



Lean DevOps

A Practical Guide to
On Demand Service Delivery



Robert Benefield

FREE SAMPLE CHAPTER |



Lean DevOps

This page intentionally left blank

Lean DevOps

A Practical Guide to On Demand Service Delivery

Robert Benefield

◆◆ Addison-Wesley

Boston • Columbus • 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 author 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: 2022937306

Copyright © 2023 Pearson Education, Inc.

Cover image: Vinap/Shutterstock

All rights reserved. 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.pearson.com/permissions.

ISBN-13: 978-0-13-384750-5

ISBN-10: 0-13-384750-0

ScoutAutomatedPrintCode

Author: To my beautiful love, Gabrielle, my “little people” Aidan and Talia, and all of those who have enriched my life and, in the process, made the world a better place.

This page intentionally left blank

Contents

<i>Foreword</i>	<i>xv</i>
<i>Acknowledgments</i>	<i>xviii</i>
<i>About the Author</i>	<i>xx</i>
Introduction	1
Chapter 1: The Problem with IT Service Delivery	7
Approach #1: Reduce Delivery Friction	9
<i>The Downsides of Targeting Delivery Friction</i>	11
Approach #2: Managing Service Delivery Risk	12
<i>The Downsides of Targeting Service Delivery Risk</i>	14
<i>The Essence of Delivery</i>	15
Beginning the DevOps Journey	17
Summary	18
Chapter 2: How We Make Decisions	21
Examining the Decision-Making Process	22
Boyd and the Decision Process	23
The OODA Loop	26
The Ingredients of Decision Making	29
<i>Ingredient 1: The Target Outcome</i>	30
<i>Delivering Measures over Outcomes</i>	36
<i>Ingredient 2: Friction Elimination</i>	39
<i>Ingredient 3: Situational Awareness</i>	42
<i>The Challenge of Trust</i>	44
<i>The Fragility of Mental Models and Cognitive Biases</i>	45
<i>Ingredient 4: Learning</i>	48
<i>Failing to Learn</i>	48

The Pathway to Improved Decision Making	53
Summary	54
Chapter 3: Mission Command	55
The Origins of Mission Command	56
<i>Learning How to Lead Effectively the Hard Way</i>	57
Managing Through Unpredictability	58
<i>Knowledge and Awareness Weaknesses</i>	59
<i>Misalignments</i>	60
<i>Misjudgment of Ecosystem Complexity</i>	61
The Anatomy of Mission Command	62
Commander's Intent	63
<i>Brief</i>	66
<i>Situational Overview</i>	67
<i>Statement of the Desired Outcome or Overall</i>	
<i>Mission Objective</i>	67
<i>Execution Priorities</i>	67
<i>Anti-Goals and Constraints</i>	68
Backbriefing	69
Einheit: The Power of Mutual Trust	71
<i>Creating Einheit in DevOps</i>	74
<i>Continual Improvement</i>	75
<i>Staff Rides</i>	78
<i>After Action Reviews</i>	79
Organizational Impacts of Mission Command	80
Summary	81
Chapter 4: Friction	83
Understanding Ohno's Forms of Waste	84
<i>Muda (Pure Waste)</i>	86
<i>Muri (Overburden)</i>	109
<i>Mura (Fluctuation and Irregularity)</i>	113
See the Whole	125
Summary	126

Chapter 5: Risk	127
Cynefin and Decision Making	128
<i>Ordered Systems</i>	131
<i>Unordered Systems</i>	134
Reimagining Risk Management	143
<i>Have Clear and Understood Target Outcomes</i>	144
<i>Make the Best Choice the Easiest Choice</i>	145
<i>Continually Improve Ecosystem Observability</i>	147
Summary	151
Chapter 6: Situational Awareness	153
Making Sense of Our Ecosystem	154
The Mental Model	157
<i>The Problems with Mental Models</i>	158
Cognitive Bias	161
Gaining Better Situational Awareness	163
Framing	164
<i>Finding and Fixing Framing Problems</i>	165
Information Flow	169
<i>Why Ecosystem Dynamics Matter</i>	169
<i>Meeting Your Information Flow Needs</i>	172
Analysis and Improvement	181
Summary	182
Chapter 7: Learning	183
The Emergence of Skills Attainment Learning	184
<i>The Rise of the One Right Way</i>	186
Outcome-Directed Learning	188
Creating a Learning Culture	191
<i>Day-to-Day Kata</i>	191
<i>Improvement and Problem-Solving Kata</i>	192
<i>The Coaching Practice</i>	193
Summary	195

Chapter 8: Embarking on the DevOps Journey	197
The Service Delivery Challenge	204
<i>Traditional Delivery Fog in the Service World</i>	205
<i>The Challenge of the “ilities”</i>	207
The Path to Eliminating Service Delivery Fog	209
<i>The Role of Managers in Eliminating Service Delivery Fog</i>	210
<i>Identifying What You Can or Cannot Know</i>	214
<i>Ways the Team Can Eliminate Service Delivery Fog</i>	219
Summary	220
Chapter 9: Service Delivery Maturity and the Service Engineering Lead	221
Modeling Service Delivery Maturity	223
<i>The Example of Measuring Code Quality</i>	224
<i>Service Delivery Maturity Model Levels</i>	225
<i>Service Delivery Maturity Areas of Interest</i>	228
<i>Configuration Management and Delivery Hygiene</i>	232
<i>Supportability</i>	235
<i>Single Point of Failure Mitigation and Coupling Management</i>	239
<i>Engagement</i>	241
The Service Engineering Lead	243
<i>Why Have a Separate Rotating Role?</i>	244
<i>How the SE Lead Improves Awareness</i>	246
<i>Organizational Configurations with the SE Lead</i>	248
Challenges to Watch Out For	250
<i>Incentivizing Collaboration and Improvement</i>	251
<i>Developers Running Production Services</i>	253
<i>Overcoming the Operational Experience Gap</i>	254
Summary	256

Chapter 10: Automation	257
Tooling and Ecosystem Conditions	258
Building Sustainable Conditions	260
5S	261
<i>Seeing Automation 5S in Action</i>	278
Tools & Automation Engineering	283
<i>Organizational Details</i>	285
<i>Workflow and Sync Points</i>	285
Summary	287
Chapter 11: Instrumentation and Observability	289
Determining the “Right” Data	291
<i>Know the Purpose and Value</i>	293
<i>Know the Audience</i>	297
<i>Know the Source</i>	302
Making the Ecosystem Observable	307
<i>Instrumenting for Observability</i>	310
<i>Instrumenting Development</i>	310
<i>Instrumenting Packaging and Dependencies</i>	314
<i>Instrumenting Tooling</i>	316
<i>Instrumenting Environment Change and</i>	
<i>Configuration Management</i>	317
<i>Instrumenting Testing</i>	319
<i>Instrumenting Production</i>	320
<i>Queryable/Reportable Live Code and Services</i>	321
<i>Presenting Task, Change, Incident, and</i>	
<i>Problem Records Together</i>	321
<i>Environment Configuration</i>	322
<i>Logging</i>	323
<i>Monitoring</i>	324
<i>Security Tracking and Analysis</i>	325
<i>Service Data</i>	326

Pulling It All Together	327
<i>Instrumenting a Wastewater Ecosystem</i>	328
<i>Instrumenting an IT Ecosystem</i>	331
Summary	333
Chapter 12: Workflow	335
Workflow and Situational Awareness	336
Managing Work Through Process	337
Managing Work Organically	339
The Tyranny of Dark Matter	340
<i>Learning to See the Disconnects in Action</i>	343
<i>Resolving Disconnects by Building Context</i>	347
Visualizing the Flow	349
<i>Workflow Board Basics</i>	351
<i>State Columns</i>	352
<i>State Columns for Operations</i>	353
<i>Swim Lanes</i>	355
Task Cards	358
Preventing Dark Matter	359
Using the Board	362
Seeing the Problems	363
Limiting Work in Progress	365
The Limits of a Workflow Board	367
<i>Managing the Board</i>	367
<i>Managing Flow and Improvement</i>	368
Summary	368
Chapter 13: Queue Master	371
An Introduction to the Queue Master	372
<i>Role Mechanics</i>	374
<i>“Follow the Sun” Queue Mastering</i>	384
Queue Master Rollout Challenges	389

<i>Team Members Don't See the Value</i>	389
<i>More Traditionally Minded Managers</i>	
<i>Thwarting Rollout</i>	390
<i>Pushy Queue Masters</i>	391
<i>Junior Team Members as Queue Masters</i>	391
<i>Queue Masters Who Struggle to Lead Sync Points</i>	394
Summary	394
Chapter 14: Cycles and Sync Points	395
Inform, Align, Reflect, and Improve	396
<i>Top-Down Alignment Control Approach</i>	397
<i>Alignment Through Iterative Approaches</i>	397
Service Operations Synchronization and Improvement	400
<i>The Tactical Cycle</i>	400
Important Differences Between Kickoffs and	
Sprint Planning	404
<i>Daily Standup</i>	408
<i>Retrospective</i>	411
<i>General Meeting Structure</i>	413
The Learning and Improvement Discussion	415
<i>The Strategic Cycle</i>	421
Strategic Review	424
<i>General Review Structure</i>	426
<i>A3 Problem Solving for the Strategic Review</i>	427
Summary	432
Chapter 15: Governance	433
Factors for Successful Governance	434
<i>Meeting Intent</i>	435
<i>No Target Outcome Interference</i>	437
<i>Maintain Situational Awareness and Learning</i>	438
Common Governance Mistakes	440
<i>Poor Requirement Drafting and Understanding</i>	440

<i>Using Off-the-Shelf Governance Frameworks</i>	445
<i>Out-of-the-Box Process Tooling and Workflows</i>	450
Tips for Effective DevOps Governance	453
<i>Understand Governance Intent</i>	454
<i>Make It Visible</i>	454
<i>Propose Reasonable Solutions</i>	456
<i>Automation and Compliance</i>	458
<i>Be Flexible and Always Ready to Improve</i>	458
Summary	460
<i>Appendix</i>	461
<i>Index</i>	463

Foreword

Just about everyone can remember being inconvenienced by a computer outage at sometime during their lives. I particularly remember walking into the Melbourne airport on a sunny Saturday morning and discovering that the computer which ran all the check-in terminals was down. It stayed down all morning, exposing just how time-consuming and awkward the manual check-in process was. Some of us made it to our destinations many hours late; others were not so lucky. And I remember the hours-long check-in lines at the Baltimore airport many years later; the computer running Southwest's check-in kiosks was out and this time there was no manual backup process. After multiple long lines, we raced to our gate just in time to board before the door closed on a half-empty plane.

As I stood in those lines, I imagined the scene in the data center where “emergency responders” were no doubt scrambling to bring the system back up; after all, I had been there. Some years ago, I managed a factory data center, and we knew that any outage longer than a half hour would seriously curtail pack-out. It only happened once, and believe me, that was one time too many.

At last, we have a book written for the people on the emergency response side of a computer outage. It's not a book focused on being agile or lean; it's a book that focuses on being ready: ready to prevent serious service delivery problems, ready to limit the damage when they happen, and ready to uncover the cause of any incident and keep it from happening again. You would think, given our increasing reliance on technology, that there would be a lot of information about how to do these things well. But I haven't seen a lot of guidance for those who hold this challenging job, at least until now.

We have decades, if not centuries, of experience handling emergency situations in fire departments, hospitals, and military organizations. But rarely have we thought to apply that experience to the technical world, probably because we don't consider a technical failure to be a matter of life or death. But as Robert demonstrates throughout this book, lessons from other domains can be surprisingly relevant for those seeking to prevent computer

outages and, when they do (inevitably) occur, limit their damage and ensure a rapid, safe recovery.

Service delivery teams do a lot more than minimize outages and smoothly handle any that may arise; they provide a critical interface between their organization and its consumers. This book is based on a simple premise—the purpose of a technical system is to bring about expected outcomes for its consumers, and the purpose of service delivery teams is to ensure those expectations are met. They do this by understanding consumers’ intent when using the service, discovering what prevents consumers from achieving the outcomes they expect, and learning how to improve the system and close the gaps.

This book is not about meeting service delivery targets; it’s about paying attention: paying attention to consumers, being aware of their expectations, and being attentive to their frustrations. It’s about paying attention to the ways in which complex technical systems interact, how information flows through organizations, and how decisions are made and executed. It’s about paying attention to and improving the way work gets done and the outcomes that are delivered.

Be warned that focusing on delivering outcomes is a novel concept for IT teams, because historically, there was a large time and distance gap between technical teams and their consumers. Most “standard” IT practices, including Agile practices, presume that an intermediary translates consumer intent into proxy features and goals, which then guide the decisions of the technical team. This book eliminates such proxies, and in doing so, it may contradict some of your views on appropriate roles, responsibilities, and processes. Therefore, you might want to check your confirmation bias at the door as you enter, because, as the saying goes, “It ain’t what you don’t know that gets you into trouble. It’s what you know for sure that just ain’t so.”

This is an important book. It is well written and engaging, with great stories and easy-to-understand analogies. It simultaneously challenges readers and offers very practical advice. I highly recommend it.

Mary Poppendieck
January, 2022

Register your copy of *Lean DevOps* on the InformIT site for convenient access to updates and/or 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 (9780133847505) and click Submit. Look on the Registered Products tab for an Access Bonus Content link next to this product, and follow that link to access any available bonus materials. If you would like to be notified of exclusive offers on new editions and updates, please check the box to receive email from us.

Figure Credits

Cover image:

Vinap/Shutterstock

Figures 2.1, 2.2, 2.4–2.11,
3.1–3.9, 4.1–4.7, 4.9–4.17,
5.1–5.11, 6.1–6.10, 7.1–7.4,
8.1–8.8, 9.2–9.8, 10.1–10.6,
11.1–11.8, 11.10, 12.1–12.9,
13.1–13.8, 14.1–14.7,
15.1–15.4:

Courtesy of Gabrielle Benefield

Figure 2.3:

John R. Boyd

Acknowledgments

I am far from a natural at writing a book. Its linear format feels overly constraining to the way that I think. Books also feel like a permanent record, while life to me feels more like a constant quest to test ideas, gain knowledge, and correct mistakes to improve yourself and, if you are lucky, maybe make the world a little bit better along the way.

This made creating this book rather arduous. In my desire to help you, the reader, get the most coherent message, I spent an inordinate amount of time rewriting and shuffling content around. I probably would still be at it today if it weren't for the immense help and patience of several key people along the way.

The person most central to helping turn this book into reality, and the one who drew the great illustrations in this book, is my wife and partner-in-crime Gabrielle Benefield. With her product innovation background and my delivery experience, we constantly challenge each other to understand what so often goes wrong with teams and organizations in order to find the best way to help others sustainably deliver the outcomes that really matter. This book, along with the Mobius Loop innovation framework, have come out of years of deep discussion and collaboration to bring out the best in people in order to overcome the problems and dysfunction that so often thwart organizational success.

I also cannot thank Mary and Tom Poppendieck enough for their help and inspiration throughout this process. Both so generously spent the time challenging me to dig deep into my experiences to try to best convey to the world *why* the key pillars of this book are so critical to DevOps success.

I want to extend a special thanks to Colonel David A Hopley OBE RM (Retd). I learned and adopted many of the principles of Mission Command by working alongside many talented military veterans early on in my career. Despite my hazy awareness of their origins and understanding of the nuances that make them so effective, the approach not only resonated with

me but always seemed to work better than traditionally run teams. Being the formal commanding officer of the UK Special Forces (SBS) and as a member of the teaching staff at the Joint Services' Defence College, Colonel Hopley helped improve upon my lay description of Mission Command.

I also want to thank those who hung in through thick and thin reviewing the many iterations of the manuscript, some who also spent several hours helping me clarify my often-jumbled thoughts. This includes: Steve Freeman, James Duncan Davidson, David McNab, Jeff Sutherland, and Tim Beattie. I am also grateful for the patience of the editorial and production team at Addison-Wesley, including: Haze Humbert, Chris Guzikowski, and Menka Mehta. Your help turned this book into something that I am proud of.

Finally, there are a number of people who were instrumental as sounding boards to help clarify many of the ideas in this book, including (in no particular order): Peter Webb, Priit Kaasik, Heikko Ellermaa, Martin Sulg, Ainar Sepp, Sarah-Jane Mason, Chris Matts, Ademar Aguiar, and Henrik Kniberg. You helped me express *why* the approaches in this book not only work, but can help others on their own journey to learn, improve, and maybe make the world a little bit better for all of us.

About the Author

Robert Benefield is an experienced technical leader who has decades of experience delivering robust on-demand services to solve hard problems in demanding ecosystems including banking and securities trading, medical and pharmaceutical, energy, telecom, government, and Internet services. His continual eagerness to learn and work with others to make a difference has taken him from building computers and writing code in the early days of the Internet at Silicon Valley startups to the executive suite in large multinational companies. He shares his unique experience in the hopes that others can continue to build on it without having to collect quite as many scars along the way.

Introduction

Delivering on-demand services well is never easy. Your success hinges on having both the capability and capacity to deliver what your customer needs while doing so at high speed with the consistency, reliability, security, privacy, and cost effectiveness that they expect. This is just as true whether you are providing an IT service or a more traditional courier or electric utility service.

However, unlike more traditional services, IT service providers are far less restricted by organizational size or physical location. With so many quickly deployable tools and cloud capabilities available, even the smallest IT service providers can now instantly scale to address nearly any identified market need globally.

Where IT service delivery providers do struggle is predictably and reliably delivering services that match customer expectations. This, of course, matters. No one wants the frustration and disappointment of a service that falls short of what is needed. What makes this particularly frustrating is that such shortfalls are not caused by misunderstanding the market need or the functionality customers are looking for. If anything, IT is flooded by tools and techniques that allow businesses to analyze and validate ideas quickly. Instead, the problems arise from awareness gaps caused by the way organizations deliver and manage the services themselves.

As IT service stacks grow in complexity, it becomes far more difficult to determine, let alone ensure, that the dynamics between service components and the delivery ecosystem match what the customer expected. Rather than put measures in place to improve their awareness and understanding of these dynamics, delivery teams have focused on other factors like delivering more faster, using the latest cloud technologies and architectural approaches, or adopting the process or methodology most in fashion. Unfortunately, in the process the delivery teams unknowingly create further disconnects that fragment the information flow and context necessary to understand those dynamics.

As the resulting gap between what delivery teams *believe* they are providing and what is *actually* delivered grows, the team's ability to maintain sufficient context to make effective decisions steadily degrades. Even when disconnects are found, organizations often double down on more processes and misunderstood tooling that do little to effectively bridge the gaps. This creates a vicious loop that creates more frustration as the team drifts further away from being able to deliver to meet customer expectations.

Learning How to See

It is not inevitable that delivery teams have to fall into such a dysfunctional spiral. To break the cycle, you first need to understand the many ways you can lose your situational awareness, from deeply entrenched bad habits that fragment information flow to biases and perceptions that distort your understanding of a situation and what is important. Only then can you begin to put measures in place to counteract these tendencies and improve everyone's situational awareness.

Sharpening your situational awareness is like gaining a new sense or superpower you never knew you had. I like to think of it as learning how to see.

The primary objective of this book is to help you on that journey so that you and your organization can close the awareness gap and deliver services that your customers can use to reach their target outcomes. This book is geared primarily for two audiences. The first comprises the individual contributors, like software developers and IT Operations staff, who are in the trenches delivering the services. The other key audience comprises the managers and leaders who are responsible for building and directing those delivery teams.

For individual contributors the journey begins by looking at the delivery process itself. The first step is how you determine the objective of the work you are performing. Can it be used to check how well what is delivered aligns with the target outcomes of the customer, or are the measures more output-focused, such as the number of features or service uptime? Then there are the ways you acquire, understand, learn, and improve your ability to deliver. There are a number of misperceptions that inject flaws into our decision making, and ultimately the effectiveness of our actions.

To break this cycle, in this book you will find various techniques to help you measure and improve your situational awareness and the quality of information flow across your organization so that you are able to make better delivery decisions that move your services closer toward meeting your customer's target outcomes. Along the way many of the excuses people make for not changing their behavior and way of working, from managing work to governance procedures, will be debunked so that you and the team can continue to make progress.

Here, managers and delivery leadership will find strategies to help delivery teams spot and eliminate awareness gaps and misalignments that hinder effective delivery. This begins by identifying the various problems that arise from many of the management styles, requirements management techniques, processes, communication styles, and incentive structures that have traditionally been relied upon to direct and control people. These lead to poor decision making, conflict, reduced learning and improvement, and ultimately failure to deliver in a way that meets customer expectations. You will also learn about the power of Mission Command, as well as ways to communicate, inspire, and support the members of your team to effectively deliver to the organization's vision and the outcomes customers are trying to achieve.

Those who do not fit neatly in one of the two audiences described likely will also find value within these pages. For example, you might uncover and correct your own misperceptions about service delivery. This can help you better understand and more effectively interact with service delivery teams.

The thinking and techniques in this book are part of the larger Mobius outcome delivery approach, which you can find at <https://mobiusloop.com>. Mobius has been developed by a community dedicated to harnessing the power of innovation and delivery excellence to more effectively achieve the outcomes that matter.

How to Use This Book

This book can be divided into three parts. The first part introduces the key dynamics that underlie the service delivery challenge. It sets the scene for how those of us in IT service delivery are constantly in danger of focusing

far too much on removing delivery friction or reducing perceived delivery risk, often at the cost of maintaining situational awareness and ensuring teams have the ability to learn and improve.

Understanding these dynamics is important for any IT service delivery organization, and especially for those that wish to pursue the promise of DevOps. Overlooking them and missing their effects are what causes so many who pursue DevOps and Agile delivery approaches to fail to meet their promise from the start. This lack of awareness and appreciation of the way they can distort how we perceive the delivery ecosystem is also where many automation tooling and artificial intelligence/machine learning (AI/ML) approaches so often fail.

The second part of the book dives into each of the key elements and the role they play in service delivery. It explores their importance, how they are so often misapplied, and the repercussions to service delivery and the team. I personally feel that this is the most important part of the book, and the one that is so often missing from most guides out there.

The third and last part of the book is a practical guide to help you improve your own service delivery effectiveness. It includes ways to determine the maturity of your team to ensure you have the key elements in place to deliver consistently and effectively. It also has a number of suggestions for how to organize and manage the flow of work, build and deploy instrumentation and automation solutions, and deal with governance associated with internal controls and those required to meet legal and regulatory requirements.

My Own Journey

This book draws from my own journey working in the trenches as an individual contributor and, later, a technical leader to build great IT services and improve the effectiveness of the teams delivering them. I know from firsthand experience that every ecosystem has different challenges, and what I have learned from my own missteps along the way has kept me level-headed and practical. More than anything, I want this book to be a useful addition to your bookshelf for a long time. This is why the focus of this

book is to help you better understand your circumstances so that you and your organization can deliver more effectively, not to talk about some specific process or set of technology that will quickly be supplanted by the next big fad.

I have been blessed throughout my career to have met and worked alongside a number of people far smarter than I am who early on in my career exposed me to revolutionary concepts and ways of working. Some had worked alongside John Boyd's "Fighter Mafia." Others were Training Within Industry (TWI) veterans, or had to come up with ways to deliver highly reliable services long before the existence of concepts like cloud computing or continuous delivery. Only later did I realize that what I learned along the way is what has allowed me to quickly cut through delivery ecosystem noise to help teams overcome seemingly intractable problems. At times it has felt like a superpower, one that I hope I can share to help you reduce your own delivery pain and frustration to secure success.

This page intentionally left blank

Chapter 2

How We Make Decisions

Prepare for the unknown by studying how others in the past have coped with the unforeseeable and the unpredictable.

George S. Patton

Decisions act as our steering wheel that guides us through the pathways of life. This is true whether a decision is simple, like deciding whether to have a cup of coffee, or complex, such as choosing the best architectural design for a new critical service. Decisions are also how we guide our own actions when delivering IT services, whether it is in how we approach coding a new feature or troubleshooting a production problem.

However, being effective at decision making isn't an innate skill. In fact, it is one that is surprisingly difficult to learn. For starters, effectiveness is more than how fast you decide or how adeptly you execute that decision. It must also achieve its objective while accounting for any conditions that might change the dynamics and thereby the outcome trajectory of your actions in the executing ecosystem.

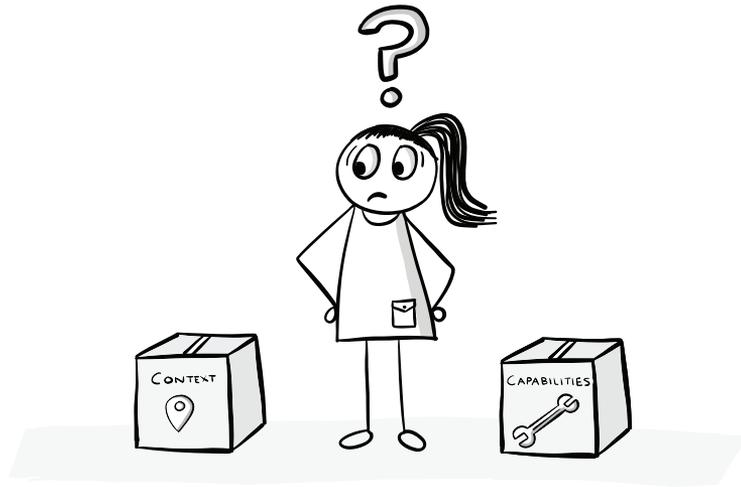
In this chapter, we take a deeper look at the decision-making process itself. We will also explore the ingredients necessary for effective decision making, how they impact the decision process, and how they can become impaired.

Examining the Decision-Making Process

Decision making is the process of pulling together any *information* and *context* about our situation and evaluating it against the *capabilities* available to progress toward the *desired outcome*. This is an iterative, rapid process. To work well we have to determine with each cycle whether the decision progressed us toward the outcome. If so, how effective was it, and if not, why not? We also have to look to see if anything unexpected occurred that can tell us more about the situation and the efficacy of our current capabilities that we can use to adjust and adapt.

Figure 2.1

Figuring out the right mix of context and capabilities for decision making can be challenging.



Even though we make decisions all the time, the process for making them can be surprisingly complex and fraught with mistakes. Consider the example of taking a friend to a coffee shop. While the task is inherently simple, there are all sorts of elements involved that, without the right level of scrutiny, can cause problems. You may find that you have the wrong address (*wrong or incomplete information*), that you took a wrong turn because you thought you were on a different street (*flawed situational context*), or that your car is having engine trouble and the coffee shop is too far to walk to

(mismatched capabilities). It is also possible that your capabilities, context, and information to get to the coffee shop are all fine, but your friend is angry because the shop has no Internet service and she only agreed to go there with you because she assumed she would be able to get online (*misunderstood target outcome*).

Spotting and rectifying mistakes under such simple conditions is easy. However, as the setting becomes far more complex, particularly as decisions take the form of large chains like those needed in IT service delivery, mistakes can far more easily hide under several layers of interactions, where they can remain undiscovered all while causing seemingly intractable problems. As these mistakes mount, they steadily degrade our understanding of our delivery ecosystem in ways that, unless found, undermine the overall effectiveness of future decisions.

The military strategist John Boyd became captivated by the importance of decision making while trying to understand what factors determined the likelihood of success in combat. He studied how simple mistakes could cascade and destroy any advantage a unit might have, and sought ways to improve decision-making processes in order to create a strategic advantage over the enemy. His work soon came to revolutionize how elite units, and many Western militaries, began to approach warfare.

Boyd and the Decision Process

Like many of us, John Boyd began his search for what factors increased the likelihood for success by looking at the tools (in this case weapons) of the victor. Boyd was an American fighter pilot who had served in the Korean War, where he flew the highly regarded F-86 fighter jet that dominated the skies against Soviet-designed MiG-15s. He knew firsthand there were differences in each aircraft model's capabilities. He theorized that there must be a way to quantitatively calculate an aircraft's performance so that it could be used to compare the relative performance of different types. If this were possible, one could then determine both the optimal design and the combat maneuvers that would be the most advantageous against the enemy.

His work led to the discovery of Energy-Maneuverability (EM) theory. It modeled an aircraft's specific energy to determine its relative performance against an enemy. The formula was revolutionary and is still used today in the design of military aircraft.

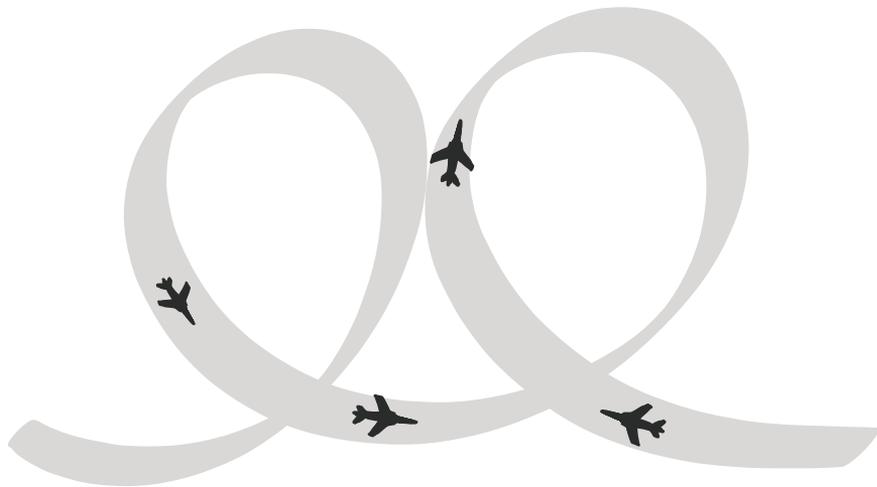
From there, Boyd looked to combine his theory with the optimal processes to fully exploit an aircraft's capabilities. He used his time at the Fighter Weapons School in Nevada to develop *Aerial Attack Study*, a book of aerial combat maneuvers first released inside the US military in 1961. It is considered so comprehensive that it is still used by combat pilots as the definitive source today.

Having both a means to create the best tools and processes to use them, most of us would figure that Boyd now possessed the formula for success in warfare. Despite all of this, Boyd was still troubled.

When he ran his own formula against some of the most successful weapons of World War II and the Korean War, he found many instances where the "successful" ones were far less capable than those of the enemy. Particularly disturbing to him, this included the highly regarded F-86.

Figure 2.2

Searching for the ingredients to air superiority was difficult.



As Boyd went back to earlier wars, he found that this was hardly unique. In fact, throughout history, having superior weaponry seemed to rarely be a relevant factor in determining the victor. As Boyd continued to research, he repeatedly found instances where numerically and technically superior

forces lost spectacularly to their poorly equipped opponents. That meant that despite his revolutionary work on EM theory, combat success couldn't be determined by any one formula or set of maneuvers.

Boyd studied great military tacticians from Sun Tzu and Alexander the Great to the Mongols, Clausewitz, and Helmut von Moltke. He also interviewed surviving officers of the most successful German Army units during World War II to understand what made them different. He soon realized that battlefield success hinged on which side could make the right decisions more quickly and accurately to reach a given objective. This was true even in cases where the victor possessed inferior weapons, fewer soldiers, poorer training, and battlefield terrain disadvantages. Not only that, but he noticed that this decision-making advantage could be gained just as well by either optimizing your own decision-making abilities or by thwarting those of your opponent.

This realization led Boyd to examine more deeply how the decision-making process works and what can make it more or less effective. In the process, he invented what is now known as the *OODA loop*.

Operation Millennium Challenge

A good recent demonstration of outmaneuvering superior enemy forces happened in the *Operation Millennium Challenge 2002* war game event held by the US Armed Forces in the Persian Gulf. On one side was the "Blue" side with the latest US weapons, while on the other was an unknown "Red" adversary armed with little more than light weaponry.

The Blue side had built elaborate plans to make the most of its superior firepower. What they had not counted on was the savvy leadership of the Red team under Lt. General Paul Van Riper. Van Riper knew there was no way he could win head-to-head against his more capable opponent. Instead, he used his ability to adapt rapidly to outmaneuver and overwhelm the Blue side. He used motorcycle messengers and World War II-style light signal communications to avoid Blue's sophisticated electronic surveillance. He then launched a surprise raid on the Blue fleet by using lightly armed speedboats to locate the fleet followed by a large salvo of missiles. This approach not only eliminated any advantages Blue had, but also used Blue's size and rigidity of command as friction against itself.

Despite Red's seemingly long odds, the Blue side's ability to react was subsequently overwhelmed, resulting in 16 ships being "sunk," including an

aircraft carrier, ten cruisers, and five of the six amphibious ships involved. In all, 20,000 Blue service personnel were “lost” by the surprise maneuver, greatly embarrassing the US military.

The OODA Loop

Boyd hypothesized that all intelligent creatures and organizations undergo decision loops continuously as they interact within their environment. Boyd described this as four interrelated and overlapping processes that are cycled through continuously and which he called the OODA loop, depicted in Figure 2.3. These processes include the following:

- **Observe:** The collection of all practically accessible current information. This can be anything from observed activity, unfolding conditions, known capabilities and weaknesses present, available intelligence data, and whatever else is at hand. While having lots of information can be useful, information quality is often more important. Even understanding what information you do not have can improve the efficacy of decisions.
- **Orient:** The analysis and synthesis of information to form a mental perspective. The best way to think of this is as the context needed for making your decision. Becoming oriented to a competitive situation means bringing to bear not only previous training and experiences, but also the cultural traditions and heritage of whoever is doing the orienting—a complex interaction that each person and organization handles differently. Together with observe, orient forms the foundation for situational awareness.
- **Decide:** Determining a course of action based upon one’s mental assessment of how likely the action is to move toward the desired outcome.
- **Act:** Following through on a decision. The results of the action, as well as how well these results adhere to the mental model of what was expected, can be used to adjust key aspects of our orientation toward both the problem and our understanding of the greater world. This is the core of the *learning* process.

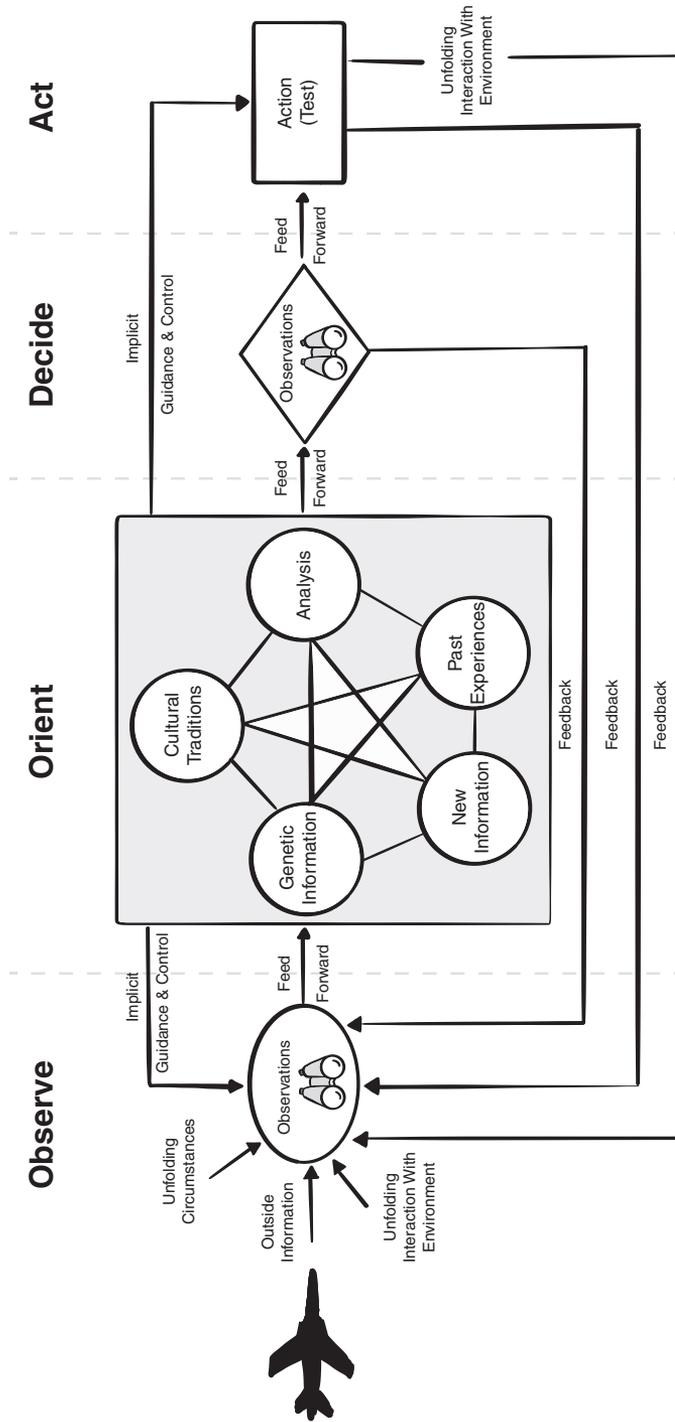


Figure 2.3
OODA loop.

Many who compare OODA to the popular PDCA cycle¹ by W. Edwards Demming miss the fact that OODA makes clear that the decision process is rarely a simple one-dimensional cycle that starts with observation and ends with action. The complex lines and arrows in Boyd's original diagram visualize the hundreds of possible loops through these simple steps in order to get to the desired outcome. As such, the most suitable path is not always the next in the list. Instead, it is the one that ensures there is sufficient alignment to the situation. That means that there will be many cases where steps are skipped, repeated, or even reversed before moving on. In some situations all steps occur simultaneously.

To help illustrate this nonlinear looping, let's take a very simplistic example of a typical process failure.

You get an alert that appears to be a failed production service (Observe->Orient).

You decide to investigate (Decide->Observe).

Before you can act, someone points out (Observe) that the alert came from a node that was taken out of production.

As you change your investigation (Orient) you then may go to see what might have been missed to cause the spurious alert (Observe).

Then you fix it (Decide->Act).

The action will likely involve changing the way people approach pulling nodes from production (Orient).

This solution may need to be checked and tuned over time to ensure it is effective without being too cumbersome (Observe, with further possible Decide->Act->Orient->Observe cycles).

What is important to recognize is that rapidly changing conditions and information discovered along the way can change the decision maker's alignment, leading to a new orientation. This can necessitate canceling or revising a plan of action, seeking out new information, or even throwing out decisions you may have *thought* you needed to make and replacing them with different and more appropriate ones. It may even require knowing when to throw away a once well-used practice in order to incorporate new learning.

1. Known as "Plan-Do-Check-Act," it is a control and continuous improvement cycle pioneered by Demming and Walter Shewart and now used heavily in Lean Manufacturing.

With the loop in hand and understanding how changing conditions can affect the way it is traversed, Boyd became interested in exploring the ways that one could not only out-decide their opponent, but also disrupt their decision process. Was there a way to overwhelm the enemy by changing the dynamics on the battlefield beyond the enemy's ability to decide effectively?

Boyd ultimately called this “getting inside” your opponent's decision cycle. Increasing the rate of change beyond the enemy's ability to adjust effectively can overwhelm their decision making enough to render them vulnerable to a nimbler opponent. He realized that the path to do this started with traversing the OODA decision loop to the target outcome faster than your opponent can.

With this knowledge, Boyd tried to identify the ingredients necessary to drive effective decision making.

The Ingredients of Decision Making

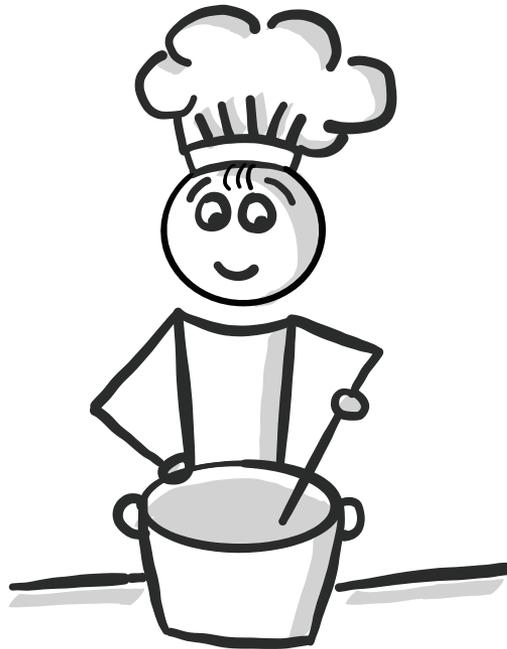


Figure 2.4

Success is having sufficient amounts of the right ingredients.

Just like any recipe, decisions are only as good as the quality of the ingredients on hand and the skill used putting them together. While the importance of any one ingredient can often differ from one decision to the next, they all play an important role in determining the efficacy of the decision-making process.

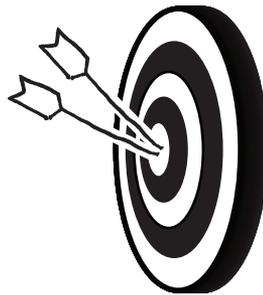
In order to understand more, let's take a look at each of these ingredients.

Ingredient 1: The Target Outcome

An effective decision is one that helps you in your pursuit of your *target outcome*. Target outcomes are the intended purpose of the decision. The better any target outcomes are understood by those making decisions, the better the person making the decision can choose the option that is most likely to lead to progress toward achieving those outcomes.

Figure 2.5

The target outcome.



Besides helping to highlight the better option in a decision, awareness of the target outcome also helps to determine decision efficacy by providing a means to measure any progress made toward it. This aids in learning and improvement, making it far easier to investigate cases where progress did not match expectations. This can help answer questions like:

- Was the decision executed in a timely and accurate way? If not, why not?
- Did the decision maker and others involved in executing it have sufficient situational awareness of current conditions to make an optimal match? If not, why not?

- Was the decision the best one to execute, or were there other, more suitable decisions that could have been made that might have had a more positive impact on outcome progress? If better alternatives existed, what can be learned to make it more likely to choose the more optimal decisions in the future?

The third point is the most important of these to note. While there are other ways to find awareness and execution improvement areas, it is far more difficult to determine whether the decision chosen to pursue, as well as the process for choosing it, could itself be improved without an adequately defined target outcome to measure it against. Remaining unaware that the target premise used to decide which decisions you need to make is flawed is incredibly corrosive to the efficacy of your decision-making abilities.

Despite their importance, target outcomes rarely get the attention they deserve. The problem begins with the fact that most decision frameworks spend far more time focusing on how quickly you execute a decision and not whether the decision is the one you need to make. This is especially true in IT, where far more value is often placed on team responsiveness and delivery speed. What information IT teams do get about outcome intent is usually a description of an output, such as a new feature, or performance target, like mean time to recover (MTTR).

There are also challenges around the communication of the outcomes themselves. Some people simply do not think delivery teams need to know anything about the desired outcome. More frequently, however, target outcomes are not clearly communicated because the requestor may not have a clear idea themselves of what they are. This “not knowing” your target outcome may sound strange but is surprisingly common.

To illustrate this, let’s go back to our coffee shop example.

When we say we want a cup of coffee, what exactly is our underlying intent? It probably isn’t to simply have one in our possession. In fact, most of us would be pretty upset if all we had were the disappointingly tepid remains of one.

Figure 2.6

Tepid coffee rarely sparks joy.



Typically, what we are looking for are the conditions that we associate with obtaining or consuming a cup of coffee. We might like the taste, the warmth, or the fact that it helps us feel more alert. It is possible that we don't actually care about the coffee but are more interested in the opportunity to relax or chat with friends or colleagues. Getting a coffee may be an excuse to find a convenient place to finalize a business opportunity. It might even be as simple as creating an opportunity to get Internet access, as it was with our friend earlier.

What sometimes makes achieving target outcomes tricky is that changing conditions can alter its intent or invalidate it altogether. For instance, stepping out of the office into a hot day can cause someone to change their order to an iced coffee. Suddenly finding yourself late for a meeting can throw out the idea of relaxing at the coffee shop, possibly eliminating the trip entirely. If we are with our Internet-obsessed friend, we might go to a different coffee shop where the coffee is less enjoyable and choose a smaller size or a different blend than we would normally have.

As a result, not only is all the hard work put in to enact the wrong decision a waste, but any measures used to gauge progress can at best be meaningless and at worst perpetuate the problem. Even if by chance you reach your desired outcome, the initial flaw means that your ability to learn from

your actions to improve and replicate the success has been undermined, opening you up for surprise and disappointment when it eventually fails.

The military faces very similar challenges. It is easy for a commander to order his troops to march up a hill or send a squadron to bomb a target. But if the commander's intent is not known, conditions might change and destroy any advantage the commander thought the action might have attained at best, or at worst lead to the destruction of the unit.

For this reason, Boyd and others found that it was better to turn the problem on its head, focusing instead on communicating to troops the target outcome rather than the actions they should use to try to achieve it. He became a strong proponent of Mission Command as a means for commanders and subordinates to communicate and build a true understanding of the target outcomes desired from a given mission. Mission Command, as discussed later in Chapter 3, "Mission Command," is an approach to communicate the intent and target outcomes desired that still gives those executing to achieve them the ability to adjust their actions as circumstances require them to do so.

This Mission Command method is useful in the service delivery world of DevOps. Not only does it enable teams to adjust to changing delivery conditions more quickly than more top-down approaches, it also enables teams to better discover and correct any flaws found in their understanding of the target outcomes themselves. Anyone who has worked in a delivery team knows that such flaws happen all the time, mostly due to the fact that few ever get to meet, let alone discuss, target outcomes with the actual customer. Most delivery teams instead have to rely upon proxies from the business or make intelligent guesses based upon their own experience and the data available around them.

It is in these proxies and guesses where things can go very wrong. The worst of these is when the focus is not on the actual outcome itself but on the solution itself, the method used to deliver the solution, or measures that have little connection to the target outcome.

To understand better, let's go through some of the common patterns of dysfunction to see why they occur and how they lead us astray.

Solution Bias

How many times have you been absolutely certain you knew exactly what was needed in order to solve a problem only to find out afterward that the solution was wrong? You are far from alone. It is human nature to approach a problem with a predetermined solution before truly understanding the problem itself. Sometimes it is the allure of a solution that misguides us, while at other times we fail to take the necessary time and effort to fully understand the expected outcome.

We commonly encounter solution bias when vendors push their products and services. Rather than spending any effort figuring out the problems you have and the outcomes you desire, vendors try to create artificial needs that are satisfied by their solution. Some do it by flashing fancy feature “bling.” Others list off recognized established companies or competitors who are their customers in the hopes that this will entice you to join the bandwagon. Some look to become your answer for your “strategy” in a particular area, be it “cloud,” “mobile,” “Agile,” “offshoring,” “DevOps,” or some other industry flavor of the month.

Solution providers are notorious for assuming their offering is the answer for everything. They are hardly the only one who suffers from solution bias, however. Even without the influence of slick marketing literature and sales techniques, customers are just as likely to fall into this trap. Whatever the cause, having a solution bias puts those responsible for delivery in an unenviable spot.

The School Bell

Figure 2.7

School bell.



Architect Alastair Parvin provided an illustrative example of this very problem in a story he gave at a TED talk in 2013.² A school had approached his architecture firm to redesign the school building. It was Victorian era with narrow hallways so cramped that it was difficult for students to get through between classes. The school administration had accepted that the construction was going to be expensive, but needed the problem to be solved.

At first glance the problem looked straightforward. The school wanted a new school building with spacious halls to accommodate the students moving between classes. Most of us would take that as a signal to start roughing out a design and a project plan for the school administration to look over.

The architects in Parvin's firm took a different approach. They started by looking more closely at the problem (crowded hallways) and the desired but only indirectly requested outcome (allowing students to move easily between classes). Anyone who has worked with flow problems knows that short, intense bursts can be a real mess. It is better to find ways to have lower and more even flows. This gave the architects an idea.

They came back to the school and recommended that instead of a new building they should redesign the bell system so that smaller school bells went off in different places at different times. This would allow for student traffic to be distributed more evenly, effectively eliminating congestion at a much lower cost. The solution worked brilliantly. By moving away from a pre-determined solution the school managed to save millions.

Execution Bias



Figure 2.8

Bill let it be known he was completely focused on the execution.

2. "Architecture for the People by the People," Alistair Parvin: Ted Talk, https://www.ted.com/talks/alastair_parvin_architecture_for_the_people_by_the_people/transcript?language=en

Requesting and delivering inappropriate solutions is not the only way we stray from the target outcome. We all carry any number of personal biases about *how* we, or others on our behalf, execute. There is nothing in itself wrong with having a favorite technology, programming language, process, or management approach. Our previous personal experience, the advice of people we trust, and the realities of the ecosystem we are working in are inevitably going to create preferences.

The problem occurs when the method of execution eclipses the actual target outcome. As Boyd discovered, knowing all the possible maneuvers is of little benefit if there is no target to aim for.

In technology, execution bias can take the form of a favored process framework, technology, vendor, or delivery approach regardless of their appropriateness. When execution bias occurs, there is so much focus on how a particular process, technology, or approach is executed that any target outcomes get overlooked. This is why there can be so many excellent by-the-book Agile, Prince2, and ITIL implementations out there that, despite their process excellence, still fail to deliver what the customer really needs.

Delivering Measures over Outcomes

Good metrics measuring the effectiveness in moving toward achieving a target outcome can be very useful. They can help spot problem areas that require addressing, guide course correcting, and ultimately help teams learn and improve. However, it is far too easy to place more focus on achieving a metric target than the outcomes it is supposed to be aiding. There are two rather common causes for this.

The first is when the target outcomes are poorly understood or, as in our coffee example, more qualitative in nature. It may take a number of attempts with the customer to provide enough clues to really understand what they are looking for. This can seem arduous, especially when multiple teams who have little to no direct interaction with the customer are delivering components that together provide the solution. In many cases teams might find it too hard or time consuming to even try. Instead, they opt to track more easy-to-measure metrics that either follow some standard set of industry best practices or at least might *seem* like a reasonable proxy to measure customer value.

The other cause is the more general and long-held practice by managers to create measures to evaluate the individual performance of teams and team members. In Frederick Wilson Taylor's book *The Principles of Scientific Management*, management's role was to tell their workers what to do and how to do it, using the measures of output and method compliance as the means to gauge, and even incentivize, productivity.

Both cases encourage the use of measures like number of outputs per period of time, found defects, mean time to recover, delivery within budget, and the like. All are easy to measure and have the added benefit of being localized to a given person or team. However, unless a customer is paying only for outputs, they rarely have anything but the most tenuous link to the actual desired outcome.

Encouraging more outputs may not sound like such a bad thing on its face. Who wouldn't want more work done in a given time, or higher uptime? The former nicely measures delivery friction reductions, while the latter provides a sense that risk is being well managed. However, the problem is two-fold. The first and most obvious is that, as we saw in our coffee story, few outcomes are met simply by possessing an output. This is especially true if an output is only a localized part of the delivery. Likewise, few customers would be happy to have a great database on the backend or a high uptime on a service that doesn't help them achieve what they need.

Another challenge comes from human behavior when they know they are being assessed on how well they meet a target. Most will often ruthlessly chase after the target, even if doing so means stripping out everything but the absolute minimum required to achieve it. This is extremely problematic, especially if there is only at best a loose connection between the measure and the desired outcome. I have personally observed teams that have sacrificed service quality and maintainability by reclassifying defects as feature requests to meet defect targets. The defects were still there, and often even worse than before, but as the count of what was being measured was dropping, the team was rewarded for their deceit.

The Objectives and Key Results (OKR) framework tries to overcome this by using a common set of clearly defined qualitative objectives that are intended to provide targeted measures. However, many find doing this well to be difficult. The best objectives are directly tied to the pursuit of the outcomes the customer wants, with key result measures that are both tied to those objectives and have sufficiently difficult targets (to encourage

innovation) that they are reachable only some of the time. In this way the targets form more of a direction to go in and a means to measure progress toward rather than some definitive targets that staff and teams feel they are being evaluated against. Unfortunately, this thinking goes against many of the bad assessment habits many of us grew up with in school. Instead, organizations tend to turn objectives into the same traditional output performance criteria with outputs becoming the key results to be quantitatively measured, thereby losing the value and original intent of OKRs.

Lost Bills

The dysfunction caused by output-based metrics is bad enough when everyone involved is in the same organization. It is even worse when such metrics form the basis of compensation between two different firms.

One large company was looking to streamline the operational support of various backend systems used to support the business. The old operational approach had led to many small teams that felt very protective of their space. Rather than fight with them, the CIO decided to wipe the problem away by outsourcing the entire mess.

The systems, much like the business, were stable and mature. The CIO decided that the most effective way to structure the outsourcing contract would be to pay for the number of incident tickets handled.

One of the outsourcing targets was the billing system. Even though billing was one of the most critical parts of the company, and one of the most interconnected, it was generally very stable. The billing process ran every night from 1 to 3 a.m. As it required locking customer records in both the CRM and billing databases, it was important that the process complete well before the Sales and Customer Support teams needed to access the system during the business day.

After a few months, the billing runs started to fail. The outsourced support investigated and noticed a read error of data in the cache containing customer records being processed. They cleared the cache, restarted the billing process, and everything continued seamlessly. The error began to happen more and more frequently. One of the support people decided to automate the recovery with a script. This eliminated the manual work, sped up recovery, and as it opened and closed incident tickets associated with the error, the outsourcer would reap income from the work. Everyone seemed to win.

Another few months went by. Customer Support started noticing that the number of complaints about incorrect bills had been growing significantly. The problems

were very strange, with customers being billed for services they had never had. Eventually, it got the attention of the executive staff and a team of engineers was sent in to investigate. What they found was startling.

A CRM system upgrade had turned on the ability for Sales and Customer Support staff to insert double byte and other special characters in the customer database. The billing batch processes could not handle these and would throw an exception when encountered, halting the billing process. Such failures are never good and indicate that there are likely awareness and quality issues in the delivery process.

But this particular problem was far worse than that.

The script that the outsourced support had written “solved” the offending error by dropping the customer identification database records that had been loaded before the special character was hit. The now orphaned billable assets would then be appended to subsequent customers as the process was restarted. This caused three types of errors:

Some customers were billed for products and services they never used.

Some customers were not billed for products and services. Some proportion of these customers also disappeared entirely from the customer and billing databases.

As the outsourced operational support staff were paid by the number of tickets they handled, they had little incentive to find the root cause of the problem, let alone fix it. Instead, they were measured and rewarded for efforts that ultimately ruined the integrity of the customer and billing databases and the overall billing process.

Ingredient 2: Friction Elimination



Figure 2.9

Friction elimination needs to be strategic, not random.

As mentioned in Chapter 1, friction is anything that increases the time and effort needed to reach a target outcome. It can arise at any point along the journey. This includes the decision cycle itself. We can be affected by it from the time and effort to gather and understand the information surrounding a situation through one or more of the various steps needed to make the decision, act upon it, and review its result. Friction can be so potent as to prevent you from reaching an outcome entirely.

Eliminating delivery friction, whether in the form of provisioning infrastructure, building and deploying code, or consuming new capabilities quickly, is what attracts people to DevOps and on-demand service delivery. There is also a lot of value in trying to mitigate many of the sources of wasteful friction whenever possible. But as Boyd found, and what so many of us miss, is that eliminating friction only provides value if it actually results in helping you reach your target outcome.

Many confuse friction elimination with increasing response and delivery speed rather than reaching the outcome. Teams dream of instantly spinning up more capacity and delivering tens or even hundreds of releases a day. Organizations will tout having hundreds of features available at a push of a button. We do all this believing, with little strong supporting evidence, that being fast and producing more will somehow get us to the outcome. It is like buying a fast race car with lots of fancy features and expecting that merely having it, regardless of the type of track or knowing how to drive effectively on it, is enough to win every race.

The gap between the lure of possessing a “potential ability” and effective outcome delivery is often most pronounced when there is friction in the decision cycle itself. “Red” side’s win over “Blue” in *Operation Millennium Challenge 2002* is a great demonstration of this. The “Blue” side’s superior weaponry was not sufficient to overcome its inferior decision loop agility against the “Red” side.

This same decision cycle friction was likely at work in Boyd’s analyses of Korean War era fighter aircraft. The superior abilities of enemy MiG fighter jets over the F-86 that Boyd found was often made irrelevant in the battlefield due to the relatively higher friction in communication flow and command structures of Communist forces compared to those of the US. This friction made it far more difficult for Communist forces to quickly understand and adjust to changing battlefield dynamics, as well as to catch

and learn from mistakes. American pilots did not have such problems and exploited these differences to their own advantage.

While the IT service delivery ecosystem does not face quite the same adversarial challenges, these same friction factors do have a significant impact on an organization's success. Teams can have an abundance of capabilities to build and deploy code quickly, yet still have so much decision-making friction that they are prevented from delivering effectively. Such friction can arise from such sources as defects and poorly understood target outcomes to poor technical operations skills. Interestingly, symptoms of decision friction do not necessarily manifest as problems of delivery agility. Often they take the form of misconfigurations, fragile and irreproducible “snowflake” configurations,³ costly rework, or ineffective troubleshooting and problem resolution. All of these not only impact the service quality that the customer experiences but also can consume more time and resources than former traditional delivery methods did.

Even when delivery agility is a problem, it is not always caused by poor delivery capabilities. I regularly encounter technical delivery teams with nearly flawless Agile practices, continuous integration (CI)/continuous delivery (CD) tooling, and low friction release processes that stumble because the process for deciding what to deliver is excessively slow. At one company, it typically took *17 months* for a six-week technical delivery project to traverse the approval and business prioritization process in order to get into the team's work queue.

Another frequent cause of delivery friction occurs when work that contains one or more key dependencies is split across different delivery teams. The best case in such a situation is that one team ends up waiting days or weeks for others to finish their piece before they can start. However, if the dependencies are deep or even cycle back and forth between teams, such poor planning can drag out for months. At one such company this problem was not only a regular occurrence but in places was *seven* dependencies deep. This meant that even when teams had two-week-long sprints, it would take a *minimum* of 14 weeks for affected functionality to make it through the delivery process. If a bug in an upstream library or module was encountered or if the critical functionality was deprioritized by an upstream team, this delay could (and often did) stretch for additional weeks or months.

3. <https://martinfowler.com/bliki/SnowflakeServer.html>

Delivery friction can also come from behaviors that are often perceived to improve performance. Long work weeks and aggressive velocity targets can burn out teams, increasing defect rates and reducing team efficiency. Multitasking is often a method used to keep team members fully utilized. But having large amounts of work in progress (WIP), along with the resulting unevenness and predictably unpredictable delivery of work, as well as the added cost of constant context switching, usually results in slower and lower-quality service delivery.

Friction in feedback and information flow can also reduce decision, and thereby delivery, effectiveness. I have encountered many companies that have fancy service instrumentation, data analytics tools, and elaborate reports and yet continually miss opportunities to exploit the information effectively because it takes them far too long to collect, process, and understand the data. One company was trying to use geolocation data in order to give customers location-specific offers like a coupon for 20 percent off a grocery item or an entrée at a restaurant. However, they soon found that it took three days to gather and process the information required, thus making it impossible to execute their plan for just-in-time offers.

These are just a small sampling of all the types of delivery friction. Others can be found in Chapter 4, “Friction,” which goes into considerably more depth. Understanding the various friction patterns that exist and their root causes can help you start your own journey to uncover and eliminate the sources of friction in your delivery ecosystem. The next decision-making ingredient, *situational awareness*, not only can aid in this search, but is the key ingredient for building enough context and understanding of your delivery ecosystem to make effective decisions.

Ingredient 3: Situational Awareness

Figure 2.10

Never overestimate your level of situational awareness.



Situational awareness is your ability to understand what is going on in your current ecosystem, and combine it with your existing knowledge and skills to make the most appropriate decisions to progress toward your desired outcome. For Boyd, all this gathering and combining takes place in the OODA loop's *Orient* step.

It makes sense that having information about your delivery ecosystem provides a level of insight into its dynamics to improve the speed and accuracy of the decisions you make in it. It is gaining this edge that has been the draw for everyone from investment banks and marketing companies to IT services organizations to invest in large-scale Big Data and business intelligence initiatives as well as advanced IT monitoring and service instrumentation tools.

But while the desire to collect as much information as possible seems logical, collecting and processing information that is unsuitable or has no clear purpose can actually harm your decision-making abilities. It can distract, misdirect, and slow down your decision making.

What makes information suitable depends heavily upon the right mix of the following factors:

- **Timeliness:** Relevant information needs to be both current and available in a consumable form at the time the decision is being made to have any positive impact. Information that arrives too slowly due to ecosystem friction, is now out of date, or is buried in extraneous data will reduce decision-making effectiveness. Collecting just enough of the right data is just as important.
- **Accuracy:** Accurate information means fewer misdirected or inaccurate decisions that need to be corrected to get back on the path to the target outcome.
- **Context:** Context is the framing of information in relation to the known dynamics of the ecosystem you are executing in. It is how we make sense of our current situation and is an important prerequisite for gaining situational awareness. We use context to not only understand the relevance of information to a given situation, but also to gauge the level of risk various options have based upon our grasp of the orderliness and predictability of the ecosystem. How all of this works, and how it often goes wrong, is covered in considerable depth in Chapter 5, “Risk.”

- **Knowledge:** The more you know about the relative appropriateness of the options and capabilities available to you, the greater the chance you will choose the option that will provide the most positive impact toward achieving the target outcome. How much knowledge you have and how effectively it can aid your decision making depends heavily upon the decision-making ingredient of learning effectiveness.

Assuming you are able to obtain suitable information, you then need to combine it effectively to gain sufficient situational awareness. There are many ways this process can fall apart, making it the most fragile component of the decision process. Boyd was intrigued by this, and spent the rest of his life trying to understand what conditions could damage or improve it.

Understanding these conditions is both important and discussed more fully in Chapter 6, “Situational Awareness.” For the purposes of this chapter, it is important to point out the two most common places where situational awareness can silently deteriorate, often with chronic if not catastrophic consequences. These are the challenge with trust, and the fragility of mental models and cognitive biases.

The Challenge of Trust

Most of us cannot acquire all the information directly ourselves and instead rely upon other people, either directly or from tools they build and run. If any of those people do not trust those who consume the information or what they will do with it, they may withhold, hide, or distort it. This can be done either directly, if they are handling the data themselves, or by suboptimizing the data collection and processing mechanisms they are asked to build and manage.

Lack-of-trust issues are surprisingly common. The vast majority of the time they arise not out of malice but out of fear of blame, micro-management, loss of control, or simply fear of the unknown. They are also just as likely to be caused by problems between peers as well as between managers and individual contributors in both directions. Many of these issues can develop out of anything from simple misunderstandings caused by insufficient contact or communication to larger organizational culture issues. Whatever the cause, the resulting damage to decision making is very real and can be far reaching.

The Fragility of Mental Models and Cognitive Biases

Human brains have evolved to rely upon two particular shortcuts to both speed up and reduce the amount of mental effort required to make decisions. These shortcuts, the mental model and cognitive bias, are generally quite useful when we are in stable, well-known environments. However, as you will see, these mechanisms are particularly susceptible to flaws that can severely degrade our decision-making abilities.

Mental models are sets of patterns that our brains build that are internal representations of an ecosystem and the relationships between its various parts to anticipate events and predict the probable behaviors of elements within it. Mental models allow us to quickly determine what actions are likely the most optimal to take, dramatically reducing the amount of information that needs to be gathered and analyzed.

Mental models are valuable, which is why people with experience are often quicker and more effective than a novice in situations that are similar to what they have been through previously. However, this mental model mechanism has several serious flaws. For one, a mental model is only as good as our understanding of the information we use to build and tune it. This information may be partial, misconstrued, or even factually incorrect. Boyd knew that this would inevitably lead to faulty assumptions that damage decision making.

The most obvious threat to mental model accuracy is having a rival try to intentionally seed inaccurate information and mistrust of factual information and knowledge. Studies have shown that disinformation, even when knowingly false or quickly debunked, can have a lingering effect that damages decision making.^{4,5} In recent years such behavior has moved from the battlefield and propagandist's toolbox to the political sphere to discredit

4. "Debunking: A Meta-Analysis of the Psychological Efficacy of Messages Countering Misinformation"; Chan, Man-pui Sally; Jones, Christopher R.; Jamieson, Kathleen Hall; Albarracín, Dolores. *Psychological Science*, September 2017. DOI: 10.1177/0956797617714579.
5. "Displacing Misinformation about Events: An Experimental Test of Causal Corrections"; Nyhan, Brendan; Reifler, Jason. *Journal of Experimental Political Science*, 2015. DOI: 10.1017/XPS.2014.22.

opponents through overly simplistic and factually incorrect descriptions of policies and positions.

But incorrect information, let alone active disinformation campaigns, is far from the most common problem mental models face. More often they become flawed through limited or significantly laggy information flow. Without sufficient feedback, we are unable to fill in important details or spot and correct errors in our assumptions. When an ecosystem is dynamic or complex, the number and size of erroneous assumptions can become so significant that the affected party becomes crippled.

Like mental models, cognitive biases are a form of mental shortcut. A cognitive bias trades precision for speed and a reduction of the cognitive load needed to make decisions. This “good enough” approach takes many forms. The following are just a handful of examples:

- **Representativeness biases** are used when making judgments about the probability of an event or subject of uncertainty by judging its similarity to a prototype in their mind. For instance, if you regularly receive a lot of frivolous monitoring alerts, you are likely to assume that the next set of monitoring alerts you get are also going to be frivolous and do not need to be investigated, despite the fact that they may be important. Similarly, a tall person is likely to be assumed to be good at playing basketball despite only sharing the characteristic of height with people who are good at the sport.
- **Anchoring biases** occur when someone relies upon the first piece of information learned when making a choice regardless of its relevancy. I have seen developers spend hours troubleshooting a problem based on the first alert or error message they see, only to realize much later that it was an extraneous effect of the underlying root cause. This led them on a wild goose chase that slowed down their ability to resolve the issue.
- **Availability biases** are the estimation of the probability of an event occurring based on how easily the person can imagine an event occurring. For instance, if they recently heard about a shark attack somewhere on the other side of the world, they will overestimate the chances they will be attacked by a shark at the beach.

- **Satisficing** is choosing the first option that satisfies the necessary decision criteria regardless of whether there are other better options available. An example is choosing a relational database to store unstructured data.

When they are relied upon heavily, it is extremely easy for cognitive biases to result in lots of bad decisions. In their worst form they can damage your overall decision-making ability.

Unfortunately, we in IT regularly fall under the sway of any number of biases, from hindsight bias (seeing past events as predictable and the correct actions obvious even when they were not at the time) and automation bias (assuming that automated aids are more likely to be correct even when contradicting information exists) to confirmation bias, the IKEA effect (placing a disproportionately high value on your own creation despite its actual level of quality), and loss aversion (reluctance to give up on an investment even when not doing so is more costly). But perhaps the most common bias is *optimism bias*.

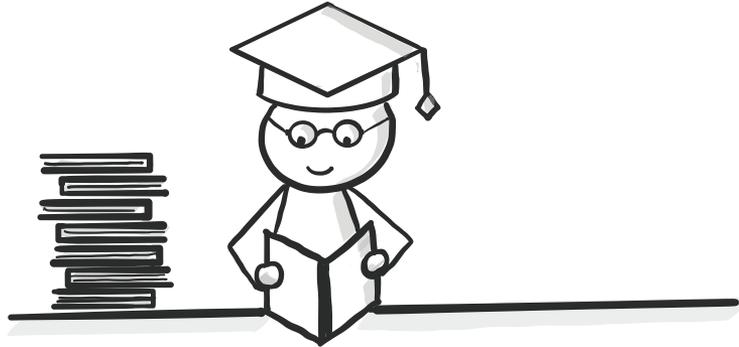
IT organizations have grown used to expecting that software and services will work and be used as we imagine they will be, regardless of any events or changes in the delivery ecosystem. A great example of this is the way that people view defects. Defects only exist in the minds of most engineers if they are demonstrated to be both real *and* having a direct negative impact, while risks are those situations that engineers view as likely events. This extreme optimism is so rampant that the entire discipline of Chaos Engineering has been developed to try to counteract it by actively seeding service environments with regular failure situations to convince engineers to design and deliver solutions that are resilient to them.

Another serious problem with mental models and cognitive biases is their stickiness in our minds. When they are fresh they are relatively malleable, though if you were ever taught an incorrect fact you likely know that it will still take repeated reinforcement to replace it with a more accurate one learned later. The longer that a mental model or bias remains unchallenged in our minds, the more embedded it becomes and the harder it is to change.

Ingredient 4: Learning

Figure 2.11

Learning isn't just school and certifications.



The final and frequently misunderstood piece of the decision-making process is *learning*. Learning is more than the sum of the facts you know or the degrees and certifications you hold. It is how well you can usefully capture and integrate key knowledge, skills, and feedback into improving your decision-making abilities.

Learning serves several purposes in the decision process. It is the way we acquire the knowledge and skills we need to improve our situational awareness. It is also the process we use to analyze the efficacy of our decisions and actions toward reaching our target outcomes so that we can improve them.

Learning is all about fully grasping the answer to the question *why*. This could be why a decision did or did not work as expected, where and why our understanding of the target outcome was or was not accurate, or how and why we had the level of situational awareness we did. It is the core of various frameworks, including Lean, and it is how we improve every part of the decision-making process.

The best part about learning is that it doesn't need to occur in a formal or structured way. In fact, our most useful learning happens spontaneously as we go through life. This should help us constantly improve the quality of our decisions—that is, if we remain willing and able to learn.

Failing to Learn

Learning is a lifelong pursuit, one we like to think we get far better at it with experience. However, it is actually really easy for us to get *worse* at it

over time. We even joke that “you can’t teach an old dog new tricks.” What causes us to get worse at doing something so fundamental?

A common belief is that any problems we may have are merely a result of having less dedicated time for learning. It is true that learning requires time to try things out, observe what happens, and reflect to really understand the reasons we got the results we did so that we can apply that knowledge to future decision making. However, the problem is far deeper.

For learning to occur we also need to be amenable to accepting new knowledge. That is pretty easy to do if you have no prior thoughts or expectations on a given topic. But as we gain experience, we start to collect various interpretations and expectations on a whole raft of things that we can use to build and tune our mental models. When they accurately reflect our situation, we can speed up and improve the accuracy of our decision making. But, unfortunately, not everything we pick up is accurate. Sometimes we learn the wrong lesson, or conditions change in important ways that invalidate our previous experiences.

As we discussed earlier, these inaccuracies can create serious flaws in our situational awareness. They also can damage our learning ability. When confronted with new knowledge that conflicts with our current view, we will sometimes simply ignore or overlook the new information because it does not meet our preset expectations. Other times we might actively contest the new evidence unless it can irrefutably disprove some interpretation we hold, even if the view we hold cannot itself stand up to such a test. This high bar creates friction that can slow or prevent us from learning.

The causes of this resistance have roots that go back to our formal schooling days.

“But That’s Just a PC”

There are a lot of scenarios where long-held habits and beliefs can get in the way of our ability to learn. Sometimes the outcome we are targeting becomes so ingrained that we fail to learn when our customers’ needs change. Other times we get so used to the idiosyncrasies of our environment that we fail to see important problems or risks in the state of our code and repository health, the suitability of our long-held processes, or challenges in skill and relationship dynamics in our teams. These can all

make life far more difficult than it needs to be and even put the business in jeopardy.

An energy company fell into such a predicament. Only after struggling to overcome their established beliefs did they discover that these very beliefs were endangering the core of its business.

The company had long relied upon a critical internally built application written in Fortran. It was designed to calculate how much electricity each power plant should generate across each service area, hour by hour each day over the entire 24-hour period. This information was critical for managing the health of the generation plants and the power grid, as well as the quality and availability of the electricity supply. Even more importantly, the generated results also had to be provided to government regulators. While generation and distribution had some buffer to allow things to coast along for a couple of hours with no data, failing the regulator's requirement meant millions in immediate fines.

Even though this application and the data it provided were so important to the business, few company employees were aware of its existence. For decades the application had run with little fanfare on a large mainframe in the basement of a building at one of the power stations. It only came to light when the building it was in was slated to be demolished.

In the process of trying to figure out what to do, management discovered that all of the original application developers had long since retired. The mainframe itself, while impressively large, was so ancient that the manufacturer had long gone out of business.

The situation seemed dire, but the business didn't seem to fully recognize the risk it had. The setup had been running so seamlessly for the most part that, at first glance, all that was required was to move the equipment. It didn't help that there was a strong belief that mainframes are by their nature failure-proof. That this one had been running since the 1970s only seemed to solidify that thinking.

The team I was leading at the time was tasked with coming up with a workable solution to the problem. What confidence others had quickly faded when we began to understand the complexities of the situation. We quickly swarmed the problem, with one part of the team looking into the problem of moving the existing system, while the rest of us looked for other alternatives.

Along the way we managed to dig up the source code and, while it was ancient enough to require some effort to port to a more contemporary

platform, doing so would be fairly straightforward. We were also fortunate enough that some of us had long ago become well versed in the intricacies of Fortran.

After a few weeks we managed to port the entire application to Linux. Eager to see if it would work, we ran it alongside the existing system. The hourly process took on average 50 minutes or so end-to-end, so we wanted to make sure the ported application would be not only just as accurate but could also fit in the same execution window.

That was when the real fun began.

At its first execution, the new setup executed the entire process in less than one second. After the main system completed its work we compared and found identical results. In disbelief, we continued running both in parallel for several weeks, each time coming up with the same results. Decades of technology innovation meant that what had once required a huge mainframe and nearly an hour of calculation time could now run on a tiny virtual machine instance in almost no time.

Convincing the business that the new setup was a far superior approach that was both cheaper and had the side effect of removing a massive business risk was extremely difficult. People had long grown to believe that mainframes are far superior in every way to a standard PC. Even though there was ample proof that doing so was safe and expedient, moving such an important application to an ordinary computer, let alone a virtual machine, seemed in many people's minds to be far more dangerous than the existing setup. They were simply too set in their ways to easily accept this new knowledge.

After yet more testing and a lot of pressure to decommission the mainframe, the management team finally gave in and migrated to the ported application.

The Shortcomings of the Formal Education System

Overcoming preexisting beliefs about a solution is not the only problem that can negatively affect our ability to learn. Preconceptions about learning and the learning process can also play a major factor.

Most of us spend a large chunk of our childhood being formally educated in school. While this is useful for kickstarting our life of learning, many of the techniques used can lead to unhealthy habits and expectations

that can impede our ability to learn later in life. The worst of these is creating the expectation that there is only one right way of working.

Having only one right method is easy to teach and test for in a standardized way. The problem is that life is not so black and white. Oftentimes, especially in IT, either there is no single right answer or the right answer no longer is right later down the line. However, the expectation of singular permanence makes people resistant to learning new and potentially better ways of working.

This problem is compounded by the fact that standardized teaching is done in a top-down way. It creates an expectation that there is an all-knowing authority that is there to judge our ability to follow instructions. Such an approach is not only unrealistic and noncollaborative, it also discourages critical thinking and innovation.

Hindsight Bias and Motivated Forgetting

Failing to learn is never good. But the problem is far more than simply missing out on knowing some information. It also can make it a great deal harder to improve our overall decision-making abilities. Ironically, the places where learning is often the most dysfunctional are those officially designed to encourage it. This includes everything from incident postmortems to sprint and project retrospectives, end project reports, and lessons learned reviews.

The reason for this goes back to some of the dysfunctions covered earlier. Being subconsciously predisposed to believe that the world operates in a particular way makes it incredibly easy to believe that events are more predictable than they are. This *hindsight bias* causes us to overlook the facts as they were known at the time and misremember events as they happened. This creates false narratives that can make any missteps appear to be due to a lack of skill or a total disregard for what seem to be obvious facts.

When this is coupled with a lack of trust and fear of failure, people can be motivated to avoid blame by actively forgetting or obfuscating useful details. This creates a culture that is not conducive to learning.

The Pathway to Improved Decision Making

We have discussed how the quality of decisions is dependent upon the strength of our understanding of the problems and desired outcome, the friction in our decision process, our level of situational awareness, and our ability to learn. We have also covered many of the ways that each of these key elements can degrade. How can we begin to improve our own decision-making abilities?

As you probably could have guessed by now, the journey is not as simple as following a set of predefined practices. We have to first dig deeper into how we go about making decisions and answer the following questions:

- Do you and your teammates understand your customer's target outcomes? Do you know why they are important. Do you have a good idea for how to measure that you are on track to reaching them?
- What factors in your ecosystem are causing rework, misalignments, and other issues that add friction that slows down and makes your decision-making and delivery processes inefficient?
- How does the information you and your teammates need to gain situational awareness flow across your ecosystem? Where is it at risk of degrading, and how do you know when it is happening and its potential impact on your decision making?
- How effectively does your organization learn and improve? Where are the weak points, and what effect are they having on your decision-making effectiveness?

Effective DevOps implementations not only can answer these questions with ease, but can tell you what is being done to fix any challenge the team faces. Each implementation journey is different, and there is no one answer for what is right.

This book takes the same approach.

If you think you have a good grasp of the service delivery problem space and are ready to dive straight into practical application, you can jump ahead to Section 3. In those chapters I have written a number of techniques and

approaches that I have used in numerous organizations large and small and found extremely helpful for everyone across the organization.

However, if you are like me, you may feel that you need to understand a bit more about the thinking behind this approach to feel more comfortable with it before making any serious commitment. This deep dive begins with Chapter 3, “Mission Command,” in the next section.

Summary

Even though we make decisions every day, becoming an effective decision-maker is a skill that takes a lot of time and effort to master. John Boyd’s OODA loop shows that the process is iterative and continual, with success dependent upon the following:

- Understanding the target outcome desired
- Gaining situational awareness through observing and understanding your capabilities and the dynamics of the ecosystem you are operating in
- Identifying the impacts of friction areas to gauge what decision options will most likely help you progress toward the outcome desired
- Following through with the action, and observing its impact so that you can determine how far you progressed, and if there are new details or changes that must be learned and incorporated into your situational awareness

Boyd showed that continually honing situational awareness and learning are at least as important as the speed of decision-making itself. Damage to any ingredient harms decision quality, and with it the ability to succeed.

Index

Numerics

- 5S, 261–262
 - for large Internet service companies, 281–283
 - seiketsu*, 273–277
 - seiso*, 270–273
 - seiton*, 266–270
 - shitsuke*, 277–278
 - sort (*'seiri'*), 262–266
 - for start-ups, 278–281

A

- A3 problem solving, 427–432
- accountability, governance and, 435–436
- accuracy, 43, 169–172
- action, 26
- after action reviews, 79–80
- Agile, 4, 10, 41, 77, 102, 447, 451
- agility, 10
- Alexander the Great, 57
- alignment
 - service operations synchronization and improvement, 400
 - strategic cycle, 421–423, 424–432
 - tactical cycle, 400–402
 - daily standup, 408–411
 - kickoff, 403–408
 - Queue Master brief, 402–403
 - retrospective, 411–421
 - through iterative approaches, 397
 - kanban*, 399
 - Scrum Sprint model, 398
 - top-down approach, 397
- all-in-one tools, 275–277
- ALM (Application Lifecycle Management), 277
- analysis paralysis, 134
- analytics, data collection and, 292
- anchoring bias, 46, 162

- Anderson, D., 351
- Andon, 350–351
- anti-goals, 68–69
- APIs, 121
- assessment. *See also* maturity model
 - review processes, 187–188
 - standardized testing, 185–187
- automation, 257–258, 261, 318. *See also* Tools & Automation Engineering
 - bias, 47, 162
 - build and CI system, 311
 - ecosystem conditions, 258–260
 - governance and, 458
 - implementing 5S principles
 - for large Internet service companies, 281–283
 - for start-ups, 278–281
 - tools, 268
- availability, 223–224
 - bias, 46
- awareness, 95–96, 99, 115, 214. *See also* situational awareness
 - development, 312–314
 - fog, 205–207, 222
 - eliminating, 209–214, 219–220
 - identifying what you can or cannot know, 214
 - operational, 206–207
 - SE Lead and, 246–248
- AWS, 124

B

- Babel, 11–12
- backbriefing, 63, 69–71, 194
- Battle of Kasserine Pass, 176
- best practices, 273–274, 451
- bias, 161–162
 - anchoring, 46
 - automation, 47

availability, 46
cognitive, 46–47
confirmation, 290
execution, 35–36
hindsight, 47, 51–52
representativeness, 46
solution, 34, 35
Big Data, 43, 289, 326
Blockbuster Entertainment, 7
Boyd, J., 4, 18, 23, 24–25, 28, 33, 36, 40–41, 74, 77, 83. *See also* OODA loop
Aerial Attack Study, 24
“Destruction and Creation,” 77
bugs, 271, 305, 314
bullwhip effect, 116–118
business continuity planning, 204, 236
Business Intelligence, 43, 69, 214, 289

C
CABs (Change Advisory Boards), 14–15
CEO, 68
change management, 317–319, 447–449
Chaos Engineering, 47, 79, 216–217
chaotic contexts, 137–139
dictators, 139–141
firefighter arsonists, 141–142
clear contexts, 131–133
cloud environments, 96–97, 124, 202
Clousewitz, C., 62
coaching practice, 193–195
code(ing), 10, 13, 41, 48–51, 92, 95–96, 112–113, 136, 233, 241, 251, 252, 268, 271, 319. *See also* software defects, 89
metrics, 311
Power Peg, 266
quality, measuring, 224–225
queryable, 321
cognitive bias, 46–47, 161–162. *See also* bias
collaboration, incentivizing, 251–252
color coding, 90
Commander’s Intent, 63–66, 211
communication, 95. *See also* information disconnects and, 348–349

implicit, 73
Queue Master role and, 385
competitive objectives, 96
complexity, 135–137, 351
environment, 128–129
kanban board and, 360–362
unpredictability and, 61
complicated contexts, 133–134
configuration management, 274, 277, 317–319
confirmation bias, 162, 290
confusion, 143
constraints, 68–69, 345
context, 43, 141–142, 261, 306
chaotic, 137–139
dictators, 139–141
firefighter arsonists, 141–142
clear, 131–133
complex, 135–137
complicated, 133–134
of confusion, 143
disconnects and, 347–349
target outcomes and, 144
continual improvement, 75–77, 84
coaching practice, 193–195
kata, 191
day-to-day, 191–192
improvement and problem solving, 192–193
outcome-directed learning, 188–190
continuous delivery, 4, 10, 41, 218–219
continuous integration, 233, 259, 280, 283, 312
contracts, 442–445
coupling management, measuring, 239–241
CRM, 39
C-suite, 116
culture, information flow and, 177–178
customer engagement, 242
cycles. *See* strategic cycle; tactical cycle
Cynefin, 129–130, 210. *See also* context

D
daily standup, 408–411
Dark Matter, 340–342. *See also* friction preventing, 359–362

- Queue Master role and, 379–380
- dark pools, 266
- dashboards, 312, 314
- data
 - collection, 291–292
 - determining the purpose and value of, 293–294
 - buying furniture online, 294–295
 - getting successfully treated at the hospital, 295–297
 - knowing the consumer, 297–299
 - capture and presentation distortions, 301
 - language mismatches, 299–301
 - monitoring, 324–325
 - presenting, 321–322
 - service, 326–327
 - sources, 302–303
 - timeliness, 305–306
 - trustworthiness and consistency, 303–305
- day-to-day *kata*, 191–192
- DBAs (database administrators), specialization, 100–101
- DBE (Database Engineering), 101
- decision-making, 21, 22–23, 144, 221–222, 396. *See also* Mission Command; situational awareness on the battlefield, 25
- cognitive bias and, 46–47
- contexts
 - chaotic, 137–142
 - clear, 131–133
 - complex, 135–137
 - complicated, 133–134
 - of confusion, 143
- execution bias and, 35–36
- friction elimination, 39–42
- improving, 53–54
- information
 - accuracy, 43
 - context, 43
 - dis-, 45–46
 - timeliness of, 43
- knowledge and, 44
- learning and, 48–51
- making the best choice the easiest choice, 145–147
- OODA loop, 25
 - action, 26
 - decide, 26
 - nonlinear looping, 28
 - observe, 26
 - orient, 26
 - PDCA cycle and, 28
- risk, 129
- situational awareness, 42–44
- solution bias and, 34, 35
- target outcome, 30–33
- target outcome and, 16
- trust and, 44
- unpredictability and, 58
- defects, 47, 88–91. *See also* failure(s)
 - causes of, 89
 - in code, 89
 - spotting, 90–91
- “Defend against the Madman” approach, 121–122
- delivery, 221–222
 - “as code,” 284
 - continuous, 4, 10, 41, 218–219
 - ecosystem, 197
 - fog, 205–207
 - eliminating, 209–214, 219–220
 - identifying what you can or cannot know, 214
- friction, 37, 57
 - dependencies and, 41
 - eliminating, 39–42
 - feedback and, 42
 - multitasking and, 42
 - reducing, 9–11
 - targeting, 11–12
- hygiene, measuring, 232–234
- ilities*, 207–209
- metrics, 36–39
- “mistake proof,” 90
- outcome-directed approach, 190
- pipeline, 238
- speed, 201
- teams, 15

- demand
 - bullwhip effect, 116–118
 - unexpected variability in, 114–116
 - Demming, E., 28
 - dependency(ies), 41, 357
 - failures and, 237
 - instrumenting, 314–316
 - management, 269–270
 - serial loopback, 102
 - team, 313
 - deployment, automation and, 260
 - development environment,
 - instrumenting, 310–314
 - DevOps, 4, 9, 10, 15, 18, 33, 40, 53, 76, 151, 168, 197, 199, 251
 - Einbeit*, 74–75
 - governance, 453
 - flexibility and, 458–460
 - intent, 454
 - proposals, 456–457
 - requirements compliance, 458
 - transparency, 454–456
 - production services and, 253
 - retrospectives, 192–193
 - Strategic Reviews, 192–193
 - uptime, 65
 - waste. *See also* waste
 - defects, 88–91
 - excessive processes, 106–109
 - handoffs, 101–103
 - overproduction, 91–93
 - overspecialization, 97–101
 - systemic impediments, 94–97
 - task switching, 105–106
 - WIP (“work in progress”), 103–104
 - dictators, 139–141
 - direct interference, 438–439
 - directive briefing, 66
 - anti-goals and constraints, 68–69
 - execution priorities, 67–68
 - situational overview, 67
 - statement of desired outcome, 67
 - disaster recovery, 236
 - disconnects, 343–346
 - communication, 348–349
 - resolving, 347–349
 - disinformation, 45–46
 - Docker, 97
 - Donaldson, L., 296
 - downtime, 123
- E**
- ecosystem, 183, 204, 221, 222
 - 5S, 261–262
 - seiketsu*, 273–277
 - seiso*, 270–273
 - seiton*, 266–270
 - shitsuke*, 277–278
 - sort (‘*seiri*’), 262–266
 - automation and, 258–260
 - complexity, 61
 - delivery, 197
 - development environment,
 - instrumenting, 310–314
 - dynamics, 169–172
 - hygiene, 232–234
 - IT, instrumenting, 331–333
 - logging, 323–324
 - monitoring, 324–325
 - observability, 147–151, 307–309
 - rebuilding, 323
 - situational awareness and, 154–156
 - transparency, 243–244
 - unknowns, 310
 - education, formal, shortcomings of, 51–52
 - effectiveness, 57–58, 198
 - Einbeit*, 71–75
 - EM (Energy-Maneuverability) theory, 24, 83
 - engagement, measuring, 241–243
 - entry management, Queue Master, 378–379
 - Equifax, 9
 - execution bias, 35–36
 - execution priorities, 67–68
 - experimentation, 135–137
 - exploits, 8
- F**
- failing to learn, 48–51
 - failure(s), 181, 211

- dependencies and, 237
- single point of, 239–241, 314
- feedback, 46, 135–136
 - friction, 42
 - timeliness of, 306
- Fingerspitzengefühl*, 74
- firefighter arsonists, 141–142
- flexibility, governance and, 458–460
- “fog of war,” 58. *See also* awareness
- “follow the sun” model, 384–385
 - other work when on duty, 386–389
 - sync point management, 385
- formal education, shortcomings of, 51–52
- Fortran, 48–51
- Fowler, M., 119
- frameworks. *See also* Mission Command
 - Cynefin, 129–130
 - governance, 445–449
 - maturity model, 222, 223–224
 - areas of interest, 228–229
 - configuration management and delivery hygiene, measuring, 232–234
 - engagement, measuring, 241–243
 - incentivizing collaboration and improvement, 251–252
 - information flow and instrumentation, measuring, 229–231
 - levels, 225–228
 - measuring code quality, 224–225
 - single point of failure mitigation and coupling management, measuring, 239–241
 - supportability, measuring, 235–239
 - OKR (Objectives and Key Results), 37
 - service management, 13–14
- framing, 164, 204
- finding and fixing problems in organizations
 - customer engagement, 165–167
 - locally focused metrics, 168–169

- output metrics, 167–168
- Fredendall, L., 176
- friction, 9, 14, 37, 57, 83, 84, 459–460.
 - See also* waste
 - configuration, 204
 - dependencies and, 41
 - elimination, 39–42
 - feedback and, 42
 - governance and, 437–438
 - multitasking and, 42
 - reducing, 9–11, 359–362
 - speed and, 40
 - targeting, 11–12
 - waste and, 85–86

G

- “gemba walk,” 307–309, 430–431
- Gilb, T., 208
- Gitlab.com, 12
- goals
 - anti-, 68–69
 - competitive objectives, 96
- governance, 13, 112
 - accountability and, 435–436
 - common mistakes, 440
 - out-of-the-box process tooling and workflows, 450–453
 - poor requirement drafting and understanding, 440–445
 - using off-the-shelf frameworks, 445–447
- DevOps, 453
- factors for successful, 434–435
 - maintaining situational awareness and learning, 438–440
 - meeting intent, 435–437
 - no target outcome interference, 437–438
- flexibility and, 458–460
- intent, 454
- process and, 447–449
- proposals, 456–457
- requirements compliance, 458
- transparency and, 454–456

H

handoffs, 101–103
hindsight bias, 47, 51–52, 162
Hölzle, U., 123
hooks, queryable, 321
human brain, 154–156. *See also* mental models, situational awareness and

I

IDE (integrated development environment), 90
IKEA effect, 47
ilities, 207–209, 217–218, 221–222, 223, 224
 known manageables, 214–215
 known unmanageables, 216–217
 unknowns, 217–219
implicit communication, 73
improvement and problem solving *kata*, 192–193
incident management, 179, 238, 338
indirect interference, 439–440
industry best practices, 13–14
information, 396. *See also* data
 accuracy, 43
 capturing, 290
 context, 43
 flow, 169, 198–199
 cultural disconnects, 177–178
 disconnects, 347–349
 ecosystem dynamics and, 169–172
 measuring, 229–231
 SE Lead and, 246–248
 tracers, 172–173
 transmission mismatches, 173–176
 trust and, 178–180
 visualizing, 349–351
 timeliness of, 43
innovation, 7, 8, 48–51
instrumentation, 291, 293–294, 295, 305, 306
 configuration management, 317–319
 dependency, 314–316

 development, 310–314
 interpreting, 300
 IT ecosystem, 331–333
 measuring, 229–231
 observability and, 310
 packaging, 314–316
 production, 320
 security, 325–326
 testing, 319–320
 tool, 316
 wastewater management, 328–331
intent, governance and, 434, 435–437, 454
interference
 direct, 438–439
 indirect, 439–440
 target outcome, 437–438
internal cloud, 97
intuition, 73, 74
Invisible Gorilla, The, 155
irregularity, 83
IT, 197
 cloud environments, friction and, 96–97
 ecosystem, instrumenting, 331–333
 failures, 8, 13
 friction, 86
 governance, 433
 information flow, 175
 mura (“irregularity”), 113–114
 overproduction, 92
 risk, 127
 service management, 337
 service providers, data resilience and recovery, 204
 snowflakes, 119–120, 121
 specialization and, 100–101
 task switching, 105–106
ITIL, 13, 452–453

J-K

JM (Job Methods), 84
JMX (Java Management Extensions), 321
kaizen, 272, 356
kanban, 10, 349–352, 399
 flow and, 368

- kaizen*, 356
 - limits of a workflow board, 367
 - managing the board, 367–368
 - simplicity and, 360–362
 - state columns, 352–355
 - swim lanes, 355–358
 - tally boards, 359–360
 - task cards, 358–359
 - timestamping tasks, 364
 - using the board, 362–363
 - WIP (work in progress), 365–367
 - work blockages, 364–365
 - kata*, 191, 272
 - day-to-day, 191–192
 - improvement and problem solving, 192–193
 - kickoff, 403–408
 - Klein, G., 158
 - Knight Capital Group, 9, 263–266, 301
 - knowledge. *See also* learning
 - decision-making and, 44
 - execution mismatches, 108
 - operational, 206–207
 - “single points of failure,” 120
 - skills and, 185
 - unpredictability and, 59–60
 - known
 - knowns, 131
 - manageables, 214–215
 - unknowns, 133
 - unmanageables, 216–217
 - Koçulu, A., 11–12
 - Korean War, 23, 24, 40–41
 - Kubernetes, 10, 97
- L**
- large Internet service companies,
 - implementing 5S principles, 281–283
 - leadership, 55
 - Commander’s Intent, 63–66
 - effective, 57–58
 - service delivery, 3
 - top-down approach, 58
 - Lean, 77, 80, 83, 125–126, 145, 147–148, 350
 - 5S, 261–262
 - seiketsu*, 273–277
 - seiso*, 270–273
 - seiton*, 266–270
 - shitsuke*, 277–278
 - sort (‘*seiri*’), 262–266
 - kanban*, 349–351
 - flow and, 368
 - kaizen*, 356
 - limits of a workflow board, 367
 - managing the board, 367–368
 - simplicity and, 360–362
 - state columns, 352–355
 - swim lanes, 355–358
 - tally boards, 359–360
 - task cards, 358–359
 - timestamping tasks, 364
 - using the board, 362–363
 - WIP (work in progress), 365–367
 - work blockages, 364–365
 - kata*, 272
 - day-to-day, 191–192
 - improvement and problem solving, 192–193
 - mura* (“irregularity”), 113–114.
 - See also* unpredictability
 - bullwhip effect, 116–118
 - minimizing, 120–122
 - snowflakes, 119–120
 - unexpected demand variability, 114–116
 - unmanaged variability, 118
 - Muri* (“overburden”), 109–110
 - overloading releases, 111–113
 - people and, 110–111
 - outcome-directed learning, 188–190
 - “stop the line,” 90
 - “walking the gemba,” 307–309
 - waste and, 84, 90
 - learning, 48–51, 183
 - continual improvement and, 75–77
 - effective leadership, 57–58
 - formal education, 51–52
 - hindsight bias and, 51–52
 - kata*, 191
 - day-to-day, 191–192
 - improvement and problem solving, 192–193

- outcome-directed, 188–190
- skills and, 184–185
- standardized tests and, 185–187
- Legal and Compliance teams, 300–301, 442–445, 454
- locally focused metrics, 168–169
- logging, 323–324

M

- MacArthur, D., 85
- managers, 210–214, 277–278, 371–372.
 - See also* Queue Master; Service Engineering Lead; teams
- managing
 - micro-, 278
 - risk, service delivery, 12–15
 - waste, overproduction, 91–93
- manufacturing
 - waste. *See also* waste
 - defects, 88–91
 - excessive processes, 106–109
 - handoffs, 101–103
 - overspecialization, 97–101
 - task switching, 105–106
 - waiting, 94–97
 - WIP (“work in progress”), 103–104
- maturity model, 222, 223–224
 - areas of interest, 228–229
 - incentivizing collaboration and improvement, 251–252
 - levels, 225–228
 - measuring code quality, 224–225
 - metrics
 - configuration management and delivery hygiene, 232–234
 - engagement, 241–243
 - information flow and instrumentation, 229–231
 - single point of failure mitigation and coupling management, 239–241
 - supportability, 235–239
- Maven, 270
- Meltdown, 8
- mental models, 45, 157–158
 - cognitive bias and, 161–162
 - data interpretation issues, 160
 - disinformation and, 45–46
 - problems with, 158–160
 - resilience, 160–161
- mentoring, 194–195, 211
- method-driven management, 340, 446–447
- metrics
 - availability, 223–224
 - code, 224–225, 311
 - configuration management and delivery hygiene, 232–234
 - delivery, 36–39
 - engagement, 241–243
 - information flow and instrumentation, 229–231
 - interpreting, 300
 - locally focused, 168–169
 - output, 167–168
 - single point of failure mitigation and coupling management, 239–241
 - for success, 294–295
 - supportability, 235–239
- Microsoft Azure, 13
- misalignment, 60–61, 396. *See also* alignment
- Mission Command, 3, 33, 55–56
 - after action reviews, 79–80
 - anatomy of, 62–63
 - backbriefing, 69–71
 - Commander’s Intent, 63–66, 211
 - continual improvement and, 75–77
 - directive briefing, 66
 - anti-goals and constraints, 68–69
 - execution priorities, 67–68
 - situational overview, 67
 - statement of desired outcome, 67
 - Einheit*, 71–74
 - organizational impacts of, 80–81
 - origins of, 56–57
 - staff rides, 78–79
 - target outcome, 62–63
 - unpredictability, 58–59
 - ecosystem complexity and, 61
 - misalignments and, 60–61
 - through knowledge and awareness weaknesses, 59–60
 - unpredictability and, 56–57

Mobius outcome delivery approach, 3
Moltke, H., 62, 66, 72, 77, 78, 80, 83
monitoring, 324–325
MTBF (Mean Time Between Failure),
123
Muda, 86–88, 115. *See also* waste
multitasking, 42
mura (“irregularity”), 113–114. *See also*
unpredictability
bullwhip effect, 116–118
minimizing, 120–122
snowflakes, 119–120
unexpected demand variability,
114–116
unmanaged variability, 118
Muri (“overburden”), 109–110, 115
overloading releases, 111–113
people and, 110–111

N

Napoleon, 57–58, 72
negative events, 75
Netflix, Simian Army, 122–124, 216–217
Node, 11–12
nonlinear looping, 28

O

observability, 295
ecosystem, 147–151, 307–309
instrumenting for, 310
observation, 26
office hours, Queue Master, 382–383
off-the-shelf tools, 303–304, 305
Ohno, T., 83, 85
OKR (Objectives and Key Results)
framework, 37
on-demand service delivery, 205
OODA loop, 25, 135
act, 26
decide, 26
nonlinear looping, 28
observe, 26
orient, 26, 43–44
PDCA cycle and, 28
open source software, 10

Operation Millennium Challenge,
25, 40
optimism bias, 162
ordered systems
clear contexts, 131–133
complicated contexts, 133–134
organizational impacts of Mission
Command, 80–81
organizations
finding and fixing framing problems
customer engagement, 165–167
locally focused metrics, 168–169
output metrics, 167–168
overburden and, 110–111
specialization and, 98–99, 100–101
orientation, 26
origins of Mission Command, 56–57
outages, 122, 272
outcome-directed learning, 188–190
output metrics, 167–168
overburden, 83, 109–110
overloading releases, 111–113
people and, 110–111
overproduction, 91–93
overspecialization, 97–101
ownership, building a sense of, 199

P

PaaS (platform-as-a-service), 10
partially done work, 103–104
Parvin, A., 35
PDCA cycle, 28
performance, 207, 208. *See also* *ilities*
PII (personally identifiable information),
438–439
“playing it safe,” 76
POMs (project object models), 270
Pony Test, 64–65
Power Peg, 266
pride, building a sense of, 199
process(es), 201. *See also* governance
excessive, 106–109
failure, 48–51
CRM, 39
OODA loop and, 28
governance and, 447–449
review, 187–188

- standardization, 275
- workflow and, 337–339

production, instrumenting, 320

productivity, 84

Puppet, 10, 97

pure waste, 86–88

Q

QA (quality assurance), 319

queryable hooks, 321

Queue Master, 75, 91, 104, 106, 147–148, 175, 181, 220, 372–373, 404–405

daily standup and, 408–411

“follow the sun” model, 384–385

- other work when on duty, 386–389
- sync point management, 385

retrospective and, 414, 415

role mechanics, 374

- Dark Matter handling, 379–380
- entry management, 376–377
- maintaining flow, 380–381
- office hours, 382–383
- pattern recognition, 381–382
- rotation, 374–376, 386
- sorting and dependency discovery, 378–379

rollout challenges, 389

- junior team members as Queue Masters, 391–393
- pushy Queue Masters, 391
- sync points, 394
- team members not seeing the value, 389–390
- traditional managers thwarting rollout, 390–391

tactical cycle and, 400–403

typical day of, 386–389

R

rebuilding the ecosystem, 323

reflection, 396

releases, overloading, 111–113

reliability, 207

representativeness bias, 46

resilience, 160–161, 204

retrospective, 192–193, 411–413

- general meeting structure, 413–415
- learning and improvement discussion, 415–421

review, 187–188

- kaizen*, 356
- strategic, 424–432

risk management, 127, 143–144

- decision-making, 129
- chaotic contexts, 137–142
- clear contexts, 131–133
- complex contexts, 135–137
- complicated contexts, 133–134
- confusion and, 143
- making the best choice the easiest choice, 145–147

ecosystem observability, improving, 147–151

ilities, 208

service delivery, 9

- downsides of targeting, 14–15
- managing, 12–14

rotation of duties, 142, 240

- Queue Master, 374–376, 386
- Service Engineering Lead, 244–246

S

SaaS (Software-as-a-Service), 200

satisficing, 47

Scharnhorst, D., 72

- On War*, 58

Scrum, 10, 77, 214, 314, 398, 404–408

SDK (software development kit), 445

security

- breaches

 - Equifax, 9
 - SolarWinds, 128

- tracking and analysis, 325–326

seiketsu, 273–277

seiso, 270–273

seiton, 266–270

“Selective Attention Test,” 155

separation of duties, 435–436, 439

serial loopback, 102

serverless architecture, 136

- service delivery, 3, 4, 15, 17, 204–205, 221
 - agility, 10
 - availability, 223–224
 - data, 326–327
 - on-demand, 205
 - disconnects, 343–346, 347–349
 - fog, 205–207
 - eliminating, 209–214, 219–220
 - identifying what you can or cannot know, 214
 - friction, 57, 83
 - dependencies and, 41
 - eliminating, 39–42
 - feedback and, 42
 - reducing, 9–11
 - waste and, 85–86
 - ilities*, 207–209
 - known manageables, 214–215
 - known unmanageables, 216–217
 - leadership, 3
 - maturity model, 223–224
 - areas of interest, 228–229
 - configuration management and delivery hygiene, measuring, 232–234
 - engagement, measuring, 241–243
 - incentivizing collaboration and improvement, 251–252
 - information flow and instrumentation, measuring, 229–231
 - levels, 225–228
 - measuring code quality, 224–225
 - single point of failure mitigation and coupling management, measuring, 239–241
 - supportability, measuring, 235–239
 - measuring, 17
 - metrics, 36–39
 - Mission Command, 33
 - outcome-directed approach, 190
 - risk
 - downsides of targeting, 14–15
 - managing, 12–14
 - situational awareness, 205–207
 - speed, 9–10, 201
 - systems thinking, 80
 - target outcomes, 16, 18, 293–294
 - unknowns, 217–219
 - uptime and, 65
- Service Engineering Lead, 243–244, 380
 - challenges, 250–251
 - on the delivery team, 250, 253
 - duty rotation and, 244–246
 - operational, 249–250
 - organizational configurations and, 248–250
 - overcoming the operational experience gap, 254–255
 - retrospective and, 414–415
 - tactical cycle and, 400–402
 - team awareness and, 246–248
- Shingo, S., 87
- shitsuke*, 277–278
- silos, 136–137
- Simian Army, 124, 216–217
- single point of failure, 120, 239–241
- situational awareness, 2–3, 4, 16, 42–44, 111, 114, 115–116, 120, 206. *See also kanban*
 - backbriefing and, 69–71
 - cognitive bias and, 161–162
 - Dark Matter, 340–342
 - disconnects, 343–346, 347–349
 - Fingerspitzengefühl*, 74
 - framing, 164, 165–169, 204
 - governance and, 438–440
 - human brain and, 154–156
 - information flow, 169, 206
 - cultural disconnects, 177–178
 - ecosystem dynamics and, 169–172
 - measuring, 229–231
 - tracers, 172–173
 - transmission mismatches, 173–176
 - trust and, 178–180
- mental models, 157–158
 - data interpretation issues, 160
 - problems with, 158–160
 - resilience, 160–161

- method-driven management and, 340
 - operational knowledge, 206–207
 - strategies for improving, 163, 181–182
 - unpredictability and, 59–60
 - workflow and, 336–337
 - Skillman, P., 183
 - skills, 184–185
 - SLAs (service-level agreements), 202, 203, 445
 - Snowden, D., 129, 131
 - snowflakes, 119–120, 121, 136
 - software
 - buggy, 271
 - customization, 279
 - development, overloading releases, 111–113
 - open source, 10
 - packaging, 234
 - SMARS, 266, 301
 - SolarWinds, Sunburst exploit, 128
 - solution bias, 34, 35
 - sort (*'seiri'*), 262–266
 - specialists, 98
 - specialization, 98–99, 100–101, 102–104
 - Spectre, 8
 - speed
 - friction and, 40
 - service delivery, 9–10, 201
 - Spotify, 99
 - sprints, 214, 404–408
 - staff rides, 77, 78–79
 - standardization, 273–275, 281
 - all-in-one tools, 275–277
 - test, 185–187
 - start-ups
 - development awareness, 312–314
 - implementing 5S principles, 278–281
 - state columns, 352–355
 - strategic cycle, 421–423
 - review, 192–193, 424–425
 - A3 problem solving, 427–432
 - structure, 426–427
 - success, 24, 294–295
 - Sun Tzu, 25
 - sunk cost fallacy, 162
 - supply chain friction, bullwhip effect, 116–118
 - supportability, measuring, 235–239
 - swim lanes, 355–358
 - sync points, 400
 - synthetic transaction monitoring, 304
 - systemic impediments, 94–97
 - systems thinking, 80–81
- ## T
- tactical cycle, 400–402
 - daily standup, 408–411
 - kickoff, 403–408
 - Queue Master brief, 402–403
 - retrospective, 411–413
 - general meeting structure, 413–415
 - learning and improvement discussion, 415–421
 - tally boards, 359–360
 - target outcomes, 16, 17, 18, 30–33, 62–63, 67, 214, 293–294
 - Commander's Intent and, 63–66
 - context and, 144
 - current state and, 188–189
 - governance and, 437–438
 - metrics, 36–39
 - Pony Test, 64–65
 - task(s), 371. *See also* *kanban*; workflow(s)
 - black holes, 346
 - cards, 358–359
 - Dark Matter, 340–342
 - identifiers, 311
 - structured approach, 183
 - switching, 105–106
 - team-based approach, 189
 - Taylor, F. W., *The Principles of Scientific Management*, 37
 - TCO (total cost of ownership), 146–147
 - TDD (test-driven development), 122
 - teaching to the test, 186–187
 - teams, 198, 199. *See also* Queue Master; Service Engineering Lead dependencies, 313

- eliminating fog, 219–220
- gauging maturity of, 238–239
- Legal and Compliance, 300–301
- operational experience gap, 254–255
- rotation of duties, 142, 240
- seiketsu* and, 274
- separation of duties, 435–436, 439
- Service Engineering Lead, 243–244, 253
 - awareness and, 246–248
 - duty rotation and, 244–246
 - organizational configurations and, 248–250
- telecommunications, legacy voice products, 92–93
- Terraform, 97
- testing, instrumenting, 319–320
- Three Mile Island, 159–160
- tiger teams, 194
- tight coupling, 96, 97, 102
- tools, 24, 43, 44, 90, 95–96, 203, 207, 217, 221–222, 250, 251, 274.
 - See also* Cynefin
 - A3 problem solving, 427–432
 - all-in-one, 275–277
 - automation, 258–260, 268
 - Dark Matter and, 342
 - instrumenting, 316
 - off-the-shelf, 303–304, 305
 - out-of-the-box, 450–453
 - problem-solving, 193
 - seiketsu* and, 274
 - Simian Army, 124
 - target outcome and, 16
- Tools & Automation Engineering, 283–284
 - organizational details, 285
 - retrospective and, 435
 - workflow and sync points, 285–287
- top-down
 - alignment, 397
 - leadership, 58
- Toyota, 74, 83, 85, 87, 125–126
- tracers, 172–173, 309
- trading systems, 272–273

- transparency, 121, 136, 207, 214, 243–244
 - governance and, 454–456
 - standardization and, 273–277
- trust, 52, 73
 - decision-making and, 44
 - information flow and, 178–180
- TWI (“Traning Within Industry”) program, 4, 84

U

- unknown unknowns, 135
- unordered systems, 134–135
 - chaotic contexts, 137–139
 - dictators, 139–141
 - firefighter arsonists, 141–142
 - complex contexts, 135–137
 - confusion, 143
 - observability and, 148
- unpredictability, 56–57
 - continual improvement and, 75–77
 - ecosystem complexity and, 61
 - “fog of war,” 58
 - managing through, 58–59
 - misalignments and, 60–61
 - through knowledge and awareness weaknesses, 59–60
 - workload, 335
- uptime, 65, 202, 207, 445

V

- variability, 123
 - demand, 114–116
 - minimizing, 120–122
 - standardization and, 273–277
 - unmanaged, 116–118
- visualizing, workflow, 349–351.
 - See also* *kanban*

W

- “walking the gemba,” 307–309, 430–431
- war-gaming, 78–79, 182
- waste, 83, 84–85, 106–109
 - defects, 88–91
 - causes of, 89

- in code, 89
 - spotting, 90–91
- friction and, 85–86
- handoffs, 101–103
- mura* (“irregularity”), 113–114
 - bullwhip effect, 116–118
 - minimizing, 120–122
 - snowflakes, 119–120
 - unexpected demand variability, 114–116
 - unmanaged variability, 118
- Muri* (“overburden”), 109–110
- overproduction, 91–93
- overspecialization, 97–101
- partially done work, 103–104
- process, 106–109
- pure, 86–88
- systemic impediments, 94–97
- task switching, 105–106
- waiting, 94–97
- wastewater management ecosystem, instrumenting, 328–331
- WIP (work in progress), 103–104, 365–367
- Woolworths Australia, 17
- work environment, 95
- workflow(s). *See also kanban*
 - board, 362–363
 - flow and improvement, managing, 368
 - limits of, 367
 - managing, 367–368
 - simplifying, 360–362
 - state columns, 352–355
 - swim lanes, 355–358
 - task cards, 358–359
 - timestamping tasks, 364
 - WIP (work in progress), 365–367
 - work blockages, 364–365
 - Dark Matter and, 340–342
 - disconnects, 343–346, 347–349
 - managing organically, 339–340
 - out-of-the-box, 450–453
 - patterns, 381–382
 - processes and, 337–339
 - Queue Master role and, 380–381
 - situational awareness and, 336–337
 - visualizing, 349–351
- World War II, 24, 79, 85, 176