 215
 - van Vliet, H., 247
 - Variability, 231
 - interfaces, 276, 284
 - Variability guides, 232, 340
 - Variation points, 231, 232
 - documenting, 235–238, 238
 - dynamism and dynamic architectures, 234–235
 - variation mechanisms, 232–233
 - Versions, architecture documentation for, 21
 - Vertical slices, 73
 - Vicknair, W., 260
 - View packets, 344–347
 - View selection, 315–333
 - examples, 329–335
 - prioritizing and staging, 329, 332
 - stakeholder documentation needs, 316–328, 317–322, 324–326
 - stakeholder/view tables, 326–327
 - View-to-view associations in documentation
 - beyond views, 353–354, 354
 - Viewpoints
 - ISO/IEC 42010, 402–403, 403, 405
 - requirements, 360–361, 361
 - Rozanski and Woods viewpoint sets, 408–413, 409, 413
 - Views, 22–23
 - associations between, 251
 - combining. *See* Combining views
 - component-and-connector. *See* Component-and-connector (C&C) views
 - coordination, 209–213
 - documentation beyond views, 350–356, 351, 354–356
 - in documentation packages, 337, 363–365
 - repetition across, 349–350
 - standard organization, 337–344, 338
 - standard organization variations, 344–349, 348
 - documentation roadmaps, 352
 - history of, 23–25
 - ISO/IEC 42010, 401
 - mapping between, 353–354, 354
 - module. *See* Module views
 - notations, 53
 - and published styles and patterns, 343

- Views and Beyond approach, 19–20, 339
 - in Agile development environment, 20, 415–418, 418
 - architecture decisions, 239
 - architecture documentation, 19–20
 - context diagrams, 341
 - compared to DoDAF, 421–425
 - compared to ISO/IEC 42010, 400–405
 - compared to Rozanski and Woods, 411–413
 - compared to RUP, 406–408
 - Virtual buses in AADL, 474
 - Virtual machines, 99
 - layered style, 90
 - Virtual processors in AADL, 474
 - Visibility of interface property in module views, 58, 58
 - Visual styles in presentations, 369
 - Vitruvius, 399
 - Vocabulary of view for context diagrams, 226–228, 226–228
 - Voelter, Markus, 356
 - von Goethe, Johann Wolfgang, 289
- W**
- Web-based documentation, 365–366. See also Wikis.
 - Web modules, ATIA server-side, 69
 - Web Services Description Language (WSDL), 279
 - Weeks, Edward, 375
 - Weiss, David M., 121, 122, 380, 396, 516
 - Weiss, M., 247
 - Wikis, 365–368
 - Windowsapps, 69
 - Wolf, Alexander L., 4, 11, 24, 33, 48
 - Wolff, Bobby, 187
 - Woods, Eoin, 3, 13, 25, 37, 341–343, 350, 408–413, 409, 413, 430, 516
 - Work assignment style, 205–206
 - elements, relations, and properties, 202–203
 - notations, 203–204, 204
 - overview, 202
 - purpose, 203
 - relation to other styles, 204–205
 - specializing, 208
 - UML, 446–447, 447
 - Work assignment views in combined views, 255
 - Work breakdown structure (WBS), 202
 - World Wide Web in client-server style, 164–165
 - Wright, Frank Lloyd, 14, 55, 316
 - Writer roles in C&C views, 128
 - Wuerges, H., 286
- X**
- XMI model in AADL, 473
 - XZip component, 281
- Y**
- Yahoo! Pipes, 160–161, 161
 - Yang, Chen Ning, 315
- Z**
- Z language, 304–305, 311
 - Zachman, J. A., 404
 - Zip component API, 281–285