



SCALABILITY RULES

PRINCIPLES FOR SCALING WEB SITES
S E C O N D E D I T I O N

MARTIN L. ABBOTT

MICHAEL T. FISHER

"Whether you're taking on a role as a technology leader in a new company or you simply want to make great technology decisions, Scalability Rules will be the go-to resource on your bookshelf."

—Chad Dickerson, CTO, Etsy

FREE SAMPLE CHAPTER

SHARE WITH OTHERS



Praise for the First Edition of *Scalability Rules*

“Once again, Abbott and Fisher provide a book that I’ll be giving to our engineers. It’s an essential read for anyone dealing with scaling an online business.”

—**Chris Lalonde**, GM of Data Stores, Rackspace

“Abbott and Fisher again tackle the difficult problem of scalability in their unique and practical manner. Distilling the challenges of operating a fast-growing presence on the Internet into 50 easy-to-understand rules, the authors provide a modern cookbook of scalability recipes that guide the reader through the difficulties of fast growth.”

—**Geoffrey Weber**, VP, Internet Operations, Shutterfly

“Abbott and Fisher have distilled years of wisdom into a set of cogent principles to avoid many nonobvious mistakes.”

—**Jonathan Heiliger**, VP, Technical Operations, Facebook

“In *The Art of Scalability*, the AKF team taught us that scale is not just a technology challenge. Scale is obtained only through a combination of people, process, *and* technology. With *Scalability Rules*, Martin Abbott and Michael Fisher fill our scalability toolbox with easily implemented and time-tested rules that once applied will enable massive scale.”

—**Jerome Labat**, VP, Product Development IT, Intuit

“When I joined Etsy, I partnered with Mike and Marty to hit the ground running in my new role, and it was one of the best investments of time I have made in my career. The indispensable advice from my experience working with Mike and Marty is fully captured here in this book. Whether you’re taking on a role as a technology leader in a new company or you simply want to make great technology decisions, *Scalability Rules* will be the go-to resource on your bookshelf.”

—**Chad Dickerson**, CTO, Etsy

“*Scalability Rules* provides an essential set of practical tools and concepts anyone can use when designing, upgrading, or inheriting a technology platform. It’s very easy to focus on an immediate problem and overlook issues that will appear in the future. This book ensures strategic design principles are applied to everyday challenges.”

—**Robert Guild**, Director and Senior Architect, Financial Services

“An insightful, practical guide to designing and building scalable systems. A must-read for both product building and operations teams, this book offers concise and crisp insights gained from years of practical experience of AKF principals. With the complexity of modern systems, scalability considerations should be an integral part of the architecture and implementation process. Scaling systems for hypergrowth requires an agile, iterative approach that is closely aligned with product features; this book shows you how.”

—**Nanda Kishore**, CTO, ShareThis

“For organizations looking to scale technology, people, and processes rapidly or effectively, the twin pairing of *Scalability Rules* and *The Art of Scalability* is unbeatable. The rules-driven approach in *Scalability Rules* not only makes this an easy reference companion, but also allows organizations to tailor the Abbott and Fisher approach to their specific needs both immediately and in the future!”

—**Jeremy Wright**, CEO, BNOTIONS.ca, and Founder, b5media

Scalability Rules

Second Edition

This page intentionally left blank

Scalability Rules

Principles for Scaling Web Sites

Second Edition

Martin L. Abbott
Michael T. Fisher

◆ Addison-Wesley

Boston • Columbus • Indianapolis • New York • San Francisco • Amsterdam • Cape Town
Dubai • London • Madrid • Milan • Munich • Paris • Montreal • Toronto • Delhi • Mexico City
São Paulo • Sydney • Hong Kong • Seoul • Singapore • Taipei • Tokyo

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

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

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

For government sales inquiries, please contact governmentsales@pearsoned.com.

For questions about sales outside the U.S., please contact intlcs@pearson.com.

Visit us on the Web: informit.com/aw

Library of Congress Control Number: 2016944687

Copyright © 2017 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, request forms and the appropriate contacts within the Pearson Education Global Rights & Permissions Department, please visit www.pearsoned.com/permissions/.

ISBN-13: 978-0-13-443160-4

ISBN-10: 0-13-443160-X

Text printed in the United States on recycled paper at RR Donnelley in Crawfordsville, Indiana.

1 16

Editor-in-Chief

Mark L. Taub

Executive Editor

Laura Lewin

Development Editor

Songlin Qiu

Managing Editor

Sandra Schroeder

Full-Service**Production Manager**

Julie B. Nahil

Project Editor

Dana Wilson

Copy Editor

Barbara Wood

Indexer

Jack Lewis

Proofreader

Barbara Lasoff

Technical Reviewers

Camille Fournier

Chris Lalonde

Mark Uhrmacher

Editorial Assistant

Olivia Basegio

Cover Designer

Chuti Prasertsith

Compositor

The CIP Group



*This book is dedicated to our friend and partner “Big” Tom Keeven.
“Big” refers to the impact he’s had in helping countless companies scale
in his nearly 30 years in the business.*



This page intentionally left blank

Contents

Preface xv

Acknowledgments xxi

About the Authors xxiii

1 Reduce the Equation 1

Rule 1—Don't Overengineer the Solution **3**

Rule 2—Design Scale into the Solution (D-I-D Process) **6**

Design **6**

Implement **7**

Deploy **8**

Rule 3—Simplify the Solution Three Times Over **8**

How Do I Simplify My Scope? **9**

How Do I Simplify My Design? **9**

How Do I Simplify My Implementation? **10**

Rule 4—Reduce DNS Lookups **10**

Rule 5—Reduce Objects Where Possible **12**

Rule 6—Use Homogeneous Networks **15**

Summary **15**

Notes **16**

2 Distribute Your Work 19

Rule 7—Design to Clone or Replicate Things (X Axis) **22**

Rule 8—Design to Split Different Things (Y Axis) **24**

Rule 9—Design to Split Similar Things (Z Axis) **26**

Summary **28**

Notes **28**

3 Design to Scale Out Horizontally 29

Rule 10—Design Your Solution to Scale Out, Not Just Up **31**

Rule 11—Use Commodity Systems (Goldfish Not Thoroughbreds) **33**

Rule 12—Scale Out Your Hosting Solution **35**

Rule 13—Design to Leverage the Cloud **40**

Summary **42**

Notes **42**

4 Use the Right Tools	43
Rule 14—Use Databases Appropriately	47
Rule 15—Firewalls, Firewalls Everywhere!	52
Rule 16—Actively Use Log Files	55
Summary	58
Notes	58
5 Get Out of Your Own Way	59
Rule 17—Don't Check Your Work	61
Rule 18—Stop Redirecting Traffic	64
Rule 19—Relax Temporal Constraints	68
Summary	70
Notes	70
6 Use Caching Aggressively	73
Rule 20—Leverage Content Delivery Networks	75
Rule 21—Use Expires Headers	77
Rule 22—Cache Ajax Calls	80
Rule 23—Leverage Page Caches	84
Rule 24—Utilize Application Caches	86
Rule 25—Make Use of Object Caches	88
Rule 26—Put Object Caches on Their Own “Tier”	90
Summary	91
Notes	92
7 Learn from Your Mistakes	93
Rule 27—Learn Aggressively	95
Rule 28—Don't Rely on QA to Find Mistakes	100
Rule 29—Failing to Design for Rollback Is Designing for Failure	102
Summary	105
Notes	106
8 Database Rules	107
Rule 30—Remove Business Intelligence from Transaction Processing	109
Rule 31—Be Aware of Costly Relationships	111
Rule 32—Use the Right Type of Database Lock	114
Rule 33—Pass on Using Multiphase Commits	116
Rule 34—Try Not to Use <code>Select for Update</code>	118

Rule 35—Don't Select Everything	120
Summary	121
Notes	122
9 Design for Fault Tolerance and Graceful Failure	123
Rule 36—Design Using Fault-Isolative “Swim Lanes”	124
Rule 37—Never Trust Single Points of Failure	130
Rule 38—Avoid Putting Systems in Series	132
Rule 39—Ensure That You Can Wire On and Off Features	135
Summary	138
10 Avoid or Distribute State	139
Rule 40—Strive for Statelessness	140
Rule 41—Maintain Sessions in the Browser When Possible	142
Rule 42—Make Use of a Distributed Cache for States	144
Summary	146
Notes	146
11 Asynchronous Communication and Message Buses	147
Rule 43—Communicate Asynchronously as Much as Possible	149
Rule 44—Ensure That Your Message Bus Can Scale	151
Rule 45—Avoid Overcrowding Your Message Bus	154
Summary	157
12 Miscellaneous Rules	159
Rule 46—Be Wary of Scaling through Third Parties	161
Rule 47—Purge, Archive, and Cost-Justify Storage	163
Rule 48—Partition Inductive, Deductive, Batch, and User Interaction (OLTP) Workloads	166
Rule 49—Design Your Application to Be Monitored	169
Rule 50—Be Competent	172

Summary **174**

Notes **174**

13 Rule Review and Prioritization 177

A Risk-Benefit Model for Evaluating Scalability
Projects and Initiatives **177**

50 Scalability Rules in Brief **180**

Rule 1—Don't Overengineer the Solution **180**

Rule 2—Design Scale into the Solution
(D-I-D Process) **181**

Rule 3—Simplify the Solution Three
Times Over **181**

Rule 4—Reduce DNS Lookups **182**

Rule 5—Reduce Objects Where Possible **182**

Rule 6—Use Homogeneous Networks **182**

Rule 7—Design to Clone or Replicate Things
(X Axis) **183**

Rule 8—Design to Split Different Things
(Y Axis) **183**

Rule 9—Design to Split Similar Things (Z Axis) **184**

Rule 10—Design Your Solution to Scale Out,
Not Just Up **184**

Rule 11—Use Commodity Systems (Goldfish Not
Thoroughbreds) **185**

Rule 12—Scale Out Your Hosting Solution **185**

Rule 13—Design to Leverage the Cloud **185**

Rule 14—Use Databases Appropriately **186**

Rule 15—Firewalls, Firewalls Everywhere! **186**

Rule 16—Actively Use Log Files **187**

Rule 17—Don't Check Your Work **187**

Rule 18—Stop Redirecting Traffic **188**

Rule 19—Relax Temporal Constraints **188**

Rule 20—Leverage Content Delivery Networks **188**

Rule 21—Use Expires Headers **189**

Rule 22—Cache Ajax Calls **189**

Rule 23—Leverage Page Caches **189**

Rule 24—Utilize Application Caches **190**

Rule 25—Make Use of Object Caches **190**

Rule 26—Put Object Caches on Their Own “Tier”	190
Rule 27—Learn Aggressively	191
Rule 28—Don’t Rely on QA to Find Mistakes	191
Rule 29—Failing to Design for Rollback Is Designing for Failure	191
Rule 30—Remove Business Intelligence from Transaction Processing	192
Rule 31—Be Aware of Costly Relationships	192
Rule 32—Use the Right Type of Database Lock	193
Rule 33—Pass on Using Multiphase Commits	193
Rule 34—Try Not to Use <code>Select for Update</code>	194
Rule 35—Don’t Select Everything	194
Rule 36—Design Using Fault-Isolative “Swim Lanes”	194
Rule 37—Never Trust Single Points of Failure	195
Rule 38—Avoid Putting Systems in Series	195
Rule 39—Ensure That You Can Wire On and Off Features	195
Rule 40—Strive for Statelessness	196
Rule 41—Maintain Sessions in the Browser When Possible	196
Rule 42—Make Use of a Distributed Cache for States	196
Rule 43—Communicate Asynchronously as Much as Possible	197
Rule 44—Ensure That Your Message Bus Can Scale	197
Rule 45—Avoid Overcrowding Your Message Bus	198
Rule 46—Be Wary of Scaling through Third Parties	198
Rule 47—Purge, Archive, and Cost-Justify Storage	198
Rule 48—Partition Inductive, Deductive, Batch, and User Interaction (OLTP) Workloads	199
Rule 49—Design Your Application to Be Monitored	199
Rule 50—Be Competent	200

A Benefit/Priority Ranking of the Scalability Rules	200
Very High—1	200
High—2	201
Medium—3	201
Low—4	202
Very Low—5	202
Summary	202
Index	205

Preface

Thanks for your interest in the second edition of *Scalability Rules*! This book is meant to serve as a primer, a refresher, and a lightweight reference manual to help engineers, architects, and managers develop and maintain scalable Internet products. It is laid out in a series of rules, each of them bundled thematically by different topics. Most of the rules are technically focused, and a smaller number of them address some critical mind-set or process concern, each of which is absolutely critical to building scalable products. The rules vary in their depth and focus. Some rules are high level, such as defining a model that can be applied to nearly any scalability problem; others are specific and may explain a technique, such as how to modify headers to maximize the “cacheability” of content. In this edition we’ve added stories from CTOs and entrepreneurs of successful Internet product companies from startups to Fortune 500 companies. These stories help to illustrate how the rules were developed and why they are so important within high-transaction environments. No story serves to better illustrate the challenges and demands of hyper-scale on the Internet than Amazon. Rick Dalzell, Amazon’s first CTO, illustrates several of the rules within this book in his story, which follows.

Taming the Wild West of the Internet

From the perspective of innovation and industry disruption, few companies have had the success of Amazon. Since its founding in 1994, Amazon has contributed to redefining at least three industries: consumer commerce, print publishing, and server hosting. And Amazon’s contributions go well beyond just disrupting industries; they’ve consistently been a thought leader in service-oriented architectures, development team construction, and a myriad of other engineering approaches. Amazon’s size and scale along all dimensions of its business are simply mind-boggling; the company has consistently grown at a rate unimaginable for traditional brick-and-mortar businesses. Since 1998, Amazon grew from \$600 million (no small business at all) in annual revenue to an astounding \$107 billion (that’s “billion” with a *B*) in 2015.¹ Walmart, the world’s largest retailer, had annual sales of \$485.7 billion in 2015.² But Walmart has been around since 1962, and it took 35 years to top \$100 billion in sales compared to Amazon’s 21 years. No book professing to codify the rules of scalability from the mouths of CTOs who have created them would be complete without one or more stories from Amazon.

Jeff Bezos incorporated Amazon (originally Cadabra) in July of 1994 and launched Amazon.com as an online bookseller in 1995. In 1997, Bezos hired Rick Dalzell, who was then the VP of information technology at Walmart. Rick spent the next ten years

at Amazon leading Amazon's development efforts. Let's join Rick as he relays the story of his Amazon career:

“When I was at Walmart, we had one of the world's largest relational databases running the company's operations. But it became clear to the Amazon team pretty quickly that the big, monolithic database approach was simply not going to work for Amazon. Even back then, we handled more transactions in a week on the Amazon system than the Walmart system had to handle in a month. And when you add to that our incredible growth, well, it was pretty clear that monoliths simply were not going to work. Jeff [Bezos] took me to lunch one day, and I told him we needed to split the monolith into services. He said, ‘That's great—but we need to build a moat around this business and get to 14 million customers.’ I explained that without starting to work on these splits, we wouldn't be able to make it through Christmas.”

Rick continued, “Now keep in mind that this is the mid- to late nineties. There weren't a lot of companies working on distributed transaction systems. There weren't a lot of places to go to find help in figuring out how to scale transaction processing systems growing in excess of 300% year on year. There weren't any rulebooks, and there weren't any experts who had ‘been there and done that.’ It was a new frontier—a new Wild, Wild West. But it was clear to us that we had to distribute this thing to be successful. Contrary to what made me successful at Walmart, if we were going to scale our solution and our organization, we were going to need to split the solution and the underlying database up into a number of services.” (The reader should note that an entire chapter of this book, Chapter 2, “Distribute Your Work,” is dedicated to such splits.)

“We started by splitting the commerce and store engine from the back-end fulfillment systems that Amazon uses. This was really the start of our journey into the services-oriented architecture that folks have heard about at Amazon. All sorts of things came out of this, including Amazon's work on team independence and the API contracts. Ultimately, the work created a new industry [infrastructure as a service] and a new business for Amazon in Amazon Web Services—but that's another story for another time. The work wasn't easy; some components of the once-monolithic database such as customer data—what we called ‘the Amazon customer database or ACB’—took several years to figure out how to segment. We started with services that were high in transaction volumes and could be quickly split in both software and data, like the front- and back-end systems that I described. Each split we made would further distribute the system and allow additional scale. Finally, we got back to solving the hairy problem of ACB and split it out around 2004.

“The team was incredibly smart, but we also had a bit of luck from time to time. It's not that we never failed, but when we would make a mistake we would quickly correct it and figure out how to fix the associated problems. The lucky piece is that none of our failures were as large and well publicized as those of some of the other companies struggling through the same learning curve. A number of key learnings in building these distributed services came out of these splits, learnings such as the need to limit session and state, stay away from distributed two-phase commit transactions, communicating asynchronously whenever possible, and so on. In fact, without a strong bias toward asynchronous communication through a publish-and-subscribe message bus, I don't

know if we could have ever split and scaled the way we did. We also learned to allow things to be eventually consistent where possible, in just about everything except payments. Real-time consistency is costly, and wherever people wouldn't really know the difference, we'd just let things get 'fuzzy' for a while and let them sync up later. And of course there were a number of 'human' or team learnings as well such as the need to keep teams small³ and to have specific contracts between teams that use the services of other teams."

Rick's story of how he led Amazon's development efforts in scaling for a decade is incredibly useful. From his insights we can garner a number of lessons that can be applied to many companies' scaling challenges. We've used Rick's story along with those of several other notable CTOs and entrepreneurs of successful Internet product companies ranging from startups to Fortune 500 companies to illustrate how important the rules in this book are to scaling high-transaction environments.

Quick Start Guide

Experienced engineers, architects, and managers can read through the header sections of all the rules that contain the what, when, how, and why. You can browse through each chapter to read these, or you can jump to Chapter 13, "Rule Review and Prioritization," which has a consolidated view of the headers. Once you've read these, go back to the chapters that are new to you or that you find more interesting.

For less experienced readers we understand that 50 rules can seem overwhelming. We do believe that you should eventually become familiar with all the rules, but we also understand that you need to prioritize your time. With that in mind, we have picked out five chapters for managers, five chapters for software developers, and five chapters for technical operations that we recommend you read before the others to get a jump start on your scalability knowledge.

Managers:

- Chapter 1, "Reduce the Equation"
- Chapter 2, "Distribute Your Work"
- Chapter 4, "Use the Right Tools"
- Chapter 7, "Learn from Your Mistakes"
- Chapter 12, "Miscellaneous Rules"

Software developers:

- Chapter 1, "Reduce the Equation"
- Chapter 2, "Distribute Your Work"
- Chapter 5, "Get Out of Your Own Way"
- Chapter 10, "Avoid or Distribute State"
- Chapter 11, "Asynchronous Communication and Message Buses"

Technical operations:

- Chapter 2, “Distribute Your Work”
- Chapter 3, “Design to Scale Out Horizontally”
- Chapter 6, “Use Caching Aggressively”
- Chapter 8, “Database Rules”
- Chapter 9, “Design for Fault Tolerance and Graceful Failure”

As you have time later, we recommend reading all the rules to familiarize yourself with the rules and concepts that we present no matter what your role. The book is short and can probably be read in a coast-to-coast flight in the United States.

After the first read, the book can be used as a reference. If you are looking to fix or re-architect an existing product, Chapter 13 offers an approach to applying the rules to your existing platform based on cost and the expected benefit (presented as a reduction of risk). If you already have your own prioritization mechanism, we do not recommend changing it for ours unless you like our approach better. If you don't have an existing method of prioritization, our method should help you think through which rules you should apply first.

If you are just starting to develop a new product, the rules can help inform and guide you as to best practices for scaling. In this case, the approach of prioritization represented in Chapter 13 can best be used as a guide to what's most important to consider in your design. You should look at the rules that are most likely to allow you to scale for your immediate and long-term needs and implement those.

For all organizations, the rules can serve to help you create a set of architectural principles to drive future development. Select the 5, 10, or 15 rules that will help your product scale best and use them as an augmentation of your existing design reviews. Engineers and architects can ask questions relevant to each of the scalability rules that you select and ensure that any new significant design meets your scalability standards. While these rules are as specific and fixed as possible, there is room for modification based on your system's particular criteria. If you or your team has extensive scalability experience, go ahead and tweak these rules as necessary to fit your particular scenario. If you and your team lack large-scale experience, use the rules exactly as is and see how far they allow you to scale.

Finally, this book is meant to serve as a reference and handbook. Chapter 13 is set up as a quick reference and summary of the rules. Whether you are experiencing problems or simply looking to develop a more scalable solution, Chapter 13 can be a quick reference guide to help pinpoint the rules that will help you out of your predicament fastest or help you define the best path forward in the event of new development. Besides using this as a desktop reference, also consider integrating this into your organization by one of many tactics such as taking one or two rules each week and discussing them at your technology all-hands meeting.

Why a Second Edition?

The first edition of *Scalability Rules* was the first book to address the topic of scalability in a rules-oriented fashion. Customers loved its brevity, ease of use, and convenience. But time and time again readers and clients of our firm, AKF Partners, asked us to tell the stories behind the rules. Because we pride ourselves in putting the needs of our clients first, we edited this book to include stories upon which the rules are based.

In addition to telling the stories of multiple CTOs and successful entrepreneurs, editing the book for a second edition allowed us to update the content to remain consistent with the best practices in our industry. The second edition also gave us the opportunity to subject our material to another round of technical peer reviews and production editing. All of this results in a second edition that's easier to read, easier to understand, and easier to apply.

How Does *Scalability Rules* Differ from *The Art of Scalability*?

The Art of Scalability, Second Edition (ISBN 0134032802, published by Addison-Wesley), our first book on the topic of scalability, focused on people, process, and technology, whereas *Scalability Rules* is predominantly a technically focused book. Don't get us wrong; we still believe that people and process are the most important components of building scalable solutions. After all, it's the organization, including both the individual contributors and the management, that succeeds or fails in producing scalable solutions. The technology isn't at fault for failing to scale—it's the people who are at fault for building it, selecting it, or integrating it. But we believe that *The Art of Scalability* adequately addresses the people and process concerns around scalability, and we wanted to go into greater depth on the technical aspects of scalability.

Scalability Rules expands on the third (technical) section of *The Art of Scalability*. The material in *Scalability Rules* is either new or discussed in a more technical fashion than in *The Art of Scalability*. As some reviewers on Amazon point out, *Scalability Rules* works well as both a standalone book and as a companion to *The Art of Scalability*.

Notes

1. "Net Sales Revenue of Amazon from 2004 to 2015,"
www.statista.com/statistics/266282/annual-net-revenue-of-amazoncom/.
2. Walmart, Corporate and Financial Facts,
http://corporate.walmart.com/_news_/news-archive/investors/2015/02/19/walmart-announces-q4-underlying-eps-of-161-and-additional-strategic-investments-in-people-e-commerce-walmart-us-comp-sales-increased-15-percent.
3. Authors' note: Famously known as Amazon's Two-Pizza Rule—no team can be larger than that which two pizzas can feed.

Register your copy of *Scalability Rules, Second Edition*, at informit.com for convenient access to downloads, updates, and corrections as they become available. To start the registration process, go to informit.com/register and log in or create an account. Enter the product ISBN (9780134431604) and click Submit. Once the process is complete, you will find any available bonus content under “Registered Products.”

Acknowledgments

The rules contained within this book weren't developed by our partnership alone. They are the result of nearly 70 years of work with clients, colleagues, and partners within nearly 400 companies, divisions, and organizations. Each of them contributed, in varying degrees, to some or all of the rules within this book. As such, we would like to acknowledge the contributions of our friends, partners, clients, coworkers, and bosses for whom or with which we've worked over the past several (combined) decades. The CTO stories from Rick Dalzell, Chris Lalonde, James Barrese, Lon Binder, Brad Peterson, Grant Klopper, Jeremy King, Tom Keeven, Tayloe Stansbury, Chris Schremser, and Chuck Geiger included herein are invaluable in helping to illustrate the need for these 50 rules. We thank each of you for your time, thoughtfulness, and consideration in telling us your stories.

We would also like to acknowledge and thank the editors who have provided guidance, feedback, and project management. The technical editors from both the first and second editions—Geoffrey Weber, Chris Lalonde, Camille Fournier, Jeremy Wright, Mark Urmacher, and Robert Guild—shared with us their combined decades of technology experience and provided invaluable insight. Our editors from Addison-Wesley, Songlin Qiu, Laura Lewin, Olivia Basegio, and Trina MacDonald, provided supportive stylistic and rhetorical guidance throughout every step of this project. Thank you all for helping with this project.

Last but certainly not least, we'd like to thank our families and friends who put up with our absence from social events as we sat in front of a computer screen and wrote. No undertaking of this magnitude is done single-handedly, and without our families' and friends' understanding and support this would have been a much more arduous journey.

This page intentionally left blank

About the Authors

Martin L. Abbott is a founding partner of AKF Partners, a growth consulting firm focusing on meeting the needs of today's fast-paced and hyper-growth companies. Marty was formerly the COO of Quigo, an advertising technology startup acquired by AOL in 2007. Prior to Quigo, Marty spent nearly six years at eBay, most recently as SVP of technology and CTO and member of the CEO's executive staff. Prior to eBay, Marty held domestic and international engineering, management, and executive positions at Gateway and Motorola. Marty has served on a number of boards of directors for public and private companies. He spent a number of years as both an active-duty and reserve officer in the US Army. Marty has a BS in computer science from the United States Military Academy, an MS in computer engineering from the University of Florida, is a graduate of the Harvard Business School Executive Education Program, and has a Doctorate of Management from Case Western Reserve University.

Michael T. Fisher is a founding partner of AKF Partners, a growth consulting firm focusing on meeting the needs of today's fast-paced and hyper-growth companies. Prior to cofounding AKF Partners, Michael held many industry roles including CTO of Quigo, acquired by AOL in 2007, and VP of engineering and architecture for PayPal. He served as a pilot in the US Army. Michael received a PhD and MBA from Case Western Reserve University's Weatherhead School of Management, an MS in information systems from Hawaii Pacific University, and a BS in computer science from the United States Military Academy (West Point). Michael is an adjunct professor in the Design and Innovation Department at Case Western Reserve University's Weatherhead School of Management.

This page intentionally left blank

Distribute Your Work

In 2004 the founding team of ServiceNow (originally called Glidesoft), built a generic workflow platform they called “Glide.” In looking for an industry in which they could apply the Glide platform, the team felt that the Information Technology Service Management (ITSM) space, founded on the Information Technology Infrastructure Library (ITIL), was primed for a platform as a service (PaaS) player. While there existed competition or potentially substitutes in this space in the form of on-premise software solutions such as Remedy, the team felt that the success of companies like Salesforce for customer relationship management (CRM) solutions was a good indication of potential adoption for online ITSM solutions.

In 2006 the company changed its name to ServiceNow in order to better represent its approach to the needs of buyers in the ITSM solution space. By 2007 the company was profitable. Unlike many startups, ServiceNow appreciated the value of designing, implementing, and deploying for scale early in its life. The initial solutions design included the notions of both fault isolation (covered in Chapter 9, “Design for Fault Tolerance and Graceful Failure”) and Z axis customer splits (covered in this chapter). This fault isolation and customer segmentation allowed the company to both scale to profitability early on and to avoid the noisy-neighbor effect common to so many early SaaS and PaaS offerings. Furthermore, the company valued the cost effectiveness afforded by multitenancy, so while they created fault isolation along customer boundaries, they still designed their solution to leverage multitenancy within a database management system (DBMS) for smaller customers not requiring complete isolation. Finally, the company also valued the insight offered by outside perspectives and the value inherent to experienced employees.

ServiceNow contracted with AKF Partners over a number of engagements to help them think through their future architectural needs and ultimately hired one of the founding partners of AKF, Tom Keeven, to augment their already-talented engineering staff. “We were born with incredible scalability from the date of launch,” indicated Tom. “Segmentation along customer boundaries using the AKF Z axis of scale went a long way to ensuring that we could scale into our early demand. But as our customer base grew and the average size of our customer increased beyond small early adopters to much larger Fortune 500 companies, the characterization of our workload changed and the average number of seats per customer dramatically increased. All of these led to each customer performing more transactions and storing more data. Furthermore, we were extending our scope of functionality, adding significantly greater value to

our customer base with each release. This functionality extension meant even greater demand was being placed on the systems for customers both large and small. Finally, we had a small problem with running multiple schemas or databases under a single DBMS within MySQL. Specifically, the catalog functionality within MySQL [sometimes technically referred to as the `information_schema`] was starting to show contention when we had 30 high-volume tenants on each DBMS instance.”

Tom Keeven’s unique experience building Web-based products from the high-flying days of Gateway Computer, to the Wild West startup days of the Internet at companies like eBay and PayPal, along with his experience across a number of clients at AKF, made him uniquely suited to helping to solve ServiceNow’s challenges. Tom explained, “The database catalog problem was simple to solve. For very large customers we simply had to dedicate a DBMS per customer, thereby reducing the burst radius of the fault isolation zone. Medium-size customers may have tenants below 30, and small customers could continue to have a high degree of multitenancy [for more on this see Chapter 9]. The AKF Scale Cube was helpful in offsetting both the increasing size of our customers and the increased demands of rapid functionality extensions and value creation. For large customers with heavy transaction processing demands we incorporated the X axis by replicating data to read-only databases. With this configuration, reports, which are typically computationally and I/O intensive but read-only, could be run without impact to the scale of the lighter-weight transaction (OLTP) requests. While the report functionality also represented a Y axis (service/function or resource-based) split, we added further Y axis splits by service to enable additional fault isolation by service, significantly greater caching of data, and faster developer throughput. All of these splits, the X, Y, and Z axes, allowed us to have consistency within the infrastructure and purchase similar commodity systems for any type of customer. Need more horsepower? The X axis allows us to increase transaction volumes easily and quickly. If data is starting to become unwieldy on databases, our architecture allows us to reduce the degree of multitenancy (Z axis) or split discrete services off (Y axis) onto similarly sized hardware.”

This chapter discusses scaling databases and services through cloning and replication, separating functionality or services, and splitting similar data sets across storage and application systems. Using these three approaches, you will be able to scale nearly any system or database to a level that approaches infinite scalability. We use the word *approaches* here as a bit of a hedge, but in our experience across hundreds of companies and thousands of systems these techniques have yet to fail. To help visualize these three approaches to scale we employ the AKF Scale Cube, a diagram we developed to represent these methods of scaling systems. Figure 2.1 shows the AKF Scale Cube, which is named after our partnership, AKF Partners.

At the heart of the AKF Scale Cube are three simple axes, each with an associated rule for scalability. The cube is a great way to represent the path from minimal scale (lower left front of the cube) to near-infinite scalability (upper right back corner of the cube). Sometimes, it’s easier to see these three axes without the confined space of the cube. Figure 2.2 shows the axes along with their associated rules. We cover each of the three rules in this chapter.

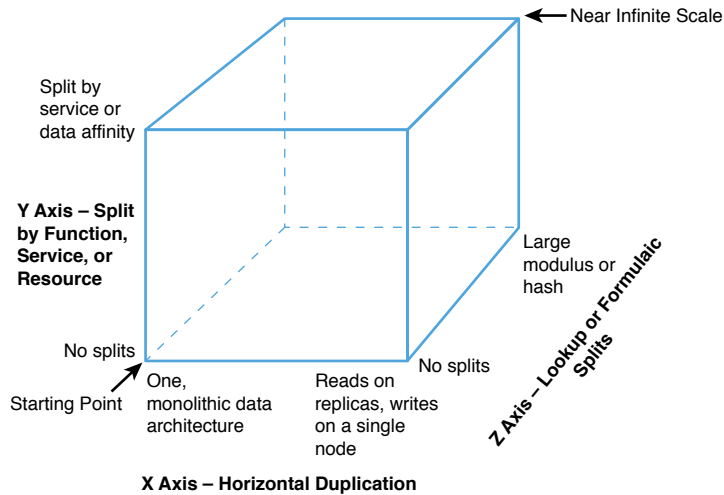


Figure 2.1 AKF Scale Cube

Not every company will need all of the capabilities (all three axes) inherent to the AKF Scale Cube. For many of our clients, one of the types of splits (X, Y, or Z) meets their needs for a decade or more. But when you have the type of viral success achieved by the likes of ServiceNow, it is likely that you will need two or more of the splits identified within this chapter.




Figure 2.2 Three axes of scale

Rule 7—Design to Clone or Replicate Things (X Axis)

Rule 7: What, When, How, and Why

What: Typically called horizontal scale, this is the duplication of services or databases to spread transaction load.

When to use:

- Databases with a very high read-to-write ratio (5:1 or greater—the higher the better).
- Any system where transaction growth exceeds data growth.

How to use:

- Simply clone services and implement a load balancer.
- For databases, ensure that the accessing code understands the difference between a read and a write.

Why: Allows for fast scale of transactions at the cost of duplicated data and functionality.

Key takeaways: X axis splits are fast to implement, are low cost from a developer effort perspective, and can scale transaction volumes nicely. However, they tend to be high cost from the perspective of operational cost of data.

Often, the hardest part of a solution to scale is the database or persistent storage tier. The beginning of this problem can be traced back to Edgar F. Codd’s 1970 paper “A Relational Model of Data for Large Shared Data Banks,”¹ which is credited with introducing the concept of the relational database management system (RDBMS). Today’s most popular RDBMSs, such as Oracle, MySQL, and SQL Server, just as the name implies, allow for relations between data elements. These relationships can exist within or between tables. The tables of most OLTP systems are normalized to third normal form,² where all records of a table have the same fields, nonkey fields cannot be described by only one of the keys in a composite key, and all nonkey fields must be described by the key. Within the table each piece of data is related to other pieces of data in that table. Between tables there are often relationships, known as foreign keys. Most applications depend on the database to support and enforce these relationships because of its ACID properties (see Table 2.1). Requiring the database to maintain and enforce these relationships makes it difficult to split the database without significant engineering effort.

Table 2.1 ACID Properties of Databases

Property	Description
Atomicity	All of the operations in the transaction will complete, or none will.
Consistency	The database will be in a consistent state when the transaction begins and ends.
Isolation	The transaction will behave as if it is the only operation being performed upon the database.
Durability	Upon completion of the transaction, the operation will not be reversed.

One technique for scaling databases is to take advantage of the fact that most applications and databases perform significantly more reads than writes. A client of ours that handles booking reservations for customers has on average 400 searches for a single booking. Each booking is a write and each search a read, resulting in a 400:1 read-to-write ratio. This type of system can be easily scaled by creating read-only copies (or replicas) of the data.

There are a couple of ways that you can distribute the read copy of your data depending on the time sensitivity of the data. Time (or temporal) sensitivity is how fresh or completely correct the read copy has to be relative to the write copy. Before you scream out that the data has to be instant, real time, in sync, and completely correct across the entire system, take a breath and appreciate the costs of such a system. While perfectly in-sync data is ideal, it costs . . . a lot. Furthermore, it doesn't always give you the return that you might expect or desire for that cost. Rule 19, "Relax Temporal Constraints" (see Chapter 5, "Get Out of Your Own Way"), will delve more into these costs and the resulting impact on the scalability of products.

Let's go back to our client with the reservation system that has 400 reads for every write. They're handling reservations for customers, so you would think the data they display to customers would have to be completely in sync. For starters you'd be keeping 400 sets of data in sync for the one piece of data that the customer wants to reserve. Second, just because the data is out of sync with the primary transactional database by 3 or 30 or 90 seconds doesn't mean that it isn't correct, just that there is a chance that it isn't correct. This client probably has 100,000 pieces of data in their system at any one time and books 10% of those each day. If those bookings are evenly distributed across the course of a day, they are booking one reservation just about every second (0.86 second). All things being equal, the chance of a customer wanting a particular booking that is already taken by another customer (assuming a 90-second sync of data) is 0.104%. Of course even at 0.1% some customers will select a booking that is already taken, which might not be ideal but can be handled in the application by doing a final check before allowing the booking to be placed in the customer's cart. Certainly every application's data needs are going to be different, but from this discussion we hope you will get a sense of how you can push back on the idea that all data has to be kept in sync in real time.

Now that we've covered the time sensitivity, let's start discussing the ways to distribute the data. One way is to use a caching tier in front of the database. An object cache can be used to read from instead of going back to the application for each query. Only when the data has been marked expired would the application have to query the primary transactional database to retrieve the data and refresh the cache. We highly recommend this as a first step given the availability of numerous excellent, open-source key-value stores that can be used as object caches.

The next step beyond an object cache between the application tier and the database tier is replicating the database. Most major relational database systems allow for some type of replication "out of the box." Many databases implement replication through some sort of *master-slave* concept—the master database being the primary transactional database that gets written to, and the slave databases being read-only copies of the

master database. The master database keeps track of updates, inserts, deletes, and so on in a binary log. Each slave requests the binary log from the master and replays these commands on its database. While this is asynchronous, the latency between data being updated in the master and then in the slave can be very low, depending on the amount of data being inserted or updated in the master database. In our client's example, 10% of the data changed each day, resulting in one update per second. This is likely a low enough volume of change to maintain the slave databases with low latency. Often this implementation consists of several slave databases or read replicas that are configured behind a load balancer. The application makes a read request to the load balancer, which passes the request in either a round-robin or least-connections manner to a read replica. Some databases further allow replication using a master-master concept in which either database can be used to read or write. Synchronization processes help ensure the consistency and coherency of the data between the masters. While this technology has been available for quite some time, we prefer solutions that rely on a single write database to help eliminate confusion and logical contention between the databases.

We call the type of split (replication) an X axis split, and it is represented on the AKF Scale Cube in Figure 2.1 as the X axis—Horizontal Duplication. An example that many developers familiar with hosting Web applications will recognize is on the Web or application tier of a system, running multiple servers behind a load balancer all with the same code. A request comes in to the load balancer which distributes it to any one of the many Web or application servers to fulfill. The great thing about this distributed model on the application tier is that you can put dozens, hundreds, or even thousands of servers behind load balancers all running the same code and handling similar requests.

The X axis can be applied to more than just the database. Web servers and application servers typically can be easily cloned. This cloning allows the distribution of transactions across systems evenly for horizontal scale. Cloning of application or Web services tends to be relatively easy to perform and allows us to scale the number of transactions processed. Unfortunately, it doesn't really help us when trying to scale the data we must manipulate to perform these transactions. In memory, caching of data unique to several customers or unique to disparate functions might create a bottleneck that keeps us from scaling these services without significant impact on customer response time. To solve these memory constraints we'll look to the Y and Z axes of our scale cube.

Rule 8—Design to Split Different Things (Y Axis)

Rule 8: What, When, How, and Why

What: Sometimes referred to as scale through services or resources, this rule focuses on scaling by splitting data sets, transactions, and engineering teams along verb (services) or noun (resources) boundaries.

When to use:

- Very large data sets where relations between data are not necessary.
- Large, complex systems where scaling engineering resources requires specialization.

How to use:

- Split up actions by using verbs, or resources by using nouns, or use a mix.
- Split both the services and the data along the lines defined by the verb/noun approach.

Why: Allows for efficient scaling of not only transactions but also very large data sets associated with those transactions. Also allows for the efficient scaling of teams.

Key takeaways: Y axis or data/service-oriented splits allow for efficient scaling of transactions, large data sets, and can help with fault isolation. Y axis splits help reduce the communication overhead of teams.

When you put aside the religious debate around the concepts of services- (SOA) and resources- (ROA) oriented architectures and look deep into their underlying premises, they have at least one thing in common. Both concepts force architects and engineers to think in terms of separation of responsibilities within their architectures. At a high and simple level, they do this through the concepts of verbs (services) and nouns (resources). Rule 8, and our second axis of scale, takes the same approach. Put simply, Rule 8 is about scaling through the separation of distinct and different functions and data within a site. The simple approach to Rule 8 tells us to split up our product by either nouns or verbs or a combination of both nouns and verbs.

Let's split up our site using the verb approach first. If our site is a relatively simple e-commerce site, we might break it into the necessary verbs of signup, login, search, browse, view, add to cart, and purchase/buy. The data necessary to perform any one of these transactions can vary significantly from the data necessary for the other transactions. For instance, while it might be argued that signup and login need the same data, they also require some data that is unique and distinct. Signup, for instance, probably needs to be capable of checking whether a user's preferred ID has been chosen by someone else in the past, whereas login might not need to have a complete understanding of every other user's ID. Signup likely needs to write a fair amount of data to some permanent data store, but login is likely a read-intensive application to validate a user's credentials. Signup may require that the user store a fair amount of personally identifiable information (PII) including credit card numbers, whereas login does not likely need access to all of this information at the time that a user would like to establish a login.

The differences and resulting opportunities for this method of scale become even more apparent when we analyze obviously distinct functions like search and login. In the case of login we are mostly concerned with validating the user's credentials and potentially establishing some notion of session (we've chosen the word *session* rather than *state* for a reason we explore in Rule 40 in Chapter 10, "Avoid or Distribute State"). Login is concerned with the user and as a result needs to cache and interact with data about that user. Search, on the other hand, is concerned with the hunt for an item and is most concerned with user intent (vis-à-vis a search string, query, or search terms typically typed into a search box) and the items that we have in stock within our catalog. Separating these sets of data allows us to cache more of them within the confines of memory available on our system and process transactions faster as a result of higher cache hit ratios. Separating this data within our back-end persistence systems (such as a database) allows

us to dedicate more “in memory” space within those systems and respond faster to the clients (application servers) making requests. Both systems respond faster as a result of better utilization of system resources. Clearly we can now scale these systems more easily and with fewer memory constraints. Moreover, the Y axis adds transaction scalability by splitting up transactions in the same fashion as Rule 7, the X axis of scale.

Hold on! What if we want to merge information about the user and our products such as in the case of recommending products? Note that we have just added another verb—*recommend*. This gives us another opportunity to perform a split of our data and our transactions. We might add a recommendation service that asynchronously evaluates past user purchase behavior against users who have similar purchase behaviors. This in turn may populate data in either the login function or the search function for display to the user when he or she interacts with the system. Or it can be a separate synchronous call made from the user’s browser to be displayed in an area dedicated to the result of the recommend call.

Now how about using nouns to split items? Again, using our e-commerce example, we might identify certain resources upon which we will ultimately take actions (rather than the verbs that represent the actions we take). We may decide that our e-commerce site is made up of a product catalog, product inventory, user account information, marketing information, and so on. Using our noun approach, we may decide to split up our data into these categories and then define a set of high-level primitives such as create, read, update, and delete actions on these primitives.

While Y axis splits are most useful in scaling data sets, they are also useful in scaling code bases. Because services or resources are now split, the actions we perform and the code necessary to perform them are split up as well. This means that very large engineering teams developing complex systems can become experts in subsets of those systems and don’t need to worry about or become experts on every other part of the system. Teams that own each service can build the interface (such as an API) into their service and own it. Assuming that each team “owns” its own code base, we can cut down on the communication overhead associated with Brooks’ Law. One tenet of Brooks’ Law is that developer productivity is reduced as a result of increasing team sizes.³ The communication effort within any team to coordinate team efforts is a square of the number of participants in the team. Therefore, with increasing team size comes decreasing developer productivity as more developer time is spent on coordination. By segmenting teams and enabling ownership, such overhead is decreased. And of course because we have split up our services, we can also scale transactions fairly easily.

Rule 9—Design to Split Similar Things (Z Axis)

Rule 9: What, When, How, and Why

What: This is very often a split by some unique aspect of the customer such as customer ID, name, geography, and so on.

When to use: Very large, similar data sets such as large and rapidly growing customer bases or when response time for a geographically distributed customer base is important.

How to use: Identify something you know about the customer, such as customer ID, last name, geography, or device, and split or partition both data and services based on that attribute.

Why: Rapid customer growth exceeds other forms of data growth, or you have the need to perform fault isolation between certain customer groups as you scale.

Key takeaways: Z axis splits are effective at helping you to scale customer bases but can also be applied to other very large data sets that can't be pulled apart using the Y axis methodology.

Often referred to as *sharding* and *podding*, Rule 9 is about taking one data set or service and partitioning it into several pieces. These pieces are often equal in size but may be of different sizes if there is value in having several unequally sized chunks or shards. One reason to have unequally sized shards is to enable application rollouts that limit your risk by affecting first a small customer segment, and then increasingly large segments of customers as you feel you have identified and resolved major problems. It also serves as a great method for allowing discovery—as you roll out first to smaller segments, if a feature is not getting the traction you expect (or if you want to expose an “early” release to learn about usage of a feature), you can modify the feature before it is exposed to everybody.

Often sharding is accomplished by separating something we know about the requestor or customer. Let's say that we are a time card and attendance-tracking SaaS provider. We are responsible for tracking the time and attendance for employees of each of our clients, who are in turn enterprise-class customers with more than 1,000 employees each. We might determine that we can easily partition or shard our solution by company, meaning that each company could have its own dedicated Web, application, and database servers. Given that we also want to leverage the cost efficiencies enabled by multitenancy, we also want to have multiple small companies exist within a single shard. Really big companies with many employees might get dedicated hardware, whereas smaller companies with fewer employees could cohabit within a larger number of shards. We have leveraged the fact that there is a relationship between employees and companies to create scalable partitions of systems that allow us to employ smaller, cost-effective hardware and scale horizontally (we discuss horizontal scale further in Rule 10 in the next chapter).

Maybe we are a provider of advertising services for mobile phones. In this case, we very likely know something about the end user's device and carrier. Both of these create compelling characteristics by which we can partition our data. If we are an e-commerce player, we might split users by their geography to make more efficient use of our available inventory in distribution centers, and to give the fastest response time on the e-commerce Web site. Or maybe we create partitions of data that allow us to evenly distribute users based on the recency, frequency, and monetization of their purchases. Or, if all else fails, maybe we just use some modulus or hash of a user identification (userid) number that we've assigned the user at signup.

Why would we ever decide to partition similar things? For hyper-growth companies, the answer is easy. The speed with which we can answer any request is at least partially determined by the cache hit ratio of near and distant caches. This speed in turn indicates how many transactions we can process on any given system, which in turn determines

how many systems we need to process a number of requests. In the extreme case, without partitioning of data, our transactions might become agonizingly slow as we attempt to traverse huge amounts of monolithic data to come to a single answer for a single user. Where speed is paramount and the data to answer any request is large, designing to split different things (Rule 8) and similar things (Rule 9) becomes a necessity.

Splitting similar things obviously isn't just limited to customers, but customers are the most frequent and easiest implementation of Rule 9 within our consulting practice. Sometimes we recommend splitting product catalogs, for instance. But when we split diverse catalogs into items such as lawn chairs and diapers, we often categorize these as splits of different things. We've also helped clients shard their systems by splitting along a modulus or hash of a transaction ID. In these cases, we really don't know anything about the requestor, but we do have a monotonically increasing number upon which we can act. These types of splits can be performed on systems that log transactions for future reference as in a system designed to retain errors for future evaluation.

Summary

We maintain that three simple rules can help you scale nearly everything. Scaling along the X, Y, and Z axes each has its own set of benefits. Typically X axis scaling has the lowest cost from a design and software development perspective; Y and Z axis scaling is a little more challenging to design but gives you more flexibility to further fully separate your services, customers, and even engineering teams. There are undoubtedly more ways to scale systems and platforms, but armed with these three rules, few if any scale-related problems will stand in your way:

- **Scale by cloning**—Cloning or duplicating data and services allows you to scale transactions easily.
- **Scale by splitting different things**—Use nouns or verbs to identify data and services to separate. If done properly, both transactions and data sets can be scaled efficiently.
- **Scale by splitting similar things**—Typically these are customer data sets. Set customers up into unique and separated shards or swim lanes (see Chapter 9 for the definition of *swim lane*) to enable transaction and data scaling.

Notes

1. Edgar F. Codd, "A Relational Model of Data for Large Shared Data Banks," 1970, www.seas.upenn.edu/~zives/03f/cis550/codd.pdf.
2. Wikipedia, "Third Normal Form," http://en.wikipedia.org/wiki/Third_normal_form.
3. Wikipedia, "Brooks' Law," https://en.wikipedia.org/wiki/Brooks'_law.

Index

Numbers

- 2PC (two-phase commit), 117
- 9/11, market sensitivity to, 159
- 24.0 disaster, PayPal, 60, 107
- 80-20 rule (Pareto Principle), 9, 87
- 1992 W3 Project page, 13

A

- A/B customer testing, 96-97
- ACB (Amazon customer database), xvi
- Access logs, make use of, 56
- Accidents, learn from, 97-98
- ACID properties (atomicity, consistency, isolation, and durability)
 - of databases, 22, 107
 - relax temporal constraints, 68-70
 - of reliable transactions, 60
 - use relational databases for, 47
- Actions, learning from our mistakes, 98
- Active/passive configuration, removing
 - single point of failure, 131
- Actively use log files rule, 55-58, 187
- Aggregation, log, 56, 57
- Ajax (Asynchronous JavaScript and XML)
 - cache calls, 82-84, 190
 - example of, 73
 - overview of, 81-82
- AKF Scale Cube, 20-21
- Alerts
 - design application to be monitored, 57
 - in modern coupled systems, 97
 - monitor for events, 160
 - monitor logs through, 57
 - set up automatic, 170
- Aliasing, redirecting traffic for, 65
- Amazon
 - AWS's regions, 39
 - at edge of OLTP in 2001, 1-3
 - fault isolation in, 129
 - success story, xv-xvii
- Amazon customer database (ACB), xvi
- Amazon DynamoDB, 48-49
- Amdahl's Law, 32
- Apache
 - object cache, 89
 - use of log files, 56
- Apache Hadoop, 51
- Apache Web server, 66-67
- Application caches, utilize, 86-88, 190
- Application servers
 - avoid state/replication servers within, 144-145
 - competency in, 173
 - design to scale out, 33
 - put object cache on its own tier vs., 90-91
 - scale by cloning, 24
- Archive storage data, 163-166, 198-199
- Art of Scalability, Second Edition, Scalability Rules* vs., xix
- Asynchronous calls, across swim lanes, 128-129
- Asynchronous communication and message buses, rules
 - Avoid overcrowding message bus, 154-157, 198
 - Communicate asynchronously, 149-151, 197
 - Ensure message bus can scale, 151-154, 197
 - overview of, 147-148
- Asynchronous completion, eliminate
 - checking your work for, 64
- Asynchronous JavaScript and XML. *See* Ajax (Asynchronous JavaScript and XML)
- Asynchronous transfer of data, for business systems, 109, 111

- Atomicity
 - as ACID property of databases, 22
 - defined, 107
 - overview of, 60
 - Attributes, data model design, 112
 - Auction format, eBay, 1
 - Authorization cookie, 143
 - Availability. *See also* HA (high availability)
 - CAP Theorem, 69
 - fault isolation benefits, 127
 - impact of firewalls on, 53–55
 - number of components in series and, 134–135
 - Avoid or Distribute State, rules
 - Maintain sessions in browser when possible, 142–143, 196
 - Make use of distributed cache for states, 144–146, 196–197
 - overview of, 139–140
 - Strive for statelessness, 140–142, 196
 - Avoid overcrowding message bus rule, 154–157, 198
 - Avoid putting systems in series rule, 132–135, 195
- B**
- Barrese, James, 43–44
 - BASE architecture, database consistency, 69
 - Batch workloads
 - induction and, 168
 - partition, 166–169, 199
 - Bayesian belief networks, 172
 - Be aware of costly relationships rule, 111–114, 192–193
 - Be competent rule, 172–174, 200
 - Be wary of scaling through third parties rule, 161–163, 198
 - Benefit ranking, scalability rules, 200–202
 - Bezos, Jeff, xv–xvi
 - BigTable, Google, 49
 - Binder, Lon, 73
 - Black Friday, capacity planning/monitoring for, 159
 - Brewer (or CAP) Theorem, 69
 - Broken Windows Theory
 - Freakonomics* comparison to, 168–169
 - New York experiment, 167
 - Browsers
 - leverage Ajax cache in, 82–84
 - maintain session data in, 142–143
 - Reduce DNS lookups rule and, 11–12
 - simultaneous connection feature of, 13–14
 - store session data in distributed cache vs., 144–146
 - Business
 - cache objects for, 74–75
 - learn aggressively from operations, 97
 - monitor metrics for, 170–172
 - optimize SQL queries, 114
 - Business intelligence, remove from transaction processing, 109–111, 192
- C**
- Cache Ajax calls rule, 80–84, 189
 - Cache-Control header
 - leverage Ajax cache, 82–84
 - leverage page caches, 84–85
 - overview of, 78–79
 - Cache is King, in technology, 73
 - Cache miss, page caches, 84
 - Caches
 - aggressive use of. *See* Use Caching
 - Aggressively, rules
 - challenges of adding, 75
 - store data with distributed, 144–146
 - store writes needed soon in, 63
 - Callbacks, asynchronous communication, 148
 - CAP (or Brewer) Theorem, 69
 - Capacity on demand, for unpredictable events, 160
 - Cassandra, 49
 - CDNs (content delivery networks)
 - example of using, 74
 - leveraging, 75–77, 188
 - offload traffic via, 75–77
 - Ceph, 48
 - Charles Schwab, 159–160
 - Circuit breakers, 136
 - Click tracking, 65, 67–68
 - Clone
 - design to. *See* X axis (Design to Clone or Replicate things) rule
 - as solution to single point of failure, 131

- Cloud
 - design to leverage, 40–42, 185–186
 - make use of while you grow, 36
 - Cluster, 125–126, 162–163
 - Codd, Edgar F., 47
 - Code
 - failing to design for rollback of, 102–105
 - redirection with, 66
 - session/state require complexity of, 141
 - set HTTP Expires header in, 79
 - Code bases, splitting different things, 26
 - Colocation providers, 36
 - Columns
 - design for rollback, 105
 - specify for Select * or Insert, 121
 - Commodity systems, 33–35, 185
 - Communicate asynchronously when possible
 - rule, 149–151, 197
 - Competency, for each component, 172–174, 200
 - Competitive differentiation, don't check
 - work for, 64
 - Complexity
 - of Apache Web server redirects, 67
 - design to scale out and, 30
 - don't overengineer the solution, 3–5
 - scale out your hosting solution, 39
 - session and state requiring code, 141
 - simplify solution three times over vs., 8–10
 - Components
 - competency in, 172–174, 200
 - reducing number of in series, 132–135
 - Concurrency, 115–116
 - CONF, PayPal, 103
 - Config file markdown, feature wire-on/
 - wire-off approach, 136
 - Confinity, PayPal, 59–61
 - Conflict, resolving in rules, 63–64
 - Consistency
 - as ACID property of databases, 22
 - database locks facilitating, 115–116
 - defined, 107
 - relaxing temporal constraints for, 69
 - understanding, 60
 - Constraint satisfaction problems (CSPs), 68
 - Constraints. *See* Temporal constraints
 - Content delivery networks. *See* CDNs
 - (content delivery networks)
 - Continuous integration, stored procedures
 - impeding, 110
 - Cookies, 142–143
 - Cordrey, Tanya, 123
 - Cost
 - database, 109–110
 - database locks, 115
 - design to leverage cloud, 41
 - design to scale out, 32
 - doubling activity and, 63
 - fault isolation benefits, 127
 - firewall, 54
 - logging, 57–58
 - message buses, 155–157
 - multiple live sites and, 36–40
 - scale out with commodity systems, 34–35
 - scale through third parties, 162
 - of session and state, 141
 - Cost-justify storage data, 163–166, 198–199
 - Cost-Value Data Dilemma, 52
 - Costly relationships, 111–114, 192–193
 - Couchbase, 49
 - CouchDB, 49
 - CPUs, scale out via commodity systems, 34
 - Craigslist, rivalry with eBay, 141
 - Crime rates, Broken Windows Theory for,
 - 167–168
 - CSPs (Constraint satisfaction problems), 68
 - Cursors, 118–119
 - Customers, ways to watch, 96–97
 - Cyber Monday, capacity planning/
 - monitoring for, 159
- D**
- D-I-D (Design-Implement-Deploy) process,
 - 6–8, 181
 - Dalzell, Rick, 128, 139, xv–xvii
 - Data
 - cost/value of message bus, 155–157
 - design for roll back of, 105
 - mapping issues, 120–121
 - monitor scope vs. amount of, 171–172
 - normal forms and integrity of, 113–114
 - Data centers, 35–40
 - Data definition language (DDL) statements,
 - 112
 - Data sets, 24–28

- Data stores, object caches as, 89
- Database as a service (DBaaS), 44
- Database locks, 114–116, 119
- Database management system (DBMS), 19–20
- Database-only implementations, distributed session/state and, 145
- Database Rules
 - Be aware of costly relationships, 111–114, 192–193
 - Don't select everything, 120–121, 194
 - overview of, 107–108
 - Pass on using multiphase commits, 116–118, 193
 - Remove business intelligence from transaction processing, 109–111, 192
 - Try not to use "Select for Update," 118–119, 194
 - Use right type of database lock, 114–116, 193
- Databases
 - ACID properties of, 22
 - alternative storage strategies to, 48–51
 - competency in, 173
 - design for rollback, 105
 - design to scale out, 33
 - solve slow load times, 74
 - use appropriately, 47–48, 186
- DBaaS (Database as a service), 44
- DBMS (database management system), 19–20
- DDL (data definition language) statements, 112
- Deadlock, from 2PC protocol, 107
- Debug errors, with log files, 57
- Deductive workloads
 - defined, 167
 - partitioning, 166–169, 199
- Deployment, design for, 8
- Design
 - overengineering, 3–5
 - simplifying, 9–10
- Design for Fault Tolerance/Graceful Failure, rules
 - Avoid putting systems in series, 132–135, 195
 - Design using fault-isolative swim lanes, 124–130, 194
 - Ensure you can wire on/wire off features, 135–138, 195
 - Never trust single points of failure, 130–132, 195
 - overview of, 123–124
- Design-Implement-Deploy (D-I-D) process, 6–8, 181
- Design scale into the solution rule, 6–8, 181
- Design to leverage cloud rule, 40–42, 185–186
- Design to Scale Out Horizontally, rules
 - Design to leverage cloud, 40–42, 185–186
 - Design your solution to scale out, not just up, 31–33, 184
 - overview of, 29–31
 - Scale out your hosting solution, 35–40, 185
 - Use commodity systems, 33–35, 185
- Design using fault-isolative swim lanes rule, 124–130, 194
- Design your application to be monitored rule, 169–172, 199–200
- Design your solution to scale out, not just up rule, 31–33, 184
- Disaster recovery, lowering costs by scaling out, 36
- Distribute Your Work, rules
 - Design to Clone or Replicate things (X axis), 22–24, 183
 - Design to Split Different things (Y axis), 24–26, 184–185
 - Design to Split Similar things (Z axis), 26–28, 185
 - overview of, 19–21
- Distributed cache, 144–146
- DNS (Domain Name System)
 - reduce lookups, 10–12, 182
 - use content delivery networks, 75–77
- Document stores, 49
- Domain Name System. *See* DNS (Domain Name System)
- Domains, redirect traffic for misspelled/changed, 65
- Don't check your work rule, 61–64, 187
- Don't overengineer the solution rule, 3–5, 180–181
- Don't rely on QA to find mistakes rule, 100–102, 191

- Don't select everything rule, 120–121, 194
- Double-checking yourself, stop, 61–64, 187
- Dubner, Stephen J., 168–169
- Durability
 - as ACID property of databases, 22
 - defined, 107
 - distributed session/state considerations, 145
 - understanding, 60
- Dynamic content
 - adding CDN for, 78
 - use caches to scale, 73
- Dynamic Site Accelerator, Akamai, 78
- E**
- E-commerce, capacity planning/monitoring
 - for, 159
- EBay
 - Brad Peterson at, 160
 - at edge of OLTP in 2001, 1
 - James Barrese at, 43
 - June 1999 outages at, 1–2
 - on PayPal failure, 104–105
 - power of simple and easy and, 141
 - redesigned architectural principles, 2–3
- Edge servers, content delivery networks as, 76
- Elasticsearch, ELK framework, 57
- Electrical circuits, reduce number of
 - components in series, 132–135
- ELK (Elasticsearch, Logstash, Kibana)
 - framework, 57
- Employees, multiple live site consideration, 40
- Engineering resources, scale by splitting
 - different things, 26
- Ensure message bus can scale rule, 151–154, 197
- Ensure you can wire on/wire off features
 - rule, 135–138, 195
- Enterprise resource planning (ERP), 110–111
- Enterprise service bus. *See* Message bus
- Entities
 - data model design, 112
 - relationship between, 113
- Entity relationship diagrams (ERDs), 112
- ERDs (entity relationship diagrams), 112
- ERP (enterprise resource planning), 110–111
- Error logs, making use of, 56
- ETag header, leverage page caches, 85
- Ethernet, collision domains in, 126
- Expertise, defer to, 99
- Expires headers
 - leveraging Ajax cache, 82–84
 - leveraging page caches, 84–85
 - use, 77–80, 189
- Explicit database locks, 115
- Extensible record stores (ERSs), 49
- Extent database locks, 115
- External API, asynchronous calls to, 149
- F**
- F-graph, Facebook, 173
- Failing to design for rollback rule, 102–105, 191–192
- Failure. *See also* Design for Fault Tolerance/Graceful Failure, rules
 - asynchronous calls prevent spreading of, 149
 - be preoccupied with, 99
 - components in series subject to, 132–135
 - design swim lanes for, 126–127
 - ensure message bus can scale, 151–154
 - learn from mistakes. *See* Learn From Your Mistakes, rules
 - never trust single points of, 130–132, 195
- Fault isolation
 - achieving, 127–130
 - benefits of, 127
 - design swim lanes for, 124–127, 194
 - example of swim lanes, 123–124
 - partition inductive, deductive, batch, and user interactive/OLTP workloads, 168
 - scale by splitting similar things, 27
 - scale out your hosting solution, 36
- Fault isolation domains, 126
- Fault tolerance. *See also* Design for Fault Tolerance/Graceful Failure, rules
 - implement asynchronous calls, 149
 - remove business intelligence from transaction processing, 111
- Features, wire on/wire off, 135–138, 194
- Fifth normal form, 113
- File markdown, feature wire-on/wire-off, 137
- File systems, as overlooked storage systems, 48
- Financial services solutions, 159–160
- Firesheep, 143

- Firewalls
 - both inside and outside network, 134
 - design to scale out, not up, 29–31
 - homogeneous networks for, 15
 - Firewalls, firewalls, everywhere rule, 52–55, 186
 - First normal form, 113
 - Flash crash, 160
 - FOR UPDATE cursors, minimize, 118–119
 - Foreign keys, enforcing referential integrity, 112
 - Fourth normal form, 113
 - Freakonomics*, 167
 - Frequency of data access, RFM analysis, 164–165
- G**
- Garrett, Jesse James, 81
 - Geiger, Chuck, 103–104
 - Get Out of Your Own Way, rules
 - Don't check your work, 61–64, 187
 - overview of, 59–61
 - Relax temporal constraints, 68–70, 188
 - Stop redirecting traffic, 64–68, 188
 - GFS (Google File System), storage, 48
 - Giuliani, Mayor, 167
 - Glide workflow program, 19
 - Goldfish vs. thoroughbreds concept, 35
 - Google
 - BigTable, 49
 - at edge of OLTP in 2001, 1–3
 - MapReduce data storage, 50–51
 - no state/session in, 141
 - Google File System (GFS), storage, 48
 - Granularity, of database locks, 115
 - Gray, Jim, 60
- H**
- HA (high availability)
 - design to scale out, not up, 29–30
 - protect against data corruption via, 62
 - scalability and, 124
 - scale out your hosting solution, 36
 - with swim lanes, 126
 - Hardware, provisioning in cloud, 41
 - Hash function, 89
 - HBase, extensible record store, 49
 - Header() command, Expires header, 79
 - Headers, using Expire, 77–80
 - Health Insurance Portability and Accountability Act (HIPAA), 29
 - HIPAA (Health Insurance Portability and Accountability Act), 29
 - Hit ratio, object caches, 90–91
 - Homogeneous networks rule, 15
 - Horizontal duplication, scaling by cloning, 24
 - Horizontal scale
 - design to scale out. *See* Design to Scale Out
 - Horizontally, Rules
 - distribute your work through, 20–21
 - remove single point of failure, 132
 - scale by cloning, 22–24
 - Hosting solution, scale out your, 35–40, 185
 - Hot/cold configuration, remove single point of failure, 131
 - HTML, redirection with, 66
 - HTTP 3xx status codes, redirection with, 65–66
 - HTTP headers
 - control caching, 78
 - Expires header, 79–80
 - leverage Ajax cache, 82–84
 - leverage page caches, 84–85
 - HTTPS, protect cookies from sidejacking, 143
 - Hubs, homogeneous networks and, 15
- I**
- IaaS (Infrastructure as a service)
 - created by Amazon, xvi
 - design to leverage cloud, 41–42
 - rent capacity from, 160
 - Images, reduce objects where possible, 13
 - Implementation
 - as actual coding of solution, 10
 - design scale into the solution, 7
 - simplify solution three times over, 10
 - Implicit database locks, 115
 - Incident detection, fault isolation benefits, 127
 - Inconsistent demand, cloud for, 41
 - Inductive workloads, partition, 166–169, 199
 - Infrastructure as a service. *See* IaaS (Infrastructure as a service)
 - Infrastructure, competency in, 173–174, 200
 - Insert, specify columns with, 121
 - Intuit Inc., 93–95
 - Isolation
 - as ACID property of databases, 22
 - database locks facilitating, 115–116
 - defined, 107
 - understanding, 60

Issue identification, learning from mistakes, 98
 ITSM (Information Technology Service Management), 19

J

JIT (just-in-time) scalability, 6–8

K

Keep-alives, for performance/scale, 79
 Keeven, Tom, 19–20
 Key value stores, storage strategy, 48–49
 Kibana, ELK framework, 57
 King, Jeremy, 1–3, 4
 Klopper, Grant, 123–124

L

Lalonde, Chris, 44–47
 Last-Modified header, 82–85
 Latency
 CDNs solving issues with, 73
 keep-alives reducing, 79–80
 Law of the Instrument (Maslow's Hammer), 43, 107
 Learn aggressively rule, 95–100, 191
 Learn From Your Mistakes, rules
 Don't rely on QA to find mistakes, 100–102, 191
 Failing to design for roll back, 102–105
 Failing to design for rollback, 191–192
 Learn aggressively, 95–100, 191
 overview of, 93–95
 Learning culture, importance of, 95–100
 Legal requirements, don't check work for, 64
 Leverage content delivery networks rule, 75–77, 188
 Leverage page caches rule, 84–86, 189
 Levitt, Steven D., 168–169
 Load balancers
 competency in, 173
 in database replication, 24
 removing single point of failure, 132
 Locks, using right type of database, 114–116, 193
 Log files, actively using, 55–58
 Login, using verbs to split items, 25
 Logstash, ELK framework, 57
 Long-running processes, asynchronous calls to, 149
 Lookups, reducing DNS, 10–12

M

Maintain sessions in browser when possible rule, 142–143, 196
 Make use of distributed cache for states rule, 144–146, 196–197
 Make use of object caches rule, 88–90, 190
 Managers, quick start guide for, xvii
 Manual markdown command, wire-on/wire-off features, 136
 Map coloring problem, 68
 Market sensitivity to events, 159–160
 Market storm, 160
 Maslow, Abraham, 43
 Maslow's Hammer (Law of the Instrument), 43, 107
 Master-slave configuration
 in database replication, 23–24
 removing single point of failure, 131–132
 Mean time to failure (MTTF), failed writes, 61
 Memcached, 48, 89–90
 Message bus
 avoid overcrowding, 154–157, 198
 ensuring scalability of, 151–154, 197
 implementation of asynchronous communication, 148
 Meta tags, misconceptions about, 77–78
 Methods, asynchronous calls to overly complex, 149–150
 Minimum viable product, simplifying scope, 9
 Miscellaneous rules
 Be competent, 172–174, 200
 Be wary of scaling through third parties, 161–163, 198
 Design your application to be monitored, 169–172, 199–200
 overview of, 159–161
 Partition inductive, deductive, batch, and user interactive/OLTP workloads, 166–169, 199
 Purge, archive, and cost-justify storage, 163–166, 198–199
 Mod_alias module, 67
 Mod_expires module, 79
 Mod_rewrite module, 67
 MogileFS, 48
 Monetization, RFM analysis, 164
 MongoDB, 44–45, 49

- Monitoring
 - caches, 75
 - design application for, 169–172
 - events vs. abnormal changes in market, 160
 - log files, 57–58
 - object cache for hit ratio, 90
- Moore's Law, 33, 103
- MQ server, ZirMed system, 147–148
- MTTF (Mean time to failure), failed writes, 61
- Multiphase commit protocols, avoid, 116–118, 193
- Multiplicative effect, of items in series, 134–135
- Multitenancy, session/state destroying value of, 139–140

- N**
- NASDAQ, 159–160
- NCache, 89
- Network transit costs, scaling hosting solution, 39
- Neural nets, application monitoring, 172
- Never trust single points of failure rule, 130–132, 195
- Nodes, content delivery networks as, 76
- Non-normalized data models, 112
- Nonpersistent object caches, distributed session/state, 145
- Nonrelational databases, 45–46
- Normal Accident Theory, 97
- Normal forms
 - data model design, 112
 - most common, 113
 - relationship between data integrity and, 113–114
- Normalization, data model design, 112
- NoSQL database, 48–51
- Nouns (resources)
 - entities as, 112
 - scaling, 25–26
- NOWAIT keyword, Oracle database, 119
- Nuclear power generation, learning aggressively from, 97

- O**
- Object caches
 - making use of, 88–90, 190
 - nonpersistent, 145
 - putting on own tier, 90–91, 190
 - scaling by cloning, 23
- Object-relational mapping (ORM), 74
- ObjectRocket, 44–47
- Objects
 - Reduce DNS lookups rule and, 11–12
 - Reduce objects where possible rule, 12–14
- Online transaction processing (OLTP)
 - overview of, 20
 - partition workloads, 166–169, 199
 - relational structure between tables, 47
- Open-source solutions, simplify implementation with, 10
- Optimizer, ensure maximum concurrency, 115–116
- Organizations, learning culture, 95–96
- ORM (object-relational mapping), 74
- Overcrowding, avoid message bus, 154–157
- Overengineering
 - resist, 3–5, 180–181
 - simplify design, 9–10, 162
- Overuse, of tools, 44
- Ownership, component, 174

- P**
- Paas (platform as a service), 148, 154
- Page caches, leveraging, 84–85, 190
- Page database locks, 115
- Page weight, reduce, 14
- Parallel circuits, 132–135
- Pareto Principle (80–20 rule), 9, 87
- Partition inductive, deductive, batch, and user interaction (OLTP) workloads rule, 166–169
- Partition inductive, deductive, batch, and user interactive/OLTP workloads rule, 199
- Partitioning
 - by splitting similar things, 27–28
 - tolerance of CAP Theorem for, 69
- Pass on multiphase commits rule, 116–118, 193
- Payment Card Industry (PCI) compliance, firewalls for, 55
- PayPal
 - 24.0 incident, 60, 107–108
 - engineering and architecture, 59–61

- failure to design rollback in, 103–105
 - James Barrese at, 43
 - temporal constraints in, 70
 - PCI (Payment Card Industry) compliance, firewalls for, 55
 - Pending transactions, PayPal engineering, 59–61
 - Performance
 - database locks and statistics for, 116
 - reduce DNS lookups for, 10–12
 - reduce objects where possible for, 12–14
 - use of log files for, 56
 - Perimeter security devices, firewalls as, 53
 - Petrow, Charles, 97
 - Persistence, distributed session/state, 145
 - Persistence tier, 88–90, 144–145
 - Personally identifiable information (PII), firewalls for, 54–55
 - Peterson, Brad, 159–160, 170
 - Petopia.com, 1
 - PHP function, asynchronous communication, 148
 - Platform as a service (PaaS), 148, 154
 - Pod, 27–28, 125–126
 - Pool, 125–126
 - Post/Redirect/Get (PRG) pattern, redirect traffic, 65
 - Postmortem process, learning from mistakes, 98
 - PRG (Post/Redirect/Get) pattern, redirect traffic, 65
 - Primary keys, entity integrity, 112
 - Priority ranking, scalability rules, 200–202
 - Processes, asynchronous calls to long-running, 149
 - Product detail pages, content management service, 74
 - Protocols, use homogeneous networks, 15
 - Proxy headers, 78
 - Purge, archive, and cost-justify storage rule, 163–166, 198–199
 - Put object caches on their own tier rule, 90–91, 190
- Q**
- QA (Quality Assurance) personnel
 - don't rely on to find mistakes, 100–102, 191
 - important role of, 100–101
 - leverage cloud for, 42
 - when to hire, 101
 - Queries
 - don't use Select * in, 120–121
 - joining tables and optimizing, 113–114
 - Quick start guide, to this book, xvii–xviii
- R**
- Rackspace, 45–47
 - Rate of growth, choosing data storage, 51–52
 - RDBMSs (Relational database management systems)
 - nonrelational vs., 45–46
 - overuse of, 107
 - relax temporal constraints in, 68–70
 - scale by cloning, 22–24
 - use appropriately, 47–52
 - Read and write ratios, choosing data storage, 51–52
 - Read-only replicas of databases
 - across swim lanes, 129
 - removing single point of failure, 132
 - Real Application Clusters (RAC), 48
 - Recency of access, RFM analysis, 164
 - Redirects
 - reasons for use of, 65
 - stop traffic, 66–68, 188
 - Redis, 48
 - Reduce DNS lookups rule, 10–12, 182
 - Reduce objects where possible rule, 12–14, 182
 - Reduce the Equation, rules
 - Design scale into the solution, 6–8, 181
 - Don't overengineer the solution, 3–5, 180–181
 - overview of, 1–3
 - Reduce DNS lookups, 10–12, 182
 - Reduce objects where possible, 12–14, 182
 - Simplify solution three times over, 8–10, 181
 - Use homogeneous networks, 15, 182–183
 - Redundancy, don't check your work, 61
 - Refresh, redirects with HTML via, 66
 - Regulatory requirements, don't check work for, 64
 - Relational database management systems (RDBMSs). *See* RDBMSs (relational database management systems)

- Relationships
 - be aware of costly database, 111–114
 - choose data storage solution, 51–52
 - increase concurrency by changing entity, 116
 - Relax temporal constraints rule, 68–70, 188
 - Reliability, distributed session/state considerations, 145
 - Remote procedure calls (RPCs), application monitoring, 172
 - Remove business intelligence from transaction processing rule, 109–111, 192
 - Renting
 - data centers in various locations, 39
 - risk, by leveraging cloud, 42
 - Replication (split)
 - key-value stores and, 48–49
 - relational databases and synchronous, 47–48
 - scale out hosting solution, 36–40
 - scale using, 20, 22–24
 - session, 141
 - Requirements, overengineering, 3–5
 - Resilience, commitment to, 99
 - Resource Oriented Architecture (ROA), 25–26
 - Resources (nouns)
 - entities as, 112
 - scaling, 25–26
 - Reverse proxy cache, install page cache in front of, 84–85
 - RFM (recency, frequency, and monetization) analysis, of business value, 164–166
 - Risk
 - leverage cloud to rent, 42
 - QA for mitigation of, 102
 - use firewalls for significant reduction of, 52
 - ROA (Resource Oriented Architecture), 25–26
 - Roe v. Wade, 168
 - Rollback, failure to design for, 102–105, 191–192
 - Round trips, Ajax eliminating browser, 81–82
 - Rows
 - database locks for, 115
 - entities as, 112
 - RPCs(remote procedure calls), application monitoring, 172
 - Rule review and prioritization
 - 50 scalability rules in brief. *See* Scalability rules in brief
 - overview of, 177
 - risk-benefit model for evaluation, 177–180
 - Rules, resolve conflict in, 63–64
 - Runtime, feature wire-on/wire-off, 137
- S**
- SaaS (Software as a Service) applications
 - design application for monitoring, 170–172
 - learn from your mistakes, 93–95
 - relax temporal constraints, 68–70
 - Sampling
 - ensure your message bus can scale, 156
 - log, 56
 - Scalability
 - benefit/priority ranking of rules for, 200–202
 - fault isolation benefits, 127
 - firewall issues with, 52
 - of message bus, 152
 - Scalability Rules*, xix
 - Scalability rules in brief
 - Actively use log files, 187
 - Avoid overcrowding message bus, 198
 - Avoid putting systems in series, 195
 - Be aware of costly relationships, 192–193
 - Be competent, 200
 - Be wary of scaling through third parties, 198
 - Cache Ajax calls, 189
 - cloning or replicating things. *See* X axis (Design to Clone or Replicate things) rule
 - Communicate asynchronously, 197
 - Design scale into the solution, 181
 - Design solution to scale out, 184
 - Design to leverage cloud, 185–186
 - Design using fault-isolative swim lanes, 194
 - Design your application to be monitored, 199–200
 - Don't check your work, 187
 - Don't overengineer the solution, 180–181
 - Don't rely on QA to find mistakes, 100–102

- Don't select everything, 194
- Ensure message bus can scale, 197
- Ensure you can wire on/wire off features, 195
- Failing to design for roll back, 102–105
- Firewalls, firewalls, everywhere rule, 186–187
- Learn aggressively, 95–100
- Leverage content delivery networks, 188
- Leverage page caches, 189
- Maintain sessions in browser when possible rule, 196
- Make use of distributed cache for states rule, 196–197
- Make use of object caches, 190
- Never trust single points of failure, 195
- Partition inductive, deductive, batch, and user interactive/OLTP workloads, 199
- Pass on multiphase commits, 193
- Purge, archive, and cost-justify storage, 198–199
- Put object caches on their own tier, 190
- Reduce DNS lookups, 182
- Reduce objects where possible, 182
- Relax temporal constraints, 188
- Remove business intelligence from transaction processing, 192
- Scale out your hosting solution, 185
- Simplify solution three times over, 181
- splitting different things. *See* Y axis (Design to Split Different things) rule
- splitting similar things. *See* Z axis (Design to Split Similar things) rule
- Stop redirecting traffic rule, 188
- Strive for statelessness rule, 196
- Try not to use "Select for Update," 194
- Use commodity systems, 185
- Use databases appropriately, 186
- Use Expires headers, 189
- Use homogeneous networks, 182–183
- Use right type of database lock, 193
- Utilize application caches, 190
- Scale
 - along X axis, 27–28
 - Design into the solution, 6–8, 181
 - overengineering as enemy of, 5
- Scale out
 - defined, 31
 - design solution to, 31–33, 185
 - to leverage cloud, 40–42
 - overview of, 29–31
 - use commodity systems, 33–35
- Scale out your hosting solution rule, 35–40, 185
- Schremser, Chris, 29–31, 147–148
- Scope
 - monitor amount of data vs., 171–172
 - Simplify solution three times over, 9, 181
- SDNs (software-defined networks), 55
- Search engine rankings, traffic redirects affecting, 64–68
- Search, use verbs to split items, 25
- Second normal form, 113
- Security
 - manage with firewalls, 53–55
 - reduce risk via, 53
 - store session data in cookies and, 142–143
- Select *, do not use in queries, 120–121
- "Select for Update," try not to use, 118–119, 194
- Semantic changes of data, design for roll back, 105
- Sensitivity to operations, 99
- Series circuits, avoid putting systems in, 132–135, 194
- Servers
 - locate distributed cache away from working, 144–145
 - redirects via embedded modules, 66–67
- Service Oriented Architecture (SOA), 25–26
- ServiceNow, 19–21
- Services
 - asynchronous communication between applications and, 148
 - competency in delivery of, 173–174
 - make asynchronous calls to changing, 149–150
- Session cookies, 132, 143
- Session data, storing, 142–146
- Sessions
 - scale with replication technologies, 141
 - and state destroying value of multitenancy, 139–140
 - strive for statelessness, 140–142
- Setcookie, storing cookies in browsers, 143

- Sharding
 - definition of, 126
 - difficulty in relational databases, 48
 - with extensible record stores, 49
 - role of relationships between entities in, 113
 - scale by splitting similar things, 27–28
 - as solution for PayPal, 60
 - Share
 - across swim lanes with read replicas, 129
 - nothing between swim lanes, 127–129
 - Shared libraries, feature wire-on/wire-off with, 137
 - Signal intelligence, 160
 - Signup, using verbs to split items, 25
 - Simplification
 - is better, 162
 - reluctance to interpret failures using, 99
 - Simplify solution three times over rule, 8–10, 181
 - Simultaneous connection feature, browsers, 13–14
 - Simultaneous reconciliation, PayPal
 - constraints, 59–60
 - Single points of failure (SPOF), never trust, 130–132, 194
 - Singleton antipattern, 130
 - Singleton pattern, as single point of failure, 130
 - Sixth normal form, 113
 - SOA (Service Oriented Architecture), 25–26
 - Social construction, 95–100
 - Social contagion (viral growth), in learning cultures, 95–96
 - Software as a Service. *See* SaaS (Software as a Service) applications
 - Software-defined networks (SDNs), 55
 - Software developers, quick start guide, xvii
 - Software, overly complex, 5
 - Splits
 - design using fault-isolative swim lanes, 125–130
 - types of, 125–126
 - utilize application caches, 86–88
 - X axis (clone or replicate things), 26–28
 - Y axis (splitting different things), 24–26
 - Z axis (splitting similar things), 26–28
 - Splunk, 57
 - SPOF (single points of failure), never trust, 130–132, 194
 - SQL queries
 - design for roll back, 105
 - optimize, 114
 - Stand-in service, feature wire-on/wire-off, 136
 - Stansbury, Tayloe, 93–95
 - State
 - avoid or distribute rules for. *See* Avoid or Distribute State, rules
 - decisions about implementing, 140
 - distributed cache for, 143
 - Statelessness, strive for, 140–142, 196
 - Static content
 - add CDN for, 77
 - content delivery network, 76
 - use caches to scale, 73
 - Stop redirecting traffic rule, 64–68, 188
 - Storage solutions
 - alternatives to databases, 48
 - Apache Hadoop, 51
 - document stores, 49
 - extensible record stores, 49
 - Google's MapReduce, 50–51
 - key-value stores, 48–49
 - purge, archive, and cost-justify, 163–166
 - Stored procedures, 109–110
 - Strive for statelessness rule, 140–142, 196
 - Sudoku, 68
 - Swim lanes
 - creating fault isolation domains with, 126
 - design using fault-isolative, 124–130, 194
 - fault isolation example, 123–124
 - implement asynchronous calls, 149
 - isolate firewalls with, 55
 - scale by splitting similar things, 28
 - scale message bus, 154
 - Switches, homogeneous networks for, 15
 - Synchronous calls
 - asynchronous calls vs., 148
 - duplicate service/put in swim lane, 129
 - none between swim lanes, 128–129
- T**
- Table database locks, 115
 - Table-style DBMSs, 49

- Tables
 - competency in, 173
 - design for rollback, 105
 - entity sets as, 112
 - query optimization/joining, 113–114
 - using databases appropriately, 47–48
 - Tackley, Graham, 123–124
 - Teams, separate product/business intelligence, 111
 - Technical operations, quick start guide, xviii
 - Technology, learn aggressively from, 97
 - Temporal constraint satisfaction problem (TCSP), 68
 - Temporal constraints
 - failure of PayPal to relax, 60–61
 - make asynchronous calls for, 150
 - relax, 70, 188
 - Temporary demand, leverage cloud for, 41
 - The Guardian* newspaper, UK, 123–124
 - Third normal form, 113
 - Third parties
 - avoid state/replication servers within clustered, 144–145
 - be wary of scaling through, 161–163, 198
 - design to leverage cloud, 41
 - make asynchronous calls to, 149
 - simplify implementation with, 10
 - Thoroughbreds vs. goldfish concept, 35
 - Three Mile Island, 97
 - Three-phase commit (3PC), 117
 - Three-site configuration, scale out hosting solution, 39–40
 - Tier, object cache, 90–91
 - Time-outs
 - asynchronous calls across swim lanes, 128–129
 - feature wire-on/wire-off with, 136
 - Time sensitivity, scale by cloning and, 23
 - Time to market, fault isolation benefits, 127
 - Timeline, learn from mistakes, 98
 - Tipping Point* (Gladwell), 167
 - Tomcat, 56
 - Tools. *See* Use Right Tools, rules
 - Top SQL report, object caches, 89
 - Trade-offs
 - data integrity through normal forms, 113–114
 - firewalls, 55
 - flexibility vs. scalability, 49–50
 - store session data in browser, 142–143
 - understand, 46–47
 - Traffic
 - design application for monitoring, 170
 - stop redirects of, 64–68, 188
 - use content delivery networks to offload, 75–77
 - Transactions
 - ACID properties of reliable, 60
 - cost of logging, 57–58
 - design solution to scale out, 31–33
 - remove business intelligence from processing, 109–111, 192
 - scale by splitting different things, 24–26
 - session and state using longer-running, 141–142
 - Try not to use "Select for Update" rule, 118–119, 194
 - Two-phase commit (2PC), PayPal 24.0 disaster, 60
 - avoid multiphase commits, 117–118
 - database rules and, 107–108
- ## U
- URLs
 - avoid traffic redirects, 67
 - redirect traffic for shortened, 65
 - Use Caching Aggressively, rules
 - Cache Ajax calls, 80–84, 189
 - Leverage content delivery networks, 75–77, 188
 - Leverage page caches, 84–86, 189
 - Make use of object caches, 88–90, 190
 - overview of, 73–75
 - Put object caches on their own tier, 90–91, 190
 - Use Expires headers, 77–80, 188
 - Utilize application caches, 86–88, 190
 - Use commodity systems rule, 33–35, 185
 - Use databases appropriately rule, 47–48, 186
 - Use Expires headers rule, 77–80, 189
 - Use homogeneous networks rule, 15, 182–183
 - Use Right Tools, rules
 - Actively use log files, 55–58, 187
 - Firewalls everywhere! 52–55
 - Firewalls, firewalls everywhere! 186–187
 - overview of, 43–47
 - Use databases appropriately, 47–52, 186

Use right type of database lock rule, 114–116, 193

User interactive (OLTP) workloads, partition, 166–169, 199

User interfaces, Ajax for interactive, 81–84

Users, make something overly complex for, 4–5

Utilize application caches rule, 86–88, 190

V

V3 commerce platform, eBay, 1–3

Value

- data on message buses, 155–157
- data storage and, 164
- purge data low in, 165–166
- use RFM analysis to calculate business, 164–166

Vendors

- be vendor-neutral, 162
- be wary of scaling through third parties, 161–163, 198
- design to leverage cloud, 40–42
- homogeneous networks and, 15

Verbs (services)

- relationships as, 112
- scale by splitting different things, 25–26

Viral growth (social contagion), in learning cultures, 95–96

Virtualization

- design to leverage cloud, 40–42
- scale out using commodity systems, 35
- swim lanes and, 129–130

Voting phase, two-phase commit, 107

W

Wade v. Roe, 168

Walmart, xv

Warby Parker, 73

Web pages

- 1992 W3 Project as first Web page, 13
- reduce DNS lookups for, 10–12
- reduce objects where possible for, 12–14

Web servers

- avoid state systems requiring affinity, 144–145
- install page cache in front of, 84–85
- put object cache on its own tier vs., 90–91
- scale by cloning, 24

Websphere, 56

Weight, reduce objects where possible, 14

What question, application monitoring, 171–172

Why question

- application monitoring, 172
- learn from mistakes, 98–99

Wide column stores, alternative storage, 49

WiFi network, and Firesheep, 143

Wire on/wire off frameworks

- design for fault isolation, 128–129
- design for rollback, 105
- ensure for features, 135–138, 194

Workloads, partition inductive/deductive/batch/user interactive/OLTP, 166–169

World Trade Center bombing, market sensitivity to, 159

"Write once, read many" system

- change entity relationships to increase concurrency, 116
- scale by cloning, 23
- use file system for storage, 48

Writes

- don't immediately validate, 61–64
- store in cache if needed soon, 63

X

X axis (Design to Clone or Replicate things) rule

- brief summary of, 183
- distribute your work, 22–24
- overview of, 20–21
- purge, archive, and cost-justify storage, 164–165
- scale extensible record stores, 49
- as solution to single point of failure, 131
- store session data in browser, 142
- summary, 28

Y

Y axis (Design to Split Different things) rule

- brief summary of, 183–184
- database locks and, 116
- distribute your work, 24–26
- handle session and state needs, 141
- overview of, 20–21
- scale extensible record stores, 49

- scale message bus, 152–154
- scale out your hosting solution, 36
- summary, 28
- utilize application caches, 86–88

Z

- Z axis (Design to Split Similar things) rule
 - brief summary of, 183–184
 - database locks and, 116

- distribute your work, 26–28
- handle session and state needs, 141
- overview of, 20–21
- scale message bus, 153–154
- scale out your hosting solution, 36
- summary, 28
- utilize application caches, 86–88

- Zero Tolerance Program, 167–168
- ZirMed system, 29–31, 147–148