



SECOND
EDITION

The
**DevOps
Handbook**

HOW TO CREATE WORLD-CLASS
AGILITY, RELIABILITY, & SECURITY
IN TECHNOLOGY ORGANIZATIONS

GENE KIM,
JEZ HUMBLE,
PATRICK DEBOIS,
& JOHN WILLIS

Featuring new foreword
and updated material by
NICOLE FORSGREN, PhD

Original foreword by
JOHN ALLSPAW



**SECOND
EDITION**

The
**DevOps
Handbook**

**HOW TO CREATE WORLD-CLASS
AGILITY, RELIABILITY, & SECURITY
IN TECHNOLOGY ORGANIZATIONS**

**GENE KIM,
JEZ HUMBLE,
PATRICK DEBOIS,
& JOHN WILLIS**

Featuring new foreword
and updated material by
NICOLE FORSGREN, PhD

Original foreword by
JOHN ALLSPAW

SECOND
EDITION

The
DevOps
Handbook

HOW TO CREATE WORLD-CLASS
AGILITY, RELIABILITY, & SECURITY
IN TECHNOLOGY ORGANIZATIONS

**SECOND
EDITION**

The
**DevOps
Handbook**

**HOW TO CREATE WORLD-CLASS
AGILITY, RELIABILITY, & SECURITY
IN TECHNOLOGY ORGANIZATIONS**

**GENE KIM, JEZ HUMBLE, PATRICK DEBOIS,
& JOHN WILLIS**

**FEATURING NEW FOREWORD AND UPDATED
MATERIALS FROM NICOLE FORSGREN, PHD**

ORIGINAL FOREWORD BY JOHN ALLSPAW

IT Revolution
Portland, Oregon



25 NW 23rd Pl, Suite 6314

Portland, OR 97210

All rights reserved. For information about permission to reproduce
selections from this book,

write to Permissions, IT Revolution Press, LLC,
25 NW 23rd Pl, Suite 6314, Portland, OR 97210

The DevOps Handbook, Second Edition © 2021 by Gene Kim,
Matthew “Jez” Humble,
Patrick Debois, John Willis, and Nicole Forsgren

First edition © 2016 by Gene Kim, Jez Humble, Patrick Debois, and
John Willis

Printed in the United States of America

27 26 25 24 23 22 21 1 2 3 4 5 6 7 8 9 10

Cover design by Devon Smith Creative

Cover illustration by eboy

Book design by Devon Smith Creative

ISBN: 9781950508402

eBook ISBN: 9781950508433

Web PDF ISBN: 9781942788867

Audio ISBN: 9781950508440

The author of the 18F case study has dedicated the work to the
public domain by

waiving all of his or her rights to the work worldwide under
copyright law, including all related

and neighboring rights, to the extent allowed by law. You can copy,
modify, distribute, and

perform case study 18F, even for commercial purposes, all without
asking permission.

For information about special discounts for bulk purchases

or for information on booking authors for an event,

please visit ITRevolution.com.

THE DEVOPS HANDBOOK, SECOND EDITION

CONTENTS

Figures & Tables

Note from the Publisher on the Second Edition

Foreword to the Second Edition: Nicole Forsgren

Foreword to the First Edition: John Allspaw

Preface

Introduction

Part I—The Three Ways

Part I Introduction

01 Agile, Continuous Delivery, and the Three Ways

NEW Case Study: Approaching Cruising
Altitude: American Airlines' DevOps
Journey (Part 1) (2020)

02 The First Way: The Principles of Flow

NEW Case Study: Flow and Constraint
Management in Healthcare (2021)

03 The Second Way: The Principles of Feedback

NEW Case Study: Pulling the Andon Cord at
Excella (2018)

04 The Third Way: The Principles of Continual
Learning and Experimentation

NEW Case Study: The Story of Bell Labs (1925)

Part 1 Conclusion

Part II—Where to Start

Part II Introduction

05 Selecting Which Value Stream to Start With

Case Study: Nordstrom's DevOps Transformation (2014–2015)

NEW Case Study: Kessel Run: The Brownfield Transformation of a Mid-Air Refueling System (2020)

NEW Case Study: Scaling DevOps Across the Business: American Airlines' DevOps Journey (Part 2) (2020)

NEW Case Study: Saving the Economy From Ruin (With a Hyperscale PaaS) at HMRC (2020)

06 Understanding the Work in Our Value Stream, Making it Visible, and Expanding it Across the Organization

Case Study: Nordstrom's Experience with Value Stream Mapping (2015)

Case Study: Operation InVersion at LinkedIn (2011)

07 How to Design Our Organization and Architecture with Conway's Law in Mind

Case Study: Conway's Law at Etsy (2015)

Case Study: API Enablement at Target (2015)

08 How to Get Great Outcomes by Integrating Operations into the Daily Work of Development

Case Study: Big Fish Games (2014)

NEW Case Study: Better Ways of Working at Nationwide Building Society (2020)

Part II Conclusion

Part III—The First Way: The Technical Practices of Flow

Part III Introduction

09 Create the Foundations of Our Deployment Pipeline

Case Study: Enterprise Data Warehouse (2009)

NEW Case Study: How a Hotel Company Ran \$30B of Revenue in Containers (2020)

10 Enable Fast and Reliable Automated Testing

Case Study: Google Web Server (2005)

11 Enable and Practice Continuous Integration

Case Study: HP LaserJet Firmware (2006)

Case Study: Continuous Integration of Bazaarvoice (2012)

12 Automate and Enable Low-Risk Releases

Case Study: Daily Deployments at CSG International (2013)

Case Study: Etsy—Self-Service Developer Deployment: An Example of Continuous Deployment (2014)

Case Study: Dixons Retail—Blue-Green Deployment for Point-of-Sale System (2008)

Case Study: Dark Launch of Facebook Chat (2008)

NEW Case Study: Creating a Win-Win for Dev & Ops at CSG (2016)

13 Architect for Low-Risk Releases

Case Study: Evolutionary Architecture at Amazon (2002)

Case Study: Strangler Fig Pattern at Blackboard Learn (2011)

Part III Conclusion

Part IV—The Second Way: The Technical Practices of Feedback

Part IV Introduction

14 Create Telemetry to Enable Seeing and Solving Problems

Case Study: DevOps Transformation at Etsy (2012)

Case Study: Creating Self-Service Metrics at LinkedIn (2011)

15 Analyze Telemetry to Better Anticipate Problems and Achieve Goals

Case Study: Telemetry at Netflix (2012)

Case Study: Auto-Scaling Capacity at Netflix (2012)

Case Study: Advanced Anomaly Detection (2014)

16 Enable Feedback So Development and Operations Can Safely Deploy Code

Case Study: Right Media (2006)

Case Study: The Launch and HandOff Readiness Review Google (2010)

17 Integrate Hypothesis-Driven Development and A/B Testing into Our Daily Work

Case Study: Hypothesis-Driven Development at Intuit, Inc. (2012)

Case Study: Doubling Revenue Growth through Fast Release Cycle Experimentation at Yahoo! Answers (2010)

18 Create Review and Coordination Processes to Increase Quality of Our Current Work

Case Study: Peer Review at GitHub (2011)

NEW Case Study: From Six-Eye Principle to Release at Scale at Adidas (2020)

Case Study: Code Reviews at Google (2010)

Case Study: Pair Programming Replacing Broken Code Review Processes at Pivotal Labs (2011)

Part IV Conclusion

Part V—The Third Way: The Technical Practices of Continual Learning and Experimentation

Part V Introduction

19 Enable and Inject Learning into Daily Work

Case Study: AWS US-East and Netflix (2011)

NEW Case Study: Turning an Outage into a Powerful Learning Opportunity at CSG (2020)

20 Convert Local Discoveries into Global Improvements

Case Study: Standardizing a New Technology Stack at Etsy (2010)

NEW Case Study: Crowdsourcing Technology Governance at Target (2018)

21 Reserve Time to Create Organizational Learning and Improvement

Case Study: Thirty-Day Challenge at Target (2015)

Case Study: Internal Technology Conferences at Nationwide Insurance, Capital One, and Target (2014)

Part V Conclusion

Part VI—The Technological Practices of Integrating Information Security, Change Management, and Compliance

Part VI Introduction

22 Information Security Is Everyone's Job Every Day

Case Study: Static Security Testing at Twitter (2009)

Case Study: 18F Automating Compliance for the Federal Government with Compliance Masonry (2016)

Case Study: Instrumenting the Environment at Etsy (2010)

NEW Case Study: Shifting Security Left at Fannie Mae (2020)

23 Protecting the Deployment Pipeline

Case Study: Automated Infrastructure Changes
as Standard Changes at Salesforce.com
(2012)

Case Study: PCI Compliance and a Cautionary
Tale of Separating Duties at Etsy (2014)

NEW Case Study: Biz and Tech Partnership
toward Ten “No Fear Releases” Per Day at
Capital One (2020)

Case Study: Proving Compliance in Regulated
Environments (2015)

Case Study: Relying on Production Telemetry
for ATM Systems (2013)

Part VI Conclusion

A Call to Action: Conclusion to *The DevOps
Handbook*

Afterword to the Second Edition

Appendices

Bibliography

Notes

Index

Acknowledgments

About the Authors

FIGURES & TABLES

Table 0.1: The Ever Accelerating Trend toward Faster, Cheaper, Lower Risk Delivery of Software

Figure 0.1: Deployments per Day vs. Number of Developers

Figure 1.1: Lead Time vs. Process Time of a Deployment Operation

Figure 1.2: A Technology Value Stream with a Deployment Lead Time of Three Months

Figure 1.3: A Technology Value Stream with a Lead Time of Minutes

Figure 1.4: The Three Ways

Figure 1.5: American Airlines' DevOps Transformation Journey

Figure 2.1: An Example Kanban Board Spanning Requirements, Dev, Test, Staging, and In Production

Figure 2.2: Simulation of "Envelope Game"

Figure 3.1: Feedback Cycle Times

Figure 3.2: Cycle Time vs. Andon Pulls at Excella

Table 4.1: The Westrum Organizational Typology Model

Figure 5.1: The Technology Adoption Curve

Figure 5.2: **American Airlines' Delivery Transformation**

Table 5.1: **American Airlines' New Vocabulary**

Figure 6.1: **An Example Value Stream Map**

Figure 6.2: **Invest 20% of Capacity in those Who Create Positive, User-Invisible Value**

Figure 7.1: **Functional vs. Market Orientation**

Table 7.1: **Specialists vs. Generalists vs. "E-shaped" Staff**

Figure 8.1: **Functional Teams in Silos vs. Long-Lived, Multi-Skilled Teams**

Figure **The Deployment Pipeline**
10.1:

Figure **The Ideal and Non-Ideal Automated Testing**
10.2: **Pyramids**

Figure **Running Automated and Manual Tests in**
10.3: **Parallel**

Figure **Number of Developers Deploying per Week at**
12.1: **Facebook**

Figure **Daily Deployments at CSG International**
12.2:

Figure **Elite and High Performers Achieve Faster**
12.3: **Deployment Lead Times and MTTR (2019)**

Figure **The Deployinator Console at Etsy**

12.4:

Figure **Blue-Green Deployment Patterns**

12.5:

Figure **The Canary Release Pattern**

12.6:

Figure **How Structure Influences Behavior and**

12.7: **Quality**

Figure **From Siloed Approach to Cross-Functional**

12.8: **Teams**

Figure **Conventional vs. Cross-Functional Structure**

12.9:

Figure **Google Cloud Datastore**

13.1:

Table 13.1: **Architectural Archetypes**

Figure **Blackboard Learn Code Repository: Before**

13.2: **Building Blocks**

Figure **Blackboard Learn Code Repository: After**

13.3: **Building Blocks**

Figure **Incident Resolution Time for Elite, High,**

14.1: **Medium, and Low Performers (2019)**

Figure **Monitoring Framework**

14.2:

Figure **One Line of Code to Generate Telemetry using**

14.3: StatsD and Graphite at Etsy

Figure User Excitement of New Features in User
14.4: Forum Posts after Deployments

Figure Standard Deviations (σ) Mean (μ) with
15.1: Gaussian Distribution

Figure Downloads per Minute: Over-Alerting when
15.2: Using “Three Standard Deviation” Rule

Figure Downloads per Minute: Histogram of Data
15.3: Showing Non-Gaussian Distribution

Figure Netflix Customer Viewing Demand for Five
15.4: Days

Figure Netflix Stryer Forecasting Customer Traffic
15.5: and the Resulting AWS Schedule of Computer
Resources

Figure Autodesk Share Price and Thirty-Day Moving
15.6: Average Filter

Figure Transaction Volume: Under-Alerting Using
15.7: “Three Standard Deviation” Rule

Figure Transaction Volume: Using Kolmogorov-
15.8: Smirnov Test to Alert on Anomalies

Figure Deployment to Etsy.com Causes PHP Run-
16.1: Time Warnings and Is Quickly Fixed

Figure The “Service Handback” at Google
16.2:

Figure **The LRR and HRR at Google**
16.3:

Figure **Comments and Suggestions on a GitHub Pull**
18.1: **Request**

Figure **Size of Change vs. Lead Time for Reviews at**
18.2: **Google**

Figure **The ASREDS Learning Loop**
21.1:

Figure **Jenkins Running Automated Security Testing**
22.1:

Figure **Number of Brakeman Security Vulnerabilities**
22.2: **Detected**

Figure **Time to Remediate vs. Time to Update**
22.3: **Dependencies (TTU)**

Figure **Five Behavioral Clusters for Open-Source**
22.4: **Projects**

Figure **Developers See SQL Injection Attempts in**
22.5: **Graphite at Etsy**

Figure **Average Development Window by Day of Week**
AF.1: **per User**

Figure A.1: **The Core, Chronic Conflict Facing Every IT
Organization**

Table A.1: **The Downward Spiral**

Figure A.2: Queue Size and Wait Times as Function of Percent Utilization

Table A.2: Two Stories

Figure A.3: The Toyota Andon Cord

NOTE FROM THE PUBLISHER ON THE SECOND EDITION

Impact of the First Edition

Since the original publication of *The DevOps Handbook*, data from the *State of DevOps Reports* and other research continue to show that DevOps improves time to value for businesses and increases productivity and worker well-being. It also helps create nimble, agile businesses that can adjust to overwhelming change, as witnessed in the COVID-19 pandemic of 2020 and beyond.

“I think 2020 has been illuminating in showing what technology can do in a time of incredible crisis,” Gene Kim said in his “State of DevOps: 2020 and Beyond” article. “The crisis provided a catalyst for rapid change. And I’m thankful we were able to rise and meet it.”¹

One of the underpinnings of DevOps and *The DevOps Handbook* is that it shows—and is indeed written for—the horses not the unicorns of the business and technology world. DevOps was never, and still is not, only effective at technology giants—the FAANGs—or startups. This book and the DevOps community as a whole have shown time and time again that DevOps practices and processes can take even the most legacy-riddled, old “horse” enterprise organization and turn it into a nimble technology organization.

In 2021, it is clearer than ever before that every business is a technology business and every leader is a technology leader. Not only can technology no longer be ignored or relegated to the basements; it must also be considered a vital part of the entire strategic endeavor of the business.

Changes to the Second Edition

In this expanded edition of *The DevOps Handbook*, the authors have updated the main text where new research, learnings, and experiences have developed and shaped our understanding of DevOps and how it is used in the industry. Additionally we are pleased to include renowned researcher Dr. Nicole Forsgren as co-author to help update and expand the text with new research and supporting metrics.

CONTINUOUS LEARNING

We've added some additional insights and resources we've learned since the first edition came out. These "Continuous Learning" sections are highlighted throughout the book as you see here and include new supporting data and additional resources, tools, and techniques to use on your DevOps journey.

We've also expanded the book with additional case studies to illustrate how far DevOps has spread throughout all industries, especially how it has spread beyond the IT department and into the C-suite itself. In addition, at the end of each case study we have added a key takeaway or two that highlight the most important, though not exclusive, lessons learned. Finally, we've updated the conclusion to each part with new resources to continue your learning journey.

What's Next for DevOps and the Age of Software

If the past five years have taught us anything, it is how important technology is and how much we can achieve when IT and the business speak openly and honestly, as DevOps facilitates.

Perhaps nothing illustrates this more than the rapid changes that were necessary due to the COVID-19 pandemic

of 2020 and beyond. Through the use of DevOps, organizations mobilized technology to get services to customers, internal and external, in a moment of rapid, unprecedented change. These large, complex organizations, known for their inability to pivot and adapt quickly, suddenly had no other choice.

American Airlines also was able to take advantage of their ongoing DevOps transformation to build big wins quickly, as you can read about in [Chapters 1 and 5](#).

Dr. Chris Strear relates his experiences using the Theory of Constraints to optimize flow in hospitals, as you can read about in [Chapter 2](#).

In 2020 Nationwide Building Society, the world's largest mutual financial institution, was able to respond in weeks to customer needs versus the typical years, thanks to their ongoing DevOps transformation. You can read more about their experience in [Chapter 8](#).

But while technology is a piece of a successful transformation into future ways of working, business leadership must lead the charge. The bottleneck of today is no longer just technical practices (though they still exist); the biggest challenge and necessity is getting business leadership on board. Transformation must be co-created between the business and technology, and the theories presented here can lead that change.

The enterprise can no longer sustain a binary thought process: top down or tech only. We must achieve true collaboration. Ninety percent of that work involves getting the right people engaged, onboard, and aligned. Start there and we can maintain the resulting motivation into the future.

—IT Revolution
Portland, OR

June 2021

FOREWORD TO THE SECOND EDITION

It has been five years since the first edition of *The DevOps Handbook* was released. While so much has changed, so much has also remained the same. Some of our tools and technologies are no longer in vogue or might not exist, but that shouldn't detract any readers. Although the technology landscape has shifted, the underlying principles presented in this book remain as relevant as ever.

In fact, the need for DevOps is even greater today, as organizations need to deliver value quickly, safely, securely, and reliably to their customers and users. To do this, they need to transform their internal processes and leverage technology to deliver value—using DevOps practices. This is true for organizations around the world and across all industry verticals.

Over the past several years, I've led research in the annual *State of DevOps Reports* (first with Puppet and later with DORA and Google), with co-authors Jez Humble and Gene Kim. The research has confirmed that many of the practices described in this book lead to improved outcomes like speed and stability of software releases; reductions in deployment pain and burnout of our engineers; and contributions to organizational performance, including profitability, productivity, customer satisfaction, effectiveness, and efficiency.

For the second edition of *The DevOps Handbook*, we have refreshed the text with updated data based on the latest research and best practices, and included new case studies to share even more stories about what transformations look like “on the ground.” Thanks for joining us on this journey of continuous improvement.

—Nicole Forsgren, PhD
Partner at Microsoft Research
2021

FOREWORD TO THE FIRST EDITION

In the past, many fields of engineering have experienced a sort of notable evolution, continually “leveling up” its understanding of its own work. While there are university curriculums and professional support organizations situated within specific disciplines of engineering (civil, mechanical, electrical, nuclear, etc.), the fact is, modern society needs all forms of engineering to recognize the benefits of and work in a multidisciplinary way.

Think about the design of a high-performance vehicle. Where does the work of a mechanical engineer end and the work of an electrical engineer begin? Where (and how, and when) should someone with domain knowledge of aerodynamics (who certainly would have well-formed opinions on the shape, size, and placement of windows) collaborate with an expert in passenger ergonomics? What about the chemical influences of fuel mixture and oil on the materials of the engine and transmission over the lifetime of the vehicle? There are other questions we can ask about the design of an automobile, but the end result is the same: success in modern technical endeavors absolutely requires multiple perspectives and expertise to collaborate.

In order for a field or discipline to progress and mature, it needs to reach a point where it can thoughtfully reflect on its origins, seek out a diverse set of perspectives on those reflections, and place that synthesis into a context that is useful for how the community pictures the future.

This book represents such a synthesis and should be seen as a seminal collection of perspectives on the (I will argue, still emerging and quickly evolving) field of software engineering and operations.

No matter what industry you are in, or what product or service your organization provides, this way of thinking is paramount and necessary for survival for every business and technology leader.

—John Allspaw, CTO, Etsy
Brooklyn, NY, August 2016

PREFACE

Aha!

The journey to complete *The DevOps Handbook* has been a long one—it started with weekly working Skype calls between the co-authors in February of 2011, with the vision of creating a prescriptive guide that would serve as a companion to the as-yet-unfinished book *The Phoenix Project: A Novel About IT, DevOps, and Helping Your Business Win*.

More than five years later, with over two thousand hours of work, *The DevOps Handbook* is finally here. Completing this book has been an extremely long process, although one that has been highly rewarding and full of incredible learning, with a scope that is much broader than we originally envisioned. Throughout the project, all the co-authors shared a belief that DevOps is genuinely important, formed in a personal “aha” moment much earlier in each of our professional careers, which I suspect many of our readers will resonate with.

Gene Kim

I’ve had the privilege of studying high-performing technology organizations since 1999, and one of the earliest findings was that boundary-spanning between the different functional groups of IT Operations, Information Security, and Development was critical to success. But I still remember the first time I saw the magnitude of the downward spiral that would result when these functions worked toward opposing goals.

It was 2006, and I had the opportunity to spend a week with the group that managed the outsourced IT Operations of a large airline reservation service. They described the downstream consequences of their large,

annual software releases: each release would cause immense chaos and disruption for the outsourcer as well as customers; there would be SLA (service level agreement) penalties because of the customer-impacting outages; there would be layoffs of the most talented and experienced staff because of the resulting profit shortfalls; there would be much unplanned work and firefighting so that the remaining staff couldn't work on the ever-growing service request backlogs coming from customers; the contract would be held together by the heroics of middle management; and everyone felt that the contract would be doomed to be put out for re-bid in three years.

The sense of hopelessness and futility that resulted created, for me, the beginnings of a moral crusade. Development seemed to always be viewed as strategic, but IT Operations was viewed as tactical, often delegated away or outsourced entirely, only to return in five years in worse shape than it was first handed over.

For many years, many of us knew that there must be a better way. I remember seeing the talks coming out of the 2009 Velocity Conference, describing amazing outcomes enabled by architecture, technical practices, and cultural norms that we now know as DevOps. I was so excited because it clearly pointed to the better way that we had all been searching for. And helping spread that word was one of my personal motivations to co-author *The Phoenix Project*. You can imagine how incredibly rewarding it was to see the broader community react to that book, describing how it helped them achieve their own “aha” moments.

Jez Humble

My DevOps “aha” moment was at a startup in 2000—my first job after graduating. For some time, I was one of two technical staff. I did everything: networking, programming, support, systems administration. We deployed software to production by FTP directly from our work-stations.

Then in 2004 I got a job at ThoughtWorks, a consultancy where my first gig was working on a project involving about seventy people. I was on a team of eight engineers whose full-time job was to deploy our software into a production-like environment. In the beginning, it was really stressful. But over a few months, we went from manual deployments that took two weeks to an automated deployment that took one hour, where we could roll forward and back in milliseconds using the blue-green deployment pattern during normal business hours.

That project inspired a lot of the ideas in both the *Continuous Delivery* book and this one. A lot of what drives me and others working in this space is the knowledge that, whatever your constraints, we can always do better, and the desire to help people on their journey.

Patrick Debois

For me, it was a collection of moments. In 2007 I was working on a data center migration project with some Agile teams. I was jealous that they had such high productivity—able to get so much done in so little time.

For my next assignment, I started experimenting with Kanban in Operations and saw how the dynamic of the team changed. Later, at the Agile Toronto 2008 conference, I presented my IEEE paper on this, but I felt

it didn't resonate widely in the Agile community. We started an Agile system administration group, but I overlooked the human side of things.

After seeing the 2009 Velocity Conference presentation "10 Deploys per Day" by John Allspaw and Paul Hammond, I was convinced others were thinking in a similar way. So I decided to organize the first DevOpsDays, accidentally coining the term DevOps.

The energy at the event was unique and contagious. When people started to thank me because it changed their life for the better, I understood the impact. I haven't stopped promoting DevOps since.

John Willis

In 2008, I had just sold a consulting business that focused on large-scale, legacy IT operations practices around configuration management and monitoring (Tivoli) when I first met Luke Kanies (the founder of Puppet Labs). Luke was giving a presentation on Puppet at an O'Reilly open-source conference on configuration management (CM).

At first I was just hanging out at the back of the room, killing time and thinking, "What could this twenty-year-old tell me about configuration management?" After all, I had literally been working my entire life at some of the largest enterprises in the world, helping them architect CM and other operations management solutions. However, about five minutes into his session, I moved up to the first row and realized everything I had been doing for the last twenty years was wrong. Luke was describing what I now call second-generation CM.

After his session, I had an opportunity to sit down and have coffee with him. I was totally sold on what we now call infrastructure as code. However, while we met for coffee, Luke started going even further, explaining his ideas. He started telling me he believed that Operations was going to have to start behaving like software developers. They were going to have to keep their configurations in source control and adopt CI/CD delivery patterns for their workflow. Being the old IT Operations person at the time, I think I replied to him with something like, “That idea is going to sink like Led Zeppelin with Ops folk.” (I was clearly wrong.)

Then about a year later, in 2009, at another O’Reilly conference, Velocity, I saw Andrew Clay Shafer give a presentation on Agile infrastructure. In his presentation, Andrew showed this iconic picture of a wall between developers and Operations with a metaphorical depiction of work being thrown over the wall. He coined this “the wall of confusion.” The ideas he expressed in that presentation codified what Luke was trying to tell me a year earlier. That was the light bulb for me. Later that year, I was the only American invited to the original DevOpsDays in Ghent. By the time that event was over, this thing we call DevOps was clearly in my blood.

DevOps Myths

Clearly, the co-authors of this book all came to a similar epiphany, even if they came there from very different directions. But there is now an overwhelming weight of evidence that the problems described above happen almost everywhere, and that the solutions associated with DevOps are nearly universally applicable.

The goal of writing this book is to describe how to replicate the DevOps transformations we've been a part of or have observed, as well as dispel many of the myths of why DevOps won't work in certain situations. Below are some of the most common myths we hear about DevOps.

Myth—DevOps Is Only for Startups: While DevOps practices have been pioneered by the web-scale, internet “unicorn” companies such as Google, Amazon, Netflix, and Etsy, each of these organizations has, at some point in their history, risked going out of business because of the problems associated with more traditional “horse” organizations: highly dangerous code releases that were prone to catastrophic failure, inability to release features fast enough to beat the competition, compliance concerns, an inability to scale, high levels of distrust between Development and Operations, and so forth.

However, each of these organizations was able to transform their architecture, technical practices, and culture to create the amazing outcomes that we associate with DevOps. As Dr. Branden Williams, an information security executive, said, “Let there be no more talk of DevOps unicorns or horses but only thoroughbreds and horses heading to the glue factory.”¹

Myth—DevOps Replaces Agile: DevOps principles and practices are compatible with Agile, with many observing that DevOps is a logical continuation of the Agile journey that started in 2001. Agile often serves as an effective enabler of DevOps because of its focus on small teams continually delivering high-quality code to customers.

Many DevOps practices emerge if we continue to manage our work beyond the goal of “potentially shippable code” at the end of each iteration, extending it to having our code

always in a deployable state, with developers checking into trunk daily, and if we demonstrate our features in production-like environments.

Myth—DevOps Is Incompatible with ITIL: Many view DevOps as a backlash to ITIL or ITSM (IT Service Management), which was originally published in 1989. ITIL has broadly influenced multiple generations of Ops practitioners, including one of the co-authors, and is an ever-evolving library of practices intended to codify the processes and practices that underpin world-class IT Operations, spanning service strategy, design, and support.

DevOps practices can be made compatible with ITIL processes. To support the shorter lead times and higher deployment frequencies associated with DevOps, many areas of ITIL become fully automated, solving many problems associated with the configuration and release management processes (e.g., keeping the configuration management database and definitive software libraries up to date). And because DevOps requires fast detection and recovery when service incidents occur, the ITIL disciplines of service design, incident, and problem management remain as relevant as ever.

Myth—DevOps Is Incompatible with Information Security and Compliance: The absence of traditional controls (e.g., segregation of duty, change approval processes, manual security reviews at the end of the project) may dismay information security and compliance professionals.

However, that doesn't mean that DevOps organizations don't have effective controls. Instead of security and compliance activities only being performed at the end of the project, controls are integrated into every stage of daily work

in the software development life cycle, resulting in better quality, security, and compliance outcomes.

Myth—DevOps Means Eliminating IT Operations, or “NoOps”: Many misinterpret DevOps as the complete elimination of the IT Operations function. However, this is rarely the case. While the nature of IT Operations work may change, it remains as important as ever. IT Operations collaborates far earlier in the software life cycle with Development, who continues to work with IT Operations long after the code has been deployed into production.

Instead of IT Operations doing manual work that comes from work tickets, it enables developer productivity through APIs and self-serviced platforms that create environments, test and deploy code, monitor and display production telemetry, and so forth. By doing this, IT Operations becomes more like Development (as do QA and Infosec), engaged in product development, where the product is the platform that developers use to safely, quickly, and securely test, deploy, and run their IT services in production.

Myth—DevOps Is Just “Infrastructure as Code” or Automation: While many of the DevOps patterns shown in this book require automation, DevOps also requires cultural norms and an architecture that allows for the shared goals to be achieved throughout the IT value stream. This goes far beyond just automation. As Christopher Little, a technology executive and one of the earliest chroniclers of DevOps, wrote, “DevOps isn’t about automation, just as astronomy isn’t about telescopes.”²

Myth—DevOps Is Only for Open-Source Software: Although many DevOps success stories take place in organizations using software such as the LAMP stack (Linux, Apache, MySQL, PHP), achieving DevOps outcomes is

independent of the technology being used. Successes have been achieved with applications written in [Microsoft.NET](#), COBOL, and mainframe assembly code, as well as with SAP and even embedded systems (e.g., HP LaserJet firmware).

Spreading the Aha! Moment

Each of the authors has been inspired by the amazing innovations happening in the DevOps community and the outcomes they are creating, including safe systems of work and enabling small teams to quickly and independently develop and validate code that can be safely deployed to customers. Given our belief that DevOps is a manifestation of creating dynamic, learning organizations that continually reinforce high trust cultural norms, it is inevitable that these organizations will continue to innovate and win in the marketplace.

It is our sincere hope that *The DevOps Handbook* will serve as a valuable resource for many people in different ways:

- a guide for planning and executing DevOps transformations
- a set of case studies to research and learn from
- a chronicle of the history of DevOps
- a means to create a coalition that spans Product Owners, Architecture, Development, QA, IT Operations, and Information Security to achieve common goals
- a way to get the highest levels of leadership support for DevOps initiatives, as well as a moral imperative to change the way we manage technology organizations to enable better effectiveness and efficiency, as well as enabling a happier and more humane work environment, helping everyone become lifelong

learners—this not only helps everyone achieve their highest goals as human beings, but also helps their organizations win

INTRODUCTION

Imagine a World Where Dev and Ops Become DevOps

Imagine a world where product owners, Development, QA, IT Operations, and Infosec work together, not only to help each other, but also to ensure that the overall organization succeeds. By working toward a common goal, they enable the fast flow of planned work into production (e.g., performing tens, hundreds, or even thousands of code deploys per day), while achieving world-class stability, reliability, availability, and security.

In this world, cross-functional teams rigorously test their hypotheses of which features will most delight users and advance the organizational goals. They care not just about implementing user features, but also about actively ensuring their work flows smoothly and frequently through the entire value stream without causing chaos and disruption to IT Operations or any other internal or external customer.

Simultaneously, QA, IT Operations, and Infosec are always working on ways to reduce friction for the team, creating the work systems that enable developers to be more productive and get better outcomes. By adding the expertise of QA, IT Operations, and Infosec into delivery teams and automated self-service tools and platforms, teams are able to use that expertise in their daily work without being dependent on other teams.

This enables organizations to create a safe system of work, where small teams are able to quickly and independently develop, test, and deploy code and value quickly, safely, securely, and reliably to customers. This allows organizations to maximize developer productivity, enable

organizational learning, create high employee satisfaction, and win in the marketplace.

These are the outcomes that result from DevOps. For most of us, this is not the world we live in. More often than not, the system we work in is broken, resulting in extremely poor outcomes that fall well short of our true potential. In our world, Development and IT Operations are adversaries; testing and Infosec activities happen only at the end of a project, too late to correct any problems found; and almost any critical activity requires too much manual effort and too many handoffs, leaving us always waiting. Not only does this contribute to extremely long lead times to get anything done, but the quality of our work, especially production deployments, is also problematic and chaotic, resulting in negative impacts to our customers and our business.

As a result, we fall far short of our goals, and the whole organization is dissatisfied with the performance of IT, resulting in budget reductions and frustrated, unhappy employees who feel powerless to change the process and its outcomes.* The solution? We need to change how we work; DevOps shows us the best way forward.

To better understand the potential of the DevOps revolution, let us look at the manufacturing revolution of the 1980s. By adopting Lean principles and practices, manufacturing organizations dramatically improved plant productivity, customer lead times, product quality, and customer satisfaction, enabling them to win in the marketplace.

Before the revolution, average manufacturing plant order lead times were six weeks, with fewer than 70% of orders shipped on time. By 2005, with the widespread implementation of Lean practices, average product lead times

had dropped to less than three weeks, and more than 95% of orders were shipped on time.¹ Organizations that did not implement Lean practices lost market share, and many went out of business entirely.

Similarly, the bar has been raised for delivering technology products and services—what was good enough in previous decades is not good enough now. For each of the last four decades, the cost and time required to develop and deploy strategic business capabilities and features have dropped by orders of magnitude. During the 1970s and 1980s, most new features required one to five years to develop and deploy, often costing tens of millions of dollars.

By the 2000s, because of advances in technology and the adoption of Agile principles and practices, the time required to develop new functionality had dropped to weeks or months but deploying into production still required weeks or months, often with catastrophic outcomes.

And by 2010, with the introduction of DevOps and the never-ending commoditization of hardware, software, and now the cloud, features (even entire startup companies) could be created in weeks, quickly being deployed into production in just hours or minutes—for these organizations, deployment finally became routine and low risk. These organizations are able to perform experiments to test business ideas, discovering which ideas create the most value for customers and the organization as a whole, and which are then further developed into features that can be rapidly and safely deployed into production.

Table 0.1: The Ever-Accelerating Trend toward Faster, Cheaper, Lower Risk Delivery of Software

	1970s–1980s	1990s	2000s–Present
Era	Mainframes	Client/Server	Commoditization and Cloud
Representative technology of era	COBOL, DB2 on MVS, etc.	C++, Oracle, Solaris, etc.	Java, MySQL, Red Hat, Ruby on Rails, PHP, etc.
Cycle time	1–5 years	3–12 months	2–12 weeks
Cost	\$1M–\$100M	\$100k–\$10M	\$10k–\$1M
At risk	The whole company	A product line or division	A product feature
Cost of failure	Bankruptcy, sell the company, massive layoffs	Revenue miss, CIO's job	Negligible

Source: Adrian Cockcroft, “Velocity and Volume (or Speed Wins),” presentation at FlowCon, San Francisco, CA, November 2013.

Today, organizations adopting DevOps principles and practices often deploy changes hundreds or even thousands of times per day. In an age where competitive advantage requires fast time to market and relentless experimentation, organizations that are unable to replicate these outcomes are destined to lose in the marketplace to more nimble competitors and could potentially go out of business entirely, much like the manufacturing organizations that did not adopt Lean principles.

These days, regardless of what industry we are competing in, the way we acquire customers and deliver value to them is dependent on the technology value stream. Put even more succinctly, as Jeffrey Immelt, CEO of General Electric, stated, “Every industry and company that is not bringing software to the core of their business will be disrupted.”² Or as Jeffrey Snover, Technical Fellow at Microsoft, said, “In previous economic eras, businesses created value by moving atoms. Now they create value by moving bits.”³

It’s difficult to overstate the enormity of this problem—it affects every organization, independent of the industry we operate in, the size of our organization, whether we are profit or nonprofit. Now more than ever, how technology work is managed and performed predicts whether our organizations will win in the marketplace or even survive. In many cases, we will need to adopt principles and practices that look very different from those that have successfully guided us over the past decades. (See [Appendix 1](#).)

Now that we have established the urgency of the problem that DevOps solves, let us take some time to explore in more detail the symptomatology of the problem, why it occurs, and why, without dramatic intervention, the problem worsens over time.

The Problem: Something in Your Organization Must Need Improvement (Or You Wouldn’t Be Reading This Book)

Most organizations are not able to deploy production changes in minutes or hours, instead requiring weeks or months. Nor are they able to deploy hundreds or thousands of changes into production per day; instead, they struggle to deploy monthly or even quarterly. Nor are production deployments

routine, instead involving outages and chronic firefighting and heroics.

In an age where competitive advantage requires fast time to market, high service levels, and relentless experimentation, these organizations are at a significant competitive disadvantage. This is in large part due to their inability to resolve a core, chronic conflict within their technology organization.

The Core, Chronic Conflict

In almost every IT organization, there is an inherent conflict between Development and IT Operations that creates a downward spiral, resulting in ever-slower time to market for new products and features, reduced quality, increased outages, and, worst of all, an ever-increasing amount of technical debt.

The term “technical debt” was first coined by Ward Cunningham. Analogous to financial debt, technical debt describes how decisions we make lead to problems that get increasingly more difficult to fix over time, continually reducing our available options in the future—even when taken on judiciously, we still incur interest.

One factor that contributes to this is the often competing goals of Development and IT Operations. IT organizations are responsible for many things. Among them are the two following goals, which must be pursued simultaneously:

- Respond to the rapidly changing competitive landscape.
- Provide stable, reliable, and secure service to the customer.

Frequently, Development will take responsibility for responding to changes in the market and for deploying

features and changes into production as quickly as possible. IT Operations will take responsibility for providing customers with IT service that is stable, reliable, and secure, making it difficult or even impossible for anyone to introduce production changes that could jeopardize production. Configured this way, Development and IT Operations have diametrically opposed goals and incentives.

Dr. Eliyahu M. Goldratt, one of the founders of the manufacturing management movement, called these types of configuration “the core, chronic conflict”—when organizational measurements and incentives across different silos prevent the achievement of global, organizational goals.^{4†}

This conflict creates a downward spiral so powerful it prevents the achievement of desired business outcomes, both inside and outside the IT organization. These chronic conflicts often put technology workers into situations that lead to poor software and service quality and bad customer outcomes, as well as a daily need for workarounds, firefighting, and heroics, whether in Product Management, Development, QA, IT Operations, or Information Security. (See [Appendix 2](#).)

Downward Spiral in Three Acts

The downward spiral in IT has three acts that are likely familiar to most IT practitioners. The first act begins in IT Operations, where our goal is to keep applications and infrastructure running so that our organization can deliver value to customers. In our daily work, many of our problems are due to applications and infrastructure that are complex, poorly documented, and incredibly fragile. This is the technical debt and daily workarounds that we live with

constantly, always promising that we'll fix the mess when we have a little more time. But that time never comes.

Alarmingly, our most fragile artifacts support either our most important revenue-generating systems or our most critical projects. In other words, the systems most prone to failure are also our most important and are at the epicenter of our most urgent changes. When these changes fail, they jeopardize our most important organizational promises, such as availability to customers, revenue goals, security of customer data, accurate financial reporting, and so forth.

The second act begins when somebody has to compensate for the latest broken promise—it could be a product manager promising a bigger, bolder feature to dazzle customers with or a business executive setting an even larger revenue target. Then, oblivious to what technology can or can't do, or what factors led to missing our earlier commitment, they commit the technology organization to deliver upon this new promise.

As a result, Development is tasked with another urgent project that inevitably requires solving new technical challenges and cutting corners to meet the promised release date, further adding to our technical debt—made, of course, with the promise that we'll fix any resulting problems when we have a little more time.

This sets the stage for the third and final act, where everything becomes just a little more difficult, bit by bit—everybody gets a little busier, work takes a little more time, communications become a little slower, and work queues get a little longer. Our work becomes more tightly coupled, smaller actions cause bigger failures, and we become more fearful and less tolerant of making changes. Work requires more communication, coordination, and approvals; teams

must wait just a little longer for their dependent work to get done; and our quality keeps getting worse. The wheels begin grinding slower and require more effort to keep turning. (See [Appendix 3](#).)

Although it's difficult to see in the moment, the downward spiral is obvious when one takes a step back. We notice that production code deployments are taking ever-longer to complete, moving from minutes to hours to days to weeks. And worse, the deployment outcomes have become even more problematic, resulting in an ever-increasing number of customer-impacting outages that require more heroics and firefighting in Operations, further depriving them of their ability to pay down technical debt.

As a result, our product delivery cycles continue to move slower and slower, fewer projects are undertaken, and those that are, are less ambitious. Furthermore, the feedback on everyone's work becomes slower and weaker, especially the feedback signals from our customers. And, regardless of what we try, things seem to get worse—we are no longer able to respond quickly to our changing competitive landscape, nor are we able to provide stable, reliable service to our customers. As a result, we ultimately lose in the marketplace.

Time and time again, we learn that when IT fails, the entire organization fails. As Steven J. Spear noted in his book *The High-Velocity Edge*, whether the damages “unfold slowly like a wasting disease” or rapidly “like a fiery crash ... the destruction can be just as complete.”⁵

Why Does This Downward Spiral Happen Everywhere?

For over a decade, the authors of this book have observed this destructive spiral occur in countless organizations of all types and sizes. We understand better than ever why this downward spiral occurs and why it requires DevOps

principles to mitigate. First, as described earlier, every IT organization has two opposing goals, and second, every company is a technology company, whether they know it or not.

As Christopher Little, a software executive and one of the earliest chroniclers of DevOps, said, “Every company is a technology company, regardless of what business they think they’re in. A bank is just an IT company with a banking license.”^{6†} To convince ourselves that this is the case, consider that the vast majority of capital projects have some reliance on IT. As the saying goes, “It is virtually impossible to make any business decision that doesn’t result in at least one IT change.”

In the business and finance context, projects are critical because they serve as the primary mechanism for change inside organizations. Projects are typically what management needs to approve, budget for, and be held accountable for; therefore, they are the mechanism that achieves the goals and aspirations of the organization, whether it is to grow or even shrink.[§]

Projects are typically funded through capital spending (e.g., factories, equipment, and major projects, and expenditures are capitalized when payback is expected to take years), of which 50% is now technology related. This is even true in “low tech” industry verticals with the lowest historical spending on technology, such as energy, metal, resource extraction, automotive, and construction.⁸ In other words, business leaders are far more reliant upon the effective management of IT in order to achieve their goals than they think.[¶]

The Costs: Human and Economic

When people are trapped in this downward spiral for years, especially those who are downstream of Development, they often feel stuck in a system that preordains failure and leaves them powerless to change the outcomes. This powerlessness is often followed by burnout, with the associated feelings of fatigue, cynicism, and even hopelessness and despair.

Many psychologists assert that creating systems that cause feelings of powerlessness is one of the most damaging things we can do to fellow human beings—we deprive other people of their ability to control their own outcomes and even create a culture where people are afraid to do the right thing because of fear of punishment, failure, or jeopardizing their livelihood. This type of culture can create the conditions for learned helplessness, where people become unwilling or unable to act in a way that avoids the same problem in the future.

For our employees, it means long hours, working on weekends, and a decreased quality of life, not just for the employee, but for everyone who depends on them, including family and friends. It is not surprising that when this occurs, we lose our best people (except for those who feel like they can't leave because of a sense of duty or obligation).

In addition to the human suffering that comes with the current way of working, the opportunity cost of the value that we could be creating is staggering—the authors believe that we are missing out on approximately \$2.6 trillion of value creation per year, which is, at the time of this writing, equivalent to the annual economic output of France, the sixth-largest economy in the world.

Consider the following calculation—both IDC and Gartner estimated that in 2011 approximately 5% of the worldwide gross domestic product (\$3.1 trillion) was spent

on IT (hardware, services, and telecom).¹⁰ If we estimate that 50% of that \$3.1 trillion was spent on operating costs and maintaining existing systems, and that one-third of that 50% was spent on urgent and unplanned work or rework, approximately \$520 billion was wasted.

If adopting DevOps could enable us—through better management and increased operational excellence—to halve that waste and redeploy that human potential into something that’s five times the value (a modest proposal), we could create \$2.6 trillion of value per year.

The Ethics of DevOps: There Is a Better Way

In the previous sections, we described the problems and the negative consequences of the status quo due to the core, chronic conflict, from the inability to achieve organizational goals to the damage we inflict on fellow human beings. By solving these problems, DevOps astonishingly enables us to simultaneously improve organizational performance, achieve the goals of all the various functional technology roles (e.g., Development, QA, IT Operations, Infosec), and improve the human condition.

This exciting and rare combination may explain why DevOps has generated so much excitement and enthusiasm with so many and in such a short time, including technology leaders, engineers, and much of the software ecosystem we reside in.

Breaking the Downward Spiral with DevOps

Ideally, small teams of developers independently implement their features, validate their correctness in production-like environments, and have their code deployed into production quickly, safely, and securely. Code deployments are routine and predictable. Instead of starting deployments at midnight

on Friday and spending all weekend working to complete them, deployments occur throughout the business day when everyone is already in the office and without our customers even noticing—except when they see new features and bug fixes that delight them. And, by deploying code in the middle of the workday, IT Operations is working during normal business hours like everyone else for the first time in decades.

By creating fast feedback loops at every step of the process, everyone can immediately see the effects of their actions. Whenever changes are committed into version control, fast automated tests are run in production-like environments, giving continual assurance that the code and environments operate as designed and are always in a secure and deployable state.

Automated testing helps developers discover their mistakes quickly (usually within minutes), which enables faster fixes as well as genuine learning—learning that is impossible when mistakes are discovered six months later during integration testing, when memories and the link between cause and effect have long faded. Instead of accruing technical debt, problems are fixed as they are found, mobilizing the entire organization if needed because global goals outweigh local goals.

Pervasive production telemetry in our code and production environments ensures that problems are detected and corrected quickly, confirming that everything is working as intended and that customers are getting value from the software we create.

In this scenario, everyone feels productive—the architecture allows small teams to work safely and architecturally decoupled from the work of other teams who use self-service platforms that leverage the collective

experience of Operations and Information Security. Instead of everyone waiting all the time, with large amounts of late, urgent rework, teams work independently and productively in small batches, quickly and frequently delivering new value to customers.

Even high-profile product and feature releases become routine by using dark launch techniques. Long before the launch date, we put all the required code for the feature into production, invisible to everyone except internal employees and small cohorts of real users, allowing us to test and evolve the feature until it achieves the desired business goal.

And, instead of firefighting for days or weeks to make the new functionality work, we merely change a feature toggle or configuration setting. This small change makes the new feature visible to ever-larger segments of customers, automatically rolling back if something goes wrong. As a result, our releases are controlled, predictable, reversible, and low stress.

It's not just feature releases that are calmer—all sorts of problems are being found and fixed early, when they are smaller and when they are cheaper and easier to correct. With every fix, we also generate organizational learnings, enabling us to prevent the problem from recurring and enabling us to detect and correct similar problems faster in the future.

Furthermore, everyone is constantly learning, fostering a hypothesis-driven culture where the scientific method is used to ensure nothing is taken for granted—we do nothing without measuring and treating product development and process improvement as experiments.

Because we value everyone's time, we don't spend years building features that our customers don't want, deploying

code that doesn't work, or fixing something that isn't actually the cause of our problem.

Because we care about achieving goals, we create long-term teams that are responsible for meeting them. Instead of project teams where developers are reassigned and shuffled around after each release, never receiving feedback on their work, we keep teams intact so they can keep iterating and improving, using those learnings to better achieve their goals. This is equally true for the product teams who are solving problems for our external customers, as well as our internal platform teams who are helping other teams be more productive, safe, and secure.

Instead of a culture of fear, we have a high-trust, collaborative culture, where people are rewarded for taking risks. They are able to fearlessly talk about problems as opposed to hiding them or putting them on the back burner—after all, we must see problems in order to solve them.

And, because everyone fully owns the quality of their work, everyone builds automated testing into their daily work and uses peer reviews to gain confidence that problems are addressed long before they can impact a customer. These processes mitigate risk, as opposed to approvals from distant authorities, allowing us to deliver value quickly, reliably, and securely—even proving to skeptical auditors that we have an effective system of internal controls.

When something does go wrong, we conduct blameless post-mortems, not to punish anyone but to better understand what caused the accident and how to prevent it. This ritual reinforces our culture of learning. We also hold internal technology conferences to elevate our skills and ensure that everyone is always teaching and learning.

Because we care about quality, we even inject faults into our production environment so we can learn how our system fails in a planned manner. We conduct planned exercises to practice large-scale failures, randomly kill processes and compute servers in production, and inject network latencies and other nefarious acts to ensure we grow ever more resilient. By doing this, we enable better resilience, as well as organizational learning and improvement.

In this world, everyone has ownership in their work, regardless of their role in the technology organization. They have confidence that their work matters and is meaningfully contributing to organizational goals, proven by their low-stress work environment and their organization's success in the marketplace. Their proof is that the organization is indeed winning in the marketplace.

The Business Value of DevOps

We have decisive evidence of the business value of DevOps. From 2013 through 2016, as part of Puppet Labs' *State Of DevOps Report*, to which authors Nicole Forsgren, Jez Humble, and Gene Kim contributed, we collected data from over twenty-five thousand technology professionals with the goal of better understanding the health and habits of organizations at all stages of DevOps adoption.**

The first surprise this data revealed was how much high-performing organizations using DevOps practices were outperforming their non-high-performing peers in the following areas:¹¹

- Throughput metrics
 - code and change deployments (thirty times more frequent)

- code and change deployment lead time (two hundred times faster)
- Reliability metrics
 - production deployments (sixty times higher change success rate)
 - mean time to restore service (168 times faster)

Organizational performance metrics

- productivity, market share, and profitability goals (two times more likely to exceed)
- market capitalization growth (50% higher over three years)

In other words, high performers were both more agile and more reliable, providing empirical evidence that DevOps enables us to break the core, chronic conflict. High performers deployed code thirty times more frequently, and the time required to go from “code committed” to “successfully running in production” was two hundred times faster—high performers had lead times measured in minutes or hours, while low performers had lead times measured in weeks, months, or even quarters.

Furthermore, high performers were twice as likely to exceed profitability, market share, and productivity goals. And, for those organizations that provided a stock ticker symbol, we found that high performers had 50% higher market capitalization growth over three years. They also had higher employee job satisfaction, lower rates of employee burnout, and their employees were 2.2 times more likely to recommend their organization to friends as a great place to work.^{††} High performers also had better information security outcomes. By integrating security objectives into all stages of

the development and operations processes, they spent 50% less time remediating security issues.

DevOps Helps Scale Developer Productivity

When we increase the number of developers, individual developer productivity often significantly decreases due to communication, integration, and testing overhead.

This is highlighted in the famous book by Frederick Brook, *The Mythical Man-Month*, where he explains that when projects are late, adding more developers not only decreases individual developer productivity but also decreases overall productivity.¹³

On the other hand, DevOps shows us that when we have the right architecture, the right technical practices, and the right cultural norms, small teams of developers are able to quickly, safely, and independently develop, integrate, test, and deploy changes into production.

As Randy Shoup, formerly a director of engineering at Google and now VP Engineering at eBay, observed, large organizations using DevOps “have thousands of developers, but their architecture and practices enable small teams to still be incredibly productive, as if they were a startup.”¹⁴

The *2015 State of DevOps Report* examined not only “deploys per day” but also “deploys per day per developer.” They hypothesized that high performers would be able to scale their number of deployments as team sizes grew.¹⁵

Indeed, this is what they found. [Figure 0.1](#) shows that in low performers, deploys per day per developer go down as team size increases, stays constant for medium performers, and increases linearly for high performers. In other words, organizations adopting DevOps are able to linearly increase

the number of deploys per day as they increase their number of developers, just as Google, Amazon, and Netflix have done.##

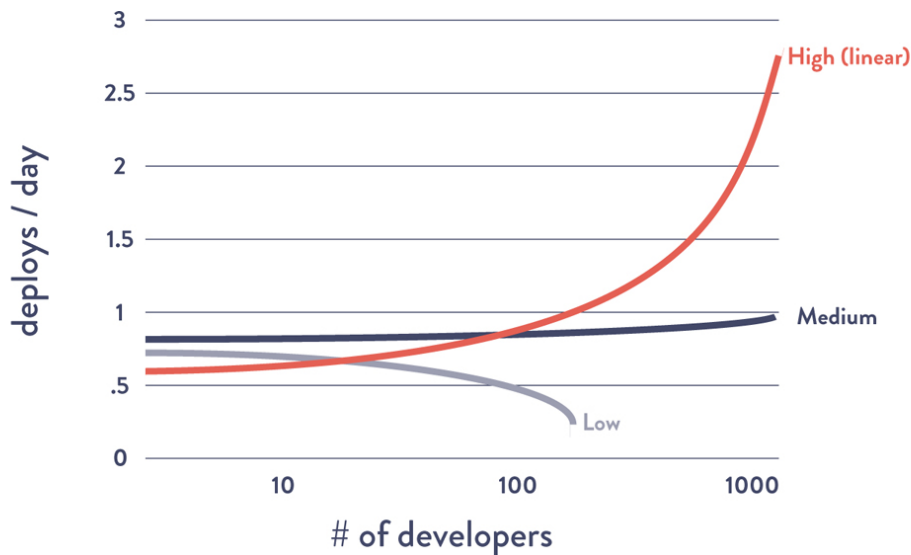


Figure 0.1: Deployments per Day vs. Number of Developers

Only organizations that are deploying at least once per day are shown.

Source: Puppet Labs, 2015 State Of DevOps Report.

The Universality of the Solution

One of the most influential books in the Lean manufacturing movement is *The Goal: A Process of Ongoing Improvement*, written by Dr. Eliyahu M. Goldratt in 1984. It influenced an entire generation of professional plant managers around the world. It was a novel about a plant manager who had to fix his cost and product due date issues in ninety days, otherwise his plant would be shut down.

Later in his career, Dr. Goldratt described the letters he received in response to *The Goal*. These letters would typically read, “You have obviously been hiding in our factory, because you’ve described my life [as a plant manager] exactly ...”¹⁷ Most importantly, these letters showed that people were able to replicate the breakthroughs in performance that were described in the book in their own work environments.

The Phoenix Project: A Novel About IT, DevOps, and Helping Your Business Win, written by Gene Kim, Kevin Behr, and George Spafford in 2013, was closely modeled after *The Goal*. It is a novel that follows an IT leader who faces all the typical problems that are endemic in IT organizations: an over-budget, behind-schedule project that must get to market in order for the company to survive. He experiences catastrophic deployments; problems with availability, security, and compliance; and so forth.

Ultimately, he and his team use DevOps principles and practices to overcome those challenges, helping their organization win in the marketplace. In addition, the novel shows how DevOps practices improved the workplace environment for the team, creating lower stress and higher satisfaction because of greater practitioner involvement throughout the process.

As with *The Goal*, there is tremendous evidence of the universality of the problems and solutions described in *The Phoenix Project*. Consider some of the statements found in the Amazon reviews: “I find myself relating to the characters in *The Phoenix Project* ... I’ve probably met most of them over the course of my career,” “If you have ever worked in any aspect of IT, DevOps, or Infosec you will definitely be able to relate to this book,” or “There’s not a character in *The Phoenix Project* that I don’t identify with myself or someone I know in real life ... not to mention the problems faced and overcome by those characters.”¹⁸

The DevOps Handbook: An Essential Guide

In the remainder of this book, we will describe how to replicate the transformation described in *The Phoenix Project*, as well as provide many case studies of how other

organizations have used DevOps principles and practices to replicate those outcomes.

The purpose of *The DevOps Handbook* is to give you the theory, principles, and practices you need to successfully start your DevOps initiative and achieve your desired outcomes. This guidance is based on decades of sound management theory, the study of high-performing technology organizations, work we have done helping organizations transform, and research that validates the effectiveness of the prescribed DevOps practices, as well as interviews with relevant subject matter experts and analyses of nearly one hundred case studies presented at the DevOps Enterprise Summit.

Broken into six parts, this book covers DevOps theories and principles using the Three Ways, a specific view of the underpinning theory originally introduced in *The Phoenix Project*. *The DevOps Handbook* is for everyone who performs or influences work in the technology value stream (which typically includes Product Management, Development, QA, IT Operations, and Information Security), as well as for business and marketing leadership, where most technology initiatives originate.

The reader is not expected to have extensive knowledge of any of these domains, or of DevOps, Agile, ITIL, Lean, or process improvement. Each of these topics is introduced and explained in the book as it becomes necessary.

Our intent is to create a working knowledge of the critical concepts in each of these domains, to serve as a primer and to introduce the language necessary to help practitioners work with all their peers across the entire IT value stream, and to frame shared goals.

This book will be of value to business leaders and stakeholders who are increasingly reliant upon the technology organization to achieve their goals.

Furthermore, this book is intended for readers whose organizations might not be experiencing all the problems described in the book (e.g., long deployment lead times or painful deployments). Even readers in this fortunate position will benefit from understanding DevOps principles, especially those relating to shared goals, feedback, and continual learning.

In [Part I](#), we present a brief history of DevOps and introduce the underpinning theory and key themes from relevant bodies of knowledge that span decades. We then present the high-level principles of the Three Ways: Flow, Feedback, and Continual Learning and Experimentation.

[Part II](#) describes how and where to start and presents concepts such as value streams, organizational design principles and patterns, organizational adoption patterns, and case studies.

[Part III](#) describes how to accelerate flow by building the foundations of our deployment pipeline: enabling fast and effective automated testing, continuous integration, continuous delivery, and architecting for low-risk releases.

[Part IV](#) discusses how to accelerate and amplify feedback by creating effective production telemetry to see and solve problems, better anticipate problems and achieve goals, enable feedback so that Development and Operations can safely deploy changes, integrate A/B testing into our daily work, and create review and coordination processes to increase the quality of our work.

[Part V](#) describes how we accelerate continual learning and experimentation by establishing a just culture, converting

local discoveries into global improvements, and properly reserving time to create organizational learning and improvements.

Finally, in [Part VI](#) we describe how to properly integrate security and compliance into our daily work by integrating preventative security controls into shared source code repositories and services, integrating security into our deployment pipeline, enhancing telemetry to better enable detection and recovery, protecting the deployment pipeline, and achieving change management objectives.

By codifying these practices, we hope to accelerate the adoption of DevOps practices, increase the success of DevOps initiatives, and lower the activation energy required for DevOps transformations.

-
- * This is just a small sample of the problems found in typical IT organizations.
 - † In the manufacturing realm, a similar core, chronic conflict existed: the need to simultaneously ensure on-time shipments to customers and to control costs. How this core, chronic conflict was broken is described in [Appendix 2](#).
 - ‡ In 2013, the European bank HSBC employed more software developers than Google.⁷
 - § For now, let us suspend the discussion of whether software should be funded as a “project” or a “product.” This is discussed later in the book.
 - ¶ For instance, Dr. Vernon Richardson and his colleagues published this astonishing finding. They studied the 10-K SEC filings of 184 public corporations and divided them into three groups: A) firms with material weaknesses with IT-related deficiencies, B) firms with material weaknesses with no IT-related deficiencies, and C) “clean firms” with no material weaknesses. Firms in Group A saw eight times higher CEO turnover than Group C, and there was four times higher CFO turnover in Group A than in Group C. Clearly, IT may matter far more than we typically think.⁹
 - ** The *State of DevOps Report* has since been repeated every year. Additionally, the key findings from the 2013–2018 reports were collected into the book *Accelerate: The Science of Lean Software and DevOps: Building and Scaling High Performing Technology Organizations*.

- †† As measured by employee Net Promoter Score (eNPS). This is a significant finding, as research has shown that “companies with highly engaged workers grew revenues two and a half times as much as those with low engagement levels. And [publicly traded] stocks of companies with a high-trust work environment outperformed market indexes by a factor of three from 1997 through 2011.”¹²

- ‡‡ Another more extreme example is Amazon. In 2011, Amazon was performing approximately seven thousand deploys per day. By 2015, they were performing 130,000 deploys per day.¹⁶



Part I

THE THREE WAYS

PART I: INTRODUCTION

In Part I of *The DevOps Handbook*, we will explore how the convergence of several important movements in management and technology set the stage for the DevOps movement. We describe value streams, how DevOps is the result of applying Lean principles to the technology value stream, and the Three Ways: Flow, Feedback, and Continual Learning and Experimentation.

Primary focuses within these chapters include:

- The principles of Flow, which accelerate the delivery of work from Development to Operations to our customers.
- The principles of Feedback, which enable us to create ever-safer systems of work.
- The principles of Continual Learning and Experimentation, which foster a high-trust culture and a scientific approach to organizational improvement and risk-taking as part of our daily work.

A Brief History

DevOps and its resulting technical, architectural, and cultural practices represent a convergence of many philosophical and management movements. While many organizations have developed these principles independently, understanding that DevOps resulted from a broad stroke of movements, a phenomenon described by John Willis (one of the co-authors of this book) as the “convergence of Dev and Ops,” shows an amazing progression of thinking and improbable connections. There are decades of lessons learned from manufacturing, high-reliability organizations, high-trust

management models, and others that have brought us to the DevOps practices we know today.

DevOps is the outcome of applying the most trusted principles from the domain of physical manufacturing and leadership to the IT value stream. DevOps relies on bodies of knowledge from Lean, Theory of Constraints, the Toyota Production System, resilience engineering, learning organizations, safety culture, human factors, and many others. Other valuable contexts that DevOps draws from include high-trust management cultures, servant leadership, and organizational change management.

The result is world-class quality, reliability, stability, and security at ever-lower cost and effort and accelerated flow and reliability throughout the technology value stream, including Product Management, Development, QA, IT Operations, and Infosec.

While the foundation of DevOps can be seen as being derived from Lean, the Theory of Constraints, and the Toyota Kata movement, many also view DevOps as the logical continuation of the Agile software journey that began in 2001.

The Lean Movement

Techniques such as value stream mapping, kanban boards, and total productive maintenance were codified for the Toyota Production System in the 1980s. In 1997, the Lean Enterprise Institute started researching applications of Lean to other value streams, such as the service industry and healthcare.

Two of Lean's central tenets include the deeply held belief that the manufacturing lead time required to convert raw materials into finished goods is the best predictor of quality, customer satisfaction, and employee happiness, and that one

of the best predictors of short lead times is small batch sizes of work.

Lean principles focus on how to create value for the customer through systems thinking by creating constancy of purpose, embracing scientific thinking, creating flow and pull (versus push), assuring quality at the source, leading with humility, and respecting every individual.

The Agile Manifesto

The Agile Manifesto was created in 2001 at an invite-only event by seventeen experts in what was then known as “lightweight methods” in software development. They wanted to create a set of values and principles that captured the advantage of these more adaptive methods, compared to the prevailing software development processes such as waterfall development and methodologies such as the Rational Unified Process.

One key principle was to “deliver working software frequently, from a couple of weeks to a couple of months, with a preference to the shorter timescale,”¹ emphasizing the desire for small batch sizes—incremental releases instead of large, big-bang releases. Other principles emphasized the need for small, self-motivated teams working in a high-trust management model.

Agile is credited for dramatically increasing the productivity and responsiveness of many development organizations. And interestingly, many of the key moments in DevOps history also occurred within the Agile community or at Agile conferences, as described below.

Agile Infrastructure and Velocity Movement

At the 2008 Agile conference in Toronto, Canada, Patrick Debois and Andrew Shafer held a “birds of a feather” session

on applying Agile principles to infrastructure as opposed to application code. (In its early days, this was referred to as “Agile system administration.”) Although they were the only people who showed up, they rapidly gained a following of like-minded thinkers, including co-author John Willis.

CONTINUOUS LEARNING

Around the same time, a few academics started studying system administrators, how they applied engineering principles to their work, and how it impacted performance. The leading experts included a group from IBM Research, with ethnographies led by Dr. Eben Haber, Dr. Eser Kandogan, and Dr. Paul Maglio. This was extended to include behavioral quantitative studies led by co-author Dr. Nicole Forsgren in 2007–2009. Nicole went on to lead the research in the 2014–2019 *State of DevOps Reports*, the industry-standard research into practices and capabilities that drive software delivery and performance; these were published by Puppet and DORA.

Later, at the 2009 Velocity conference, John Allspaw and Paul Hammond gave the seminal “10 Deploys per Day: Dev and Ops Cooperation at Flickr” presentation, where they described how they created shared goals between Dev and Ops and used continuous integration practices to make deployment part of everyone’s daily work. According to firsthand accounts, everyone attending the presentation immediately knew they were in the presence of something profound and of historic significance.

Patrick Debois was so excited by Allspaw and Hammond's idea that he created the first DevOpsDays in Ghent, Belgium, in 2009, where the term "DevOps" was coined.

The Continuous Delivery Movement

Building upon the development discipline of continuous build, test, and integration, Jez Humble and David Farley extended the concept to continuous delivery, which defined the role of a "deployment pipeline" to ensure that code and infrastructure are always in a deployable state and that all code checked into trunk can be safely deployed into production. This idea was first presented at the 2006 Agile conference and was also independently developed in 2009 by Tim Fitz in a blog post on his website titled "Continuous Deployment."^{*}

Toyota Kata

In 2009, Mike Rother wrote *Toyota Kata: Managing People for Improvement, Adaptiveness, and Superior Results*, which framed his twenty-year journey to understand and codify the Toyota Production System. He had been one of the graduate students who flew with GM executives to visit Toyota plants and helped develop the Lean toolkit, but he was puzzled when none of the companies adopting these practices replicated the level of performance observed at the Toyota plants.

He concluded that the Lean community missed the most important practice of all, which he called the *improvement kata*.² He explains that every organization has work routines, and the improvement kata requires creating structure for the daily, habitual practice of improvement work because daily practice is what improves outcomes. The constant cycle of establishing desired future states, setting target outcomes on

a cadence, and the continual improvement of daily work is what guided improvement at Toyota.

Throughout the rest of Part I, we will look at value streams, how Lean principles can be applied to the technology value stream, and the Three Ways of Flow, Feedback, and Continual Learning and Experimentation.

* DevOps also extends and builds upon the practices of infrastructure as code, which was pioneered by Dr. Mark Burgess, Luke Kanies, and Adam Jacob. In *infrastructure as code*, the work of Operations is automated and treated like application code, so that modern development practices can be applied to the entire development stream. This further enabled fast deployment flow, including continuous integration (pioneered by Grady Booch and integrated as one of the key 12 practices of Extreme Programming), continuous delivery (pioneered by Jez Humble and David Farley), and continuous deployment (pioneered by Etsy, Wealthfront, and Eric Ries's work at IMVU).

1

AGILE, CONTINUOUS DELIVERY, AND THE THREE WAYS

In this chapter, we present an introduction to the underpinning theory of Lean Manufacturing, as well as the Three Ways—the principles from which the observed DevOps behaviors can be derived.

Our focus here is primarily on theory and principles, describing many decades of lessons learned from manufacturing, high-reliability organizations, high-trust management models, and others, from which DevOps practices have been derived. The resulting concrete principles and patterns, and their practical application to the technology value stream, are presented in the remaining chapters of the book.

The Manufacturing Value Stream

One of the fundamental concepts in Lean is the value stream. We will define it first in the context of manufacturing and then extrapolate how it applies to DevOps and the technology value stream.

Karen Martin and Mike Osterling define a value stream in their book *Value Stream Mapping: How to Visualize Work and Align Leadership for Organizational Transformation* as “the sequence of activities an organization undertakes to deliver upon a customer request,” or “the sequence of activities required to design, produce, and deliver a good or service to a customer, including the dual flows of information and material.”¹

In manufacturing operations, the value stream is often easy to see and observe: it starts when a customer order is received and the raw materials are released onto the plant floor. To enable fast and predictable lead times in any value stream, there is usually a relentless focus on creating a smooth and even flow of work, using techniques such as small batch sizes, reducing work in process (WIP), preventing rework to ensure defects are not passed to downstream work centers, and constantly optimizing systems toward global goals.

The Technology Value Stream

Many principles and patterns that enable the fast flow of work in physical processes are equally applicable to technology work (and, for that matter, for all knowledge work). In DevOps, we typically define our technology value stream as the process required to convert a business hypothesis into a technology-enabled service or feature that delivers value to the customer.

The input to our process is the formulation of a business objective, concept, idea, or hypothesis, and it starts when we accept the work in Development, adding it to our committed backlog of work.

From there, Development teams that follow a typical Agile or iterative process will likely transform that idea into user stories and some sort of feature specification, which is then implemented in code into the application or service being built. The code is then checked into the version control repository, where each change is integrated and tested with the rest of the software system.

Because value is created only when our services are running in production, we must ensure that we are not only delivering fast flow, but that our deployments can also be

performed without causing chaos and disruptions, such as service outages, service impairments, or security or compliance failures.

Focus on Deployment Lead Time

For the remainder of this book, our attention will be on deployment lead time, a subset of the value stream described above. This value stream begins when any engineer* in our value stream (which includes Development, QA, IT Operations, and Infosec) checks a change into version control and ends when that change is successfully running in production, providing value to the customer and generating useful feedback and telemetry.

The first phase of work that includes design and development is akin to Lean Product Development and is highly variable and highly uncertain, often requiring high degrees of creativity and work that may never be performed again. Because of this, we expect high variability of process times. In contrast, the second phase of work, which includes testing, deployment, and operations, is akin to Lean Manufacturing. It strives to be predictable and mechanistic, with the goal of achieving work outputs with minimized variability (e.g., short and predictable lead times, near zero defects).

Instead of large batches of work being processed sequentially through the design/development value stream and then through the test/operations value stream (such as when we have a large-batch waterfall process or long-lived feature branches), our goal is to have testing, deployment, and operations happening simultaneously with design/development, enabling fast flow and high quality. This method succeeds when we work in small batches and build quality into every part of our value stream.†

Defining Lead Time vs. Processing Time

In the Lean community, lead time is one of two measures commonly used to measure performance in value streams, with the other being processing time (sometimes known as touch time or task time).[‡]

Whereas the lead time clock starts when the request is made and ends when it is fulfilled, the process time clock only starts when we begin work on the customer request—specifically, it omits the time that the work is in queue, waiting to be processed (Figure 1.1).

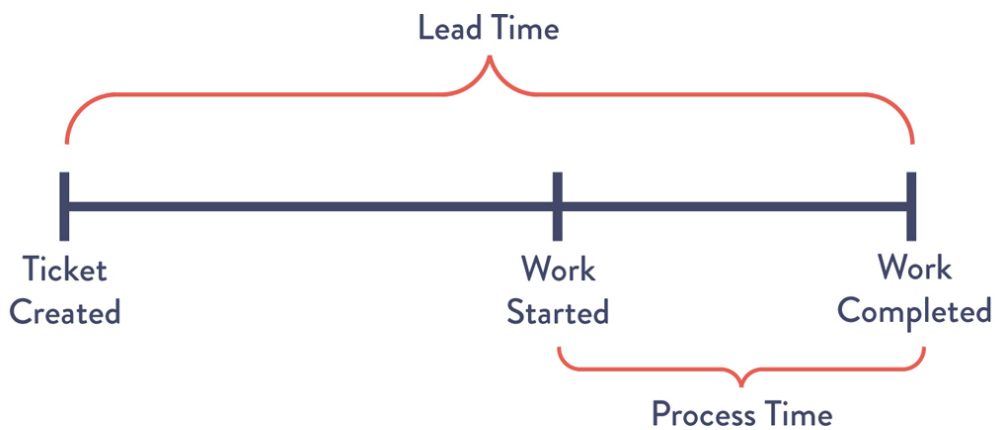


Figure 1.1: Lead Time vs. Process Time of a Deployment Operation

Because lead time is what the customer experiences, we typically focus our process improvement attention there instead of on process time. However, the proportion of process time to lead time serves as an important measure of efficiency—achieving fast flow and short lead times almost always requires reducing the time our work is waiting in queues.

The Common Scenario: Deployment Lead Times Requiring Months

Many teams and organizations find themselves in situations where deployment lead times require months. This is especially common in large, complex organizations that are

working with tightly coupled, monolithic systems, often with scarce integration test environments, long test and production environment lead times, high reliance on manual testing, and multiple required approval processes. When this occurs, our value stream may look like [Figure 1.2](#):




Figure 1.2: A Technology Value Stream with a Deployment Lead Time of Three Months

Source: Damon Edwards, “DevOps Kaizen,” 2015.

When we have long deployment lead times, heroics are required at almost every stage of the value stream. We may discover that nothing works at the end of the project when we merge all the Development team’s changes together, resulting in code that no longer builds correctly or passes any of our tests. Fixing each problem requires days or weeks of investigation to determine who broke the code and how it can be fixed, and still results in poor customer outcomes.

Our DevOps Ideal: Deployment Lead Times of Minutes

In the DevOps ideal, developers receive fast, constant feedback on their work, enabling them to quickly and independently implement, integrate, and validate their code and have the code deployed into the production environment (either by deploying the code themselves or by others).

We achieve this by continually checking small code changes into our version control repository, performing automated and exploratory testing against it and deploying it into production. This enables us to have a high degree of confidence that our changes will operate as designed in

production and that any problems can be quickly detected and corrected.

This is most easily achieved when we have architecture that is modular, well encapsulated, and loosely coupled so that small teams are able to work with high degrees of autonomy, with failures being small and contained, and without causing global disruptions.

In this scenario, our deployment lead time is measured in minutes, or, in the worst case, hours. Our resulting value stream map should look something like [Figure 1.3](#).

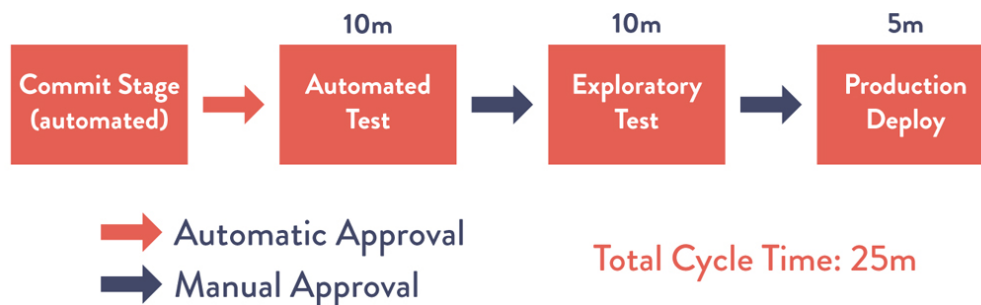


Figure 1.3: A Technology Value Stream with a Lead Time of Minutes

Observing “%C/A” as a Measure of Rework

In addition to lead times and process times, the third key metric in the technology value stream is percent complete and accurate (%C/A). This metric reflects the quality of the output of each step in our value stream.

Karen Martin and Mike Osterling state that “the %C/A can be obtained by asking downstream customers what percentage of the time they receive work that is ‘usable as is,’ meaning that they can do their work without having to correct the information that was provided, add missing information that should have been supplied, or clarify information that should have and could have been clearer.”³

CONTINUOUS LEARNING

Flow Metrics to Measure Delivery of Business Value

When measuring the end-to-end value of any value stream it is important to stay away from proxy metrics (counting the number of lines of code committed or solely the frequency of deployments). While these metrics can reveal local optimizations, they don't directly link to business outcomes such as revenue.

Using flow metrics provides a look into the end-to-end value of your software delivery, making software products and value streams as visible as widgets on a production line. In his book *Project to Product: How to Survive and Thrive in the Age of Digital Disruption with the Flow Framework*, Dr. Mik Kersten describes flow metrics as: flow velocity, flow efficiency, flow time, flow load, and flow distribution:⁴

- **Flow velocity:** number of flow items (e.g., work items) that are completed in a set time period. Helps to answer whether value delivery is accelerating.
- **Flow efficiency:** the proportion of flow items actively worked on to the total time that has elapsed. Identifies inefficiencies like long wait times and helps teams see if work upstream is in a wait state or not.
- **Flow time:** a unit of business value pulled by a stakeholder through a product's value

stream (i.e., features, defects, risks, and debts). Helps teams see if time to value is getting shorter.

- **Flow load:** number of active or waiting flow items in a value stream. This is similar to a measure of work in progress (WIP) based on flow items. High flow load leads to inefficiencies and to reduced flow velocity or increased flow time. Helps teams see if demand is outweighing capacity.
- **Flow distribution:** the proportion of each flow item type in a value stream. Each value stream can track and adjust these depending on their needs in order to maximize the business value being delivered.

The Three Ways: The Principles Underpinning DevOps

The Phoenix Project: A Novel about IT, DevOps, and Helping Your Business Win presents the Three Ways as the set of underpinning principles from which all the observed DevOps behaviors and patterns are derived ([Figure 1.4](#)).




Figure 1.4: The Three Ways

Source: Gene Kim, “The Three Ways: The Principles Underpinning DevOps,” [ITRevolution.com](http://itrevolution.com/the-three-ways-principles-underpinning-devops/) (blog), August 22, 2012, <http://itrevolution.com/the-three-ways-principles-underpinning-devops/>.

The First Way enables fast left-to-right flow of work from Development to Operations to the customer. In order to maximize flow, we need to make work visible, reduce our batch sizes and intervals of work, build in quality by preventing defects from being passed to downstream work centers, and constantly optimize for global goals.

By speeding up flow through the technology value stream, we reduce the lead time required to fulfill internal or customer requests, especially the time required to deploy code into the production environment. By doing this, we increase the quality of work as well as our throughput and boost our ability to innovate and out-experiment the competition.

The resulting practices include continuous build, integration, test, and deployment processes, creating environments on demand, limiting work in process (WIP),

and building systems and organizations that are safe to change.

The Second Way enables the fast and constant flow of feedback from right to left at all stages of our value stream. It requires that we amplify feedback to prevent problems from happening again, or that we enable faster detection and recovery. By doing this, we create quality at the source and generate or embed knowledge where it is needed—this allows us to create ever-safer systems of work where problems are found and fixed long before a catastrophic failure occurs.

By seeing problems as they occur and swarming them until effective countermeasures are in place, we continually shorten and amplify our feedback loops, a core tenet of virtually all modern process improvement methodologies. This maximizes the opportunities for our organization to learn and improve.

The Third Way enables the creation of a generative, high-trust culture that supports a dynamic, disciplined, and scientific approach to experimentation and risk-taking, facilitating the creation of organizational learning, both from our successes and failures. Furthermore, by continually shortening and amplifying our feedback loops, we create ever-safer systems of work and are better able to take risks and perform experiments that help us learn faster than our competition and win in the marketplace.

As part of the Third Way, we also design our system of work so that we can multiply the effects of new knowledge, transforming local discoveries into global improvements. Regardless of where someone performs work, they do so with the cumulative and collective experience of everyone in the organization and throughout the organization's history.

CONTINUOUS LEARNING

Research Supported: The Three Ways

The Three Ways aren't just a good idea: research has shown that adopting these strategies leads to superior outcomes for both organizations and people.

In a six-year study led by co-author Dr. Nicole Forsgren in the 2014–2019 *State of DevOps Reports*, with Puppet and then DORA and published in the book *Accelerate: The Science of Lean and DevOps*, data shows there are better outcomes by combining capabilities and practices like continuous integration, test, deployment, and working in small batches (the First Way), fast feedback and monitoring (the Second Way), and a generative culture (the Third Way).⁵

The Three Ways help teams become elite performers by shipping software faster and more reliably, helping contribute to their organization's revenue, market share, and customer satisfaction. Elite performers are twice as likely to meet or exceed their organizational performance goals. The Three Ways also improve the well-being of those doing the work. By adopting these practices, the research from the *State of DevOps Reports* shows decreased burnout and deployment pain.⁶

→ CASE STUDY: NEW TO SECOND EDITION

Approaching Cruising Altitude: American Airlines' DevOps Journey (Part 1) (2020)

American Airlines' DevOps journey grew out of a series of questions, the first being simply "What is DevOps?"

"We were really starting at the very bottom, at the very beginning," Maya Leibman, Executive Vice President and Chief Information Officer of American Airlines related at the DevOps Enterprise Summit-London 2020.⁷

To get started, the team did their research but, most importantly, they stopped making excuses. In the beginning of DevOps, most examples were coming from digital-native companies like Netflix and Spotify. It was easy for the team to discount their accomplishments—after all, they were born in the cloud. But as more traditional enterprises, companies like Target, Nordstrom, and Starbucks, got on board, American Airlines knew they didn't have any excuses left.

The team started by,

1. setting concrete goals
2. formalizing their toolchain
3. bringin in coaches and mentors from outside the company
4. experimenting and automating
5. conducting immersive practical training (to learn while they were doing)

All of this was tied to their ultimate goal, which was to deliver value faster.

As Leibman said:

There were so many times when a business counterpart would bring something to the table, a new idea, and they'd say, "Oh this is what we want to do but it's going to take IT six months or a year to get it done." And those experiences just killed me. So the impetus behind this was really "how do we not be the long tent pole." We knew there was a better way of working that would help us achieve that.⁸

Next, they decided what outputs they were going to measure:⁹

- deployment frequency
- deployment cycle time
- change failure rate
- development cycle time
- number of incidents
- mean time to recover (MTTR)

Early successes in value stream mapping helped team members better understand the end-to-end processes of the system and inspired motivation. From these successes, they built energy around how to attack issues and improve them. They also conducted immersive learning opportunities across IT.

These initial successes, learning about DevOps and starting to actually practice some elements of it, led

them to the second big question on their DevOps journey: Finance, friend or foe?

The current finance approval process was cumbersome and lengthy, with months of approval cycles. “I used to describe it as a process that’s designed to make you give up,” said Leibman.¹⁰

The process looked like this:¹¹

- No projects approved without Finance’s involvement.
- Projects were approved but no headcount added to do them (and no other priorities were stopped).
- Requests were given equal scrutiny regardless of size or risk.
- Requests were given equal scrutiny, even if the request was a top corporate priority and there was no question that it was going to be done.
- Projects were often completed before they were approved.

Even Finance knew that the process needed to change, but a lack of trust between Finance and IT caused a block. To help shed light on where the money was being spent and to build trust with Finance, the team undertook a cost mapping exercise and assigned all the costs to their products, including the costs to run them.

After this exercise, the IT team was able to better see where money was actually being invested and question whether that was the best use of it. And Finance was

able to gain the visibility they needed to trust there weren't large amounts of waste.

This visibility built the trust needed for experimentation. Finance took four product teams and gave them a set budget for the year. The teams defined the OKRs and used the budget for the top priorities they felt met those OKRs. This allowed the team to test before rollout and focus on accountability and outcomes, and Finance was able to gain even more visibility.

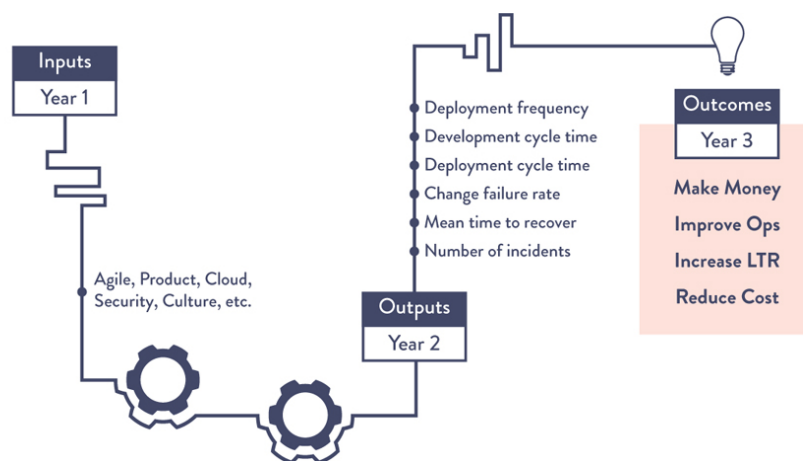


Figure 1.5: American Airlines' DevOps Transformation Journey

Source: With permission of Ross Clanton

This success allowed them to scale the new model against all of their products and define a new funding process. “This was a huge accelerator in our journey,” said Leibman.¹²

With Finance on board and new processes in place, American Airlines discovered the third question in their DevOps journey: How do we know what the score is? With each small success, the team wanted to better understand how they were doing overall. In other words, they wanted to know what the score was.

For the American Airlines team, year one of their DevOps journey was really focused on inputs: learning about Agile/DevOps, focusing on products, cloud, and security, etc. Year two of their journey focused more on outputs, including the metrics they began measuring, like deployment frequency and mean time to recover. Finally in year three they started to focus not just on inputs and outputs but on outcomes. “At the end of the day, what do we really want to do?” said Leibman.

They came up with the following outcomes: make money, improve Ops, increase LTR, and reduce cost.¹³

In year one, one of our objectives was X% of people are going to go to Agile training. That really represents an input. In year two, as we started focusing more on outputs, the objectives sort of changed to X% of teams are going to up their agile maturity from this level to this level. And by the time we got to year three, agile wasn't even an objective anymore. We realized the inputs and outputs are great, we have to measure them, but ultimately we have to be focused on the outcome.¹⁴

This finally led to the fourth question in their DevOps journey: What's a product? It was clear that it was time to flesh out their taxonomy. This proved to be one of the most challenging moments of their journey. There were lots of opinions and no single right answer. In the end, they decided to just get started, put something on paper, organize around it, and fix it as they learned. And ultimately, this all led to their fifth question: Does this feel way bigger than DevOps? To

answer that and to show some specific product success examples, we'll continue the American Airlines journey later in the book.

This study illustrates applying the Three Ways by using value stream mapping to help optimize flow, selecting outcomes to measure in order to establish fast feedback, and creating immersive learning experiences to build a culture of continual learning and experimentation.

Conclusion

In this chapter, we described the concepts of value streams, lead time as one of the key measures of effectiveness for both manufacturing and technology value streams, and the high-level concepts behind each of the Three Ways, the principles that underpin DevOps.

In the following chapters, the principles for each of the Three Ways are described in greater detail. The first of these principles is Flow, which focuses on how we create the fast flow of work in any value stream, whether it's in manufacturing or technology work. The practices that enable fast flow are described in [Part III](#).

* Going forward, *engineer* refers to anyone working in our value stream, not just developers.

† In fact, with techniques such as test-driven development, testing occurs even before the first line of code is written.

‡ In this book, the term *process time* will be favored for the same reason Karen Martin and Mike Osterling cite: “To minimize confusion, we avoid using the term cycle time as it has several definitions synonymous with processing time and pace or frequency of output, to name a few.”²

2

THE FIRST WAY: THE PRINCIPLES OF FLOW

In the technology value stream, work typically flows from Development to Operations, the functional areas between our business and our customers. The First Way requires the fast and smooth flow of work from Development to Operations in order to deliver value to customers quickly. We optimize for this global goal instead of local goals, such as Development feature completion rates, test find/fix ratios, or Operations availability measures.

We increase flow by making work visible, by reducing batch sizes, and by building quality in, preventing defects from being passed to downstream work centers. By speeding up the flow through the technology value stream, we reduce the lead time required to fulfill internal and external customer requests, further increasing the quality of our work while making us more responsive to customer and market needs and able to out-experiment the competition.

Our goal is to decrease the amount of time required for changes to be deployed into production and to increase the reliability and quality of those services. Clues on how we do this in the technology value stream can be gleaned from how Lean principles were applied to the manufacturing value stream.

Make Our Work Visible

A significant difference between technology and manufacturing value streams is that our work is invisible. Unlike physical processes, in the technology value stream we cannot easily see where flow is being impeded or when work

is piling up in front of constrained work centers. Transferring work between work centers in manufacturing is usually highly visible and slow because inventory must be physically moved.

However, in technology work the move can be done with a click of a button, such as by reassigning a work ticket to another team. Because it is so easy to move, work can bounce between teams endlessly due to incomplete information, or work can be passed onto downstream work centers with problems that remain completely invisible until we are late delivering what we promised to the customer or our application fails in the production environment.

To help us see where work is flowing well and where work is queued or stalled, we need to make our work as visible as possible. One of the best methods of doing this is using visual work boards, such as kanban boards or sprint planning boards, where work can be represented on physical or electronic cards. Work originates on the left (often being pulled from a backlog), is pulled from work center to work center (represented in columns), and finishes when it reaches the right side of the board, usually in a column labeled “done” or “in production.”

Not only does our work become visible, but we can also manage our work so that it flows from left to right as quickly as possible. This also helps surface unnecessary handoffs in our work, which can introduce errors and additional delays. Furthermore, we can measure lead time from when a card is placed on the board to when it is moved into the “done” column.

Ideally, our kanban board will span the entire value stream, defining work as completed only when it reaches the right side of the board ([Figure 2.1](#)). Work is not done when

Development completes the implementation of a feature. Rather, it is only done when our application is running successfully in production, delivering value to the customer.

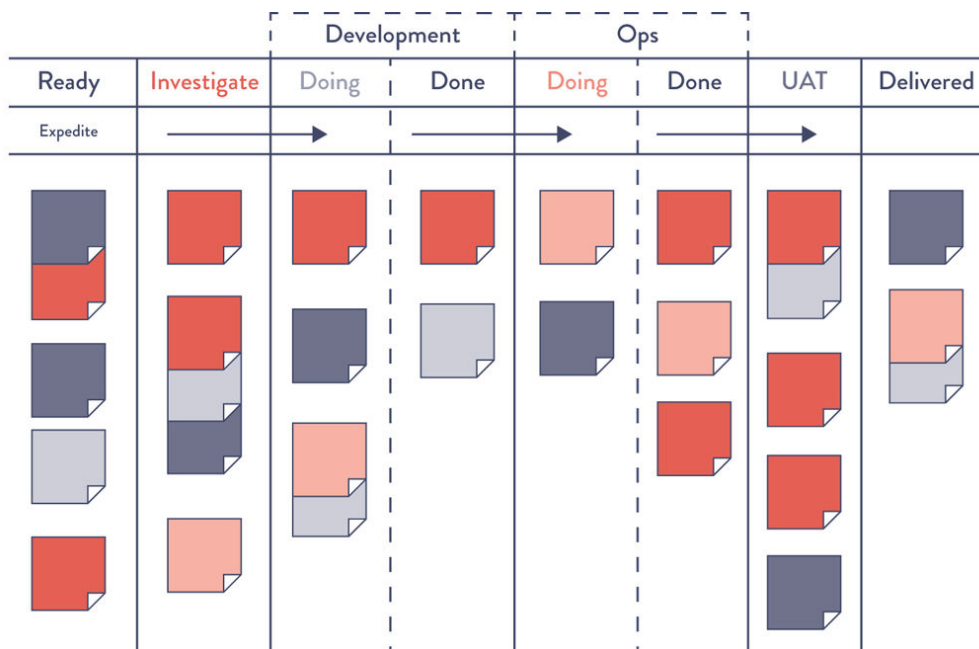


Figure 2.1: An Example Kanban Board Spanning Requirements, Dev, Test, Staging, and In Production

Source: David J. Andersen and Dominica DeGrandis, *Kanban for IT Ops*, training materials for workshop, 2012.

By putting all work for each work center in queues and making it visible, all stakeholders can more easily prioritize work in the context of global goals. Doing this enables each work center to single-task on the highest priority work until it is completed, increasing throughput.

Limit Work in Process (WIP)

In manufacturing, daily work is typically dictated by a production schedule that is generated regularly (e.g., daily, weekly), establishing which jobs must be run based on customer orders, order due dates, parts available, and so forth.

In technology, our work is usually far more dynamic—this is especially the case in shared services, where teams must

satisfy the demands of many different stakeholders. As a result, daily work becomes dominated by the priority du jour, often with requests for urgent work coming in through every communication mechanism possible, including ticketing systems, outage calls, emails, phone calls, chat rooms, and management escalations.

Disruptions in manufacturing are also highly visible and costly. They often require breaking the current job and scrapping any incomplete work in process in order to start the new job. This high level of effort discourages frequent disruptions.

However, interrupting technology workers is easy because the consequences are invisible to almost everyone, even though the negative impact to productivity may be far greater than in manufacturing. For instance, an engineer assigned to multiple projects must switch between tasks, incurring all the costs of having to reestablish context, as well as cognitive rules and goals.

Studies have shown that the time to complete even simple tasks, such as sorting geometric shapes, significantly degrades when multitasking. Of course, because our work in the technology value stream is far more cognitively complex than sorting geometric shapes, the effects of multitasking on process time is much worse.¹

We can limit multitasking when we use a kanban board to manage our work, such as by codifying and enforcing WIP (work in process) limits for each column or work center, that puts an upper limit on the number of cards that can be in a column.

For example, we may set a WIP limit of three cards for testing. When there are already three cards in the test lane, no new cards can be added to the lane unless a card is

completed or removed from the “in work” column and put back into queue (i.e., putting the card back to the column to the left). Nothing can be worked on until it is first represented in a work card, reinforcing that all work must be made visible.

Dominica DeGrandis, one of the leading experts on using kanban in DevOps value streams and author of *Making Work Visible*, notes that “controlling queue size [WIP] is an extremely powerful management tool, as it is one of the few leading indicators of lead time—with most work items, we don’t know how long it will take until it’s actually completed.”²

Limiting WIP also makes it easier to see problems that prevent the completion of work. * For instance, when we limit WIP, we find that we may have nothing to do because we are waiting on someone else. Although it may be tempting to start new work (i.e., “It’s better to be doing something than nothing”), a far better action would be to find out what is causing the delay and help fix that problem. Bad multitasking often occurs when people are assigned to multiple projects, resulting in prioritization problems. In other words, as David J. Anderson, author of *Kanban: Successful Evolutionary Change for Your Technology Business*, said, “Stop starting. Start finishing.”⁴

Reduce Batch Sizes

Another key component to creating smooth and fast flow is performing work in small batch sizes. Prior to the Lean manufacturing revolution, it was common practice to manufacture in large batch sizes (or lot sizes), especially for operations where job setup or switching between jobs was time-consuming or costly. For example, producing large car body panels requires setting large and heavy dies onto metal

stamping machines, a process that can take days. When changeover cost is so expensive, we often stamp as many panels at a time as possible, creating large batches in order to reduce the number of changeovers.

However, large batch sizes result in skyrocketing levels of WIP and high levels of variability in flow that cascade through the entire manufacturing plant. The results are long lead times and poor quality—if a problem is found in one body panel, the entire batch has to be scrapped.

One of the key lessons in Lean is that in order to shrink lead times and increase quality, we must strive to continually shrink batch sizes. The theoretical lower limit for batch size is single-piece flow, where each operation is performed one unit at a time.[†]

The dramatic differences between large and small batch sizes can be seen in the simple newsletter mailing simulation described in *Lean Thinking: Banish Waste and Create Wealth in Your Corporation* by James P. Womack and Daniel T. Jones.⁵

Suppose in our own example we have ten brochures to send, and mailing each brochure requires four steps: (1) fold the paper, (2) insert the paper into the envelope, (3) seal the envelope, and (4) stamp the envelope.

The large batch strategy (i.e., “mass production”) would be to sequentially perform one operation on each of the ten brochures. In other words, we would first fold all ten sheets of paper, then insert each of them into envelopes, then seal all ten envelopes, and then stamp them.

On the other hand, in the small batch strategy (i.e., “single-piece flow”), all the steps required to complete each brochure are performed sequentially before starting on the next brochure. In other words, we fold one sheet of paper,

insert it into the envelope, seal it, and stamp it—only then do we start the process over with the next sheet of paper.

The difference between using large and small batch sizes is dramatic (see [Figure 2.2](#) on page 24). Suppose each of the four operations takes ten seconds for each of the ten envelopes. With the large batch size strategy, the first completed and stamped envelope is produced only after 310 seconds.

Large Batches



Single-Piece Flow



Figure 2.2: Simulation of “Envelope Game”

(Fold, insert, seal, and stamp the envelope.)

Source: Stefan Luyten, “Single Piece Flow,” [Medium.com](https://medium.com/@stefanluyten/single-piece-flow-5d2c2bec845b), August 8, 2014, <https://medium.com/@stefanluyten/single-piece-flow-5d2c2bec845b>.

Worse, suppose we discover during the envelope sealing operation that we made an error in the first step of folding—in this case, the earliest we would discover the error is at two hundred seconds, and we have to refold and reinsert all ten brochures in our batch again.

In contrast, in the small batch strategy the first completed stamped envelope is produced in only forty seconds, eight times faster than the large batch strategy. And, if we made an error in the first step, we only have to redo the one brochure in our batch. Small batch sizes result in less WIP, faster lead times, faster detection of errors, and less rework.

The negative outcomes associated with large batch sizes are just as relevant to the technology value stream as in manufacturing. Consider when we have an annual schedule for software releases, where an entire year's worth of code that Development has worked on is released to production deployment.

Like in manufacturing, this large batch release creates sudden, high levels of WIP and massive disruptions to all downstream work centers, resulting in poor flow and poor quality outcomes. This validates our common experience that the larger the change going into production, the more difficult the production errors are to diagnose and fix, and the longer they take to remediate.

In a post on *Startup Lessons Learned*, Eric Ries states,

The batch size is the unit at which work-products move between stages in a development [or DevOps] process. For software, the easiest batch to see is code. Every time an engineer checks in code, they are batching up a certain amount of work. There are many techniques for controlling these batches, ranging from the tiny batches needed for continuous deployment to more traditional branch-based development, where all of the code from multiple developers working for weeks or months is batched up and integrated together.⁶

The equivalent to single piece flow in the technology value stream is realized with continuous deployment, where each change committed to version control is integrated, tested, and deployed into production. The practices that enable this are described in [Part IV](#) of this book.

Reduce the Number of Handoffs

In the technology value stream, whenever we have long deployment lead times measured in months, it is often because there are hundreds (or even thousands) of operations required to move our code from version control into the production environment. To transmit code through the value stream requires multiple departments to work on a variety of tasks, including functional testing, integration testing, environment creation, server administration, storage administration, networking, load balancing, and information security.

Each time the work passes from team to team, we require all sorts of communication: requesting, specifying, signaling, coordinating, and often prioritizing, scheduling, deconflicting, testing, and verifying. This may require using different ticketing or project management systems; writing technical specification documents; communicating via meetings, emails, or phone calls; and using file system shares, FTP servers, and Wiki pages.

Each of these steps is a potential queue where work will wait when we rely on resources that are shared between different value streams (e.g., centralized operations). The lead times for these requests are often so long that there is constant escalation to have work performed within the needed timelines.

Even under the best circumstances, some knowledge is inevitably lost with each handoff. With enough handoffs, the work can completely lose the context of the problem being solved or the organizational goal being supported. For instance, a server administrator may see a newly created ticket requesting that user accounts be created, without knowing what application or service the accounts are for, why they need to be created, what all the dependencies are, or

whether the user account creations are actually recurring work.

To mitigate these types of problems, we strive to reduce the number of handoffs, either by automating significant portions of the work, or by building platforms and reorganizing teams so they can self-service builds, testing, and deployments to deliver value to the customer themselves instead of having to be constantly dependent on others. As a result, we increase flow by reducing the amount of time that our work spends waiting in queue, as well as the amount of non-value-added time. (See [Appendix 4](#).)

Continually Identify and Elevate Our Constraints

To reduce lead times and increase throughput, we need to continually identify our system's constraints and improve its work capacity. In *Beyond the Goal*, Dr. Goldratt states, "In any value stream, there is always a direction of flow, and there is always one and only constraint; any improvement not made at that constraint is an illusion."⁷ If we improve a work center that is positioned before the constraint, work will merely pile up at the bottleneck even faster, waiting for work to be performed by the bottlenecked work center.

On the other hand, if we improve a work center positioned *after* the bottleneck, it remains starved, waiting for work to clear the bottleneck. As a solution, Dr. Goldratt defined the "five focusing steps":⁸

- Identify the system's constraint.
- Decide how to exploit the system's constraint.
- Subordinate everything else to the above decisions.
- Elevate the system's constraint.

- If a constraint has been broken in the previous steps, go back to step one but do not allow inertia to cause a system constraint.

In typical DevOps transformations, as we progress from deployment lead times measured in months or quarters to lead times measured in minutes, the constraint usually follows this progression:

- **Environment creation:** We cannot achieve deployments on demand if we always have to wait weeks or months for production or test environments. The countermeasure is to create environments that are on-demand and completely self-serviced, so that they are always available when we need them.
- **Code deployment:** We cannot achieve deployments on demand if each of our production code deployments takes weeks or months to perform (e.g., each deployment requires 1,300 manual, error-prone steps, involving up to three hundred engineers). The countermeasure is to automate our deployments as much as possible, with the goal of being completely automated so deployments can be done self-service by any developer.
- **Test setup and run:** We cannot achieve deployments on demand if every code deployment requires two weeks to set up our test environments and data sets and another four weeks to manually execute all our regression tests. The countermeasure is to automate our tests so we can execute deployments safely and to parallelize them so the test rate can keep up with our code development rate.
- **Overly tight architecture:** We cannot achieve deployments on demand if overly tight architecture

means that every time we want to make a code change we have to send our engineers to scores of committee meetings in order to get permission to make our changes. Our countermeasure is to create more loosely coupled architecture so that changes can be made safely and with more autonomy, increasing developer productivity.

After all these constraints have been broken, our constraint will likely be Development or the product owners. Because our goal is to enable small teams of developers to independently develop, test, and deploy value to customers quickly and reliably, this is where we want our constraint to be. High performers, regardless of whether an engineer is in Development, QA, Operations, or Infosec, state that their goal is to help maximize developer productivity.

When the constraint is here, we are limited only by the number of good business hypotheses we create and our ability to develop the code necessary to test these hypotheses with real customers.

The progression of constraints listed above are generalizations of typical transformations—techniques to identify the constraint in actual value streams, such as through value stream mapping and measurements, are described later in this book.

Eliminate Hardships and Waste in the Value Stream

Shigeo Shingo, one of the pioneers of the Toyota Production System, believed that waste constituted the largest threat to business viability—the commonly used definition in Lean is “the use of any material or resource beyond what the customer requires and is willing to pay for.”⁹ He defined seven major types of manufacturing waste: inventory,

overproduction, extra processing, transportation, waiting, motion, and defects.

More modern interpretations of Lean have noted that “eliminating waste” can have a demeaning and dehumanizing context; instead, the goal is reframed to reduce hardship and drudgery in our daily work through continual learning in order to achieve the organization’s goals. For the remainder of this book, the term *waste* will imply this more modern definition, as it more closely matches the DevOps ideals and desired outcomes.

In the book *Implementing Lean Software Development: From Concept to Cash*, Mary and Tom Poppendieck describe waste and hardship in the software development stream as anything that causes delay for the customer, such as activities that can be bypassed without affecting the result.¹⁰ Mary and Tom Poppendieck listed the following seven categories of waste and hardship:¹¹

- **Partially done work:** This includes any work in the value stream that has not been completed (e.g., requirement documents or change orders not yet reviewed) and work that is sitting in queue (e.g., waiting for QA review or server admin ticket). Partially done work becomes obsolete and loses value as time progresses.
- **Extra processes:** Any additional work being performed in a process that does not add value to the customer. This may include documentation not used in a downstream work center, or reviews or approvals that do not add value to the output. Extra processes add effort and increase lead times.
- **Extra features:** Features built into the service that are not needed by the organization or the customer (e.g.,

“gold plating”). Extra features add complexity and effort to testing and managing functionality.

- **Task switching:** When people are assigned to multiple projects and value streams, requiring them to context switch and manage dependencies between work, adding additional effort and time into the value stream.
- **Waiting:** Any delays between work requiring resources to wait until they can complete the current work. Delays increase cycle time and prevent the customer from getting value.
- **Motion:** The amount of effort to move information or materials from one work center to another. Motion waste can be created when people who need to communicate frequently are not colocated. Handoffs also create motion waste and often require additional communication to resolve ambiguities.
- **Defects:** Incorrect, missing, or unclear information, materials, or products create waste, as effort is needed to resolve these issues. The longer the time between defect creation and defect detection, the more difficult it is to resolve the defect.

We also add the following two categories of waste from Damon Edwards:¹²

- **Nonstandard or manual work:** Reliance on nonstandard or manual work from others, such as using non-rebuilding servers, test environments, and configurations. Ideally, any manual work that can be automated should be automated, self-serviced, and available on demand. However, some types of manual work will likely always be essential.

- **Heroics:** In order for an organization to achieve goals, individuals and teams are put in a position where they must perform unreasonable acts, which may even become a part of their daily work (e.g., nightly 2:00 AM problems in production, creating hundreds of work tickets as part of every software release).

Our goal is to make these wastes and hardships—anywhere heroics become necessary—visible, and to systematically do what is needed to alleviate or eliminate these burdens and hardships to achieve our goal of fast flow.

→ CASE STUDY: NEW TO SECOND EDITION

Flow and Constraint Management in Healthcare (2021)

DevOps and constraint management theories aren't just for software development or physical manufacturing. They can be applied to nearly any situation. Just look at this case study from the healthcare industry. At the DevOps Enterprise Summit 2021, Dr. Chris Strear, an emergency physician for more than nineteen years, related his experience improving patient outcomes by working with flow.¹³

Around 2007, our hospital was struggling. We had unbelievable problems with flow. We were boarding patients in the emergency department for hours and hours, and sometimes days, while they waited for an inpatient bed to become available.

Our hospital was so crowded and flow was so backed up that our emergency department was on ambulance diversion for sixty hours a month on average. Now that

means that for sixty hours a month, our emergency department was closed to the sickest patients in our community. One month we hit over two hundred hours of diversion.

It was horrible. We couldn't keep nurses. It was such a hard place to work that nurses would quit. And we relied on temporary nurses, on agencies for placing nurses, or traveler nurses to fill in the gaps in staffing. For the most part, these nurses weren't experienced enough to work in the kind of emergency setting where we practiced. It felt dangerous to come to work every day. It felt dangerous to take care of patients. We were just waiting around for something bad to happen.

The president of our hospital recognized how bad things were, and she put together a committee for flow, and I was lucky enough to be on that committee... .

[The change] was transformative. Within a year, we had basically eliminated ambulance diversion. We went from sixty hours a month [of ambulance diversion] to forty-five minutes a month. We improved the length of stay of all of our admitted patients in the hospital. We shortened the time patients spent in the emergency department. We virtually eliminated the patients who left the department without being seen because the waits were too long. And we did all of this in a

time when we had record volumes, record ambulance traffic, and record admissions.

[The transformation] was amazing. We took better care of patients. It was safer. And it felt so much easier to take care of patients. It was such an amazing turnaround, in fact, that we were able to stop hiring temporary nurses. We were able to fill our staff completely with dedicated emergency nurses who were qualified to work there. In fact, our department became the number one place for emergency nurses to want to work in the Portland/Vancouver area.

Honestly, I'd never been a part of anything that amazing before, and I haven't been since. We made patient care better for tens of thousands of patients, and we made life better for hundreds of healthcare workers in our hospital.¹⁴

So how did they manage this turnaround? Sometime before, Chris had been introduced to the book *The Goal*. Constraint management had a profound influence on him and the way he tackled the problem of flow at his hospital.

So I get asked a bunch of times, what's the difference? I don't have all the answers, but I've seen some trends. I've seen some recurrent themes. Flow needs to be important to leaders, not just in words, but in deeds. They need to walk the walk and not just talk the talk. And a lot of them don't do that.

Part of that is they need to create the bandwidth. The hospital leaders aren't going to be the ones who are actually going to be making the changes day to day. What they have to do is, they have to allow the people who are going to be making those changes to have enough room on their plate to put in the work. If a nurse manager, for instance, has fifteen projects, fifteen committee meetings that they have to go to day in and day out, and the leader comes along and says, "Flow's important," but now flow is their sixteenth task and the sixteenth meeting that they have to go to, really, it doesn't say that it's important. All it says is it's sixteenth most important.

And then there's managers that aren't going to have time to put in for the sixteenth project. Leaders need to figure out what really is important and what can wait, what can take a back burner, and then take an active role in clearing some of that work off of people's plates so that they can do a job. It doesn't just make those people who have to do the work more effective; it conveys to them in a very real, tangible, palpable sense that this new project, flow, is the most important task.

You have to break down silos. You're looking at flow through a system. You're not looking at flow on an inpatient unit or flow just in the emergency department, because each of these departments, when taken

individually, has competing interests. When you move a patient out of the emergency department and onto an inpatient unit, you're creating work for the inpatient unit. You incentivize people differently throughout the hospital.

When you're discussing how to make flow better, and somebody says no, it can't just stop at that. No can't be the final word. I heard time and time again, "We can't do that because that's not how we've done things." And that's ridiculous. No is okay, as long as it's followed up with another idea to try. Because if I have a lousy idea, but it's the only idea out there, then you know what? My lousy idea is the best idea we got going, and so that's the one we try.

Leaders need to make sure that they're measuring things correctly and that they're rewarding things thoughtfully. And what do I mean by that? Well, part of silos in a hospital setting is that a manager for a particular department is often measured on how things go just in that department. And they're rewarded accordingly. People behave based on how they're measured and how they're rewarded. So if improving flow in the emergency department is what's right for patients and what's right for the hospital system, but it may shift burden onto another unit, and that other unit then falls off in their metrics, that should be okay because flow

through the hospital is improved. Who cares about flow through an individual unit?

Make sure that what you're measuring is commensurate with what your overall goals are. Make sure people are rewarded appropriately, and they're not unfairly penalized for improving flow through the system. You need to think about the system, not about the department.

And finally, how we've set things up, that's all artificial, that's a constraint. It's not a natural law of physics. Keep that in mind because so much resistance comes from the uncertainty of doing something differently.

There's often this mindset that because we haven't done something a certain way before, it can't be done. But we've made all of this stuff up. How a body responds to a treatment, that's not artificial, that is a natural law. But where you put a patient, who's in charge of them, how you move a patient from one unit to another, we all just made that up and then perpetuated it. That's all negotiable.¹⁵

This case study concretely illustrates using Goldratt's Theory of Constraint and his five focusing steps to identify and illuminate that constraint and thus improve flow. In this example, the flow of people through the hospital system shows that this theory can be applied to any environment, not just manufacturing or software development.

Conclusion

Improving flow through the technology value stream is essential to achieving DevOps outcomes. We do this by making work visible, limiting WIP, reducing batch sizes and the number of handoffs, continually identifying and evaluating our constraints, and eliminating hardships in our daily work.

The specific practices that enable fast flow in the DevOps value stream are presented in [Part IV](#) of this book. In the next chapter, we present *The Second Way: The Principles of Feedback*.

* Taiichi Ohno compared enforcing WIP limits to draining water from the river of inventory in order to reveal all the problems that obstruct fast flow.³

† Also known as “batch size of one” or “1x1 flow,” terms that refer to batch size and a WIP limit of one.

3

THE SECOND WAY: THE PRINCIPLES OF FEEDBACK

While the First Way describes the principles that enable the fast flow of work from left to right, the Second Way describes the principles that enable the reciprocal fast and constant feedback from right to left at all stages of the value stream. Our goal is to create an ever-safer and more resilient system of work.

This is especially important when working in complex systems, when catastrophic outcomes, such as a manufacturing worker being hurt on the job or a nuclear reactor meltdown in progress, can result from errors.

In technology, our work happens almost entirely within complex systems with a high risk of catastrophic consequences. As in manufacturing, we often discover problems only when large failures are underway, such as a massive production outage or a security breach resulting in the theft of customer data.

We make our systems of work safer by creating fast, frequent, high-quality information flow throughout our value stream and our organization, which includes *feedback* and *feedforward* loops. This allows us to detect and remediate problems while they are smaller, cheaper, and easier to fix; to avert problems before they cause catastrophe; and to create organizational learning that we integrate into future work. When failures and accidents occur, we treat them as opportunities for learning, as opposed to causes for punishment and blame. To achieve all of the above, let us first

explore the nature of complex systems and how they can be made safer.

Working Safely within Complex Systems

One of the defining characteristics of a complex system is that it defies any single person's ability to see the system as a whole and understand how all the pieces fit together. Complex systems typically have a high degree of interconnectedness of tightly coupled components and system-level behavior that cannot be explained merely in terms of the behavior of the system components.

Dr. Charles Perrow studied the Three Mile Island partial nuclear meltdown, observing that it was impossible for anyone to understand how the reactor would behave in all circumstances or how it might fail.¹ When a problem was underway in one component, it was difficult to isolate it from other components. Failures quickly cascaded through the paths of least resistance in unpredictable ways.

Dr. Sidney Dekker, who also codified some key elements of safety culture, observed another characteristic of complex systems: doing the same thing twice will not predictably or necessarily lead to the same result.² It is this characteristic that makes static checklists and best practices, while valuable, insufficient to prevent catastrophes from occurring or manage them effectively. (See [Appendix 5](#).)

Therefore, because failure is inherent and inevitable in complex systems, we must design a safe system of work, whether in manufacturing or technology, where we can perform work without fear, confident that most errors will be detected quickly, long before they cause catastrophic outcomes, such as worker injury, product defects, or negative customer impact.

After he decoded the mechanics of the Toyota Production System as part of his doctoral thesis at Harvard Business School, Dr. Steven Spear stated that designing perfectly safe systems is likely beyond our abilities, but we can make it safer to work in complex systems when the four following conditions are met:^{3*}

- Complex work is managed so that problems in design and operations are revealed.
- Problems are swarmed and solved, resulting in quick construction of new knowledge.
- New local knowledge is exploited globally throughout the organization.
- Leaders create other leaders who continually grow these types of capabilities.

Each of these capabilities is required to work safely in a complex system. In the next sections, the first two capabilities and their importance are described, as well as how they have been created in other domains and what practices enable them in the technology value stream. (The third and fourth capabilities are described in [Chapter 4](#).)

See Problems as They Occur

In a safe system of work, we must constantly test our design and operating assumptions. Our goal is to increase information flow in our system from as many areas as possible, sooner, faster, cheaper, and with as much clarity between cause and effect as possible. The more assumptions we can invalidate, the faster we can find and fix problems, increasing our resilience, agility, and ability to learn and innovate.

We do this by creating feedback and feedforward loops into our system of work. Dr. Peter Senge, in his book *The Fifth*

Discipline: The Art & Practice of the Learning Organization, described feedback loops as a critical part of learning organizations and systems thinking.⁵ Feedback and feedforward loops cause effects within a system to reinforce or counteract each other.

In manufacturing, the absence of effective feedback often contributes to major quality and safety problems. In one well-documented case at the General Motors Fremont manufacturing plant, there were no effective procedures in place to detect problems during the assembly process, nor were there explicit procedures on what to do when problems were found. As a result, there were instances of engines being put in backward, cars missing steering wheels or tires, and cars even having to be towed off the assembly line because they wouldn't start.⁶

In contrast, in high-performing manufacturing operations there is fast, frequent, and high-quality information flow throughout the entire value stream—every work operation is measured and monitored, and any defects or significant deviations are quickly found and acted upon by the people doing the work. This is the foundation of what enables quality, safety, and continual learning and improvement.

In the technology value stream, we often get poor outcomes because of the absence of fast feedback. For instance, in a waterfall software project, we may develop code for an entire year and get no feedback on quality until we begin the testing phase—or, worse, when we release our software to customers. When feedback is this delayed and infrequent, it is too slow to enable us to prevent undesirable outcomes.

In contrast, our goal is to create fast feedback and feedforward loops wherever work is performed, at all stages of the technology value stream, encompassing Product Management, Development, QA, Infosec, and Operations. This includes the creation of automated build, integration, and test processes so that we can immediately detect when a change has been introduced that takes us out of a correctly functioning and deployable state.

We also create pervasive telemetry so we can see how all our system components are operating in testing and production environments so that we can quickly detect when they are not operating as expected. Telemetry also allows us to measure whether we are achieving our intended goals and, ideally, is radiated to the entire value stream so we can see how our actions affect other portions of the system as a whole.

Feedback loops not only enable quick detection and recovery of problems but also inform us on how to prevent these problems from occurring again in the future. Doing this increases the quality and safety of our system of work and creates organizational learning.

As Elisabeth Hendrickson, VP of Engineering at Pivotal Software, Inc. and author of *Explore It!: Reduce Risk and Increase Confidence with Exploratory Testing*, said, “When I headed up quality engineering, I described my job as ‘creating feedback cycles.’ Feedback is critical because it is what allows us to steer. We must constantly validate between customer needs, our intentions and our implementations. Testing is merely one sort of feedback.”⁷

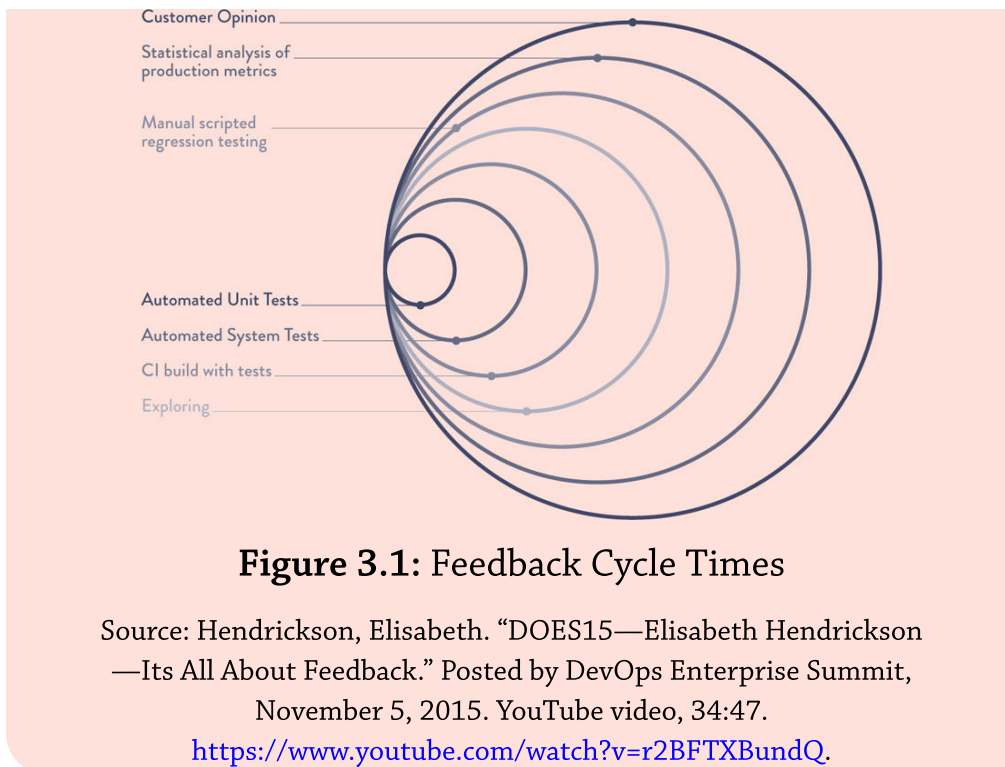
CONTINUOUS LEARNING

Feedback Types and Cycle Times

According to Elisabeth Hendrickson in her 2015 DevOps Enterprise Summit presentation, there are six types of feedback in software development:⁸

- **Dev Tests:** As a programmer, did I write the code I intended to write?
- **Continuous Integration (CI) and Testing:** As a programmer, did I write the code I intended to write without violating any existing expectations in the code?
- **Exploratory Testing:** Did we introduce any unintended consequences?
- **Acceptance Testing:** Did I get the feature I asked for?
- **Stakeholder Feedback:** As a team, are we headed in the right direction?
- **User Feedback:** Are we producing something our customers/users love?

And each type of feedback takes a different amount of time. Think of it as a series of concentric circles. The fastest loops are at the developer's station (local tests, test-driven development, etc.) and the longest are the customer or user feedback at the very end of the cycle (as seen in [Figure 3.1](#))



Swarm and Solve Problems to Build New Knowledge

Obviously, it is not sufficient to merely detect when the unexpected occurs. When problems occur, we must swarm them, mobilizing whoever is required to solve the problem.

According to Dr. Spear, the goal of swarming is to contain problems before they have a chance to spread, and to diagnose and treat the problem so that it cannot recur. “In doing so,” he says, “they build ever-deeper knowledge about how to manage the systems for doing our work, converting inevitable up-front ignorance into knowledge.”⁹

The paragon of this principle is the Toyota *Andon cord*. In a Toyota manufacturing plant, above every work center is a cord that every worker and manager is trained to pull when something goes wrong; e.g., a part is defective, a required part is not available, or even if work takes longer than documented.[†]

When the Andon cord is pulled, the team leader is alerted and immediately works to resolve the problem. If the problem cannot be resolved within a specified time (e.g., fifty-five seconds), the production line is halted so that the entire organization can be mobilized to assist with problem resolution until a successful countermeasure has been developed.

Instead of working around the problem or scheduling a fix “when we have more time,” we swarm to fix it immediately—this is nearly the opposite of the behavior at the GM Fremont plant described earlier.

Swarming is necessary for the following reasons:

- It prevents the problem from progressing downstream, where the cost and effort to repair it increases exponentially and technical debt is allowed to accumulate.
- It prevents the work center from starting new work, which will likely introduce new errors into the system.
- If the problem is not addressed, the work center could potentially have the same problem in the next operation (e.g., fifty-five seconds later), requiring more fixes and work. (See [Appendix 6](#).)

This practice of swarming seems contrary to common management practice, as we are deliberately allowing a local problem to disrupt operations globally. However, swarming enables learning. It prevents the loss of critical information due to fading memories or changing circumstances. This is especially critical in complex systems, where many problems occur because of some unexpected, idiosyncratic interaction of people, processes, products, places, and circumstances—as time passes, it becomes impossible to reconstruct exactly what was going on when the problem occurred.

As Dr. Spear notes, swarming is part of the “disciplined cycle of real-time problem recognition, diagnosis ... and treatment (countermeasures or corrective measures in manufacturing vernacular). It [is] the discipline of the Shewhart cycle—plan, do, check, act—popularized by Dr. W. Edwards Deming, but accelerated to warp speed.”¹⁰

It is only through the swarming of ever-smaller problems discovered ever-earlier in the life cycle that we can deflect problems before a catastrophe occurs. In other words, when the nuclear reactor melts down, it is already too late to avert the worst outcomes.

To enable fast feedback in the technology value stream, we must create the equivalent of an Andon cord and the related swarming response. This requires that we also create the culture that makes it safe, and even encouraged, to pull the Andon cord when something goes wrong, whether it is when a production incident occurs or when errors occur earlier in the value stream, such as when someone introduces a change that breaks our continuous build or test processes.

When conditions trigger an Andon cord pull, we swarm to solve the problem and prevent the introduction of new work until the issue has been resolved.[‡] This provides fast feedback for everyone in the value stream (especially the person who caused the system to fail), enables us to quickly isolate and diagnose the problem, and prevents further complicating factors that can obscure cause and effect.

Preventing the introduction of new work enables continuous integration and deployment, which is “single-piece flow” in the technology value stream. All changes that pass our continuous build and integration tests are deployed into production, and any changes that cause any tests to fail trigger our Andon cord and are swarmed until resolved.

→ CASE STUDY: NEW TO SECOND EDITION

Pulling the Andon Cord at Excella (2018)

Excella is an IT consulting firm. In 2019 at the DevOps Enterprise Summit, Zack Ayers, Scrum Master, and Joshua Cohen, Sr. Developer, discussed their experiments using an Andon cord to decrease cycle time, improve collaboration, and achieve a culture of higher psychological safety.¹¹

Excella noticed during a team retrospective that their cycle times were beginning to rise. They had what Joshua Cohen described as a case of the “almost dones.” He noted, “During standup, our developers would give an update on the feature they were working on the previous day. They would say, ‘Hey, I made a lot of progress. I’m almost done.’ And the next morning they would say, ‘Hey, I ran into some issues but I worked through them. I just have a few more tests to run. I’m almost done.’”¹²

This case of the “almost dones” was happening too frequently. The team decided this was an area they wanted to improve. They noticed that teammates were only bringing up issues at specific times, like during standups. They wanted the team to shift their practice to collaborating as soon as the issue was identified, not waiting until the next day’s standup or meeting.

The team decided to experiment with the idea of an Andon cord. They had two key parameters: (1) When the cord was “pulled,” everyone would stop work to identify a path toward resolution. (2) The cord would be pulled whenever someone on the team felt stuck or needed team's help.

Instead of a literal string or cord to pull, the team created a bot in Slack as a metaphorical Andon cord. When someone typed in *andon*, the bot would @here the team, notifying everyone in Slack. But they didn't want to end it there. They also created an "if/this/then/that" integration in Slack that would turn on a rotating red light, string lights, and even a dancing "tube" man in the office.

To measure their Andon cord experiment, the team decided to focus on reduction in cycle time as the key metric for success, as well as increasing the team's collaboration and getting rid of the "almost dones" by talking about issues when they arose.

In the beginning of the experiment in 2018, their cycle time hovered around three days in progress. Over the following weeks, as the Andon cord began to be pulled, they saw a slight decrease in cycle time. A few weeks later, they stopped pulling the Andon cord and saw their cycle time rise to nearly eleven days, an all-time high.¹³

They evaluated what was happening with the experiment. They realized that while pulling the cord was fun, they weren't pulling it often enough because people were afraid to ask for help and they didn't want to disturb their teammates.

In order to alleviate this, they changed how they defined when teammates should pull the Andon cord. Instead of the cord being pulled whenever a team member was stuck, they would pull the cord whenever they needed the *opinion of the team*.

With this change, they saw a huge uptick in the number of Andon cord pulls and a corresponding

decrease in cycle time.

Each time the team saw Andon cord pulls drop, they found new ways to incentivize pulls, and each time they saw their cycle times decrease with increased pulls. The team continued to iterate and eventually moved the Andon cord experiment into a practice and finally scaled it product wide, using “Andon: Code Red” to report major issues.




Figure 3.2: Cycle Time vs. Andon Pulls at Excella

Source: Zach Ayers and Joshua Cohen. “Andon Cords in Development Teams—Driving Continuous Learning,” presentation at the DevOps Enterprise Summit Las Vegas, 2019.

<https://videolibrary.doesvirtual.com/?video=504281981>.

In addition to cycle time, they found that the Andon cord promoted psychological safety. Teammates spoke up more and offered more creative solutions.

As Jeff Gallimore, Chief Technology and Innovation Officer and Cofounder of Excella, explains:

One of the counterintuitive learnings from this experiment was it challenged the generally held belief that, for developers and engineers in particular, you shouldn’t interrupt flow because it hurts individual

productivity. However, that's exactly what the Andon cord does, for the benefit of team flow and productivity.¹⁴

This case study highlights the amazing effectiveness of swarming to solve problems before a local issue can become a global issue, and how the creative integration of an Andon cord system helps decrease cycle times and improve collaboration.

Keep Pushing Quality Closer to the Source

We may inadvertently perpetuate unsafe systems of work due to the way we respond to accidents and incidents. In complex systems, adding more inspection steps and approval processes actually increases the likelihood of future failures. The effectiveness of approval processes decreases as we push decision-making further away from where the work is performed. Doing so not only lowers the quality of decisions but also increases our cycle time, thus decreasing the strength of the feedback between cause and effect and reducing our ability to learn from successes and failures.[§]

This can be seen even in smaller and less complex systems. When top-down, bureaucratic, command-and-control systems become ineffective, it is usually because the variance between “who should do something” and “who is actually doing something” is too large, due to insufficient clarity and timeliness.

Examples of ineffective quality controls, per *Lean Enterprise*, include:¹⁶

- Requiring another team to complete tedious, error-prone, and manual tasks that could be easily

automated and run as needed by the team who needs the work performed.

- Requiring approvals from busy people who are distant from the work, forcing them to make decisions without adequate knowledge of the work or the potential implications, or to merely rubber stamp their approvals.
- Creating large volumes of documentation of questionable detail, which become obsolete shortly after they are written.
- Pushing large batches of work to teams and special committees for approval and processing and then waiting for responses.

Instead, we need everyone in our value stream to find and fix problems in their area of control as part of their daily work. By doing this, we push quality and safety responsibilities and decision-making to where the work is performed, instead of relying on approvals from distant executives.

We use peer reviews of our proposed changes to gain whatever assurance is needed that our changes will operate as designed. We automate as much of the quality checking typically performed by a QA or Information Security department as possible. Instead of developers needing to request or schedule a test to be run, these tests can be performed on demand, enabling developers to quickly test their own code and even deploy those changes into production themselves.

By doing this, we truly make quality everyone's responsibility as opposed to it being the sole responsibility of a separate department. Information security is not just

Information Security's job, just as availability isn't merely the job of Operations.

Having developers share responsibility for the quality of the systems they build not only improves outcomes but also accelerates learning. This is especially important for developers, as they are typically the team that is furthest removed from the customer. As Gary Gruver observes, "It's impossible for a developer to learn anything when someone yells at them for something they broke six months ago—that's why we need to provide feedback to everyone as quickly as possible, in minutes, not months."¹⁷

Enable Optimizing for Downstream Work Centers

In the 1980s, designing for manufacturability principles sought to design parts and processes so that finished goods could be created with the lowest cost, highest quality, and fastest flow. Examples include designing parts that are wildly asymmetrical to prevent them from being put on backward and designing screw fasteners so that they are impossible to over-tighten.

This was a departure from how design was typically done, which focused on the external customers but overlooked internal stakeholders, such as the people performing the manufacturing.

Lean defines two types of customers that we must design for: the external customer (who most likely pays for the service we are delivering) and the internal customer (who receives and processes the work immediately after us). According to Lean, our most important customer is our next step downstream. Optimizing our work for them requires that we have empathy for their problems in order to better identify the design problems that prevent fast and smooth flow.

In the technology value stream, we optimize for downstream work centers by designing for operations, where operational non-functional requirements (e.g., architecture, performance, stability, testability, configurability, and security) are prioritized as highly as user features.

By doing this, we create quality at the source, resulting in a set of codified, non-functional requirements that we can proactively integrate into every service we build.

Conclusion

Creating fast feedback is critical to achieving quality, reliability, and safety in the technology value stream. We do this by seeing problems as they occur, swarming and solving problems to build new knowledge, pushing quality closer to the source, and continually optimizing for downstream work centers.

The specific practices that enable fast flow in the DevOps value stream are presented in [Part IV](#) of this book. In the next chapter, we present the Third Way, the Principles of Continual Learning and Experimentation.

* Dr. Spear extended his work to explain the long-lasting successes of other organizations, such as the Toyota supplier network, Alcoa, and the US Navy's Nuclear Power Propulsion Program.⁴

† In some of its plants, Toyota has moved to using an Andon button.

‡ Astonishingly, when the number of Andon cord pulls drop, plant managers will actually decrease the tolerances to get an increase in the number of Andon cord pulls in order to continue to enable more learnings and improvements and to detect ever-weaker failure signals.

§ In the 1700s, the British government engaged in a spectacular example of top-down, bureaucratic command and control, which proved remarkably ineffective. At the time, Georgia was still a colony, and despite the fact that the British government was three thousand miles away and lacked firsthand knowledge of local land chemistry, rockiness, topography, accessibility to water, and other conditions, it tried to plan Georgia's entire agricultural economy. The results of the

attempt were dismal and left Georgia with the lowest levels of prosperity and population in the thirteen colonies.¹⁵

4

THE THIRD WAY: THE PRINCIPLES OF CONTINUAL LEARNING AND EXPERIMENTATION

While the First Way addresses workflow from left to right and the Second Way addresses the reciprocal fast and constant feedback from right to left, the Third Way focuses on creating a culture of continual learning and experimentation. These are the principles that enable constant creation of individual knowledge, which is then turned into team and organizational knowledge.

In manufacturing operations with systemic quality and safety problems, work is typically rigidly defined and enforced. For instance, in the GM Fremont plant described in the previous chapter, workers had little ability to integrate improvements and learnings into their daily work, with suggestions for improvement “apt to meet a brick wall of indifference.”¹

In these environments, there is also often a culture of fear and low trust, where workers who make mistakes are punished, and those who make suggestions or point out problems are viewed as whistleblowers and troublemakers. When this occurs, leadership is actively suppressing, even punishing, learning and improvement, perpetuating quality and safety problems.

In contrast, high-performing manufacturing operations require and actively promote learning—instead of work being rigidly defined, the system of work is dynamic, with line workers performing experiments in their daily work to

generate new improvements enabled by rigorous standardization of work procedures and documentation of the results.

In the technology value stream, our goal is to create a high-trust culture, reinforcing that we are all lifelong learners who must take smart risks in our daily work. By applying a scientific approach to both process improvement and product development, we learn from our successes and failures, identifying which ideas don't work and reinforcing those that do. Moreover, any local learnings are rapidly turned into global improvements, so that new techniques and practices that improve the technology value stream in one area can be used by the entire organization.

CONTINUOUS LEARNING

Continual learning and experimentation do more than just improve the performance of our systems. These practices also create an inspiring, rewarding workplace where we are excited to work and collaborate with our peers.

Research led by the *State of DevOps Reports* presents compelling findings. For example, employees from organizations that embrace practices from the Third Way are 2.2 times more likely to recommend their team or organization to friends and have higher job satisfaction and lower levels of burnout.

Recent research from McKinsey also reports that culture—which includes psychological safety, collaboration, and practicing continuous improvement—is a key driver of developer velocity and organizational value.²

We reserve time for the improvement of daily work and time to further accelerate and ensure learning. We consistently introduce stress into our systems to force continual improvement. We even simulate and inject failures in our production services under controlled conditions to increase our resilience.

By creating this continual and dynamic system of learning, we enable teams to rapidly and automatically adapt to an ever-changing environment, which ultimately helps us win in the marketplace.

Enabling Organizational Learning and a Safety Culture

When we work within a complex system, by definition it is impossible for us to perfectly predict all the outcomes for any action we take. This is what contributes to unexpected, or even catastrophic, outcomes and accidents in our daily work, even when we take precautions and work carefully.

When these accidents affect our customers, we seek to understand why they happened. The root cause is often deemed to be human error, and the all-too-common management response is to “name, blame, and shame” the person who caused the problem.* And, either subtly or explicitly, management hints that the person guilty of committing the error will be punished. They then create more processes and approvals to prevent the error from happening again.

Dr. Sidney Dekker, who codified several key elements of safety culture and coined the term *just culture*, wrote, “Responses to incidents and accidents that are seen as unjust can impede safety investigations, promote fear rather than mindfulness in people who do safety-critical work, make organizations more bureaucratic rather than more careful, and cultivate professional secrecy, evasion, and self-protection.”³

These issues are especially problematic in the technology value stream—our work is almost always performed within a complex system, and how management chooses to react to failures and accidents leads to a culture of fear, which then makes it unlikely that problems and failure signals are ever reported. The result is that problems remain hidden until a catastrophe occurs.

Dr. Ron Westrum was one of the first to observe the importance of organizational culture on safety and performance. He observed that in healthcare organizations,

the presence of “generative” cultures was one of the top predictors of patient safety.⁴ Dr. Westrum defined three types of culture:^{5†}

Pathological organizations are characterized by large amounts of fear and threat. People often hoard information, withhold it for political reasons, or distort it to make themselves look better. Failure is often hidden.

Bureaucratic organizations are characterized by rules and processes, often to help individual departments maintain their “turf.” Failure is processed through a system of judgment, resulting in either punishment or justice and mercy.

Generative organizations are characterized by actively seeking and sharing information to better enable the organization to achieve its mission. Responsibilities are shared throughout the value stream, and failure results in reflection and genuine inquiry.

Just as Dr. Westrum found in healthcare organizations, a high-trust, generative culture also predicted software delivery and organizational performance in technology value streams.⁶

In the technology value stream, we establish the foundations of a generative culture by striving to create a safe system of work. When accidents and failures occur, instead of looking for human error, we look for how we can redesign the system to prevent the accident from happening again.

Table 4.1: The Westrum Organizational Typology Model

How organizations process information.

Pathological Organization	Bureaucratic Organization	Generative Organization
Information is hidden	Information may be ignored	Information is actively sought
Messengers are “shot”	Messengers are tolerated	Messengers are trained
Responsibilities are shirked	Responsibilities are compartmented	Responsibilities are shared
Bridging between teams is discouraged	Bridging between teams allowed but discouraged	Bridging between teams is rewarded
Failure is covered up	Organization is just; merciful	Failure causes injury
New ideas are crushed	New ideas create problems	New ideas are welcomed

Source: Ron Westrum, “A typology of organisation culture,” *BMJ Quality & Safety* 13, no. 2 (2004), doi:10.1136/qshc.2003.009522.

For instance, we may conduct a blameless post-mortem (also known as a retrospective) after every incident to gain the best understanding of how the accident occurred and agree upon what the best countermeasures are to improve the system, ideally preventing the problem from occurring again and enabling faster detection and recovery.

By doing this, we create organizational learning. As Bethany Macri, an engineer at Etsy who led the creation of the Morgue tool to help with recording of post-mortems,

stated, “By removing blame, you remove fear; by removing fear, you enable honesty; and honesty enables prevention.”⁷

Dr. Spear observes that the result of removing blame and putting organizational learning in its place is that “organizations become ever more self-diagnosing and self-improving, skilled at detecting problems [and] solving them.”⁸

Many of these attributes were also described by Dr. Peter Senge as attributes of learning organizations. In *The Fifth Discipline*, he wrote that these characteristics help customers, ensure quality, create competitive advantage and an energized and committed workforce, and uncover the truth.⁹

Institutionalize the Improvement of Daily Work

Teams are often not able or not willing to improve the processes they operate within. In many organizations, they are not given the capacity or authority to experiment with improvement work and change their processes based on what they discover. The result is not only that they continue to suffer from their current problems, but their suffering also grows worse over time. Mike Rother observed in Toyota Kata that in the absence of improvements, processes don't stay the same—due to chaos and entropy, processes actually degrade over time.¹⁰

In the technology value stream, when we avoid fixing our problems, relying instead on accumulated, daily workarounds, our problems and technical debt accumulate until all we are doing is performing workarounds, trying to avoid disaster, with no cycles left over for doing productive work. This is why Mike Orzen, author of *Lean IT*, observed, “Even more important than daily work is the improvement of daily work.”¹¹

We improve daily work by explicitly reserving time to pay down technical debt, fix defects, and refactor and improve problematic areas of our code and environments. We do this by reserving cycles in each development interval, or by scheduling *kaizen blitzes*, which are periods when engineers self-organize into teams to work on fixing any problem they want.

The result of these practices is that everyone finds and fixes problems in their area of control all the time, as part of their daily work. When we finally fix the daily problems that we've worked around for months (or years), we can eradicate from our systems the less obvious problems. By detecting and responding to these ever-weaker failure signals, we fix problems when it is not only easier and cheaper but also when the consequences are smaller.

Consider the following example that improved workplace safety at Alcoa, an aluminum manufacturer with \$7.8 billion in revenue in 1987. Aluminum manufacturing requires extremely high heat, high pressures, and corrosive chemicals. In 1987, Alcoa had a frightening safety record, with 2% of the ninety thousand employee workforce being injured each year—that's seven injuries per day. When Paul O'Neill started as CEO, his first goal was to have zero injuries to employees, contractors, and visitors.¹²

O'Neill wanted to be notified within twenty-four hours of anyone being injured on the job—not to punish, but to ensure and promote that learnings were being generated and incorporated to create a safer workplace. Over the course of ten years, Alcoa reduced their injury rate by 95%.¹³

The reduction in injury rates allowed Alcoa to focus on smaller problems and weaker failure signals—instead of notifying O'Neill only when injuries occurred, they started

reporting any close calls as well. ‡¹⁴ By doing this, they improved workplace safety over the subsequent twenty years and have one of the most enviable safety records in the industry.

As Dr. Spear writes,

Alcoans gradually stopped working around the difficulties, inconveniences, and impediments they experienced. Coping, fire fighting, and making do were gradually replaced throughout the organization by a dynamic of identifying opportunities for process and product improvement. As those opportunities were identified and the problems were investigated, the pockets of ignorance that they reflected were converted into nuggets of knowledge.¹⁵

This did more than reduce safety incidents; this helped give the company a greater competitive advantage in the market.

Similarly, in the technology value stream, as we make our system of work safer, we find and fix problems from ever-weaker failure signals. For example, we may initially perform blameless post-mortems only for customer-impacting incidents. Over time, we may perform them for lesser team-impacting incidents and near misses as well.

Transform Local Discoveries into Global Improvements

When new learnings are discovered locally, there must also be some mechanism to enable the rest of the organization to use and benefit from that knowledge. In other words, when teams or individuals have experiences that create expertise, our goal is to convert that tacit knowledge (i.e., knowledge that is difficult to transfer to another person by means of writing it down or verbalizing) into explicit, codified

knowledge, which becomes someone else's expertise through practice.

This ensures that when anyone else does similar work, they do so with the cumulative and collective experience of everyone in the organization who has ever done the same work. A remarkable example of turning local knowledge into global knowledge is the US Navy's Nuclear Power Propulsion Program (also known as "NR" for "Naval Reactors"), which has over 5,700 reactor-years of operation without a single reactor-related casualty or escape of radiation.¹⁶

The NR is known for their intense commitment to scripted procedures and standardized work, and the need for incident reports for any departure from procedure or normal operations to accumulate learnings, no matter how minor the failure signal—they constantly update procedures and system designs based on these learnings.

The result is that when a new crew sets out to sea on their first deployment, they and their officers benefit from the collective knowledge of 5,700 accident-free reactor years. Equally impressive is that their own experiences at sea will be added to this collective knowledge, helping future crews safely achieve their own missions.

In the technology value stream, we must create similar mechanisms to create global knowledge, such as making all our blameless post-mortem reports searchable by teams trying to solve similar problems, and by creating shared source code repositories that span the entire organization, where shared code, libraries, and configurations that embody the best collective knowledge of the entire organization can be easily utilized. All these mechanisms help convert individual expertise into artifacts that the rest of the organization can use.

Inject Resilience Patterns into Our Daily Work

Lower-performing manufacturing organizations buffer themselves from disruptions in many ways—in other words, they bulk up or add waste. For instance, to reduce the risk of a work center being idle (due to inventory arriving late, inventory that had to be scrapped, etc.), managers may choose to stockpile more inventory at each work center. However, that inventory buffer also increases WIP, which has all sorts of undesired outcomes, as previously discussed.

Similarly, to reduce the risk of a work center going down due to machinery failure, managers may increase capacity by buying more capital equipment, hiring more people, or even increasing floor space. All these options increase costs.

In contrast, high performers achieve the same results (or better) by improving daily operations, continually introducing tension to elevate performance, as well as engineering more resilience into their system.

Consider a typical experiment at one of Aisin Seiki Global's mattress factories, one of Toyota's top suppliers. Suppose they had two production lines, each capable of producing one hundred units per day. On slow days, they would send all production onto one line, experimenting with ways to increase capacity and identify vulnerabilities in their process, knowing that if overloading the line caused it to fail, they could send all production to the second line.

By relentless and constant experimentation in their daily work, they were able to continually increase capacity, often without adding any new equipment or hiring more people. The emergent pattern that results from these types of improvement rituals not only improves performance but also improves resilience, because the organization is always in a state of tension and change. This process of applying stress to

increase resilience was named *antifragility* by author and risk analyst Dr. Nassim Nicholas Taleb.¹⁷

In the technology value stream, we can introduce the same type of tension into our systems by seeking to always reduce deployment lead times, increase test coverage, decrease test execution times, and even by rearchitecting if necessary to increase developer productivity or increase reliability.

We may also perform *game day* exercises, where we rehearse large scale failures, such as turning off entire data centers. Or we may inject ever-larger scale faults into the production environment (such as the famous Netflix Chaos Monkey, which randomly kills processes and compute servers in production) to ensure that we're as resilient as we want to be.

Leaders Reinforce a Learning Culture

Traditionally, leaders are expected to be responsible for setting objectives, allocating resources for achieving those objectives, and establishing the right combination of incentives. Leaders also establish the emotional tone for the organizations they lead. In other words, leaders lead by “making all the right decisions.”

However, there is significant evidence that shows greatness is not achieved by leaders making all the right decisions—instead, the leader's role is to create the conditions so their team can discover greatness in their daily work. In other words, creating greatness requires both leaders and workers, each of whom are mutually dependent upon each other.

Jim Womack, author of *Gemba Walks*, described the complementary working relationship and mutual respect that

must occur between leaders and frontline workers. According to Womack, this relationship is necessary because neither can solve problems alone—leaders are not close enough to the work, which is required to solve any problem, and frontline workers do not have the broader organizational context or the authority to make changes outside of their area of work.^{18§}

Leaders must elevate the value of learning and disciplined problem-solving. Mike Rother formalized these methods in what he calls the *coaching kata*. The result is one that mirrors the scientific method, where we explicitly state our True North goals, such as “sustain zero accidents” in the case of Alcoa, or “double throughput within a year” in the case of Aisin.¹⁹

These strategic goals then inform the creation of iterative, shorter-term goals, which are cascaded and then executed by establishing target conditions at the value-stream or work-center level (e.g., “reduce lead time by 10% within the next two weeks”).

These target conditions frame the scientific experiment: we explicitly state the problem we are seeking to solve, our hypothesis of how our proposed countermeasure will solve it, our methods for testing that hypothesis, our interpretation of the results, and our use of learnings to inform the next iteration.

The leader helps coach the person conducting the experiment with questions that may include:

- What was your last step and what happened?
- What did you learn?
- What is your condition now?
- What is your next target condition?

- What obstacle are you working on now?
- What is your next step?
- What is your expected outcome?
- When can we check?

This problem-solving approach, in which leaders help workers see and solve problems in their daily work, is at the core of the Toyota Production System, learning organizations, the Improvement Kata, and high-reliability organizations. Mike Rother observes that he sees Toyota “as an organization defined primarily by the unique behavior routines it continually teaches to all its members.”²⁰

In the technology value stream, this scientific approach and iterative method guides all of our internal improvement processes, but also how we perform experiments to ensure that the products we build actually help our internal and external customers achieve their goals.

→ CASE STUDY: NEW TO SECOND EDITION

The Story of Bell Labs (1925)²¹

With a history spanning the development of sound motion pictures and technicolor, the transistor, Unix, electronic switching systems, and beyond, Bell Labs has been a symbol of innovation and continued success for almost a hundred years. With nine Nobel prizes and four Turing Awards, Bell Labs has applied breakthrough concepts to develop products used by nearly every human being on the planet. What was behind the creation of a culture so pervasive that it was seemingly present “in the air” to produce these types of breakthroughs?

Starting in 1925, Bell Labs was created to consolidate the research activities of Bell Systems. And while many of its products improved telecommunication systems, it never limited itself to a narrow field of focus. It was in this atmosphere that Walter Shewhart (who started at Bell Labs) developed his groundbreaking statistical-control concepts and later collaborated on them with W. Edwards Deming to create the Shewhart-Deming PDCA (Plan, Do, Check, Act) continuous improvement cycle. Their work formed the basis for the Toyota Production System.

In *The Idea Factory: Bell Labs and the Great Age of American Innovation*, Jon Gertner talks about Mervin Kelly, who envisioned an “institute of creative technology” where a cross-skilled team across multiple disciplines could openly collaborate and experiment, recognizing that any breakthrough would come from a team rather than a specific individual.²²

This aligns with the concept of *scenius*, a term coined by the pioneering composer Brian Eno. Gene Kim commonly refers to it, and Dr. Mik Kersten discusses it in his book *Project to Product*, as well as in the blog post “Project to Product: From Stories to Scenius.”²³ As Eno is credited as saying: “Scenius stands for the intelligence and the intuition of a whole cultural scene. It is the communal form of the concept of the genius.”²⁴

Gertner explains that it was clear to the researchers and engineers at Bell Labs that the ultimate aim of their organization was to transform new knowledge into new things.²⁵ In other words, the goal was to transform innovation into the delivery of something of societal value. Bell Labs had a culture that kept it continually

successful because change and challenging the status quo were its hallmarks.

A vital aspect of the culture was that there should be no fear of failures. As Kelly explained, “The odds of creating a new and popular technology were always stacked against the innovator; only where the environment allowed failure could truly groundbreaking ideas be pursued.”²⁶

Even concepts like Chaos Monkey and the SRE model have their roots in Bell Labs’ work in the hardening of telecommunication systems, which achieved five-nines availability by disrupting these systems as part of the normal testing cycle and then ensured their robustness through the automation of recovery actions.

So as we talk today about collaboration using cross-skilled teams, continuous improvement, providing psychological safety, and harnessing the ideas of our teams, take note that these concepts were present in the DNA of how Bell Labs operated. And while many people today may not know which company invented the transistor or provided the brilliant rainbow colors of Oz, the concepts behind these innovations are very much alive almost a century later.

One of the striking characteristics found in this glimpse into Bell Labs’ scenius is their commitment to building a culture that allows teams to discover greatness by engendering collaboration both vertically and horizontally.

Conclusion

The principles of the Third Way address the need for valuing organizational learning, enabling high trust and boundary-spanning between functions, accepting that failures will always occur in complex systems, and making it acceptable to talk about problems so we can create a safe system of work. It also requires institutionalizing the improvement of daily work, converting local learnings into global learnings that can be used by the entire organization, as well as continually injecting tension into our daily work.

Although fostering a culture of continual learning and experimentation is the principle of the Third Way, it is also interwoven into the First and Second Ways. In other words, improving flow and feedback requires an iterative and scientific approach that includes framing of a target condition, stating a hypothesis of what will help us get there, designing and conducting experiments, and evaluating the results. The results are not only better performance but also increased resilience and improved organizational adaptability.

* The “name, blame, shame” pattern is part of the Bad Apple Theory criticized by Dr. Sidney Dekker and extensively discussed in his book *The Field Guide to Understanding Human Error*.

† Dr. Westrum talks more extensively about generative cultures in his interview with Gene Kim on *The Idealcast* podcast.

‡ It is astonishing, instructional, and truly moving to see the level of conviction and passion that Paul O’Neill has about the moral responsibility leaders have to create workplace safety.

§ Leaders are responsible for the design and operation of processes at a higher level of aggregation where others have less perspective and authority.

PART I CONCLUSION

In Part I of *The DevOps Handbook*, we looked back at several movements in history that helped lead to the development of DevOps. We also looked at the three main principles that form the foundation for successful DevOps organizations: the principles of Flow, Feedback, and Continual Learning and Experimentation. In [Part II](#), we will begin to look at how to start a DevOps movement in your organization.

Additional Resources for Part I

Beyond The Phoenix Project by Gene Kim and John Willis is an audio series that chronicles the origins and evolution of DevOps (itrevolution.com/book/beyond-phoenix-project-audiobook/).

The *State of DevOps Reports* from DORA consistently produce excellent metrics and insights into the DevOps community (devops-research.com/research.html).

The *Octoverse Report* from GitHub, like the *State of DevOps Reports*, provides excellent metrics on the state of work in the software industry (devops-research.com/research.html).

Thoughtworks' *Tech Radar* is a great place to find upcoming and emerging tools and trends (thoughtworks.com/radar).

If you're looking to learn more about Ron Westrum's Organization Typology and Generative Cultures, check out his two part interview on *The Idealcast*, hosted by Gene Kim (itrevolution.com/the-idealcast-podcast/).

At the 2017 DevOps Enterprise Summit, Gene Kim conducted a panel on the convergence of Safety Culture and Lean with Sidney Dekker, Dr. Steven Spear, and Dr. Richard Cook (videolibrary.doesvirtual.com/?video=524027004).

The whitepaper *Getting Started with Dojos* is a great resource if you're looking to build a learning culture (itrevolution.com/resources).

The *Measuring Software Quality* whitepaper clearly breaks down some beginning metrics to track quality in software (itrevolution.com/resources).

Making Work Visible: Exposing Time Theft to Optimize Work & Flow by Dominica DeGrandis is an excellent resource to help you bring more transparency into your work (itrevolution.com/making-work-visible-by-dominica-degrandis/).

And you can learn more about the Naval Reactor program in The Idealcast episode featuring Admiral John Richardson (<https://itrevolution.com/the-idealcast-podcast/>).



Part II

WHERE TO START

PART II: INTRODUCTION

How do we decide where to start a DevOps transformation in our organization? Who needs to be involved? How should we organize our teams, protect their work capacity, and maximize their chances of success? These are the questions we aim to answer in Part II of *The DevOps Handbook*.

In the following chapters, we will walk through the process of initiating a DevOps transformation. We'll begin by evaluating the value streams in our organization, locating a good place to start, and forming a strategy to create a dedicated transformation team with specific improvement goals and eventual expansion. For each value stream being transformed, we identify the work being performed and then look at organizational design strategies and organizational archetypes that best support the transformation goals.

Primary focuses in these chapters include:

- selecting which value streams to start with
- understanding the work being done in our candidate value streams
- designing our organization and architecture with Conway's Law in mind
- enabling market-oriented outcomes through more effective collaboration between functions throughout the value stream
- protecting and enabling our teams

Beginning any transformation is full of uncertainty—we are charting a journey to an ideal end state but where virtually all the intermediate steps are unknown. These next chapters are intended to provide a thought process to guide

your decisions, provide actionable steps you can take, and illustrate case studies as examples.

5

SELECTING WHICH VALUE STREAM TO START WITH

Choosing a value stream for DevOps transformation deserves careful consideration. Not only does the value stream we choose dictate the difficulty of our transformation, but it also dictates who will be involved in the transformation. It will affect how we need to organize into teams and how we can best enable the teams and individuals in them.

Another challenge was noted by Michael Rembetsy, who helped lead the DevOps transformation as the Director of Operations at Etsy in 2009. He observed, “We must pick our transformation projects carefully—when we’re in trouble, we don’t get very many shots. Therefore, we must carefully pick and then protect those improvement projects that will most improve the state of our organization.”¹

Nordstrom’s DevOps Transformation

Let us examine how the Nordstrom team started their DevOps transformation initiative in 2013, which Courtney Kissler, their VP of E-Commerce and Store Technologies, described at the DevOps Enterprise Summit in 2014 and 2015.

Founded in 1901, Nordstrom is a leading fashion retailer focused on delivering the best possible shopping experience to its customers. In 2015, Nordstrom had annual revenue of \$13.5 billion.²

The stage for Nordstrom's DevOps journey was likely set in 2011 during one of their annual board of directors meetings.³ That year, one of the strategic topics discussed was the need for online revenue growth. They studied the plight of Blockbuster, Borders, and Barnes & Noble, which demonstrated the dire consequences when traditional retailers were late creating competitive e-commerce capabilities. These organizations were clearly at risk of losing their position in the marketplace or even going out of business entirely.*

At that time, Courtney Kissler was the senior director of Systems Delivery and Selling Technology and was responsible for a significant portion of the technology organization, including their in-store systems and online e-commerce site. As Kissler described,

In 2011, the Nordstrom technology organization was very much optimized for cost—we had outsourced many of our technology functions, we had an annual planning cycle with large batch, “waterfall” software releases. Even though we had a 97% success rate of hitting our schedule, budget, and scope goals, we were ill-equipped to achieve what the five-year business strategy required from us, as Nordstrom started optimizing for speed instead of merely optimizing for cost.⁵

Kissler and the Nordstrom technology management team had to decide where to start their initial transformation efforts. They didn't want to cause upheaval in the whole system. Instead, they wanted to focus on very specific areas of the business so they could experiment and learn. Their goal was to demonstrate early wins, which would give everyone confidence that these improvements could be replicated in

other areas of the organization. How exactly that would be achieved was still unknown.

They focused on three areas: the customer mobile application, their in-store restaurant systems, and their digital properties. Each of these areas had business goals that weren't being met; thus, they were more receptive to considering a different way of working. The stories of the first two focus areas are described below.

The Nordstrom mobile application had experienced an inauspicious start. As Kissler said, "Our customers were extremely frustrated with the product, and we had uniformly negative reviews when we launched it in the App Store."⁶ Worse, the existing structure and processes were designed so that updates could only be released twice per year. In other words, any fixes to the application would have to wait months to reach the customer.

Their first goal was to enable faster or on-demand releases, providing faster iteration and the ability to respond quickly to customer feedback. They created a product team that was solely dedicated to supporting the mobile application. They had the goal of enabling the mobile application team to be able to independently implement, test, and deliver value to customers. By doing this, the mobile application team would no longer have to depend on and coordinate with scores of other teams inside Nordstrom.

Furthermore, the mobile application team moved from planning once per year to a continuous planning process. The result was a single prioritized backlog of work for the mobile app based on customer need—gone were all the conflicting priorities when the team had to support multiple products.

Over the following year, they eliminated testing as a separate phase of work; instead, they integrated it into

everyone's daily work. † They doubled the features being delivered per month and halved the number of defects—creating a successful outcome.

Their second area of focus was the systems supporting their in-store *Café Bistro* restaurants. Unlike the mobile app value stream, where the business need was to reduce time to market and increase feature throughput, the business need here was to decrease cost and increase quality. In 2013, Nordstrom had completed eleven “restaurant re-concepts,” which required changes to the in-store applications, causing a number of customer-impacting incidents. Plus, they had planned forty-four more of these re-concepts for 2014—four times as many as in the previous year.

As Kissler stated, “One of our business leaders suggested that we triple our team size to handle these new demands, but I proposed that we had to stop throwing more bodies at the problem and instead improve the way we worked.”⁷

Kissler's team was able to identify problematic areas, such as their work intake and deployment processes, which is where they focused their improvement efforts. They were able to reduce code deployment lead times by 60% and reduce the number of production incidents 60–90%.

These successes gave the teams confidence that DevOps principles and practices were applicable to a wide variety of value streams. Kissler was promoted to VP of E-Commerce and Store Technologies in 2014.

In 2015, Kissler said that in order for the selling or customer-facing technology organization to enable the business to meet their goals, “... we needed to increase productivity in all our technology value streams, not just in a few. At the management level, we created an across-the-board

mandate to reduce cycle times by 20% for all customer-facing services.”⁸

She continued, “This is an audacious challenge. We have many problems in our current state—process and cycle times are not consistently measured across teams, nor are they visible. Our first target condition requires us to help all our teams measure, make [the work] visible, and perform experiments to start reducing their process times, iteration by iteration.”⁹ Kissler concluded,

From a high-level perspective, we believe that techniques such as value stream mapping, reducing our batch sizes toward single-piece flow, as well as using continuous delivery and microservices will get us to our desired state. However, while we are still learning, we are confident that we are heading in the right direction, and everyone knows that this effort has support from the highest levels of management.¹⁰

In this chapter, various models are presented that will enable us to replicate the thought processes that the Nordstrom team used to decide which value streams to start with. We will evaluate our candidate value streams in many ways, including whether they are a *greenfield* or a *brownfield* service, a *system of engagement* or a *system of record*. We will also estimate the risk/reward balance of transforming and assess the likely level of resistance we may get from the teams we would work with.

Greenfield vs. Brownfield Services

We often categorize our software services or products as either greenfield or brownfield. These terms were originally used for urban planning and building projects. Greenfield development is when we build on undeveloped land.

Brownfield development is when we build on land that was previously used for industrial purposes, potentially contaminated with hazardous waste or pollution. In urban development, many factors can make greenfield projects simpler than brownfield projects—there are no existing structures that need to be demolished nor are there toxic materials that need to be removed.

In technology, a greenfield project is a new software project or initiative, likely in the early stages of planning or implementation, where we build our applications and infrastructure from scratch, with few constraints. Starting with a greenfield software project can be easier, especially if the project is already funded and a team is either being created or is already in place. Furthermore, because we are starting from scratch, we can worry less about existing code bases, architectures, processes, and teams.

Greenfield DevOps projects are often pilots to demonstrate feasibility of public or private clouds, piloting deployment automation, and similar tools. An example of a greenfield DevOps project is the Hosted LabVIEW product in 2009 at National Instruments, a thirty-year-old organization with five thousand employees and \$1 billion in annual revenue.

To bring this product to market quickly, a new team was created. They were allowed to operate outside of the existing IT processes and explore the use of public clouds. The initial team included an applications architect, a systems architect, two developers, a system automation developer, an operations lead, and two offshore operations staff. By using DevOps practices, they were able to deliver Hosted LabVIEW to market in half the time of their normal product introductions.¹¹

On the other end of the spectrum are brownfield DevOps projects. These are existing products or services that are already serving customers and have potentially been in operation for years or even decades. Brownfield projects often come with significant amounts of technical debt, such as having no test automation or running on unsupported platforms. In the Nordstrom example presented earlier in this chapter, both the in-store restaurant systems and e-commerce systems were brownfield projects.

Although many believe that DevOps is primarily for greenfield projects, DevOps has been used to successfully transform brownfield projects of all sorts. In fact, over 60% of the transformation stories shared at the DevOps Enterprise Summit in 2014 were for brownfield projects.¹² In these cases, there was a large performance gap between what the customer needed and what the organization was currently delivering, and the DevOps transformations created tremendous business benefits.

Indeed, research from the *State of DevOps Reports* found that the age of the application or even the technology used was not a significant predictor of performance; instead, what predicted performance was whether the application was architected (or could be re-architected) for testability and deployability.¹³

Teams supporting brownfield projects may be very receptive to experimenting with DevOps, particularly when there is a widespread belief that traditional methods are insufficient to achieve their goals—and especially if there is a strong sense of urgency around the need for improvement.[†]

When transforming brownfield projects, we may face significant impediments and problems, especially when no automated testing exists or when there is a tightly coupled

architecture that prevents small teams from developing, testing, and deploying code independently. How we overcome these issues is discussed throughout this book.

Examples of successful brownfield transformations include:

- **American Airlines (2020):** DevOps practices can also be applied to legacy COTS (commercial off-the-shelf product). At American Airlines, their loyalty product runs on Seibel. AA moved it onto a hybrid cloud model and then invested in CI/CD pipelines to automate delivery and infrastructure end to end for their loyalty product. Since that move, teams are deploying more frequently, with more than fifty automated deployments in just a few months, plus twice as fast loyalty web service response times and 32% cost optimization in the cloud. What's more, this move caused the conversation between the business and IT to change. Instead of the business waiting on IT for changes, the teams are able to deploy more frequently and seamlessly than the business can validate and accept. Now the product teams, in partnership with the business and IT teams, are looking at how to optimize the end-to-end process to achieve higher deployment frequency.¹⁴
- **CSG (2013):** In 2013, CSG International had \$747 million in revenue and over thirty-five hundred employees, enabling over ninety thousand customer service agents to provide billing operations and customer care to over fifty million video, voice, and data customers, executing over six billion transactions, and printing and mailing over seventy million paper bill statements every month. Their initial scope of improvement was bill printing, one of

their primary businesses, and involved a COBOL mainframe application and the twenty surrounding technology platforms. As part of their transformation, they started performing daily deployments into a production-like environment and doubled the frequency of customer releases from twice annually to four times annually. As a result, they significantly increased the reliability of the application and reduced code deployment lead times from two weeks to less than one day.¹⁵

- **Etsy (2009):** In 2009, Etsy had thirty-five employees and was generating \$87 million in revenue, but after they “barely survived the holiday retail season,” they started transforming virtually every aspect of how the organization worked, eventually turning the company into one of the most admired DevOps organizations and setting the stage for a successful 2015 IPO.¹⁶
- **HP LaserJet (2007):** By implementing automated testing and continuous integration, HP created faster feedback and enabled developers to quickly confirm that their command codes actually worked. You can read the full case study in [Chapter 11](#).

→ CASE STUDY: NEW TO SECOND EDITION

Kessel Run: The Brownfield Transformation of a Mid-Air Refueling System (2020)

In October 2015, the US Air Force struck a Doctors Without Borders hospital in Afghanistan, believing it to be an enemy stronghold. Afghan commandos were under fire, and the US needed to respond quickly. Unclassified analysis later showed that a number of

failures contributed to this devastating outcome: there was no time to fully brief the crew and the aircraft didn't have the latest data to identify the hospital. As Adam Furtado, Chief of Platform for Kessel Run, put it, "basically, a failed IT ecosystem caused an AC130 gunship to attack the wrong building."¹⁷

Furtado continued to explain at the DevOps Enterprise Summit-Virtual Las Vegas 2020, "What happened here was not some kind of black swan event, it was predictable and it's going to happen again."¹⁸ They needed a solution.

Named after Han Solo's famous smuggling route, an homage to their own need to "smuggle" these new ways of working into the Department of Defense (DOD), Kessel Run is the continuing effort within the US Air Force to solve the tough business challenges that traditional defense IT isn't solving effectively. Made up of a small coalition, the group tests modern software practices, processes, and principles. Their focus is on the mission and a disregard for the status quo.

In the beginning, around 2010, walking into the DOD to work was like walking into a time machine, a completely analog environment circa 1974 where many collaborative tools, like chats and Google Docs, weren't possible. As Adam Furtado says, "You shouldn't have to go back in time to go to work."¹⁹

Eric Schmidt, Executive Chairman at Google, even testified to the US Congress in September of 2020 that "the DOD violates every rule of modern product development."²⁰

These kinds of problems weren't exclusive to the DOD. According to the US Digital Service, 94% of federal IT projects across the entire government were behind schedule or over budget, and 40% of them were never delivered.²¹

The Kessel Run coalition was watching companies like Adidas and Walmart become software companies. They wanted to transform the US Air Force into a software company that could win wars. So they turned their attention to the critical business outcome in front of them: modernizing the Air Force's Air Operation Center.

There are several physical Air Operation Centers (AOCs) around the world from which the US Air Force strategizes, plans, and executes air campaigns. Due to archaic infrastructure, all this work is done by specific people, in specific buildings, located in specific locations, to access specific data on specific hardware, a brick-and-mortar approach that has been in place for decades. The only updates they've been able to implement are Microsoft Office updates.

"You might think I'm lying, but a recent search showed 2.8 million Excel and PowerPoint files on one of the servers in one of the locations," said Adam Furtado.²²

Gall's Law states that if you want a complex system to work, build a simpler system first, then improve over time. The Kessel Run coalition did just that by applying the Strangler Fig Pattern, otherwise known as the Encasement Strategy, to incrementally and iteratively implement more modern software systems and processes at twenty-two physical locations, each with

their own software and hardware, while keeping the whole system working.

They started with a specific process: mid-air refueling. This process requires massive coordination to ensure refueling tankers are where they need to be, at the right time, at the right altitude, with the right hardware to refuel the correct aircraft. The process involved several pilots to plan every day using color pucks, an Excel macro, and lots of data entry. They had become largely efficient, but only to the point that their brains would allow, and they couldn't react quickly to changes.

Kessel Run brought in a team to digitize their process using DevOps principles, extreme programming, and balanced team models. They got their initial minimally viable product out to users in just weeks. This early program created enough efficiency that it kept one aircraft and its crew from flying every day, a fuel savings of \$214,000 a day.²³

They kept iterating. After thirty iterations, they were able to double the savings, keeping two aircraft and crew on the ground a day. The new approach saved \$13 million in fuel a month and cut the planning crew in half.²⁴

While this case study shows a great DevOps transformation in general, it also wonderfully illustrates transforming brownfield practices. Kessel Run successfully reduced the complexity and improved the reliability and stability of the mid-air refueling system, helping the US Air Force to move and change faster and more safely.

Consider Both Systems of Record and Systems of Engagement

The Gartner global research and advisory firm popularized the notion of *bimodal IT*, referring to the wide spectrum of services that typical enterprises support.²⁵ Within bimodal IT there are *systems of record*, the ERP-like systems that run our business (e.g., MRP, HR, financial reporting systems), where the correctness of the transactions and data are paramount; and *systems of engagement*, which are customer-facing or employee-facing systems, such as e-commerce systems and productivity applications.

Systems of record typically have a slower pace of change and often have regulatory and compliance requirements (e.g., SOX). Gartner calls these types of systems “Type 1,” where the organization focuses on “doing it right.”²⁶

Systems of engagement typically have a much higher pace of change to support rapid feedback loops that enable them to experiment to discover how to best meet customer needs. Gartner calls these types of systems “Type 2,” where the organization focuses on “doing it fast.”²⁷

It may be convenient to divide up our systems into these categories; however, we know that the core, chronic conflict between “doing it right” and “doing it fast” can be broken with DevOps. The data from six years of *State of DevOps Reports* show that high-performing organizations are able to simultaneously deliver higher levels of throughput and reliability.²⁸

Furthermore, because of how interdependent our systems are, our ability to make changes to any of these systems is limited by the system that is most difficult to safely change, which is almost always a system of record.

Scott Prugh, VP of Product Development at CSG, observed, “We’ve adopted a philosophy that rejects bi-modal IT because every one of our customers deserve speed and quality. This means that we need technical excellence, whether the team is supporting a thirty-year-old mainframe application, a Java application, or a mobile application.”²⁹

Consequently, when we improve brownfield systems, we should not only strive to reduce their complexity and improve their reliability and stability, but also make them faster, safer, and easier to change. Even when new functionality is added just to greenfield systems of engagement, they often cause reliability problems in the brownfield systems of record they rely on. By making these downstream systems safer to change, we help the entire organization more quickly and safely achieve its goals.

Start With the Most Sympathetic and Innovative Groups

Within every organization, there will be teams and individuals with a wide range of attitudes toward the adopting new ideas. Geoffrey A. Moore first depicted this spectrum in the form of the technology adoption life cycle in *Crossing the Chasm*, where the chasm represents the classic difficulty of reaching groups beyond the *innovators* and *early adopters* (see [Figure 5.1](#)).³⁰

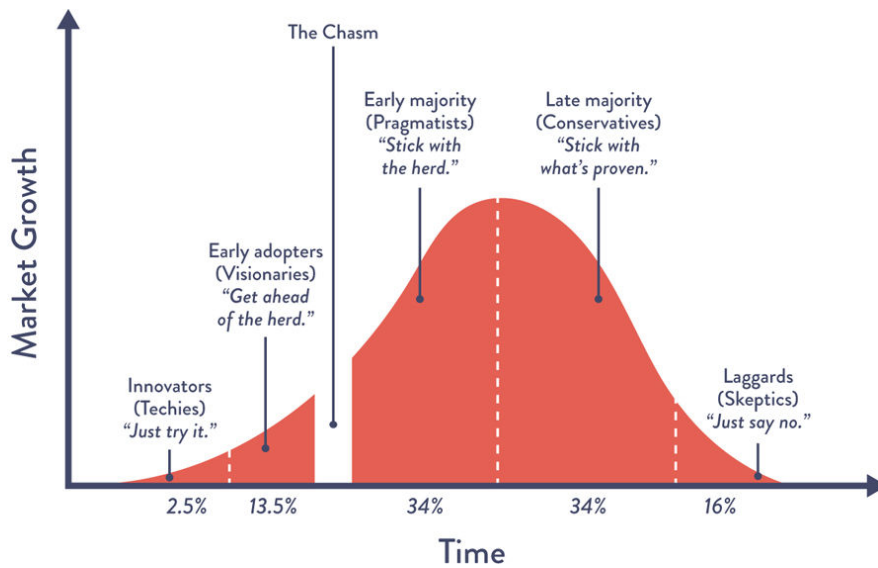


Figure 5.1: The Technology Adoption Curve

Source: Moore and McKenna, *Crossing the Chasm*, 15.

In other words, new ideas are often quickly embraced by innovators and early adopters, while others with more conservative attitudes resist them (the *early majority*, *late majority*, and laggards). Our goal is to find those teams that already believe in the need for DevOps principles and practices and that possess a desire and demonstrated ability to innovate and improve their own processes. Ideally, these groups will be enthusiastic supporters of the DevOps journey.

Especially in the early stages, we will not spend much time trying to convert the more conservative groups. Instead, we will focus our energy on creating successes with less risk-averse groups and build out our base from there (a process that is discussed further in the next section). Even if we have the highest levels of executive sponsorship, we will avoid the *big-bang approach* (i.e., starting everywhere all at once), choosing instead to focus our efforts in a few areas of the organization, ensuring that those initiatives are successful, and expanding from there.[§]

Expanding DevOps Across Our Organization

Regardless of how we scope our initial effort, we must demonstrate early wins and broadcast our successes. We do this by breaking up our larger improvement goals into small, incremental steps. This not only creates our improvements faster, it also enables us to discover when we have made the wrong choice of value stream—by detecting our errors early, we can quickly back up and try again, making different decisions armed with our new learnings.

As we generate successes, we earn the right to expand the scope of our DevOps initiative. We want to follow a safe sequence that methodically grows our levels of credibility, influence, and support. The following list, adapted from a course taught by Dr. Roberto Fernandez, a William F. Pounds Professor in Management at MIT, describes the ideal phases used by change agents to build and expand their coalition and base of support:³¹

- **Find innovators and early adopters:** In the beginning, we focus our efforts on teams who actually want to help—these are our kindred spirits and fellow travelers who are the first to volunteer to start the DevOps journey. In the ideal, these are also people who are respected and have a high degree of influence over the rest of the organization, giving our initiative more credibility.
- **Build critical mass and silent majority:** In the next phase, we seek to expand DevOps practices to more teams and value streams with the goal of creating a stable base of support. By working with teams who are receptive to our ideas, even if they are not the most visible or influential groups, we expand our coalition who are generating more successes, creating a “bandwagon effect” that further increases our

influence. We specifically bypass dangerous political battles that could jeopardize our initiative.

- **Identify the holdouts:** The “holdouts” are the high-profile, influential detractors who are most likely to resist (and maybe even sabotage) our efforts. In general, we tackle this group only after we have achieved a silent majority and have established enough successes to successfully protect our initiative.

Expanding DevOps across an organization is no small task. It can create risk to individuals, departments, and the organization as a whole. But as Ron van Kemenade, CIO of ING, who helped transform the organization into one of the most admired technology organizations, said, “Leading change requires courage, especially in corporate environments where people are scared and fight you. But if you start small, you really have nothing to fear. Any leader needs to be brave enough to allocate teams to do some calculated risk-taking.”³²

→ CASE STUDY: NEW TO SECOND EDITION

Scaling DevOps across the Business: American Airlines’ DevOps Journey ([Part 2](#)) (2020)

As we learned in [Part 1](#) of the American Airlines DevOps transformation, their journey spanned several years. By year three, they had realized that DevOps was really a bigger transformation than just a way of working in IT—it was a business transformation.

Their next challenge became how to scale these new ways of working across the entire business to accelerate and further execute their transformation and their learnings. They brought in Ross Clanton as Chief

Architect and Managing Director to help guide them to the next level.

To drive the conversation across the organization, American Airlines focused on two things:³³ the why (building a competitive advantage) and the how (business and IT teams working together to maximize business value).

To scale the vision that was set in IT, “deliver value faster,” they mapped out a transformation based on the following structure of four key pillars:³⁴

- **delivery excellence:** how we work (practices, product mindset)
- **operating excellence:** how we’re structured (product taxonomy, funding model, operating model, prioritization)
- **people excellence:** growing talent and culture (including evolving leadership behaviors)
- **technology excellence:** modernization (infrastructure and technology foundation, automation, move to the cloud, etc.)



Figure 5.2: American Airlines' Delivery Transformation

Source: Maya Leibman and Ross Clanton, “DevOps: Approaching Cruising Altitude,” presentation at DevOps Enterprise Summit-Virtual Las Vegas 2020, videolibrary.doesvirtual.com/?video=467488959.

Once they had their transformation scaling strategy in place, they had to focus on scaling the culture across the business in order to continue to drive the transformation forward.

As Clanton said in his presentation at DevOps Enterprise Summit in 2020, quoting Peter Drucker, “Culture eats strategy for breakfast.”³⁵

To scale the culture, they focused on three key attributes:³⁶

- 1. Passion:** teams focused on delighting customers, being the best at getting better, embrace failure and get stronger because of it.
- 2. Selflessness:** collaborate and share knowledge and code across organization, innersourcing, making space for others’ voices and helping others win.
- 3. Accountability:** own the outcomes even when they’re hard; how you do something is as

important as what you do.

By focusing on these three cultural pillars, the teams at American Airlines are now “empowered and going out of their way to empower others,” according to Clanton.³⁷ As the global pandemic hit in 2020, American Airlines focused on the following values to ensure their teams could still achieve success and results even amidst global change.³⁸

- action and doing over analysis
- collaboration over silos
- clarity of mission over trying to do everything
- empowerment over personal stamp on every effort (set the goals and empower teams to get there)
- getting something out (MVP) versus getting something perfect
- “We can do this” versus hierarchy (collaboration across org boundaries)
- finishing versus starting (limiting WIP and focusing on top priorities)

Instead of throwing requirements over a wall, American Airlines now has a team of stakeholders from the business, IT, Design, etc. They pivoted their planning model so leaders define clear outcomes and the teams are able to decide how they deliver on those outcomes. The teams deliver by focusing on small tasks that add value incrementally. By keeping tasks small, the teams are able to finish quickly (driving value faster) and focus on finishing not starting.

To enable all of this change, leadership needed to change their ways of working as well. Leaders have pivoted to serving the teams, removing impediments and constraints that prevent the teams from delivering value. Instead of status meetings, leadership attend playback meetings (demos) to see what the teams are doing and provide guidance in the moment.

American Airlines also realized that to help change the mindset of leaders and to get everyone aligned in thinking, talking, and acting from an Agile/DevOps perspective, they needed to provide a new vocabulary. [Table 5.1](#) shows some of the ways they helped pivot the conversation.

As Doug Parker, CEO of American Airlines related,

... [the transformation] is making us more efficient, we get projects done more quickly, the projects delivered are more designed to what the users need ... it's already making a huge difference to how to manage projects at American Airlines.

What I'm most proud of is the champion of delivery transformation is not IT anymore; it's the business leaders who embraced it. They see how much faster they're getting the work done, and they're spreading the word everywhere else. And that's making a big difference.³⁹

Table 5.1: American Airlines' New Vocabulary

Before	After
I want to create a pop-up to incentivize people to download the mobile app.	Fragile applications are prone to failure.
What did our competitors do?	What do our customers value?
When will this project be done?	When do we start seeing value?
What went wrong?	What did we learn and how can I help?
I want a completely new website.	What's the first thing we can try to experiment with this idea?

Source: Maya Leibman and Ross Clanton, "DevOps: Approaching Cruising Altitude," presentation at DevOps Enterprise Summit-Virtual Las Vegas 2020, ideolibrary.doesvirtual.com/?video=467488959.

American Airlines built a critical mass and silent majority by focusing on the how and why, as well as changing the vocabulary to achieve a common language.

→ CASE STUDY: NEW FOR SECOND EDITION

Saving the Economy From Ruin (With a Hyperscale PaaS) at HMRC (2020)

HMRC, Her Majesty's Revenue and Customs, is the tax collection agency for the UK government. In 2020, HMRC distributed hundreds of billions of pounds to UK citizens and businesses in an unprecedented financial support package that would eventually see around 25% of the entire UK workforce supported by public money. HMRC built the technology to do this in just four weeks, under conditions of incredible pressure and uncertainty.⁴⁰

HMRC's challenges went way beyond aggressive timescales. "We knew we would have millions of users, but nobody could actually tell us how many. So whatever we built had to be accessible to everyone and had to be capable of paying out billions into bank accounts within hours of launch. It also needed to be secure, with checks being conducted before money was paid out," says Ben Conrad, HMRC's Head of Agile Delivery, at the 2021 DevOps Enterprise Summit-Europe.⁴¹

And they nailed it. All the services launched on time, with most launching a week or two ahead of expectations without any issues, resulting in a 94% user satisfaction rating.⁴² HMRC went from being the least popular of all government departments to the people customers relied on to help out.

In order to achieve this amazing result, HMRC adopted some key processes and leveraged a mature digital platform, one that had evolved over the last seven years to allow teams to build digital services rapidly and deliver them at hyper-scale.

HMRC's platform (the Multichannel Digital Tax Platform or MDTP) is a collection of infrastructure

technologies that enable the organization to serve content to users over the internet. Business domains within HMRC can expose tax services to the public by funding a small cross-functional team to build a microservice or a set of microservices on the platform. MDTP removes much of the pain and complexity of getting a digital service in front of a user by providing tenants with a suite of common components that are necessary to develop and run high-quality digital products.

MDPT is the largest digital platform in the UK government and one of the largest platforms in the UK as a whole. It hosts about 1,200 microservices built by more than two thousand people split into seventy teams across eight geographic locations (although, since March 2020, the teams have all worked 100% remotely).⁴³ These teams make about one hundred deployments into production every day.

“The teams use agile methods, with deliberately lightweight governance, and they’re trusted to make changes themselves whenever and as ever they see fit,” says Matt Hyatt, Technical Delivery Manager with Equal Experts. “It only takes a few seconds to push changes through our infrastructure, so getting products and services in front of users happens really fast.”⁴⁴

Pivotal to the success of the platform has been a constant focus on three key things: culture, tooling, and practices. MDTP’s goal is to make it easy to add teams, build services, and deliver value quickly. The platform has evolved around that goal so that a cross-functional team can be spun up quickly and then use the common

tooling to design, develop, and operate a new public-facing service.

In practice, the MDPT provides a place for a digital team's code to live and automated pipelines for the code to be built and deployed through various environments and into production, where the team can get rapid user feedback. Common telemetry tooling is configured to enable a team to monitor its services via automated dashboards and alerting mechanisms so that they always know what's going on. The platform provides teams with collaboration tools to help them communicate, internally and between one another, so that they can work effectively, both remotely and in person. This is all made available to a team more or less instantly, with minimal configuration or manual steps required. The idea is to free the digital team so it can focus solely on solving business problems.

With over two thousand people making changes, potentially several times a day, things could get very messy. To avoid this, MDTP uses the concept of an *opinionated platform* (also known as a paved road platform or guardrails). For instance, if you build a microservice, it must be written in Scala and use the Play framework. If your service needs persistence, it must use Mongo. If a user needs to perform a common action like uploading a file, then the team must use a common platform service to enable that when there is one. Essentially, a little bit of governance is baked into the platform itself. The benefit to the teams of sticking to the rails is that they can deliver services extraordinarily quickly.

But the benefits don't stop there. By limiting the technology used on the platform, it's far simpler to

support. Moreover, digital teams are prevented from spending time rolling their own solutions to problems that have already been solved elsewhere.

“We can provide common services and reusable components that we know work with *all* the services. It also allows people to move between services, and, indeed, allows the services to move to new teams without worrying about whether our people have the required skills to do the job,” says Hyatt.⁴⁵

Another key differentiator is that MDTP abstracts the need to care about infrastructure away from digital teams, allowing them to focus solely on their apps. “They can still observe the infrastructure through tools like Kibana and Grafana, but none of the service teams have access to AWS accounts themselves,” Hyatt notes.⁴⁶

Importantly, the opinions of this “opinionated platform” can and do change according to user needs and demands, with a focus on self-service. “A service can be created, developed, and deployed on our platform without any direct involvement from platform teams at all,” Hyatt says.⁴⁷

MDTP plays an important role in enabling the rapid delivery of the COVID-19 financial support package, but this isn’t the whole story. A herculean effort by the team was still required. Good decision-making and fast adaptation of processes also proved vital.

“The most important aspect of delivering a system at speed is the ability for engineers to ‘just get on with it,’” says Hyatt.⁴⁸ A combination of balanced governance, established best practices, and a mandate from HMRC

teams to use proven tooling meant they weren't risking new technology problems or security concerns that would have ultimately delayed delivery.

To meet the aggressive timescales, the teams also adapted their composition and communication. A platform engineer was embedded into each COVID-19 digital service team to safely “short-circuit” existing processes. This allowed the teams to eliminate risks early and maximize collaboration, particularly on requests for new infrastructure components and for help with performance testing. Those working on the key services were also made easily identifiable via collaboration tooling, like Slack, so that requests for help could be prioritized, not just by the platform teams but by the whole two thousand-strong digital community.

The ability of the teams and business leaders to adapt so quickly in unison and then flex back into place once the services were launched is a great demonstration of the mature DevOps culture at HMRC.

The HMRC case study illustrates how much can be accomplished using a PaaS in any large organization.

Conclusion

Peter Drucker, a leader in the development of management education, observed that “little fish learn to be big fish in little ponds.”⁴⁹ By carefully choosing where and how to start, we are able to experiment and learn in areas of our organization that create value without jeopardizing the rest of the organization. By doing this, we build our base of support, earn the right to expand the use of DevOps in our

organization, and gain the recognition and gratitude of an ever-larger constituency.

- * These organizations were sometimes known as the “killer B’s that are dying.”⁴
- † The practice of relying on a stabilization phase or hardening phase at the end of a project often has very poor outcomes because it means problems are not being found and fixed as part of daily work and are left unaddressed, potentially snowballing into larger issues.
- ‡ That the services that have the largest potential business benefit are brownfield systems shouldn’t be surprising. After all, these are the systems that are most relied upon and have the largest number of existing customers or highest amount of revenue depending upon them.
- § Big-bang, top-down transformations are possible, such as the Agile transformation at PayPal in 2012 that was led by their vice president of technology, Kirsten Wolberg. However, as with any sustainable and successful transformation, this required the highest level of management support and a relentless, sustained focus on driving the necessary outcomes.

6

UNDERSTANDING THE WORK IN OUR VALUE STREAM, MAKING IT VISIBLE, AND EXPANDING IT ACROSS THE ORGANIZATION

Once we have identified a value stream to which we want to apply DevOps principles and patterns, our next step is to gain a sufficient understanding of how value is delivered to the customer: what work is performed and by whom, and what steps we can take to improve flow.

Nordstrom's Experience With Value Stream Mapping

In the previous chapter, we learned about the DevOps transformation led by Courtney Kissler and the team at Nordstrom. Over the years, they have learned that one of the most efficient ways to start improving any value stream is to conduct a workshop with all the major stakeholders and perform a value stream mapping exercise—a process (described later in this chapter) designed to help capture all the steps required to create value.¹

Kissler's favorite example of the valuable and unexpected insights that can come from value stream mapping is when her team tried to improve the long lead times associated with requests going through the Cosmetics Business Office application, a COBOL mainframe application that supported all the floor and department managers of their in-store beauty and cosmetic departments.²

This application allowed department managers to register new salespeople for various product lines carried in their

stores so they could track sales commissions, enable vendor rebates, and so forth.

Kissler explained:

I knew this particular mainframe application well—earlier in my career I supported this technology team, so I know firsthand that for nearly a decade, during each annual planning cycle, we would debate about how we needed to get this application off the mainframe. Of course, like in most organizations, even when there was full management support, we never seemed to get around to migrating it.

My team wanted to conduct a value stream mapping exercise to determine whether the COBOL application really was the problem, or maybe there was a larger problem that we needed to address. They conducted a workshop that assembled everyone with any accountability for delivering value to our internal customers, including our business partners, the mainframe team, the shared service teams, and so forth.

What they discovered was that when department managers were submitting the “product line assignment” request form, we were asking them for an employee number, which they didn’t have—so they would either leave it blank or put in something like “I don’t know.” Worse, in order to fill out the form, department managers would have to inconveniently leave the store floor in order to use a PC in the back office. The end result was all this wasted time, with work bouncing back and forth in the process.³

CONTINUOUS LEARNING

Incremental Work

It's important to note that improvements can start with incremental work. With a better understanding of the entire value stream and where the real constraints were, the team could make targeted improvements—many of which were much less expensive and much more effective than originally imagined. Even if the COBOL environment eventually needs to be migrated—maybe someday it becomes the constraint—the teams were able to take smart, targeted steps to speed up the process to value delivery along the way.

During the workshop, the participants conducted several experiments, including deleting the employee number field in the form and letting another department get that information in a downstream step. These experiments, conducted with the help of department managers, showed a four-day reduction in processing time. The team later replaced the PC application with an iPad application that allowed managers to submit the necessary information without leaving the store floor, and the processing time was further reduced to seconds.

Kissler said proudly,

With those amazing improvements, all the demands to get this application off the mainframe disappeared. Furthermore, other business leaders took notice and started coming to us with a whole list of further experiments they wanted to conduct with us in their own organizations. Everyone in the business and

technology teams was excited by the outcome because [it] solved a real business problem, and, most importantly, they learned something in the process.⁴

Identifying the Teams Supporting Our Value Stream

As this Nordstrom example demonstrates, in value streams of any complexity no one person knows all the work that must be performed in order to create value for the customer—especially since the required work must be performed by many different teams who are often far removed from each other on the organization charts, geographically, or by incentives.

As a result, after we select a candidate application or service for our DevOps initiative, we must identify all the members of the value stream who are responsible for working together to create value for the customers being served. In general, this includes:

- **Product owner:** the internal voice of the business that defines the next set of functionality in the service.
- **Development:** the team responsible for developing application functionality in the service.
- **QA:** the team responsible for ensuring that feedback loops exist to ensure the service functions as desired.
- **IT Operations/SRE:** the team often responsible for maintaining the production environment and helping ensure that required service levels are met.
- **Infosec:** the team responsible for securing systems and data.
- **Release managers:** the people responsible for managing and coordinating the production deployment and release processes.

- **Technology executives or value stream manager:** in Lean literature, someone who is responsible for “ensuring that the value stream meets or exceeds the customer [and organizational] requirements for the overall value stream, from start to finish.”⁵

Create a Value Stream Map to See the Work

After we identify our value stream members, our next step is to gain a concrete understanding of how work is performed, documented in the form of a value stream map. In our value stream, work likely begins with the product owner in the form of a customer request or the formulation of a business hypothesis. Some time later, this work is accepted by Development, where features are implemented in code and checked into our version control repository. Builds are then integrated, tested in a production-like environment, and finally deployed into production, where they (ideally) create value for our customers.

In many traditional organizations, this value stream will consist of hundreds, if not thousands, of steps, requiring work from hundreds of people. Because documenting any value stream map this complex likely requires multiple days, we may conduct a multi-day workshop, where we assemble all the key constituents and remove them from the distractions of their daily work.

Our goal is not to document every step and associated minutiae, but to sufficiently understand the areas in our value stream that are jeopardizing our goals of fast flow, short lead times, and reliable customer outcomes. Ideally, we have assembled those people with the authority to change their portion of the value stream.*

Damon Edwards, co-founder of Rundeck, observed,

In my experience, these types of value stream mapping exercises are always an eye-opener. Often, it is the first time when people see how much work and heroics are required to deliver value to the customer. For Operations, it may be the first time that they see the consequences that result when developers don't have access to correctly configured environments, which contributes to even more crazy work during code deployments. For Development, it may be the first time they see all the heroics that are required by Test and Operations in order to deploy their code into production, long after they flag a feature as "completed."⁶

Using the full breadth of knowledge brought by the teams engaged in the value stream, we should focus our investigation and scrutiny on the following areas:

- places where work must wait weeks or even months, such as getting production-like environments, change approval processes, or security review processes
- places where significant rework is generated or received

Our first pass of documenting our value stream should only consist of high-level process blocks. Typically, even for complex value streams, groups can create a diagram with five to fifteen process blocks within a few hours. Each process block should include the lead time and process time for a work item to be completed, as well as the %C/A (percent complete and accurate) as measured by the downstream consumers of the output.[†]

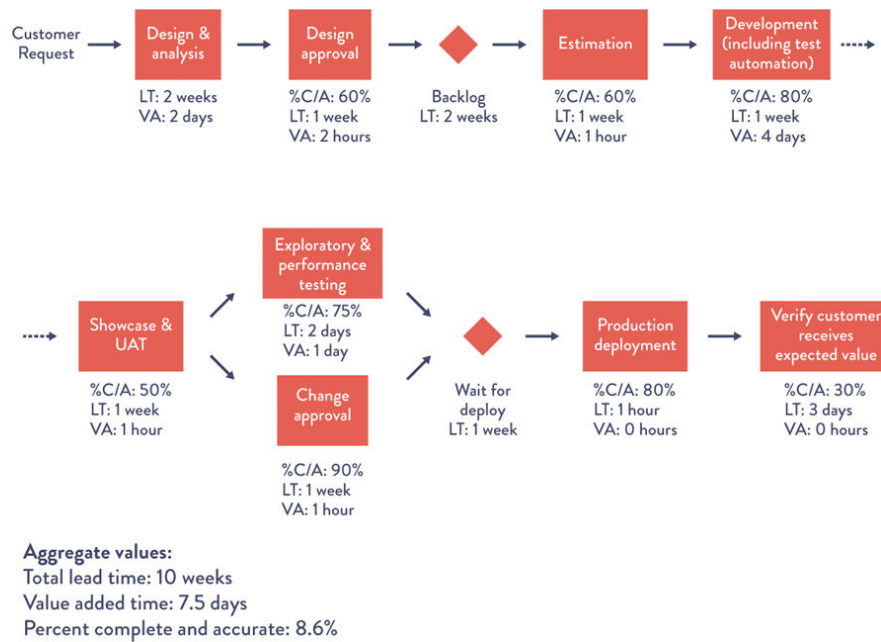


Figure 6.1: An Example of a Value Stream Map

Source: Humble, Molesky, and O'Reilly, *Lean Enterprise*, 139.

We use the metrics from our value stream map to guide our improvement efforts. In the Nordstrom example, they focused on the low %C/A rates on the request form submitted by department managers due to the absence of employee numbers. In other cases, it might be long lead times or low %C/A rates when delivering correctly configured test environments to Development teams, or it might be the long lead times required to execute and pass regression testing before each software release.

Once we identify the metric we want to improve, we should perform the next level of observations and measurements to better understand the problem and then construct an idealized, future value stream map, which serves as a target condition to achieve by some date (e.g., usually three to twelve months).

Leadership helps define this future state and then guides and enables the team to brainstorm hypotheses and countermeasures to achieve the desired improvement to that state, perform experiments to test those hypotheses, and

interpret the results to determine whether the hypotheses were correct. The teams keep repeating and iterating, using any new learnings to inform the next experiments.

Creating a Dedicated Transformation Team

One of the inherent challenges with initiatives such as DevOps transformations is that they are inevitably in conflict with ongoing business operations. Part of this is a natural outcome of how successful businesses evolve. An organization that has been successful for any extended period of time (years, decades, or even centuries) has created mechanisms to perpetuate the practices that made them successful, such as product development, order administration, and supply chain operations.

Many techniques are used to perpetuate and protect how current processes operate, such as specialization, focus on efficiency and repeatability, bureaucracies that enforce approval processes, and controls to protect against variance. In particular, bureaucracies are incredibly resilient and are designed to survive adverse conditions—one can remove half the bureaucrats, and the process will still survive.

While this is good for preserving the status quo, we often need to change how we work to adapt to changing conditions in the marketplace. Doing this requires disruption and innovation, which puts us at odds with groups who are currently responsible for daily operations and the internal bureaucracies, and who will almost always win.

In their book *The Other Side of Innovation: Solving the Execution Challenge*, Dr. Vijay Govindarajan and Dr. Chris Trimble, both faculty members of Dartmouth College's Tuck School of Business, described their studies of how disruptive innovation is achieved despite these powerful forces of daily operations. They documented how customer-driven auto

insurance products were successfully developed and marketed at Allstate, how the profitable digital publishing business was created at the *Wall Street Journal*, the development of the breakthrough trail-running shoe at Timberland, and the development of the first electric car at BMW.⁷

Based on their research, Dr. Govindarajan and Dr. Trimble assert that organizations need to create a dedicated transformation team that is able to operate outside of the rest of the organization responsible for daily operations. They called these the “dedicated team” and the “performance engine” respectively.⁸

First and foremost, we will hold this dedicated team accountable for achieving a clearly defined, measurable, system-level result (e.g., “reduce the deployment lead time from code committed into version control to successfully running in production by 50%”).

In order to execute such an initiative, we do the following:

- Assign members of the dedicated team to be solely allocated to the DevOps transformation efforts (as opposed to “maintain all your current responsibilities but spend 20% of your time on this new DevOps thing”).
- Select team members who are generalists, who have skills across a wide variety of domains.
- Select team members who have longstanding and mutually respectful relationships with key areas of the organization.
- Create a separate physical space or virtual space (such as dedicated chat channels) for the dedicated team, if possible, to maximize communication flow within the

team and to create some isolation from the rest of the organization.

If possible, we will free the transformation team from many of the rules and policies that restrict the rest of the organization, as National Instruments did (described in the previous chapter). After all, established processes are a form of institutional memory—we need the dedicated team to create the new processes and learnings required to generate our desired outcomes, creating new institutional memory.

Creating a dedicated team is not only good for the team but also good for the performance engine. By creating a separate team, we create the space for them to experiment with new practices and protect the rest of the organization from the potential disruptions and distractions associated with such experimentation.

Agree On a Shared Goal

One of the most important parts of any improvement initiative is to define a measurable goal with a clearly defined deadline, between six months and two years in the future. It should require considerable effort but still be achievable. And achievement of the goal should create obvious value for the organization as a whole and for our customers.

These goals and the time frame should be agreed upon by the executives and be known to everyone in the organization. We also want to limit the number of these types of initiatives going on simultaneously to prevent us from overly taxing the organizational change management capacity of leaders and the organization. Examples of improvement goals might include:

- Reduce the percentage of the budget spent on product support and unplanned work by 50%.

- Ensure lead time from code check-in to production release is one week or less for 95% of changes.
- Ensure releases can always be performed during normal business hours with zero downtime.
- Integrate all the required information security controls into the deployment pipeline to pass all required compliance requirements.

Once the high-level goal is made clear, teams should decide on a regular cadence to drive the improvement work. Like product development work, we want transformation work to be done in an iterative, incremental manner. A typical iteration will be in the range of two to four weeks. For each iteration, the teams should agree on a small set of goals that generates value and makes some progress toward the long-term goal. At the end of each iteration, teams should review their progress and set new goals for the next iteration.

Keep Our Improvement Planning Horizons Short

In any DevOps transformation project, we need to keep our planning horizons short, just as if we were in a startup doing product or customer development. Our initiative should strive to generate measurable improvements or actionable data within weeks (or, in the worst case, months).

By keeping our planning horizons and iteration intervals short, we achieve the following:

- flexibility and the ability to reprioritize and replan quickly
- decrease the delay between work expended and improvement realized, which strengthens our feedback loop, making it more likely to reinforce desired behaviors—when improvement initiatives are successful, it encourages more investment

- faster learning generated from the first iteration, meaning faster integration of our learnings into the next iteration
- reduction in activation energy to get improvements
- quicker realization of improvements that make meaningful differences in our daily work
- less risk that our project is killed before we can generate any demonstrable outcomes

Reserve 20% of Capacity for Non-Functional Requirements and Reducing Technical Debt

A problem common to any process improvement effort is how to properly prioritize it—after all, organizations that need it most are those that have the least amount of time to spend on improvement. This is especially true in technology organizations because of technical debt.

Organizations that struggle with financial debt only make interest payments and never reduce the loan principal, and they may eventually find themselves in situations where they can no longer service the interest payments. Similarly, organizations that don't pay down technical debt can find themselves so burdened with daily workarounds for problems left unfixed that they can no longer complete any new work. In other words, they are now only making the interest payment on their technical debt.

We will actively manage this technical debt by ensuring that we invest at least 20% of all Development and Operations capacity on refactoring and investing in automation work and architecture and non-functional requirements (NFRs). These are sometimes referred to as the “ilities,” such as maintainability, manageability, scalability, reliability, testability, deployability, and security.

After the near-death experience of eBay in the late 1990s, Marty Cagan, author of *Inspired: How to Create Products Customers Love*, the seminal book on product design and management, codified the following lesson:

The deal [between product owners and] engineering goes like this: Product management takes 20% of the team's capacity right off the top and gives this to engineering to spend as they see fit. They might use it to rewrite, re-architect, or re-factor problematic parts of the code base ... whatever they believe is necessary to avoid ever having to come to the team and say, "we need to stop and rewrite [all our code]." If you're in really bad shape today, you might need to make this 30% or even more of the resources. However, I get nervous when I find teams that think they can get away with much less than 20%.⁹

	User-visible	User-invisible
Positive Value	Feature	Architecture, non-functional requirements, process improvement
Negative Value	Defect	Technical debt

Figure 6.2: Invest 20% of Capacity in those Who Create Positive, User-Invisible Value

Source: "Machine Learning and Technical Debt with D. Sculley," Software Engineering Daily podcast, November 17, 2015, <http://softwareengineeringdaily.com/2015/11/17/machine-learning-and-technical-debt-with-d-sculley/>.

Cagan notes that when organizations do not pay their “20% tax,” technical debt will increase to the point where an organization inevitably spends all of its cycles paying down technical debt.¹⁰ At some point, the services become so fragile that feature delivery grinds to a halt because all the engineers are working on reliability issues or working around problems.

By dedicating 20% of our cycles so that Dev and Ops can create lasting countermeasures to the problems we encounter in our daily work, we ensure that technical debt doesn’t impede our ability to quickly and safely develop and operate our services in production. Elevating added pressure of technical debt from workers can also reduce levels of burnout.

CASE STUDY

Operation InVersion at LinkedIn (2011)

LinkedIn’s Operation InVersion presents an interesting case study that illustrates the need to pay down technical debt as a part of daily work. Six months after their successful IPO in 2011, LinkedIn continued to struggle with problematic deployments that became so painful that they launched Operation InVersion, where they stopped all feature development for two months in order to overhaul their computing environments, deployments, and architecture.¹¹

LinkedIn was created in 2003 to help users “connect to your network for better job opportunities.”¹² By the end of their first week of operation, they had 2,700 members; one year later, they had over one million members, and have grown exponentially since then.¹³ By

November 2015, LinkedIn had over 350 million members, who generate tens of thousands of requests per second, resulting in millions of queries per second on the LinkedIn back-end systems.¹⁴

From the beginning, LinkedIn primarily ran on their homegrown Leo application, a monolithic Java application that served every page through servlets and managed JDBC connections to various back-end Oracle databases. However, to keep up with growing traffic in their early years, two critical services were decoupled from Leo: the first handled queries around the member connection graph entirely in memory, and the second was member search, which layered over the first.

By 2010, most new development was occurring in new services, with nearly one hundred services running outside of Leo. The problem was that Leo was only being deployed once every two weeks.¹⁵

Josh Clemm, a senior engineering manager at LinkedIn, explained that by 2010, the company was having significant problems with Leo. Despite vertically scaling Leo by adding memory and CPUs, “Leo was often going down in production; it was difficult to troubleshoot and recover, and difficult to release new code... . It was clear we needed to ‘kill Leo’ and break it up into many small functional and stateless services.”¹⁶

In 2013, journalist Ashlee Vance of *Bloomberg* described how “when LinkedIn would try to add a bunch of new things at once, the site would crumble into a broken mess, requiring engineers to work long into the night and fix the problems.”¹⁷

By fall 2011, late nights were no longer a rite of passage or a bonding activity, because the problems had become intolerable. Some of LinkedIn's top engineers, including Kevin Scott, who had joined as VP of Engineering three months before their initial public offering, decided to completely stop engineering work on new features and dedicate the whole department to fixing the site's core infrastructure. They called the effort Operation InVersion.

Scott launched Operation InVersion as a way to “inject the beginnings of a cultural manifesto into his team's engineering culture. There would be no new feature development until LinkedIn's computing architecture was revamped—it's what the business and his team needed.”¹⁸

Scott described one downside “You go public, have all the world looking at you, and then we tell management that we're not going to deliver anything new while all of engineering works on this [InVersion] project for the next two months. It was a scary thing.”¹⁹

Vance described the massively positive results of Operation InVersion:

LinkedIn created a whole suite of software and tools to help it develop code for the site. Instead of waiting weeks for their new features to make their way onto LinkedIn's main site, engineers could develop a new service, have a series of automated systems examine the code for any bugs and issues the service might have interacting with existing features, and launch it right to the live LinkedIn site ... LinkedIn's engineering corps

[now] performs major upgrades to the site three times a day.²⁰

By creating a safer system of work, the value they created included fewer late-night cram sessions, with more time to develop new, innovative features.

As Josh Clemm described in his article on scaling at LinkedIn,

Scaling can be measured across many dimensions, including organizational... . [Operation InVersion] allowed the entire engineering organization to focus on improving tooling and deployment, infrastructure, and developer productivity. It was successful in enabling the engineering agility we need to build the scalable new products we have today... . [In] 2010, we already had over 150 separate services. Today, we have over 750 services.²¹

Kevin Scott stated,

Your job as an engineer and your purpose as a technology team is to help your company win. If you lead a team of engineers, it's better to take a CEO's perspective. Your job is to figure out what it is that your company, your business, your marketplace, your competitive environment needs. Apply that to your engineering team in order for your company to win.²²

By allowing LinkedIn to pay down nearly a decade of technical debt, Operation InVersion enabled stability and safety while setting the next stage of growth for the

company. However, it required two months of total focus on non-functional requirements, at the expense of all the promised features made to the public markets during an IPO. By finding and fixing problems as part of our daily work, we manage our technical debt so that we avoid these “near-death” experiences.

This case study is a good example of paying off technical debt, creating a stable and safe environment as a result. The burdens of daily workarounds were lifted and the team was able to once again focus on delivering new features to delight their customers.

Increase the Visibility of Work

To know if we are making progress toward our goal, it's essential that everyone in the organization knows the current state of work. There are many ways to make the current state visible, but what's most important is that the information we display is up to date, and that we constantly revise what we measure to make sure it's helping us understand progress toward our current target conditions.

The following section discusses patterns that can help create visibility and alignment across teams and functions.

Use Tools to Reinforce Desired Behavior

As Christopher Little, a software executive and one of the earliest chroniclers of DevOps, observed, “Anthropologists describe tools as a cultural artifact. Any discussion of culture after the invention of fire must also be about tools.”²³ Similarly, in the DevOps value stream, we use tools to reinforce our culture and accelerate desired behavior changes.

One goal is that our tooling reinforces that Development and Operations not only have shared goals but also have a common backlog of work, ideally stored in a common work system and using a shared vocabulary, so that work can be prioritized globally.

By doing this, Development and Operations may end up creating a shared work queue, instead of each silo using a different one (e.g., Development uses JIRA while Operations uses ServiceNow). A significant benefit of this is that when production incidents are shown in the same work systems as development work, it will be obvious when ongoing incidents should halt other work, especially if we have a kanban board.

Another benefit of having Development and Operations using a shared tool is a unified backlog, where everyone prioritizes improvement projects from a global perspective, selecting work that has the highest value to the organization or most reduces technical debt. As we identify technical debt, we add it to our prioritized backlog if we can't address it immediately. For issues that remain unaddressed, we can use our "20% time for non-functional requirements" to fix the top items from our backlog.

Other technologies that reinforce shared goals are chat rooms, such as IRC channels, HipChat, Campfire, Slack, Flowdock, and OpenFire. Chat rooms allow the fast sharing of information (as opposed to filling out forms that are processed through predefined workflows), the ability to invite other people as needed, and history logs that are automatically recorded for posterity and can be analyzed during post-mortem sessions.

An amazing dynamic is created when we have a mechanism that allows any team member to quickly help other team members, or even people outside their own team

—the time required to get information or needed work can go from days to minutes. In addition, because everything is being recorded, we may not need to ask someone else for help in the future—we simply search for it.

However, the rapid communication environment facilitated by chat rooms can also be a drawback. As Ryan Martens, the founder and CTO of Rally Software, observes, “In a chat room, if someone doesn’t get an answer in a couple of minutes, it’s totally accepted and expected that you can bug them again until they get what they need.”²⁴

The expectations of immediate response can, of course, lead to undesired outcomes. A constant barrage of interruptions and questions can prevent people from getting necessary work done. As a result, teams may decide that certain types of requests should go through more structured and asynchronous tools.

Conclusion

In this chapter, we identified all the teams supporting our value stream and captured in a value stream map what work is required in order to deliver value to the customer. The value stream map provides the basis for understanding our current state, including our lead time and %C/A metrics for problematic areas, and informs how we set a future state.

This enables dedicated transformation teams to rapidly iterate and experiment to improve performance. We also make sure that we allocate a sufficient amount of time for improvement, fixing known problems and architectural issues, including our non-functional requirements. The case studies from Nordstrom and LinkedIn demonstrate how dramatic improvements can be made in lead times and quality when we find problems in our value stream and pay down technical debt.

* Which makes it all the more important that we limit the level of detail being collected—everyone’s time is valuable and scarce.

† Conversely, there are many examples of using tools in a way that guarantees no behavior changes occur. For instance, an organization commits to an agile planning tool but then configures it for a waterfall process, which merely maintains the status quo.

7

HOW TO DESIGN OUR ORGANIZATION AND ARCHITECTURE WITH CONWAY'S LAW IN MIND

In the previous chapters, we identified a value stream to start our DevOps transformation and established shared goals and practices to enable a dedicated transformation team to improve how we deliver value to the customer.

In this chapter, we will start thinking about how to organize ourselves to best achieve our value stream goals. After all, how we organize our teams affects how we perform our work. Dr. Melvin Conway performed a famous experiment in 1968 with a contract research organization that had eight people who were commissioned to produce a COBOL and an ALGOL compiler. He observed, “After some initial estimates of difficulty and time, five people were assigned to the COBOL job and three to the ALGOL job. The resulting COBOL compiler ran in five phases, the ALGOL compiler ran in three.”¹

These observations led to what is now known as Conway's Law, which states that “organizations which design systems ... are constrained to produce designs which are copies of the communication structures of these organizations... . The larger an organization is, the less flexibility it has and the more pronounced the phenomenon.”²

Eric S. Raymond, author of the book *The Cathedral and the Bazaar: Musings on Linux and Open Source by an Accidental Revolutionary*, crafted a simplified (and now, more famous) version of Conway's Law in his Jargon File: “The organization

of the software and the organization of the software team will be congruent; commonly stated as ‘if you have four groups working on a compiler, you’ll get a 4-pass compiler.’”³

In other words, how we organize our teams has a powerful effect on the software we produce, as well as our resulting architectural and production outcomes. In order to get fast flow of work from Development into Operations, with high quality and great customer outcomes, we must organize our teams and our work so that Conway’s Law works to our advantage. Done poorly, Conway’s Law will prevent teams from working safely and independently; instead, they will be tightly coupled, all waiting on each other for work to be done, with even small changes creating potentially global, catastrophic consequences.

Conway’s Law at Etsy

An example of how Conway’s Law can either impede or reinforce our goals can be seen in a technology that was developed at Etsy called Sprouter. Etsy’s DevOps journey began in 2009, and its technical teams were early developers and proponents of the ideas in this book. Etsy’s 2014 revenue was nearly \$200 million, and the company had a successful IPO in 2015.⁴

Originally developed in 2007, Sprouter connected people, processes, and technology in ways that created many undesired outcomes. Sprouter, shorthand for “stored procedure router,” was originally designed to help make life easier for the developers and database teams. As Ross Snyder, a senior engineer at Etsy, said during his presentation at Surge 2011, “Sprouter was designed to allow the Dev teams to write PHP code in the application, the DBAs to write SQL inside Postgres, with Sprouter helping them meet in the middle.”⁵

Sprouter resided between their front-end PHP application and the Postgres database, centralizing access to the database and hiding the database implementation from the application layer. The problem was that adding any changes to business logic resulted in significant friction between developers and the database teams.

As Snyder observed, “For nearly any new site functionality, Sprouter required that the DBAs write a new stored procedure. As a result, every time developers wanted to add new functionality, they would need something from the DBAs, which often required them to wade through a ton of bureaucracy.”⁶

In other words, developers creating new functionality had a dependency on the DBA team, which needed to be prioritized, communicated, and coordinated, resulting in work sitting in queues, meetings, longer lead times, and so forth.⁷ This is because Sprouter created a tight coupling between the development and database teams, preventing developers from being able to independently develop, test, and deploy their code into production.

The stored procedures for the database were also tightly coupled to Sprouter. Any time a stored procedure was changed, it required changes to Sprouter too. The result was that Sprouter became an ever-larger single point of failure. Snyder explained that everything was so tightly coupled and required such a high level of synchronization as a result that almost every deployment caused a mini-outage.⁸

Both the problems associated with Sprouter and their eventual solution can be explained by Conway’s Law. Etsy initially had two teams, the developers and the DBAs, who were each responsible for two layers of the service: the

application logic layer and the stored procedure layer.⁹ Two teams working on two layers, as Conway's Law predicts.

Sprouter was intended to make life easier for both teams, but it didn't work as expected—when business rules changed, instead of changing only two layers, they now needed to make changes to three layers (in the application, in the stored procedures, and now in Sprouter). The resulting challenges of coordinating and prioritizing work across three teams significantly increased lead times which in turn caused reliability problems, as is confirmed in the 2019 *State of DevOps Report*.¹⁰

In the spring of 2009, as part of what Snyder called “the great Etsy cultural transformation,” Chad Dickerson joined as the new CTO. Dickerson put into motion many things, including a massive investment into site stability, having developers perform their own deployments into production, as well as beginning a two-year journey to eliminate Sprouter.¹¹

To do this, the team decided to move all the business logic from the database layer into the application layer, removing the need for Sprouter. They created a small team that wrote a PHP object-relational mapping (ORM) layer,^{*} enabling the front-end developers to make calls directly to the database and reducing the number of teams required to change business logic from three teams down to one team.¹²

As Snyder described, “We started using the ORM for any new areas of the site and migrated small parts of our site from Sprouter to the ORM over time. It took us two years to migrate the entire site off of Sprouter. And even though we all grumbled about Sprouter the entire time, it remained in production throughout.”¹³

By eliminating Sprouter, Etsy also eliminated the problems associated with multiple teams needing to coordinate for business logic changes, decreased the number of handoffs, and significantly increased the speed and success of production deployments, improving site stability. Furthermore, because small teams could independently develop and deploy their code without requiring another team to make changes in other areas of the system, developer productivity increased.

Sprouter was finally removed from production and Etsy's version control repositories in early 2001. As Snyder said, "Wow, it felt good."^{14†}

As Snyder and Etsy experienced, how we design our organization dictates how work is performed and, therefore, the outcomes we achieve. Throughout the rest of this chapter, we will explore how Conway's Law can negatively impact the performance of our value stream and, more importantly, how we organize our teams to use Conway's Law to our advantage.

Organizational Archetypes

In the field of decision sciences, there are three primary types of organizational structures that inform how we design our DevOps value streams with Conway's Law in mind: *functional*, *matrix*, and *market*. They are defined by Dr. Roberto Fernandez as follows:¹⁶

- **Functional-oriented organizations** optimize for expertise, division of labor, or reducing cost. These organizations centralize expertise, which helps enable career growth and skill development, and they often have tall hierarchical organizational structures. This has been the prevailing method of organization for Operations (i.e., server admins, network admins,

database admins, and so forth are all organized into separate groups).

- **Matrix-oriented organizations** attempt to combine functional and market orientation. However, as many who work in or manage matrix organizations observe, they often result in complicated organizational structures, such as individual contributors reporting to two managers or more, and sometimes achieving neither the goals of functional or market orientation.[‡]
- **Market-oriented organizations** optimize for responding quickly to customer needs. These organizations tend to be flat, composed of multiple cross-functional disciplines (e.g., marketing, engineering, etc.), which often leads to potential redundancies across the organization. This is how many prominent organizations adopting DevOps operate—in extreme examples, such as at Amazon or Netflix, each service team is simultaneously responsible for feature delivery and service support.[§]

With these three categories of organizations in mind, let's explore further how an overly functional orientation, especially in Operations, can cause undesired outcomes in the technology value stream, as Conway's Law would predict.

Problems Often Caused by Overly Functional Orientation (“Optimizing for Cost”)

In traditional IT Operations organizations, we often use functional orientation to organize our teams by their specialties. We put the database administrators in one group, the network administrators in another, the server administrators in a third, and so forth. One of the most visible consequences of this is long lead times, especially for

complex activities like large deployments where we must open up tickets with multiple groups and coordinate work handoffs, resulting in our work waiting in long queues at every step.

Compounding the issue, the person performing the work often has little visibility or understanding of how their work relates to any value stream goals (e.g., “I’m just configuring servers because someone told me to.”). This places workers in a creativity and motivation vacuum.

The problem is exacerbated when each Operations functional area has to serve multiple value streams (i.e., multiple Development teams) who all compete for their scarce cycles. In order for Development teams to get their work done in a timely manner, we often have to escalate issues to a manager or director, and eventually to someone (usually an executive) who can finally prioritize the work against the global organizational goals instead of the functional silo goals. This decision must then get cascaded down into each of the functional areas to change the local priorities, and this, in turn, slows down other teams. When every team expedites their work, the net result is that every project ends up moving at the same slow crawl.

In addition to long queues and long lead times, this situation results in poor handoffs, large amounts of rework, quality issues, bottlenecks, and delays. This gridlock impedes the achievement of important organizational goals, which often far outweigh the desire to reduce costs.[¶]

Similarly, functional orientation can also be found with centralized QA and Infosec functions, which may have worked fine (or at least, well enough) when performing less frequent software releases. However, as we increase the number of Development teams and their deployment and

release frequencies, most functionally oriented organizations will have difficulty keeping up and delivering satisfactory outcomes, especially when their work is being performed manually. Now, we'll study how market-oriented organizations work.

Enable Market-Oriented Teams (“Optimizing for Speed”)

Broadly speaking, to achieve DevOps outcomes, we need to reduce the effects of functional orientation (“optimizing for cost”) and enable market orientation (“optimizing for speed”) so we can have many small teams working safely and independently, quickly delivering value to the customer.

Taken to the extreme, market-oriented teams are responsible not only for feature development but also for testing, securing, deploying, and supporting their service in production, from idea conception to retirement. These teams are designed to be cross-functional and independent—able to design and run user experiments, build and deliver new features, deploy and run their service in production, and fix any defects without manual dependencies on other teams, thus enabling them to move faster. This model has been adopted by Amazon and Netflix and is touted by Amazon as one of the primary reasons behind their ability to move fast even as they grow.¹⁸

To achieve market orientation, we won't do a large, top-down reorganization, which often creates large amounts of disruption, fear, and paralysis. Instead, we will embed the functional engineers and skills (e.g., Ops, QA, Infosec) into each service team, or create a platform organization that provides an automated technology platform for service teams to self-serve everything they need to test, deploy, monitor,

and manage their services in testing and production environments.

This enables each service team to independently deliver value to the customer without having to open tickets with other groups, such as IT Operations, QA, or Infosec.** Research supports this approach: DORA's 2018 and 2019 *State of DevOps Reports* found that teams see superior performance in speed and stability when functional work like database change management, QA, and Infosec is integrated throughout the software delivery process.¹⁹

Making Functional Orientation Work

Having just recommended market-orientated teams, it is worth pointing out that it is possible to create effective, high-velocity organizations with functional orientation. Cross-functional and market-oriented teams are one way to achieve fast flow and reliability, but they are not the only path. We can also achieve our desired DevOps outcomes through functional orientation, as long as everyone in the value stream views customer and organizational outcomes as a shared goal, regardless of where they reside in the organization.

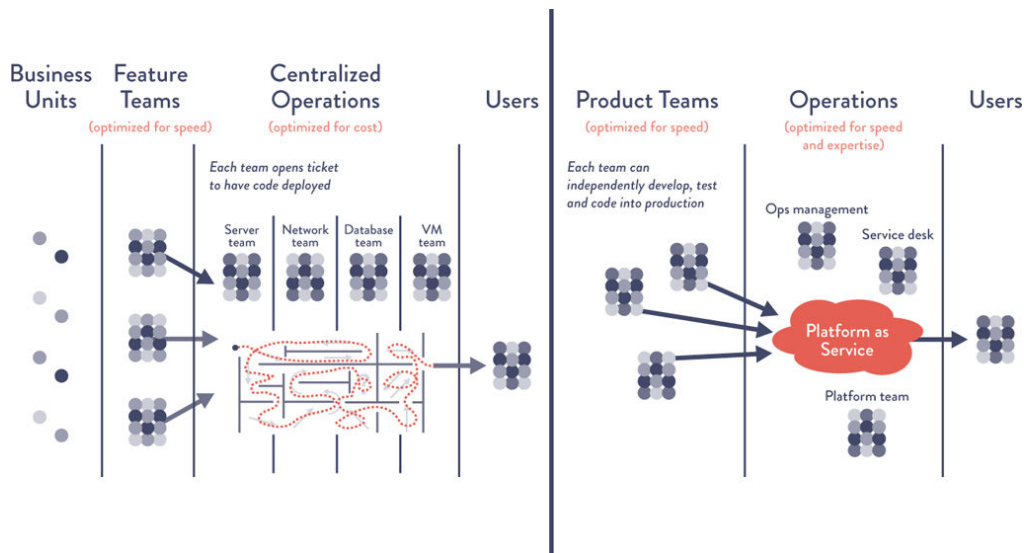


Figure 7.1: Functional vs. Market Orientation

Left: Functional orientation: all work flows through centralized IT Operations.

Right: Market orientation: all product teams can deploy their loosely coupled components self-service into production.

Source: Humble, Molesky, and O'Reilly, *Lean Enterprise*, Kindle edition, 4523 & 4592.

For example, high performance with a functional-oriented and centralized Operations group is possible, as long as service teams get what they need from Operations reliably and quickly (ideally on demand) and vice versa. Many of the most admired DevOps organizations retain functional orientation of Operations, including Etsy, Google, and GitHub.

What these organizations have in common is a high-trust culture that enables all departments to work together effectively, where all work is transparently prioritized and there is sufficient slack in the system to allow high-priority work to be completed quickly. This is, in part, enabled by automated self-service platforms that build quality into the products everyone is building.

In the Lean manufacturing movement of the 1980s, many researchers were puzzled by Toyota's functional orientation, which was at odds with the best practice of having cross-

functional, market-oriented teams. They were so puzzled it was called “the second Toyota paradox.”²⁰

As Mike Rother wrote in *Toyota Kata*,

As tempting as it seems, one cannot reorganize your way to continuous improvement and adaptiveness. What is decisive is not the form of the organization, but how people act and react. The roots of Toyota’s success lie not in its organizational structures, but in developing capability and habits in its people. It surprises many people, in fact, to find that Toyota is largely organized in a traditional, functional-department style.²¹

It is this development of habits and capabilities in people and the workforce that are the focus of our next sections.

Testing, Operations, and Security as Everyone’s Job Every Day

In high-performing organizations, everyone within a team shares a common goal—quality, availability, and security aren’t the responsibility of individual departments but are a part of everyone’s job every day.

This means that the most urgent problem of the day may be working on or deploying a customer feature or fixing a Sev 1 production incident. Alternatively, the day may require reviewing a fellow engineer’s change, applying emergency security patches to production servers, or making improvements so that fellow engineers are more productive.

Reflecting on shared goals between Development and Operations, Jody Mulkey, CTO at Ticketmaster, said, “For almost 25 years, I used an American football metaphor to describe Dev and Ops. You know, Ops is defense, who keeps the other team from scoring, and Dev is offense, trying to score goals. And one day, I realized how flawed this metaphor

was, because they never all play on the field at the same time. They're not actually on the same team!"²²

He continued, "The analogy I use now is that Ops are the offensive linemen, and Dev are the 'skill' positions (like the quarterback and wide receivers) whose job it is to move the ball down the field—the job of Ops is to help make sure Dev has enough time to properly execute the plays."²³

A striking example of how shared pain can reinforce shared goals is when Facebook was undergoing enormous growth in 2009. They were experiencing significant problems related to code deployments—while not all issues caused customer-impacting issues, there was chronic firefighting and long hours. Pedro Canahuati, their director of production engineering, described a meeting full of Ops engineers where someone asked that all people not working on an incident close their laptops, and no one could.²⁴

One of the most significant things they did to help change the outcomes of deployments was to have all Facebook engineers, engineering managers, and architects rotate through on-call duty for the services they built. By doing this, everyone who worked on the service experienced visceral feedback on the upstream architectural and coding decisions they made, which made an enormous positive impact on the downstream outcomes.

Enable Every Team Member to Be a Generalist

In extreme cases of a functionally oriented Operations organization, we have departments of specialists, such as network administrators, storage administrators, and so forth. When departments over specialize, it causes *silozation*, which Dr. Spear describes as when departments "operate more like sovereign states."²⁵ Any complex operational activity then

requires multiple handoffs and queues between the different areas of the infrastructure, leading to longer lead times (e.g., because every network change must be made by someone in the networking department).

Because we rely upon an ever-increasing number of technologies, we must have engineers who have specialized and achieved mastery in the technology areas we need. However, we don't want to create specialists who are "frozen in time," only understanding and able to contribute to that one area of the value stream.

One countermeasure is to enable and encourage every team member to be a generalist. We do this by providing opportunities for engineers to learn all the skills necessary to build and run the systems they are responsible for, and regularly rotating people through different roles. The term *full-stack engineer* is now commonly used (sometimes as a rich source of parody) to describe generalists who are familiar—at least have a general level of understanding—with the entire application stack (e.g., application code, databases, operating systems, networking, cloud).

**Table 7.1: Specialists vs. Generalists vs. “E-Shaped” Staff
(experience, expertise, exploration, and execution)**

“I-shaped” (Specialists)	“T-shaped” (Generalists)	“E-shaped”
Deep expertise in one area	Deep expertise in one area	Deep expertise in a few areas
Very few skills or experience in other areas	Broad skills across many areas	Experience across many areas Proven execution skills Always innovating
Creates bottlenecks quickly	Can step up to remove bottlenecks	Almost limitless potential
Insensitive to downstream waste and impact	Sensitive to downstream waste and impact	—
Prevents planning flexibility or absorption of variability	Helps make planning flexible and absorbs variability	—

Source: Scott Prugh, “Continuous Delivery,” [ScaledAgileFramework.com](http://scaledagileframework.com), February 14, 2013, <http://scaledagileframework.com/continuous-delivery/>.

Scott Prugh writes that CSG International has undergone a transformation that brings most resources required to build and run the product onto one team, including analysis, architecture, development, test, and operations. “By cross-training and growing engineering skills, generalists can do orders of magnitude more work than their specialist

counterparts, and it also improves our overall flow of work by removing queues and wait time.”²⁶

This approach is at odds with traditional hiring practices, but, as Prugh explains, it is well worth it. “Traditional managers will often object to hiring engineers with generalist skill sets, arguing that they are more expensive and that ‘I can hire two server administrators for every multi-skilled operations engineer.’”²⁷ However, the business benefits of enabling faster flow are overwhelming. Furthermore, as Prugh notes, “[I]nvesting in cross training is the right thing for [employees’] career growth and makes everyone’s work more fun.”²⁸

When we value people merely for their existing skills or performance in their current role rather than for their ability to acquire and deploy new skills, we (often inadvertently) reinforce what Dr. Carol Dweck describes as the *fixed mindset*, where people view their intelligence and abilities as static “givens” that can’t be changed in meaningful ways.²⁹

Instead, we want to encourage learning, help people overcome learning anxiety, help ensure that people have relevant skills and a defined career road map, and so forth. By doing this, we help foster a *growth mindset* in our engineers—after all, a learning organization requires people who are willing to learn. By encouraging everyone to learn, as well as providing training and support, we create the most sustainable and least expensive way to create greatness in our teams—by investing in the development of the people we already have.

As Jason Cox, Director of Systems Engineering at Disney, described, “Inside of Operations, we had to change our hiring practices. We looked for people who had ‘curiosity, courage, and candor,’ who were not only capable of being generalists

but also renegades... . We want to promote positive disruption so our business doesn't get stuck and can move into the future.”³⁰ As we'll see in the next section, how we fund our teams also affects our outcomes.

Fund Not Projects but Services and Products

Another way to enable high-performing outcomes is to create stable service teams with ongoing funding to execute their own strategy and road map of initiatives. These teams have the dedicated engineers needed to deliver on concrete commitments made to internal and external customers, such as features, stories, and tasks.

Contrast this to the more traditional model where Development and Test teams are assigned to a “project” and then reassigned to another project as soon as the project is completed and funding runs out. This leads to all sorts of undesired outcomes, including developers being unable to see the long-term consequences of decisions they make (a form of feedback) and a funding model that only values and pays for the earliest stages of the software life cycle—which, tragically, is also the least expensive part for successful products or services.^{††}

Our goal with a product-based funding model is to value the achievement of organizational and customer outcomes, such as revenue, customer lifetime value, or customer adoption rate, ideally with the minimum of output (e.g., amount of effort or time, lines of code). Contrast this to how projects are typically measured, such as whether they are completed within the promised budget, time, and scope.

Design Team Boundaries in Accordance with Conway's Law

As organizations grow, one of the largest challenges is maintaining effective communication and coordination between people and teams, and creating and maintaining a shared understanding and mutual trust becomes even more important. As many teams shift to embrace new patterns of work, the importance of collaboration is becoming even more apparent. It is common for teams to now include fully remote, hybrid, and distributed work configurations with team members stretched across not only office or home boundaries but time zones and sometimes even contractual boundaries (such as when work is performed by an outsourced team). Collaboration is further impeded when the primary communication mechanisms are work tickets and change requests.##

As we saw in the Etsy Sprouter example at the beginning of this chapter, the way we organize teams can create poor outcomes, a side effect of Conway's Law. These include splitting teams by function (e.g., by putting developers and testers in different locations or by outsourcing testers entirely) or by architectural layer (e.g., application, database).

These configurations require significant communication and coordination between teams, but still result in a high amount of rework, disagreements over specifications, poor handoffs, and people sitting idle waiting for somebody else.

Ideally, our software architecture should enable small teams to be independently productive, sufficiently decoupled from each other so that work can be done without excessive or unnecessary communication and coordination.

Create Loosely Coupled Architectures to Enable Developer Productivity and Safety

When we have a tightly coupled architecture, small changes can result in large-scale failures. As a result, anyone working in one part of the system must constantly coordinate with anyone else working in another part of the system they may affect, including navigating complex and bureaucratic change management processes.

Furthermore, to test that the entire system works together requires integrating changes with the changes from hundreds, or even thousands, of other developers, which may, in turn, have dependencies on tens, hundreds, or thousands of interconnected systems. Testing is done in scarce integration test environments, which often require weeks to obtain and configure. The result is not only long lead times for changes (typically measured in weeks or months) but also low developer productivity and poor deployment outcomes.

In contrast, when we have an architecture that enables small teams of developers to independently implement, test, and deploy code into production safely and quickly, we can increase and maintain developer productivity and improve deployment outcomes. These characteristics can be found in *service-oriented architectures* (SOAs) first described in the 1990s, in which services are independently testable and deployable. A key feature of SOAs is that they're composed of *loosely coupled services with bounded contexts*.^{§§}

Having architecture that is loosely coupled means that services can update in production independently, without having to update other services. Services must be decoupled from other services and, just as important, from shared databases (although they can share a database service, provided they don't have any common schemas).

Bounded contexts are described in the book *Domain-Driven Design* by Eric J. Evans. The idea is that developers should be able to understand and update the code of a service without knowing anything about the internals of its peer services. Services interact with their peers strictly through APIs and thus don't share data structures, database schemata, or other internal representations of objects. Bounded contexts ensure that services are compartmentalized and have well-defined interfaces, which also enable easier testing.

Randy Shoup, former Engineering Director for Google App Engine, observed that “organizations with these types of service-oriented architectures, such as Google and Amazon, have incredible flexibility and scalability. These organizations have tens of thousands of developers where small teams can still be incredibly productive.”³³

Keep Team Sizes Small (the “Two-Pizza Team” Rule)

Conway's Law helps us design our team boundaries in the context of desired communication patterns but it also encourages us to keep our team sizes small, reducing the amount of inter-team communication and encouraging us to keep the scope of each team's domain small and bounded.

As part of its transformation initiative away from a monolithic code base in 2002, Amazon used the *two-pizza* rule to keep team sizes small—a team only as large as can be fed with two pizzas—usually about five to ten people.³⁴

This limit on size has four important effects:

- It ensures the team has a clear, shared understanding of the system they are working on. As teams get larger, the amount of communication

required for everybody to know what's going on scales in a combinatorial fashion.

- It **limits the growth rate of the product or service being worked on**. By limiting the size of the team, we limit the rate at which their system can evolve. This also helps to ensure the team maintains a shared understanding of the system.
- It **decentralizes power and enables autonomy**. Each two-pizza team (2PT) is as autonomous as possible. The team's lead, working with the executive team, decides on the key business metric that the team is responsible for, known as the fitness function, which becomes the overall evaluation criteria for the team's experiments. The team is then able to act autonomously to maximize that metric. ¶¶
- Leading a two-pizza team is a **way for employees to gain some leadership experience** in an environment where failure does not have catastrophic consequences. An essential element of Amazon's strategy was the link between the organizational structure of a two-pizza team and the architectural approach of a service-oriented architecture.

Amazon CTO, Werner Vogels, explained the advantages of this structure to Larry Dignan of *Baseline* in 2005.

Dignan writes:

Small teams are fast ... and don't get bogged down in so-called administrivia... . Each group assigned to a particular business is completely responsible for it... . The team scopes the fix, designs it, builds it, implements it and monitors its ongoing use. This way, technology programmers and architects get direct feedback from the business people who use their code

or applications—in regular meetings and informal conversations.³⁶

CONTINUOUS LEARNING

In *Team Topologies: Organizing Business and Technology Teams for Fast Flow*, Matthew Skelton and Manuel Pais present team and organizational patterns to optimize software delivery. The book illustrates an important theme shared in this chapter: good team designs reinforce good software delivery, and good software delivery reinforces more effective teams.

Skelton and Pais also highlight best practices for teams:

- **Trust and communication take time.** They suggest it takes at least three months for team members to reach high performance, and suggest keeping teams together at least a year to benefit from their work together.
- **Just-right sizing.** They suggest eight is an ideal number, which is similar to the two-pizza team used by Amazon, and note that 150 is an upper limit (citing Dunbar's number).^{***}
- **Communication (can be) expensive.** Skelton and Pais wisely point out that while within-team communication is good, any time teams have demands or constraints on other teams, it leads to queues, context switching, and overhead.

The authors also outline four types of teams, and discuss strengths and weaknesses based on the organization, cognitive load requirements of each, and modes of team interaction.

- **Stream-aligned teams:** An end-to-end team that owns the full value stream. This is similar to the market orientation described here.
- **Platform teams:** Platform teams create and support reusable technology often used by stream-aligned teams, such as infrastructure or content management. This team may often be a third party.
- **Enabling teams:** This team contains experts who help other teams improve, such as a Center of Excellence.
- **Complicated-subsystem teams:** Teams that own development and maintenance of a subcomponent of the system that is so complicated it requires specialist knowledge.
- **Other:** The authors also touch on other team types, such as SRE (site reliability engineer) and service experience.

Another example of how architecture can profoundly improve productivity is the API Enablement program at Target, Inc.

**CASE
STUDY**

API Enablement at Target (2015)

Target is the sixth-largest retailer in the US and spends over \$1 billion on technology annually. Heather Mickman, a former director of development for Target, described the beginnings of their DevOps journey: “In the bad old days, it used to take ten different teams to provision a server at Target, and when things broke, we tended to stop making changes to prevent further issues, which of course makes everything worse.”³⁷

The hardships associated with getting environments and performing deployments created significant difficulties for development teams, as did getting access to data they needed.

As Mickman described:

The problem was that much of our core data, such as information on inventory, pricing, and stores, was locked up in legacy systems and mainframes. We often had multiple sources of truths of data, especially between e-commerce and our physical stores, which were owned by different teams, with different data structures and different priorities... .

The result was that if a new development team wanted to build something for our guests, it would take three to six months to build the integrations to get the data they needed. Worse, it would take another three to six months to do the manual testing to make sure they didn't break anything critical because of how many custom point-to-point integrations we had in a very tightly coupled system. Having to manage the interactions

with the twenty to thirty different teams, along with all their dependencies, required lots of project managers because of all the coordination and handoffs. It meant that development was spending all their time waiting in queues instead of delivering results and getting stuff done.³⁸

This long lead time for retrieving and creating data in their systems of record jeopardized important business goals, such as integrating the supply chain operations of Target's physical stores and their e-commerce site, which now required getting inventory to stores and customer homes. This pushed the Target supply chain well beyond what it was designed for, which was merely to facilitate the movement of goods from vendors to distribution centers and stores.

In an attempt to solve the data problem, in 2012 Mickman led the API Enablement team to enable development teams to “deliver new capabilities in days instead of months.”³⁹ They wanted any engineering team inside Target to be able to get and store the data they needed, such as information on their products or their stores, including operating hours, location, whether there was a Starbucks on-site, and so forth.

Time constraints played a large role in team selection. As Mickman explained:

Because our team also needed to deliver capabilities in days, not months, I needed a team who could do the work, not give it to contractors—we wanted people with kickass engineering skills, not people who knew how to manage contracts. And to make sure our

work wasn't sitting in queue, we needed to own the entire stack, which meant that we took over the Ops requirements as well. ... We brought in many new tools to support continuous integration and continuous delivery. And because we knew that if we succeeded, we would have to scale with extremely high growth, we brought in new tools such as the Cassandra database and Kafka message broker. When we asked for permission, we were told no, but we did it anyway because we knew we needed it.⁴⁰

In the following two years, the API Enablement team enabled fifty-three new business capabilities, including Ship to Store and Gift Registry, as well as their integrations with Instacart and Pinterest. As Mickman described, "Working with Pinterest suddenly became very easy, because we just provided them our APIs."⁴¹

In 2014, the API Enablement team served over 1.5 billion API calls per month. By 2015, this had grown to seventeen billion calls per month and spanning ninety different APIs. To support this capability, they routinely performed eighty deployments per week.⁴²

These changes have created major business benefits for Target—digital sales increased 42% during the 2014 holiday season and increased another 32% in Q2. During the Black Friday weekend of 2015, over 280,000 in-store pickup orders were created. By 2015, their goal was to enable 450 of their 1,800 stores to be able to fulfill e-commerce orders, up from one hundred.⁴³

"The API Enablement team shows what a team of passionate change agents can do," Mickman says. "And

it helped set us up for the next stage, which is to expand DevOps across the entire technology organization.”⁴⁴

This case study is takeaway rich, but one of the clearest pictures it paints is how architecture affects the size and organization of a team and vice versa, per Conway’s Law.

Conclusion

Through the Etsy and Target case studies, we can see how architecture and organizational design can dramatically improve our outcomes. Done incorrectly, Conway’s Law will ensure that the organization creates poor outcomes, preventing safety and agility. Done well, the organization enables developers to safely and independently develop, test, and deploy value to the customer.

-
- * Among many things, an ORM abstracts a database, enabling developers to conduct queries and data manipulation as if they were merely another object in the programming language. Popular ORMs include Hibernate for Java, SQLAlchemy for Python, and ActiveRecord for Ruby on Rails.
 - † Sprouter was one of many technologies used in development and production that Etsy eliminated as part of their transformation.¹⁵
 - ‡ For more on how to work with matrix-oriented organizations, check out the DevOps Enterprise Forum paper *Making Matrixed Organizations Successful with DevOps: Tactics for Transformation in a Less Than Optimal Organization*, which you can download at ITRevolution.com/Resources.
 - § However, as will be explained later, equally prominent organizations such as Etsy and GitHub have functional orientation.
 - ¶ Adrian Cockcroft remarked, “For companies who are now coming off of five-year IT outsourcing contracts, it’s like they’ve been frozen in time, during one of the most disruptive times in technology.”¹⁷ In other words, IT outsourcing is a tactic used to control costs through contractually enforced stasis, with firm fixed prices that schedule annual cost reductions. However, it often results in organizations being unable to respond to changing business and technology needs.

- ** For the remainder of this book, we will use *service teams* interchangeably with *feature teams*, *product teams*, *development teams*, and *delivery teams*. The intent is to specify the team primarily developing, testing, and securing the code so that value is delivered to the customer.
- †† As John Lauderbach, VP of Information Technology at Roche Bros. Supermarkets, quipped, “Every new application is like a free puppy. It’s not the upfront capital cost that kills you... . It’s the ongoing maintenance and support.”³¹
- ‡‡ When the first edition was released, most work cultures did not embrace remote or hybrid arrangements. In the intervening years, advancing technology, shifting norms, and the COVID-19 pandemic have shown teams and organizations that remote work and hybrid schedules are not only possible but productive. As such, this paragraph has been modified to account for shifting ways of working.
- §§ These properties are also found in “microservices,” which build upon the principles of SOA. One popular set of patterns for modern web architecture based on these principles is the “12-factor app.”³²
- ¶¶ In the Netflix culture, one of the seven key values is “highly aligned, loosely coupled.”³⁵
- *** Dunbar's number refers to the number of people with whom one can maintain stable social relationships with (150). It was created by British anthropologist Robin Dunbar in the 1990s.

8

HOW TO GET GREAT OUTCOMES BY INTEGRATING OPERATIONS INTO THE DAILY WORK OF DEVELOPMENT

Our goal is to enable user-oriented outcomes where many small teams can quickly and independently deliver value to customers. This can be a challenge to achieve when Operations is centralized and functionally oriented, having to serve the needs of many different development teams with potentially wildly different needs. The result can often be long lead times for needed Ops work, constant reprioritization and escalation, and poor deployment outcomes.

We can create more market-oriented outcomes by better integrating Ops capabilities into Dev teams, making both more efficient and productive. In this chapter, we'll explore many ways to achieve this, both at the organizational level and through daily rituals. By doing this, Ops can significantly improve the productivity of Dev teams throughout the entire organization, as well as enable better collaboration and organizational outcomes.

Big Fish Games

At Big Fish Games, which develops and supports hundreds of mobile and thousands of PC games and had more than \$266 million in revenue in 2013, VP of IT Operations Paul Farrall was in charge of the centralized Operations organization.¹ He was responsible for supporting many different business units that each had a great deal of autonomy.

Each of these business units had dedicated Development teams who often chose wildly different technologies. When

these groups wanted to deploy new functionality, they would have to compete for a common pool of scarce Ops resources. Furthermore, everyone was struggling with unreliable test and integration environments, as well as extremely cumbersome release processes.

Farrall thought the best way to solve this problem was by embedding Ops expertise into Development teams. He observed,

When Dev teams had problems with testing or deployment, they needed more than just technology or environments. What they also needed was help and coaching. At first, we embedded Ops engineers and architects into each of the Dev teams, but there simply weren't enough Ops engineers to cover that many teams. We were able to help more teams with what we called an *Ops liaison* model and with fewer people.²

Farrall defined two types of Ops liaisons: the business relationship manager and the dedicated release engineer.³ The business relationship managers worked with product management, line-of-business owners, project management, Dev management, and developers. They became intimately familiar with product group business drivers and product road maps, acted as advocates for product owners inside of Operations, and helped their product teams navigate the Operations landscape to prioritize and streamline work requests.

Similarly, the dedicated release engineer became intimately familiar with the product's Development and QA issues and helped product management get what they needed from the Ops organization to achieve their goals. They were familiar with the typical Dev and QA requests for Ops and would often executed the necessary work themselves. As

needed, they would also pull in dedicated technical Ops engineers (e.g., DBAs, Infosec, storage engineers, network engineers) and help determine which self-service tools the entire Operations group should prioritize building.

By doing this, Farrall was able to help Dev teams across the organization become more productive and achieve their team goals. Furthermore, he helped the teams prioritize around his global Ops constraints, reducing the number of surprises discovered mid-project and ultimately increasing the overall project throughput.

Farrall notes that both working relationships with Operations and code release velocity were noticeably improved as a result of the changes. He concludes, “The Ops liaison model allowed us to embed IT Operations expertise into the Dev and Product teams without adding new headcount.”⁴

The DevOps transformation at Big Fish Games shows how a centralized Operations team was able to achieve the outcomes typically associated with user-oriented teams. We can employ the three following broad strategies:

- Create self-service capabilities to enable developers in the service teams to be productive.
- Embed Ops engineers into the service teams.
- Assign Ops liaisons to the service teams when embedding Ops is not possible.

CONTINUOUS LEARNING

We can see how this approach at Big Fish Games follows both a platform team and enabling team approach, as outlined in *Team Topologies* by Matthew Skelton and Manuel Pais and summarized in the last chapter.

A single platform team provides infrastructure functionality for the entire organization and supports market-oriented teams. The Ops liaisons also act as members of the enabling teams.

Lastly, Ops engineers can be integrated into the Dev team's rituals that they use in their daily work, including daily standups, planning, and retrospectives.

Create Shared Services to Increase Developer Productivity

One way to enable market-oriented outcomes is for Operations to create a set of centralized platforms and tooling services that any Dev team can use to become more productive, such as getting production-like environments, deployment pipelines, automated testing tools, production telemetry dashboards, and so forth.* By doing this, we enable Dev teams to spend more time building functionality for their customer, as opposed to obtaining all the infrastructure required to deliver and support that feature in production.

All the platforms and services we provide should (ideally) be automated and available on demand without requiring a developer to open up a ticket and wait for someone to manually perform work. This ensures that Operations doesn't become a bottleneck for their customers (e.g., "We received

your work request, and it will take six weeks to manually configure those test environments”).[†]

By doing this, we enable the product teams to get what they need, when they need it, as well as reduce the need for communications and coordination. As Damon Edwards observed, “Without these self-service Operations platforms, the cloud is just Expensive Hosting 2.0.”⁶

In almost all cases, we will not mandate that internal teams use these platforms and services—these platform teams will have to win over and satisfy their internal customers, sometimes even competing with external vendors. By creating this effective internal marketplace of capabilities, we help ensure that the platforms and services we create are the easiest and most appealing choices available (the path of least resistance).

For instance, we may create a platform that provides a shared version control repository with pre-blessed security libraries, a deployment pipeline that automatically runs code quality and security scanning tools, which deploys our applications into *known, good environments* that already have production monitoring tools installed on them. Ideally, we will make life so much easier for Dev teams that they will overwhelmingly decide that using our platform is the easiest, safest, and most secure means to get their applications into production.

We build into these platforms the cumulative and collective experience of everyone in the organization, including QA, Operations, and Infosec, which helps to create an ever-safer system of work. This increases developer productivity and makes it easy for product teams to leverage common processes, such as performing automated testing and satisfying security and compliance requirements.

Creating and maintaining these platforms and tools is real product development—the customers of our platform aren't our external customers but our internal Dev teams. Like creating any great product, creating great platforms that everyone loves doesn't happen by accident. An internal platform team with poor customer focus will likely create tools that everyone will hate and quickly abandon for other alternatives, whether for another internal platform team or an external vendor.

Dianne Marsh, Director of Engineering Tools at Netflix, states that her team's charter is to “support our engineering teams' innovation and velocity. We don't build, bake, or deploy anything for these teams, nor do we manage their configurations. Instead, we build tools to enable self-service. It's okay for people to be dependent on our tools, but it's important that they don't become dependent on us.”⁷

Often, these platform teams provide other services to help their customers learn their technology, migrate off of other technologies, and even provide coaching and consulting to help elevate the state of the practice inside the organization. These shared services also facilitate standardization, enabling engineers to quickly become productive, even if they switch between teams. For instance, if every product team chooses a different toolchain, engineers may have to learn an entirely new set of technologies to do their work, putting the team goals ahead of the global goals.

In organizations where teams can only use approved tools, we can start by removing this requirement for a few teams, such as the transformation team, so that we can experiment and discover what capabilities make those teams more productive.

Internal shared services teams should continually look for internal toolchains that are widely being adopted in the organization, deciding which ones make sense to be supported centrally and made available to everyone. In general, taking something that's already working somewhere and expanding its usage is far more likely to succeed than building these capabilities from scratch.[†]

Embed Ops Engineers into Our Service Teams

Another way we can enable more user-oriented outcomes is by enabling product teams to become more self-sufficient by embedding Operations engineers within them, thus reducing their reliance on centralized Operations. These product teams may also be completely responsible for service delivery and service support.

By embedding Operations engineers into the Dev teams, their priorities are driven almost entirely by the goals of the product teams they are embedded in—as opposed to Ops focusing inwardly on solving their own problems. As a result, Ops engineers become more closely connected to their internal and external customers. Furthermore, the product teams often have the budget to fund the hiring of these Ops engineers, although interviewing and hiring decisions will likely still be done from the centralized Operations group to ensure consistency and quality of staff.

Jason Cox said,

In many parts of Disney, we have embedded Ops (system engineers) inside the product teams in our business units, along with inside Development, Test, and even Information Security. It has totally changed the dynamics of how we work. As Operations Engineers, we create the tools and capabilities that transform the way people work, and even the way they

think. In traditional Ops, we merely drove the train that someone else built. But in modern Operations Engineering, we not only help build the train, but also the bridges that the trains roll on.⁸

For new large Development projects, we may initially embed Ops engineers into those teams. Their work may include helping decide what to build and how to build it, influencing the product architecture, helping influence internal and external technology choices, helping create new capabilities in our internal platforms, and maybe even generating new operational capabilities.

After the product is released to production, embedded Ops engineers may help with the production responsibilities of the Dev team. They will take part in all of the Dev team rituals, such as planning meetings, daily standups, and demonstrations where the team shows off new features and decides which ones to ship. As the need for Ops knowledge and capabilities decreases, Ops engineers may transition to different projects or engagements, following the general pattern that the composition within product teams changes throughout its life cycle.

This paradigm has another important advantage: pairing Dev and Ops engineers together is an extremely efficient way to cross-train operations knowledge and expertise into a service team. It can also have the powerful benefit of transforming operations knowledge into automated code that can be far more reliable and widely reused.

Assign an Ops Liaison to Each Service Team

For a variety of reasons, such as cost and scarcity, we may be unable to embed Ops engineers into every product team. However, we can get many of the same benefits by assigning a designated liaison for each product team.

At Etsy, this model is called “designated Ops.”⁹ Their centralized Operations group continues to manage all the environments—not just production environments but also pre-production environments—to help ensure they remain consistent. The designated Ops engineer is responsible for understanding:

- what the new product functionality is and why we’re building it
- how it works as it pertains to operability, scalability, and observability (diagramming is strongly encouraged)
- how to monitor and collect metrics to ensure the progress, success, or failure of the functionality
- any departures from previous architectures and patterns, and the justifications for them
- any extra needs for infrastructure and how usage will affect infrastructure capacity
- feature launch plans

Furthermore, just like in the embedded Ops model, this liaison attends the team standups, integrating their needs into the Operations road map and performing any needed tasks. These liaisons escalate any resource contention or prioritization issue. By doing this, we identify any resource or time conflicts that should be evaluated and prioritized in the context of wider organizational goals.

Assigning Ops liaisons allows us to support more product teams than the embedded Ops model. Our goal is to ensure that Ops is not a constraint for the product teams. If we find that Ops liaisons are stretched too thin, preventing the product teams from achieving their goals, then we will likely need to either reduce the number of teams each liaison

supports or temporarily embed an Ops engineer into specific teams.

Integrate Ops into Dev Rituals

When Ops engineers are embedded or assigned as liaisons into our product teams, we can integrate them into our Dev team rituals. In this section, our goal is to help Ops engineers and other non-developers better understand the existing Development culture and proactively integrate them into all aspects of planning and daily work. As a result, Operations is better able to plan and radiate any needed knowledge into the product teams, influencing work long before it gets into production. The following sections describe some of the standard rituals used by Development teams using agile methods and how we would integrate Ops engineers into them. By no means are agile practices a prerequisite for this step—as Ops engineers, our goal is to discover what rituals the product teams follow, integrate into them, and add value to them.[§]

As Ernest Mueller observed, “I believe DevOps works a lot better if Operations teams adopt the same agile rituals that Dev teams have used—we’ve had fantastic successes solving many problems associated with Ops pain points, as well as integrating better with Dev teams.”¹⁰

Invite Ops to Our Dev Standups

One of the Dev rituals popularized by Scrum is the daily standup (although physically standing up has become distinctly optional in remote teams), a quick meeting where everyone on the team gets together and presents to each other three things: what was done yesterday, what is going to be done today, and what is preventing you from getting your work done.[¶]

The purpose of this ceremony is to radiate information throughout the team and to understand the work that is being done and is going to be done. By having team members present this information to each other, we learn about any tasks that are experiencing roadblocks and discover ways to help each other move our work toward completion. Furthermore, by having managers present, we can quickly resolve prioritization and resource conflicts.

A common problem is that this information is compartmentalized within the Development team. By having Ops engineers attend Dev standups, Operations can gain an awareness of the Development team's activities, enabling better planning and preparation—for instance, if we discover that the product team is planning a big feature rollout in two weeks, we can ensure that the right people and resources are available to support the rollout.

Alternatively, we may highlight areas where closer interaction or more preparation is needed (e.g., creating more monitoring checks or automation scripts). By doing this, we create the conditions where Operations can help solve our current team problems (e.g., improving performance by tuning the database instead of optimizing code) or future problems before they turn into a crisis (e.g., creating more integration test environments to enable performance testing).

Invite Ops to Our Dev Retrospectives

Another widespread agile ritual is the retrospective. At the end of each development interval, the team discusses what was successful, what could be improved, and how to incorporate the successes and improvements in future iterations or projects. The team comes up with ideas to make things better and reviews experiments from the previous

iteration. This is one of the primary mechanisms where organizational learning and the development of countermeasures occurs, with resulting work implemented immediately or added to the team's backlog.

Having Ops engineers attend our project team retrospectives means they can also benefit from any new learnings. Furthermore, when there is a deployment or release in that interval, Operations should present the outcomes and any resulting learnings, creating feedback into the product team. By doing this, we can improve how future work is planned and performed, improving our outcomes. Examples of feedback that Operations can bring to a retrospective include:

- “Two weeks ago, we found a monitoring blindspot and agreed on how to fix it. It worked. We had an incident last Tuesday, and we were able to quickly detect and correct it before any customers were impacted.”
- “Last week's deployment was one of the most difficult and lengthy we've had in over a year. Here are some ideas on how it can be improved.”
- “The promotion campaign we did last week was far more difficult than we thought it would be, and we should probably not make an offer like that again. Here are some ideas on other offers we can make to achieve our goals.”
- “During the last deployment, the biggest problem we had was that our firewall rules are now thousands of lines long, making it extremely difficult and risky to change. We need to re-architect how we prevent unauthorized network traffic.”

Feedback from Operations helps our product teams better see and understand the downstream impact of decisions they

make. When there are negative outcomes, we can make the changes necessary to prevent them in the future. Operations feedback will also likely identify more problems and defects that should be fixed—it may even uncover larger architectural issues that need to be addressed.

The additional work identified during project team retrospectives falls into the broad category of improvement work, such as fixing defects, refactoring, and automating manual work. Product managers and project managers may want to defer or deprioritize improvement work in favor of customer features.

However, we must remind everyone that improvement of daily work is more important than daily work itself, and that all teams must have dedicated capacity for this (e.g., reserving 20% of all capacity for improvement work, scheduling one day per week or one week per month, etc.). Without doing this, the productivity of the team will almost certainly grind to a halt under the weight of its own technical and process debt.

Make Relevant Ops Work Visible on Shared Kanban Boards

Often, Development teams will make their work visible on a project board or kanban board. It's far less common, however, for work boards to show the relevant Operations work that must be performed in order for the application to run successfully in production, where customer value is actually created. As a result, we are not aware of necessary Operations work until it becomes an urgent crisis, jeopardizing deadlines or creating a production outage.

Because Operations is part of the product value stream, we should put the Operations work that is relevant to product delivery on the shared kanban board. This enables us to more clearly see all the work required to move our code

into production, as well as keep track of all Operations work required to support the product. Furthermore, it enables us to see where Ops work is blocked and where work needs escalation, highlighting areas where we may need improvement.

Kanban boards are an ideal tool to create visibility, and visibility is a key component in properly recognizing and integrating Ops work into all the relevant value streams. When we do this well, we achieve user-oriented outcomes, regardless of how we've drawn our organization charts.

→ CASE STUDY: NEW TO SECOND EDITION

Better Ways of Working at Nationwide Building Society (2020)

Nationwide Building Society is the world's largest building society, with sixteen million members. In 2020, Patrick Eltridge, Chief Operating Officer, and Janet Chapman, Mission Leader, discussed their continued journey to better ways of working at the DevOps Enterprise Summit London-Virtual.

As a larger, older organization, Nationwide faces a number of challenges. As Patrick Eltridge puts it, they are part of a “hyper fluid and hyper competitive environment.”¹²

Like many organizations, they started their transformation in the IT department mainly around change activities and using Agile practices in IT delivery, seeing measurable but limited benefits.

“We deliver well and reliably, but slowly. We need to get from start to finish more quickly and to surprise and delight our members not only with the quality of our

products and services, but the speed at which they are delivered,” Janet Chapman said.¹³

In 2020, with the help of *Sooner Safer Happier* author Jonathan Smart and his team from Deloitte, Nationwide was in the middle of an organizational pivot from a functional organization to one fully aligned to member needs underpinned by stronger Agile and DevOps practices. One key objective was to bring run and change activities together into long-lived and multiskilled teams. They called it their Member Mission Operating Model.

Typical ways of working in most large organizations are the result of years of evolution. Specialist functions are gathered into departments and the work passes between them with queues at every step.

“Currently, when we process a mortgage application it gets broken into parts among all the functional teams. We all do our bits and then reassemble the outcome, test it to see what we got wrong, and then see if it fits the needs of the member and fix it if it doesn’t,” Eltridge explains. “And when we want to improve performance or reduce cost, we seek to improve the efficiency or reduce the capacity of the teams of individual specialists. What that way of working doesn’t do is optimize the flow of work to our members from beginning to end, right across and through those teams.”¹⁴

To optimize flow, Nationwide made it easy for members to tell them what they wanted. Then they brought together all the people and tools necessary into a single team to make that “want.” Everyone on the team can see all the work, and they are able to organize themselves in a way that smooths the path of the work

and optimizes the delivery of the work in a safe, controlled, and sustainable manner. If a bottleneck does appear, they add people or change the process. What they don't do is add a queuing mechanism as a first response. By moving from functional teams in silos to long-lived, multi-skilled teams, Nationwide has seen throughput improve dramatically, as well as improvements in risk and quality and lower costs.

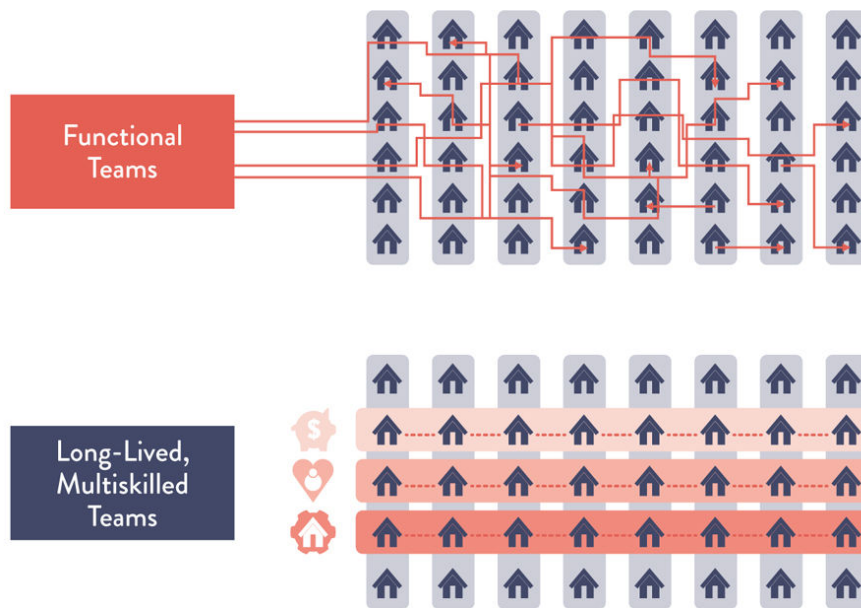


Figure 8.1: Functional Teams in Silos vs. Long-Lived, Multiskilled Teams

Source: Chapman, Janet, and Patrick Eltridge. "On A Mission: Nationwide Building Society," presentation at DevOps Enterprise Summit-Virtual London 2020. <https://videolibrary.doesvirtual.com/?video=432109857>.

They had a unique opportunity to place these new ways of working to the test with the COVID-19 pandemic. As the UK went into lockdown, Nationwide's call centers were quickly swamped due to staff absences. They needed to enable contact center staff to work from home and enable branch center staff to take calls to relieve pressure.

This was an initiative that Nationwide had been discussing for years. But it would have taken nine

months and cost more than £10 million, so it never got done.¹⁵

With an urgent need to reduce call center volumes, Nationwide gathered everyone necessary around the same “virtual” table and worked through, in real time, how to enable an agent to work from home. It took the team only four days to complete.

Next, they looked to see if they could direct calls to the branch network, but they couldn’t without the recording necessary by regulation. So instead, they only directed calls to branch networks that didn’t need the regulatory recording, which helped a bit. A small improvement on a longer path. This, again, took about four days. But over the weekend, they were able to solve the recording problem as well. Another four days.

“Afterwards, I asked the team how many corners we’d cut? How many policies we’d breached? How many security holes we’d now need to plug? They looked at me and said, ‘Well, none. We had all the specialists we needed to do it properly right there. We stuck to the policies, it’s secure, it’s fine,’” explains Eltridge.¹⁶

“When everyone you need is aligned on the most important task at the same time, you get real pace and real collaboration on solving problems. That, in essence, is what a mission is to us,” says Eltridge.¹⁷

Nationwide is now aligning people from their old, functional teams into these long-lived, multiskilled mission teams, as well as their underlying value streams. They are evolving governance and financial management to support local decision-making and continuous funding of consistent teams. They are integrating run and change activities into these long-

lived teams to enable continuous improvement of the work. And they're applying systems thinking to identify and remove failure demand from these flows.

"I think of Agile as our means and DevOps as our target. This is very much a work in progress, and we're consciously allowing the issues and opportunity to merge as we work to implement this. We're not following a templated approach," explains Eltridge. "It is most important to people to go on this journey of learning and unlearning, often with coaches, but not having the answers handed to them by a central team of experts."¹⁸

Beyond simply bringing Dev and Ops together, Nationwide brought together teams with all the skills necessary to bring value to market—moving from multiple functional teams to single, multiskilled teams. This illustrates the power of breaking down silos in order to move faster.

Conclusion

Throughout this chapter, we explored ways to integrate Operations into the daily work of Development, and we looked at how to make Dev work more visible to Operations. To accomplish this, we explored three broad strategies, including creating self-service capabilities to enable developers in service teams to be productive, embedding Ops engineers into the service teams, and assigning Ops liaisons to the service teams when embedding Ops engineers is not possible. Lastly, we described how Ops engineers can integrate with the Dev team through inclusion in their daily work, including daily standups, planning, and retrospectives.

* The terms *platform*, *shared service*, and *toolchain* will be used interchangeably in

this book.

- † Ernest Mueller observed, “At Bazaarvoice, the agreement was that these platform teams that make tools accept requirements but not work from other teams.”⁵
- ‡ After all, designing a system upfront for reuse is a common and expensive failure mode of many enterprise architectures.
- § However, if we discover that the entire Development organization merely sits at their desks all day without ever talking to each other, we may have to find a different way to engage them, such as buying them lunch, starting a book club, taking turns doing “lunch and learn” presentations, or having conversations to discover what everyone’s biggest problems are, so that we can figure out how we can make their lives better.
- ¶ Scrum is an agile development methodology, described as “a flexible, holistic product development strategy where a development team works as a unit to reach a common goal.”¹¹ It was first fully described by Ken Schwaber and Mike Beedle in the book *Agile Software Development with Scrum*. In this book, we use the term “agile development” or “iterative development” to encompass the various techniques used by special methodologies such as Agile and Scrum.

PART II CONCLUSION

In Part II: Where to Start, we explored a variety of ways to think about DevOps transformations, including how to choose where to start, relevant aspects of architecture and organizational design, and how to organize our teams. We also explored how to integrate Operations into all aspects of Development planning and daily work.

In [Part III: The First Way: The Technical Practices of Flow](#), we will now start to explore how to implement the specific technical practices to realize the principles of flow, which enable the fast flow of work from Development to Operations without causing chaos and disruption downstream.

Additional Resources

For more information on team organization or overall organizational structures, check out *Team Topologies: Organizing Business and Technology Teams for Fast Flow* from Matthew Skelton and Manuel Pais (itrevolution.com/team-topologies/).

The *Individual Contributors: From Holdouts to Holdups* paper is an excellent resource for working with individuals in organizations who are resistant to change (ITRevolution.com/Resources).

Expanding Pockets of Greatness: Spreading DevOps Horizontally in Your Organization provides clear strategies to help you build momentum and go from a few islands of DevOps goodness to a tipping point where the entire organization embraces common DevOps methods (ITRevolution.com/Resources).

In Paula Thrasher's 2020 DevOps Enterprise Summit presentation, "Interactive Virtual Value Stream Mapping: Visualizing Flow in a Virtual World," she shares her experience running value stream mapping workshops for

different organizations and teams, and walks you through the steps of setting up a virtual workshop and running one for your organization (videolibrary.doesvirtual.com/?video=466912411).



Part III

THE FIRST WAY: THE TECHNICAL
PRACTICES OF FLOW

PART III: INTRODUCTION

In Part III, our goal is to create the technical practices and architecture required to enable and sustain the fast flow of work from Development into Operations without causing chaos and disruption to the production environment or our customers. This means we need to reduce the risk associated with deploying and releasing changes into production. We will do this by implementing a set of technical practices known as *continuous delivery*.

Continuous delivery includes creating the foundations of our automated deployment pipeline, ensuring that we have automated tests that constantly validate that we are in a deployable state, having developers integrate their code into trunk daily, and architecting our environments and code to enable low-risk releases. Primary focuses within these chapters include:

- creating the foundation of our deployment pipeline
- enabling fast and reliable automated testing
- enabling and practicing continuous integration and testing
- automating, enabling, and architecting for low-risk releases

Implementing these practices reduces the lead time to get production-like environments, enables continuous testing that gives everyone fast feedback on their work, enables small teams to safely and independently develop, test, and deploy their code into production, and makes production deployments and releases a routine part of daily work.

Furthermore, integrating the objectives of QA and Operations into everyone's daily work reduces firefighting,

hardship, and toil while making people more productive and increasing joy in the work we do. We not only improve outcomes, but our organization is also better able to respond to change and achieve its mission to win in the marketplace.

9

CREATE THE FOUNDATIONS OF OUR DEPLOYMENT PIPELINE

In order to create fast and reliable flow from Dev to Ops, we must ensure that we always use production-like environments at every stage of the value stream. Furthermore, these environments must be created in an automated manner, ideally on demand from scripts and configuration information stored in version control and entirely self-serviced, without any manual work required from Operations. Our goal is to ensure that we can re-create the entire production environment based on what's in version control.

The Enterprise Data Warehouse Story (2009)

All too often, the only time we discover how our applications perform in anything resembling a production-like environment is during production deployment—far too late to correct problems without the customer being adversely impacted. An illustrative example of the spectrum of problems that can be caused by inconsistently built applications and environments is the Enterprise Data Warehouse program led by Em Campbell-Pretty at a large Australian telecommunications company in 2009. Campbell-Pretty became the general manager and business sponsor for this \$200 million program, inheriting responsibility for all the strategic objectives that relied upon this platform.

In her presentation at the 2014 DevOps Enterprise Summit, Campbell-Pretty explained,

At the time, there were ten streams of work in progress, all using waterfall processes, and all ten streams were significantly behind schedule. Only one of the ten streams had successfully reached user acceptance testing [UAT] on schedule, and it took another six months for that stream to complete UAT, with the resulting capability falling well short of business expectations. This under performance was the main catalyst for the department's Agile transformation.¹

However, after using Agile for nearly a year, they experienced only small improvements, still falling short of their needed business outcomes.

Campbell-Pretty held a program-wide retrospective and asked, "After reflecting on all our experiences over the last release, what are things we could do that would double our productivity?"²

Throughout the project, there was grumbling about the "lack of business engagement." However, during the retrospective, "improve availability of environments" was at the top of the list.³ In hindsight, it was obvious—Development teams needed provisioned environments in order to begin work and were often waiting up to eight weeks.

They created a new integration and build team that was responsible for "building quality into our processes, instead of trying to inspect quality after the fact."⁴ It was initially composed of database administrators (DBAs) and automation specialists tasked with automating their environment creation process. The team quickly made a surprising discovery: only 50% of the source code in their development and test environments matched what was running in production.⁵

Campbell-Pretty observed, “Suddenly, we understood why we encountered so many defects each time we deployed our code into new environments. In each environment, we kept fixing forward, but the changes we made were not being put back into version control.”⁶

The team carefully reverse-engineered all the changes that had been made to the different environments and put them all into version control. They also automated their environment creation process so they could repeatedly and correctly spin up environments.

Campbell-Pretty described the results, noting that “the time it took to get a correct environment went from eight weeks to one day. This was one of the key adjustments that allowed us to hit our objectives concerning our lead time, the cost to deliver, and the number of escaped defects that made it into production.”⁷

Campbell-Pretty’s story shows the variety of problems that can be traced back to inconsistently constructed environments and changes not being systematically put back into version control.

Throughout the remainder of this chapter, we will discuss how to build the mechanisms that will enable us to create environments on demand, expand the use of version control to everyone in the value stream, make infrastructure easier to rebuild than to repair, and ensure that developers run their code in production-like environments along every stage of the software development life cycle.

Enable On-Demand Creation of Dev, Test, and Production Environments

As seen in the enterprise data warehouse example above, one of the major contributing causes of chaotic, disruptive, and

sometimes even catastrophic software releases is that the first time we ever get to see how our application behaves in a production-like environment with realistic load and production data sets is during the release.* In many cases, development teams may have requested test environments in the early stages of the project.

However, when there are long lead times required for Operations to deliver test environments, teams may not receive them soon enough to perform adequate testing. Worse, test environments are often misconfigured or are so different from our production environments that we still end up with large production problems despite having performed pre-deployment testing.

In this step, we want developers to run production-like environments on their own workstations, created on demand and self-serviced. By doing this, developers can run and test their code in production-like environments as part of their daily work, providing early and constant feedback on the quality of their work.

Instead of merely documenting the specifications of the production environment in a document or on a wiki page, we create a common build mechanism that creates all of our environments, such as for development, test, and production. By doing this, anyone can get production-like environments in minutes, without opening up a ticket, let alone having to wait weeks.†

To do this requires defining and automating the creation of our known, good environments, which are stable, secure, and in a risk-reduced state, embodying the collective knowledge of the organization. All our requirements are embedded, not in documents or as knowledge in someone's

head, but codified in our automated environment build process.

Instead of Operations manually building and configuring the environment, we can use automation for any or all of the following:

- copying a virtualized environment (e.g., a VMware image, running a Vagrant script, booting an Amazon Machine Image file in EC2)
- building an automated environment creation process that starts from “bare metal” (e.g., PXE install from a baseline image)
- using “infrastructure as code” configuration management tools (e.g., Puppet, Chef, Ansible, Salt, CFEngine, etc.)
- using automated operating system configuration tools (e.g., Solaris Jumpstart, Red Hat Kickstart, Debian preseed)
- assembling an environment from a set of virtual images or containers (e.g., Docker, Kubernetes)
- spinning up a new environment in a public cloud (e.g., Amazon Web Services, Google App Engine, Microsoft Azure), private cloud (for example, using a stack based on Kubernetes), or other PaaS (platform as a service, such as OpenStack or Cloud Foundry, etc.)

Because we’ve carefully defined all aspects of the environment ahead of time, we are not only able to create new environments quickly but also ensure that these environments will be stable, reliable, consistent, and secure. This benefits everyone.

Operations benefits from this capability to create new environments quickly, because automation of the

environment creation process enforces consistency and reduces tedious, error-prone manual work. Furthermore, Development benefits by being able to reproduce all the necessary parts of the production environment to build, run, and test their code on their workstations. By doing this, we enable developers to find and fix many problems, even at the earliest stages of the project, as opposed to during integration testing or, worse, in production.

By providing developers an environment they fully control, we enable them to quickly reproduce, diagnose, and fix defects, safely isolated from production services and other shared resources. They can also experiment with changes to the environments, as well as to the infrastructure code that creates it (e.g., configuration management scripts), further creating shared knowledge between Development and Operations.[‡]

Create Our Single Repository of Truth for the Entire System

In the previous step, we enabled the on-demand creation of the development, test, and production environments. Now we must ensure that all parts of our software system can be configured and managed using a source of truth that is maintained in version control.

For decades, comprehensive use of version control has increasingly become a mandatory practice of individual developers and development teams.[§] A version control system records changes to files or sets of files stored within the system.⁹ This can be source code, assets, or other documents that may be part of a software development project. We make changes in groups called commits or revisions. Each revision, along with metadata such as who

made the change and when, is stored within the system in one way or another, allowing us to commit, compare, merge, and restore past revisions. It also minimizes risks by establishing a way to revert objects in production to previous versions.[¶]

When developers put all their application source files and configurations in version control, it becomes the single repository of truth that contains the precise intended state of the system. However, because delivering value to the customer requires both our code and the environments they run in, we need our environments in version control as well. In other words, version control is for everyone in our value stream, including QA, Operations, Infosec, as well as developers.

By putting all production artifacts into version control, our version control repository enables us to repeatedly and reliably reproduce all components of our working software system—this includes our applications and production environment, as well as all of our pre-production environments.

To ensure that we can restore production service repeatedly and predictably (and, ideally, quickly) even when catastrophic events occur, we must check in the following assets to our shared version control repository:

- all application code and dependencies (e.g., libraries, static content, etc.)
- any script used to create database schemas, application reference data, etc.
- all the environment creation tools and artifacts described in the previous step (e.g., VMware or AMI images, Puppet, Chef, or Ansible scripts.)

- any file used to create containers (e.g., Docker, Rocket, or Kubernetes definitions or composition files)
- all supporting automated tests and any manual test scripts
- any script that supports code packaging, deployment, database migration, and environment provisioning
- all project artifacts (e.g., requirements documentation, deployment procedures, release notes, etc.)
- all cloud configuration files (e.g., AWS Cloudformation templates, Microsoft Azure Stack DSC files, OpenStack HEAT)
- any other script or configuration information required to create infrastructure that supports multiple services (e.g., enterprise service buses, database management systems, DNS zone files, configuration rules for firewalls, and other networking devices)**

We may have multiple repositories for different types of objects and services, where they are labeled and tagged alongside our source code. For instance, we may store large virtual machine images, ISO files, compiled binaries, and so forth in artifact repositories (e.g., Nexus, Artifactory). Alternatively, we may put them in blob stores (e.g., Amazon S3 buckets) or put Docker images into Docker registries, and so forth. We will also create and store a cryptographic hash of these objects at build time and validate this hash at deploy time to ensure they haven't been tampered with.

It is not sufficient to merely be able to re-create any previous state of the production environment; we must also be able to re-create the entire pre-production and build processes as well. Consequently, we need to put into version

control everything relied upon by our build processes, including our tools (e.g., compilers, testing tools) and the environments they depend upon.^{††}

Research highlights the importance of version control. The 2014–2019 *State of DevOps Reports* led by co-author Dr. Nicole Forsgren show that the use of version control for all production artifacts was a higher predictor for software delivery performance, which in turn predicted organizational performance.

These findings underscore the critical role version control plays in the software development process. We now know when all application and environment changes are recorded in version control; it enables us to not only quickly see all changes that might have contributed to a problem but also provides the means to roll back to a previous known, running state, allowing us to more quickly recover from failures.

But why does using version control for our environments predict software delivery and contribute to organizational performance better than using version control for our code?

Because in almost all cases, there are orders of magnitude more configurable settings in our environment than in our code. Consequently, it is the environment that needs to be in version control the most.^{‡‡}

Version control also provides a means of communication for everyone working in the value stream—having Development, QA, Infosec, and Operations able to see each other's changes helps reduce surprises, creates visibility into each other's work, and helps build and reinforce trust. (See [Appendix 7](#).) Of course, this means that all teams must use the same version control system.

Make Infrastructure Easier to Rebuild Than to Repair

When we can quickly rebuild and re-create our applications and environments on demand, we can also quickly rebuild them instead of repairing them when things go wrong. Although this is something that almost all large-scale web operations do (i.e., operations with more than one thousand servers), we should also adopt this practice even if we have only one server in production.

Bill Baker, a distinguished engineer at Microsoft, said that we used to treat servers like pets: “You name them and when they get sick, you nurse them back to health. [Now] servers are [treated] like cattle. You number them and when they get sick, you shoot them.”¹¹

By having repeatable environment creation systems, we are able to easily increase capacity by adding more servers into rotation (i.e., horizontal scaling). We also avoid the disaster that inevitably results when we must restore service after a catastrophic failure of irreproducible infrastructure, created through years of undocumented and manual production changes.

To ensure consistency of our environments, whenever we make production changes (configuration changes, patching, upgrading, etc.), those changes need to be replicated everywhere in our production and pre-production environments, as well as in any newly created environments.

Instead of manually logging into servers and making changes, we must make changes in a way that ensures all changes are replicated everywhere automatically and that all our changes are put into version control.

Depending on the life cycle of the configuration in question, we can rely on our automated configuration systems to ensure consistency (e.g., Puppet, Chef, Ansible, Salt, Bosh, etc.), use a service mesh or configuration

management service to propagate runtime configuration (Istio, AWS Systems Manager Parameter Store etc.), or we can create new virtual machines or containers from our automated build mechanism and deploy them into production, destroying the old ones or taking them out of rotation.^{§§}

The latter pattern is what has become known as *immutable infrastructure*, where manual changes to the production environment are no longer allowed—the only way production changes can be made is to put the changes into version control and re-create the code and environments from scratch.¹³ By doing this, no variance is able to creep into production.

To prevent uncontrolled configuration variances, we may disable remote logins to production servers^{¶¶} or routinely kill and replace production instances,^{***} ensuring that manually applied production changes are removed. This action motivates everyone to put their changes in the correct way through version control. By applying such measures, we are systematically reducing the ways our infrastructure can drift from our known, good states (e.g., configuration drift, fragile artifacts, works of art, snowflakes, and so forth).

Also, we must keep our pre-production environments up to date. Specifically, we need developers run on our most current environment. Developers will often want to keep running on older environments because they fear environment updates may break existing functionality. However, we want to update them frequently so we can find problems at the earliest part of the life cycle,^{††} and research from GitHub in the 2020 *State of Octoverse* report shows that keeping your software current is the best way to secure your codebase.¹⁵

→ CASE STUDY: NEW TO SECOND EDITION

How a Hotel Company Ran \$30B of Revenue in Containers (2020)

While at one of the largest hotel companies, Dwayne Holmes, then Senior Director of DevSecOps and Enterprise Platforms, and his team containerized all of the company's revenue generating systems, which collectively supports over \$30 billion in annual revenue.¹⁶

Originally, Dwayne came from the financial sector. He was struggling to find more things to automate to increase productivity. At a local meetup on Ruby of Rails, he stumbled onto containers. For Dwayne, containers were a clear solution for accelerating business value and increasing productivity.

Containers satisfy three key things: they abstract infrastructure (the dial-tone principal—you pick up the phone and it works without needing to know how it works), specialization (Operations could create containers that developers could use over and over and over again), and automation (containers can be built over and over again and everything will just work).¹⁷

With his love of containers now fully embedded, Dwayne took a chance by leaving his comfortable position to become a contractor for one of the largest hotel companies who was ready to go all in on containers.¹⁸

With a small, cross-functional team made up of three developers and three infrastructure professionals.

Their goal was to talk about evolution versus revolution to totally change the way the enterprise worked.¹⁹

There were lots of learnings along the way, as Dwayne outlines in his 2020 DevOps Enterprise Summit presentation, but ultimately the project was successful.²⁰

For Dwayne and the hotel company, containers are the way. They're cloud portable. They're scalable. Health checks are built in. They could test for latency versus CPU, and certs are no longer in the application or managed by developers. Additionally, they are now able to focus on circuit breaking, they have APM built-in, operate zero trust, and images are very small due to good container hygiene and sidecars being used to enhance everything.²¹

During his time at the hotel company, Dwayne and his team supported over three thousand developers across multiple service providers. In 2016, microservices and containers were running in production. In 2017 \$1 billion was processed in containers, 90% of new applications were in containers, and they had Kubernetes running in production. In 2018, they were one of the top five largest production Kubernetes clusters by revenue. And by 2020, they performed thousands of builds and deployments per day and were running Kubernetes in five cloud providers.²²

Containers have become a fast-growing method of making infrastructure easier to rebuild and reuse than to repair, ultimately accelerating the delivery of business value and developer productivity.

Modify Our Definition of Development “Done” to Include Running in Production-Like Environments

Now that our environments can be created on demand and everything is checked into version control, our goal is to ensure that these environments are being used in the daily work of Development. We need to verify that our application runs as expected in a production-like environment long before the end of the project or before our first production deployment.

Most modern software development methodologies prescribe short and iterative development intervals, as opposed to the big-bang approach (i.e., the waterfall model). In general, the longer the interval between deployment, the worse the outcomes. For example, in the Scrum methodology, a sprint is a time-boxed development interval (maximum one month but typically less) within which we are required to be done, widely defined as when we have “working and potentially shippable code.”

Our goal is to ensure that Development and QA are routinely integrating the code with production-like environments at increasingly frequent intervals throughout the project. We do this by expanding the definition of “done” beyond just correct code functionality: at the end of each development interval, or more frequently, we have integrated, tested, working, and potentially shippable code, **demonstrated in a production-like environment.**

In other words, we will only accept development work as done when it can be successfully built, deployed, and confirmed that it runs as expected in a production-like environment, instead of merely when a developer believes it to be done. Ideally, it runs under a production-like load with a production-like dataset, long before the end of a sprint. This

prevents situations where a feature is called done merely because a developer can run it successfully on their laptop but nowhere else.

By having developers write, test, and run their own code in a production-like environment, the majority of the work to successfully integrate our code and environments happens during our daily work, instead of at the end of the release. By the end of our first interval, our application can be demonstrated to run correctly in a production-like environment, with the code and environment having been integrated together many times over, ideally with all the steps automated (no manual tinkering required).

Better yet, by the end of the project, we will have successfully deployed and run our code in production-like environments hundreds or even thousands of times, giving us confidence that most of our production deployment problems have been found and fixed.

Ideally, we use the same tools, such as monitoring, logging, and deployment, in our pre-production environments as we do in production. By doing this, we have familiarity and experience that will help us smoothly deploy and run, as well as diagnose and fix, our service when it is in production.

By enabling Development and Operations to gain a shared mastery of how the code and environment interact, and practicing deployments early and often, we significantly reduce the deployment risks that are associated with production code releases. This also allows us to eliminate an entire class of operational and security defects and architectural problems that are usually caught too late in the project to fix.

Conclusion

The fast flow of work from Development to Operations requires that anyone can get production-like environments on demand. By allowing developers to use production-like environments even at the earliest stages of a software project, we significantly reduce the risk of production problems later. This is one of many practices that demonstrate how Operations can make developers far more productive. We enforce the practice of developers running their code in production-like environments by incorporating it into the definition of “done.”

Furthermore, by putting all production artifacts into version control, we have a “single source of truth” that allows us to re-create the entire production environment in a quick, repeatable, and documented way, using the same development practices for Operations work as we do for Development work. And by making production infrastructure easier to rebuild than to repair, we make resolving problems easier and faster, as well as making it easier to expand capacity. Having these practices in place sets the stage for enabling comprehensive test automation, which is explored in the next chapter.

-
- * In this context, *environment* is defined as everything in the application stack except for the application, including the databases, operating systems, networking, virtualization, and all associated configurations.
 - † Most developers want to test their code, and they have often gone to extreme lengths to obtain test environments to do so. Developers have been known to reuse old test environments from previous projects (often years old) or ask someone who has a reputation of being able to find one—they just won’t ask where it came from because, invariably, someone somewhere is now missing a server.
 - ‡ Ideally, we should be finding errors before integration testing, when it is too late in the testing cycle to create fast feedback for developers. If we are unable to do so, we likely have an architectural issue that needs to be addressed. Designing our systems for testability, to include the ability to discover most defects using a

nonintegrated virtual environment on a development workstation, is a key part of creating an architecture that supports fast flow and feedback.

- § The first version control system was likely UPDATE on the CDC6600 (1969). Later came SCCS (1972), CMS on VMS (1978), RCS (1982), and so forth.⁸
- ¶ One may observe that version control fulfills some of the ITIL constructs of the Definitive Media Library (DML) and Configuration Management Database (CMDB), inventorying everything required to re-create the production environment.
- ** In future steps, we will also check into version control all the supporting infrastructure we build, such as the automated test suites and our continuous integration and deployment pipeline infrastructure.
- †† Anyone who has done a code migration for an ERP system (e.g., SAP, Oracle Financials, etc.) may recognize the following situation: When a code migration fails, it is rarely due to a coding error. Instead, it's far more likely that the migration failed due to some difference in the environments, such as between Development and QA or QA and Production.
- ‡‡ At Netflix, the average age of Netflix AWS instance is twenty-four days, with 60% being less than one week old.¹⁰
- §§ Or allow it only in emergencies, ensuring that a copy of the console log is automatically emailed to the operations team.¹²
- ¶¶ The entire application stack and environment can be bundled into containers, which can enable unprecedented simplicity and speed across the entire deployment pipeline.
- *** Kelly Shortridge has written more about this in her book *Security Chaos Engineering*.
- ††† The term *integration* has many slightly different usages in Development and Operations. In Development, integration typically refers to code integration, which is the integration of multiple code branches into trunk in version control. In continuous delivery and DevOps, integration testing refers to the testing of the application in a production-like environment or integrated test environment.¹⁴

10

ENABLE FAST AND RELIABLE AUTOMATED TESTING

At this point, Development and QA are using production-like environments in their daily work, and we are successfully integrating and running our code into a production-like environment for every feature that is accepted, with all changes checked into version control. However, we are likely to get undesired outcomes if we find and fix errors in a separate test phase, executed by a separate QA department only after all development has been completed. And, if testing is only performed a few times a year, developers learn about their mistakes months after they introduced the change that caused the error. By then, the link between cause and effect has likely faded, solving the problem requires firefighting and archaeology, and, worst of all, our ability to learn from the mistake and integrate it into our future work is significantly diminished.

CONTINUOUS LEARNING

Does better observability mean less testing?

As distributed systems have become commonplace, many organizations have rightly invested in better observability for their production systems. This has led some people to infer that better observability means less need to focus on validating software prior to deployment. This is a misconception: incidents in production are costly and hard to debug, even with excellent instrumentation and tooling. Distributed systems are sufficiently complex that, if anything, it is even more important to test individual services for correctness prior to deployment.

Automated testing addresses another significant and unsettling problem. Gary Gruver observes that “without automated testing, the more code we write, the more time and money is required to test our code—in most cases, this is a totally unscalable business model for any technology organization.”¹

The Story of Google’s Web Server (2005)

Although Google now undoubtedly exemplifies a culture that values automated testing at scale, this wasn’t always the case. In 2005, when Mike Bland joined the organization, deploying to [Google.com](https://www.google.com) was often extremely problematic, especially for the Google Web Server (GWS) team.

As Bland explains,

The GWS team had gotten into a position in the mid 2000s where it was extremely difficult to make changes to the web server, a C++ application that handled all requests to Google's home page and many other Google web pages. As important and prominent as [Google.com](https://www.google.com) was, being on the GWS team was not a glamorous assignment—it was often the dumping ground for all the different teams who were creating various search functionality, all of whom were developing code independently of each other. They had problems such as builds and tests taking too long, code being put into production without being tested, and teams checking in large, infrequent changes that conflicted with those from other teams.²

The consequences of this were large—search results could have errors or become unacceptably slow, affecting thousands of search queries on [Google.com](https://www.google.com). The potential result was not only loss of revenue but customer trust.

Bland describes how it affected developers deploying changes: “Fear became the mind-killer. Fear stopped new team members from changing things because they didn’t understand the system. But fear also stopped experienced people from changing things because they understood it all too well.”^{3*} Bland was part of the group that was determined to solve this problem.

GWS team lead Bharat Mediratta believed automated testing would help. As Bland describes,

They created a hard line: no changes would be accepted into GWS without accompanying automated tests. They set up a continuous build and religiously kept it passing. They set up test coverage monitoring and ensured that their level of test coverage went up over time. They

wrote up policy and testing guides and insisted that contributors both inside and outside the team follow them.⁵

The results were startling. As Bland notes,

GWS quickly became one of the most productive teams in the company, integrating large numbers of changes from different teams every week while maintaining a rapid release schedule. New team members were able to make productive contributions to this complex system quickly, thanks to good test coverage and code health. Ultimately, their radical policy enabled the [Google.com](https://www.google.com) homepage to quickly expand its capabilities and thrive in an amazingly fast-moving and competitive technology landscape.⁶

But GWS was still a relatively small team in a large and growing company. The team wanted to expand these practices across the entire organization. Thus, the testing grouplet was born, an informal group of engineers who wanted to elevate automated testing practices across the entire organization. Over the next five years, they helped replicate this culture of automated testing across all of Google.^{7†}

As Rachel Potvin and Josh Levenberg describe, Google's systems have evolved to automatically test thousands of commits from thousands of developers every day:

Google has an automated testing infrastructure that initiates a rebuild of all affected dependencies on almost every change committed to the repository. If a change creates widespread build breakage, a system is in place to automatically undo the change. To reduce the incidence of bad code being committed in the first place, the highly customizable Google “presubmit”

infrastructure provides automated testing and analysis of changes before they are added to the codebase. A set of global presubmit analyses are run for all changes, and code owners can create custom analyses that run only on directories within the codebase they specify.⁸

Eran Messeri, an engineer in the Google Developer Infrastructure group, notes, “Large failures happen occasionally. You’ll get a ton of instant messages and engineers knocking on your door. [When the deployment pipeline is broken,] we need to fix it right away because developers can no longer commit code. Consequently, we want to make it very easy to roll back.”⁹

What enables this system to work at Google is engineering professionalism and a high-trust culture that assumes everyone wants to do a good job, as well as the ability to detect and correct issues quickly. Messeri explains,

There are no hard policies at Google, such as, “If you break production for more than ten projects, you have an SLA to fix the issue within ten minutes.” Instead, there is mutual respect between teams and an implicit agreement that everyone does whatever it takes to keep the deployment pipeline running. We all know that one day, I’ll break your project by accident; the next day, you may break mine.¹⁰

What the teams at Google achieved has made the company one of the most productive technology organizations in the world. By 2016, automated testing and continuous integration at Google enabled over four thousand small teams to work together and stay productive, all simultaneously developing, integrating, testing, and deploying their code into production. Most of Google’s code is in a single shared repository, made up of billions of files, all

being continuously built and integrated, and in 2014 approximately 15 million lines of code were changed in approximately 250,000 files in the Google repository on a weekly basis.¹¹ Some other impressive statistics on their code infrastructure (as of 2016) include:¹²

- typically 40,000 code commits per day (16,000 from engineers and 24,000 from automated systems)
- 50,000 builds per day (on weekdays, this may exceed 90,000)
- 120,000 automated test suites
- 75 million test cases run daily
- over 99% of files stored in Google's version control system are visible to all full-time Google engineers
- codebase includes approximately one billion files and has a history of approximately 35 million commits
- the repository includes approximately two billion lines of code in nine million unique source files and is 86 TB of data

While Google-scale test automation may not be the goal for many companies, the benefits of test automation are possible for everyone. In the remainder of this chapter, we will go through the continuous integration practices required to achieve similar outcomes.

Continuously Build, Test, and Integrate Our Code and Environments

Our goal is to build quality into our product, even at the earliest stages, by having developers build automated tests as part of their daily work. This creates a fast feedback loop that

helps developers find problems early and fix them quickly when there are the fewest constraints (e.g., time, resources).

In this step, we create automated test suites that increase the frequency of integration and testing of our code and our environments from periodic to continuous. We do this by building our deployment pipeline, which will perform integration of our code and environments and trigger a series of tests every time a new change is checked into version control.[‡] (See [Figure 10.1](#).)




Figure 10.1: The Deployment Pipeline

Source: Humble and Farley, *Continuous Delivery*, 3.

The deployment pipeline, first defined by Jez Humble and David Farley in their book *Continuous Delivery: Reliable Software Releases Through Build, Test, and Deployment Automation*, ensures that all code checked into version control is automatically built and tested in a production-like environment.¹⁴ By doing this, we find any build, test, or integration errors as soon as a change is introduced, enabling us to fix them immediately. Done correctly, this allows us to always be assured that we are in a deployable and shippable state.

To achieve this, we must create automated build and test processes that run in dedicated environments. This is critical for the following reasons:

- Our build and test process can run all the time, independent of the work habits of individual engineers.

- A segregated build and test process ensures that we understand all the dependencies required to build, package, run, and test our code (i.e., removing the “it worked on the developer’s laptop, but it broke in production” problem).
- We can package our application to enable the repeatable installation of code and configurations into an environment (e.g., on Linux RPM, yum, npm; on Windows, OneGet; alternatively framework-specific packaging systems can be used, such as EAR and WAR files for Java, gems for Ruby, etc.).
- Instead of putting our code in packages, we may choose to package our applications into deployable containers (e.g., Docker, Rkt, LXD, AMIs).
- Environments can be made more production-like in a way that is consistent and repeatable (e.g., compilers are removed from the environment, debugging flags are turned off, etc.).

Our deployment pipeline validates after every change that our code successfully integrates into a production-like environment. It becomes the platform through which testers request and certify builds during acceptance testing and usability testing, and it will run automated performance and security validations.

Furthermore, it will be used to create self-service builds to UAT (user acceptance testing), integration testing, and security testing environments. In future steps, as we evolve the deployment pipeline, it will also be used to manage all activities required to take our changes from version control to deployment.

A variety of tools have been designed to provide deployment pipeline functionality, many of them open source

(e.g., Jenkins, Go.cd, Concourse, Bamboo, Microsoft Team Foundation Server, TeamCity, Gitlab CI, as well as cloud-based solutions such as CircleCi and TravisCI).[§]

We begin the deployment pipeline by running the commit stage, which builds and packages the software, runs automated unit tests, and performs additional validation such as static code analysis, duplication and test coverage analysis, and checking style.[¶] If successful, this triggers the acceptance stage, which automatically deploys the packages created in the commit stage into a production-like environment and runs the automated acceptance tests.

Once changes are accepted into version control, we want to package our code only once so that the same packages are used to deploy code throughout our entire deployment pipeline. By doing this, code will be deployed into our integrated test and staging environments in the same way that it is deployed into production. This reduces variances that can avoid downstream errors that are difficult to diagnose (e.g., using different compilers, compiler flags, library versions, or configurations).^{**}

The goal of the deployment pipeline is to provide everyone in the value stream, especially developers, the fastest possible feedback that a change has taken us out of a deployable state. This could be a change to our code, to any of our environments, to our automated tests, or even to the deployment pipeline infrastructure (e.g., a Jenkins configuration setting).

As a result, our deployment pipeline infrastructure becomes as foundational for our development processes as our version control infrastructure. Our deployment pipeline also stores the history of each code build, including information about which tests were performed on which

build, which builds have been deployed to which environment, and what the test results were. In combination with the information in our version control history, we can quickly determine what caused our deployment pipeline to break and, likely, how to fix the error. This information also helps us fulfill evidence requirements for audit and compliance purposes, with evidence being automatically generated as part of daily work.

Now that we have a working deployment pipeline infrastructure, we must create our *continuous integration* practices, which require three capabilities:

- a comprehensive and reliable set of automated tests that validate we are in a deployable state
- a culture that “stops the entire production line” when our validation tests fail
- developers working in small batches on trunk rather than long-lived feature branches

In the next section, we describe why fast and reliable automated testing is needed and how to build it.

Build a Fast and Reliable Automated Validation Test Suite

In the previous step, we started to create the automated testing infrastructure that validates that we have a *green build* (i.e., whatever is in version control is in a buildable and deployable state). To underscore why we need to perform this integration and testing step continuously, consider what happens when we only perform this operation periodically, such as during a nightly build process.

Suppose we have a team of ten developers, with everyone checking their code into version control daily, and a developer introduces a change that breaks our nightly build and test job.

In this scenario, when we discover the next day that we no longer have a green build, it will take minutes, or more likely hours, for our development team to figure out which change caused the problem, who introduced it, and how to fix it.

Worse, suppose the problem wasn't caused by a code change, but was due to a test environment issue (e.g., an incorrect configuration setting somewhere). The development team may believe that they fixed the problem because all the unit tests pass, only to discover that the tests will still fail later that night.

Further complicating the issue, ten more changes will have been checked into version control by the team that day. Each of these changes has the potential to introduce more errors that could break our automated tests, further increasing the difficulty of successfully diagnosing and fixing the problem.

In short, slow and periodic feedback kills, especially for larger development teams. The problem becomes even more daunting when we have tens, hundreds, or even thousands of other developers checking their changes into version control each day. The result is that our builds and automated tests are frequently broken, and developers even stop checking their changes into version control (“Why bother, since the builds and tests are always broken?”). Instead, they wait to integrate their code at the end of the project, resulting in all the undesired outcomes of large batch size, big-bang integrations, and production deployments.^{††}

To prevent this scenario, we need fast automated tests that run within our build and test environments whenever a new change is introduced into version control. In this way, we can find and fix any problems immediately, as the Google Web Server example demonstrated. By doing this, we ensure our

batches remain small, and, at any given point in time, we remain in a deployable state.

In general, automated tests fall into one of the following categories, from fastest to slowest:

- **Unit tests:** These typically test a single method, class, or function in isolation, providing assurance to the developer that their code operates as designed. For many reasons, including the need to keep our tests fast and stateless, unit tests often “stub out” databases and other external dependencies (e.g., functions are modified to return static, predefined values, instead of calling the real database).[#]
- **Acceptance tests:** These typically test the application as a whole to provide assurance that a higher level of functionality operates as designed (e.g., the business acceptance criteria for a user story, the correctness of an API), and that regression errors have not been introduced (i.e., we broke functionality that was previously operating correctly). Humble and Farley define the difference between unit and acceptance testing as, “The aim of a unit test is to show that a single part of the application does what the programmer intends it to... . The objective of acceptance tests is to prove that our application does what the customer meant it to, not that it works the way its programmers think it should.”¹⁵ After a build passes our unit tests, our deployment pipeline runs it against our acceptance tests. Any build that passes our acceptance tests is then typically made available for manual testing (e.g., exploratory testing, UI testing, etc.), as well as for integration testing.

- **Integration tests:** Integration tests are where we ensure that our application correctly interacts with other production applications and services, as opposed to calling stubbed out interfaces. As Humble and Farley observe, “much of the work in the SIT [system integration testing] environment involves deploying new versions of each of the applications until they all cooperate. In this situation the smoke test is usually a fully fledged set of acceptance tests that run against the whole application.”¹⁶ Integration tests are performed on builds that have passed our unit and acceptance tests. Because integration tests are often brittle, we want to minimize the number of integration tests and find as many of our defects as possible during unit and acceptance testing. The ability to use virtual or simulated versions of remote services when running acceptance tests becomes an essential architectural requirement.

When facing deadline pressures, developers may stop creating unit tests as part of their daily work, regardless of how we’ve defined “done.” To detect this, we may choose to measure and make visible our test coverage (as a function of number of classes, lines of code, permutations, etc.), maybe even failing our validation test suite when it drops below a certain level (e.g., when less than 80% of our classes have unit tests).^{§§}

Martin Fowler observes that, in general,

a ten-minute build [and test process] is perfectly within reason... . [We first] do the compilation and run tests that are more localized unit tests with the database completely stubbed out. Such tests can run very fast, keeping within the ten minute guideline. However any

bugs that involve larger scale interactions, particularly those involving the real database, won't be found. The second stage build runs a different suite of tests [acceptance tests] that do hit the real database and involve more end-to-end behavior. This suite may take a couple of hours to run.¹⁷

Catch Errors as Early in Our Automated Testing as Possible

A specific design goal of our automated test suite is to find errors as early in the testing phase as possible. This is why we run faster-running automated tests (e.g., unit tests) before slower-running automated tests (e.g., acceptance and integration tests), which are both run before any manual testing.

Another corollary of this principle is that any errors should be found with the fastest category of testing possible. If most of our errors are found in our acceptance and integration tests, the feedback we provide to developers is orders of magnitude slower than with unit tests—and integration testing requires using scarce and complex integration test environments, which can only be used by one team at a time, further delaying feedback.

Furthermore, not only are errors detected during integration testing difficult and time-consuming for developers to reproduce, even validating that it has been fixed is difficult (e.g., a developer creates a fix but then needs to wait four hours to learn whether the integration tests now pass).

Therefore, whenever we find an error with an acceptance or integration test, we should create a unit test that could find the error faster, earlier, and cheaper. Martin Fowler

described the notion of the “ideal testing pyramid,” where we are able to catch most of our errors using our unit tests.¹⁸ (See [Figure 10.2](#).) In contrast, in many testing programs the inverse is true, where most of the investment is in manual and integration testing.

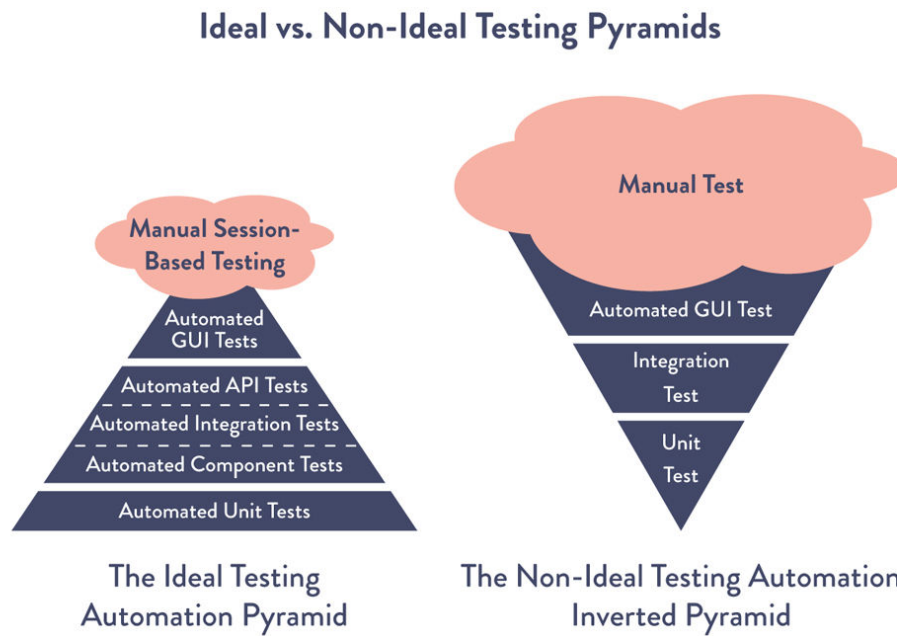


Figure 10.2: The Ideal and Non-Ideal Automated Testing Pyramids

Source: Martin Fowler, “TestPyramid,” MartinFowler.com.

If we find that unit or acceptance tests are too difficult and expensive to write and maintain, it’s likely that we have an architecture that is too tightly coupled, where strong separation between our module boundaries no longer exists (or maybe never existed). In this case, we will need to create a more loosely coupled system so modules can be independently tested without integration environments. Acceptance test suites for even the most complex applications that run in minutes are possible.

Ensure Tests Run Quickly (In Parallel, If Necessary)

Because we want our tests to run quickly, we need to design our tests to run in parallel, potentially across many different

servers. We may also want to run different categories of tests in parallel. For example, when a build passes our acceptance tests, we may run our performance testing in parallel with our security testing, as shown in [Figure 10.3](#). We may or may not allow manual exploratory testing until the build has passed all our automated tests—which enables faster feedback but may also allow manual testing on builds that will eventually fail.

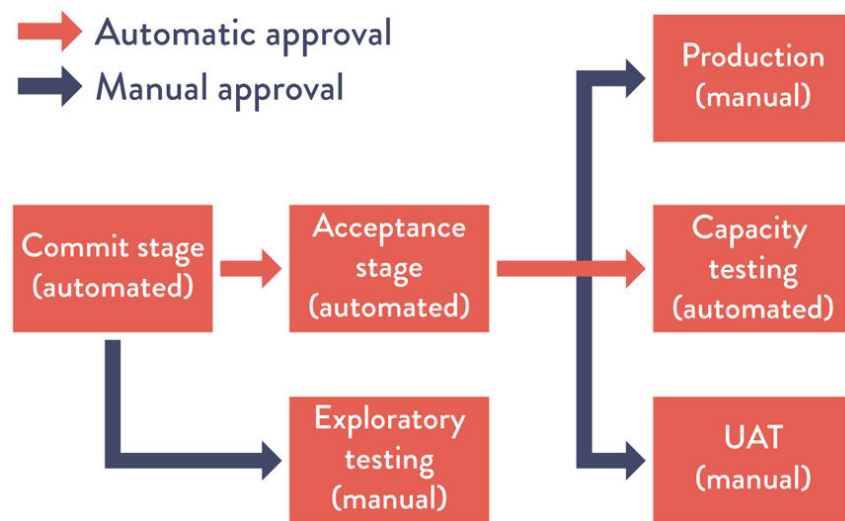


Figure 10.3: Running Automated and Manual Tests in Parallel

Source: Humble and Farley, *Continuous Delivery*, Kindle edition, location 3868.

We make any build that passes all our automated tests available to use for exploratory testing, as well as for other forms of manual or resource-intensive testing (such as performance testing). We want to do all such testing as frequently as is possible and practical, either continually or on a schedule.

Any tester (which includes all our developers) should use the latest build that has passed all the automated tests, as opposed to waiting for developers to flag a specific build as ready to test. By doing this, we ensure that testing happens as early in the process as possible.

Write Our Automated Tests before We Write the Code ("Test-Driven Development")

One of the most effective ways to ensure we have reliable automated testing is to write those tests as part of our daily work, using techniques such as *test-driven development* (TDD) and *acceptance test-driven development* (ATDD). This is when we begin every change to the system by first writing an automated test that validates the expected behavior fails, and then we write the code to make the tests pass.

This technique was developed by Kent Beck in the late 1990s as part of Extreme Programming, and has the following three steps:¹⁹

- Ensure the tests fail. "Write a test for the next bit of functionality you want to add." Check in.
- Ensure the tests pass. "Write the functional code until the test passes." Check in.
- "Refactor both new and old code to make it well structured." Ensure the tests pass. Check in again.

These automated test suites are checked into version control alongside our code, which provides a living, up-to-date specification of the system. Developers wishing to understand how to use the system can look at this test suite to find working examples of how to use the system's API.^{¶¶}

Automate as Many of Our Manual Tests as Possible

Our goal is to find as many code errors through our automated test suites, reducing our reliance on manual testing. In her 2013 presentation at Flowcon titled "On the Care and Feeding of Feedback Cycles," Elisabeth Hendrickson observed, "Although testing can be automated, creating

quality cannot. To have humans executing tests that should be automated is a waste of human potential.”²¹

By doing this, we enable all our testers (which, of course, includes developers) to work on high-value activities that cannot be automated, such as exploratory testing or improving the test process itself.

However, merely automating all our manual tests may create undesired outcomes—we do not want automated tests that are unreliable or generate false positives (i.e., tests that should have passed because the code is functionally correct but failed due to problems such as slow performance, causing timeouts, uncontrolled starting state, or unintended state due to using database stubs or shared test environments).²²

Unreliable tests that generate false positives create significant problems: they waste valuable time (e.g., forcing developers to re-run the test to determine whether there is actually a problem), increase the overall effort of running and interpreting our test results, and often result in stressed developers ignoring test results entirely or turning off the automated tests in favor of focusing on creating code.

The result is always the same: we detect the problems later, the problems are more difficult to fix, and our customers have worse outcomes, which in turn creates stress across the value stream.

To mitigate this, a small number of reliable, automated tests are almost always preferable over a large number of manual or unreliable automated tests. Therefore, we focus on automating only the tests that genuinely validate the business goals we are trying to achieve. If abandoning a test results in production defects, we should add it back to our manual test suite, with the ideal of eventually automating it.

As Gary Gruver, formerly VP of Quality Engineering, Release Engineering, and Operations for [Macys.com](https://www.macys.com), described, “For a large retailer e-commerce site, we went from running 1,300 manual tests that we ran every ten days to running only ten automated tests upon every code commit—it’s far better to run a few tests that we trust than to run tests that aren’t reliable. Over time, we grew this test suite to having hundreds of thousands of automated tests.”²³

In other words, we start with a small number of reliable automated tests and add to them over time, creating an ever-increasing level of assurance that we will quickly detect any changes to the system that take us out of a deployable state.

CONTINUOUS LEARNING

Some may think that writing automated tests isn't possible when writing features in a two-week sprint. But Angie Jones, Java champion and test automation expert, rightly points out that without automation, teams that only ship features without test automation are shipping risk and building up technical debt.²⁴

She outlines three strategies for teams to ship both features and test automation in their sprints:

- **Collaborate:** By collaborating with the business, testers, and developers, you ensure you automate the right things and have others contributing in parallel.
- **Automate strategically:** Using a hybrid approach can help teams think about test coverage, using APIs and smart design to get coverage across their scenarios.
- **Build incrementally:** Start by building what you need. As you continue to build additional features, adding more tests using a TDD (test-driven development) framework will help you think like both a tester and a developer, and will help you write more testable code.

Integrate Performance Testing into Our Test Suite

All too often, we discover that our application performs poorly during integration testing or after it has been

deployed to production. Performance problems, such as when things slow down over time, are often difficult to detect, going unnoticed until it is too late (e.g., database tables without an index). And many problems are difficult to solve, especially when they are caused by architectural decisions we made or unforeseen limitations of our networking, database, storage, or other systems.

Our goal is to write and run automated performance tests that validate our performance across the entire application stack (code, database, storage, network, virtualization, etc.) as part of the deployment pipeline so we detect problems early, when the fixes are cheapest and fastest.

By understanding how our application and environments behave under a production-like load, we can do a far better job at capacity planning, as well as detecting conditions such as:

- when our database query times grow non-linearly (e.g., we forget to create database indexes, and page load goes from one hundred minutes to thirty seconds)
- when a code change causes the number of database calls, storage use, or network traffic to increase ten-fold

When we have acceptance tests that are able to be run in parallel, we can use them as the basis of our performance tests. For instance, suppose we run an e-commerce site and have identified “search” and “checkout” as two high-value operations that must perform well under load. To test this, we may run thousands of parallel search acceptance tests simultaneously with thousands of parallel checkout tests.

Due to the large amount of compute and I/O that is required to run performance tests, creating a performance

testing environment can easily be more complex than creating the production environment for the application itself. Because of this, we may build our performance testing environment at the start of any project and ensure that we dedicate whatever resources are required to build it early and correctly. To find performance problems early, we should log performance results and evaluate each performance run against previous results. For instance, we might fail the performance tests if performance deviates more than 2% from the previous run.

Integrate Non-Functional Requirements Testing into Our Test Suite

In addition to testing that our code functions as designed and performs under production-like loads, we also want to validate every other attribute of the system we care about. These are often called non-functional requirements, which include availability, scalability, capacity, security, and so forth.

Many of these requirements are achieved through the correct configuration of our environments, so we must also build automated tests to validate that our environments have been built and configured properly. For example, we want to enforce the consistency and correctness of the following, which many non-functional requirements rely upon (e.g., security, performance, availability):

- supporting applications, databases, libraries, etc.
- language interpreters, compilers, etc.
- operating systems (e.g., audit logging enabled, etc.)
- all dependencies

When we use infrastructure as code configuration management tools (e.g., Terraform, Puppet, Chef, Ansible, Salt, Bosh), we can use the same testing frameworks that we

use to test our code to also test that our environments are configured and operating correctly (e.g., encoding environment tests into cucumber or gherkin tests). We should also run any security hardening checks as part of our automated tests to ensure that everything is configured securely and correctly (e.g., server-spec).

At any point in time, our automated tests can validate that we have a green build and that we are in a deployable state. Now, we must create an Andon cord so that when someone breaks the deployment pipeline, we take all necessary steps to get back into a green build state.

Pull Our Andon Cord When the Deployment Pipeline Breaks

When we have a green build in our deployment pipeline, we have a high degree of confidence that our code and environment will operate as designed when we deploy our changes into production.

In order to keep our deployment pipeline in a green state, we will create a virtual Andon cord, similar to the physical one in the Toyota Production System. Whenever someone introduces a change that causes our build or automated tests to fail, no new work is allowed to enter the system until the problem is fixed. And if someone needs help to resolve the problem, they can bring in whatever help they need.

When our deployment pipeline is broken, at a minimum, we notify the entire team of the failure, so anyone can either fix the problem or roll-back the commit. We may even configure the version control system to prevent further code commits until the first stage (i.e., builds and unit tests) of the deployment pipeline is back in a green state. If the problem was due to an automated test generating a false positive error, the offending test should either be rewritten or

removed.*** Every member of the team should be empowered to roll back the commit to get back into a green state.

Randy Shoup, former engineering director for Google App Engine, wrote about the importance of bringing the deployment back into a green state.

We prioritize the team goals over individual goals—whenever we help someone move their work forward, we help the entire team. This applies whether we’re helping someone fix the build or an automated test, or even performing a code review for them. And of course, we know that they’ll do the same for us, when we need help. This system worked without a lot of formality or policy—everyone knew that our job was not just “write code,” but it was to “run a service.” This is why we prioritized all quality issues, especially those related to reliability and scaling, at the highest level, treating them as Priority 0 “show-stopper” problems. From a systems perspective, these practices keep us from slipping backwards.²⁵

When later stages of the deployment pipeline fail, such as acceptance tests or performance tests, instead of stopping all new work, we will have developers and testers on call who are responsible for fixing these problems immediately. They should also create new tests that run at an earlier stage in the deployment pipeline to catch any future regressions. For example, if we discover a defect in our acceptance tests, we should write a unit test to catch the problem. Similarly, if we discover a defect in exploratory testing, we should write a unit or acceptance test.

To increase the visibility of automated test failures, we should create highly visible indicators so that the entire team can see when our build or automated tests are failing. Many

teams have created highly visible build lights that get mounted on a wall, indicating the current build status, or other fun ways of telling the team the build is broken, including lava lamps, playing a voice sample or song, klaxons, traffic lights, and so forth.

In many ways, this step is more challenging than creating our builds and test servers—those were purely technical activities, whereas this step requires changing human behavior and incentives. However, continuous integration and continuous delivery require these changes, as we explore in the next section.

Why We Need to Pull the Andon Cord

The consequence of not pulling the Andon cord and immediately fixing any deployment pipeline issues results in the all too familiar problem where it becomes ever more difficult to bring our applications and environment back into a deployable state. Consider the following situation:

- Someone checks in code that breaks the build or our automated tests, but no one fixes it.
- Someone else checks in another change onto the broken build, which also doesn't pass our automated tests—but no one sees the failing test results that would have enabled us to see the new defect, let alone to actually fix it.
- Our existing tests don't run reliably, so we are very unlikely to build new tests. (Why bother? We can't even get the current tests to run.)

When this happens, our deployments to any environment become as unreliable as when we had no automated tests or were using a waterfall method, where the majority of our problems are being discovered in production. The inevitable

outcome of this vicious cycle is that we end up where we started, with an unpredictable “stabilization phase” that takes weeks or months where our whole team is plunged into crisis, trying to get all our tests to pass, taking shortcuts because of deadline pressures, and adding to our technical debt.^{†††}

CONTINUOUS LEARNING

Data backs up the importance of automated testing. Research from DORA's 2019 *State of DevOps Report* shows that teams that use automated testing achieve superior continuous integration. Smart investments in automated testing improve our continuous integration. The report calls out that “automated tests can be a significant force-multiplier when used across several teams in an organization” and can contribute to elite performance.²⁷

The essential components of automated testing include tests that are:

- **Reliable:** A failure signals a real defect, and when tests pass, developers are confident the code will run successfully in production.
- **Consistent:** Each code commit should trigger a set of tests, providing feedback to developers.
- **Fast and reproducible:** Tests should complete in ten minutes or less so developers can quickly reproduce and fix failures in personal environments.
- **Inclusive:** Testing shouldn't just be for testers, and the best outcomes are seen when developers practice test-driven development.

The importance of exploratory and manual testing is also supported by research. DORA's

2018 *State of DevOps Report* found that testing throughout the software delivery life cycle contributes to continuous delivery outcomes and elite performance. In addition to automated testing, this includes:²⁸

- continuously reviewing and improving test suites to better find defects and keep complexity/cost under control
- allowing testers to work alongside developers throughout the software development and delivery process
- performing manual test activities such as exploratory testing, usability testing, and acceptance testing throughout the delivery process

Conclusion

In this chapter, we have created a comprehensive set of automated tests to confirm that we have a green build that is still in a passing and deployable state. We have organized our test suites and testing activities into a deployment pipeline. We have also created the cultural norm of doing whatever it takes to get back into a green build state if someone introduces a change that breaks any of our automated tests.

By doing this, we set the stage for implementing continuous integration, which allows many small teams to independently and safely develop, test, and deploy code into production, delivering value to customers.

* Bland described that at Google, one of the consequences of having so many talented developers was that it created “imposter syndrome,” a term coined by psychologists to informally describe people who are unable to internalize their accomplishments. Wikipedia states that “despite external evidence of their

competence, those exhibiting the syndrome remain convinced that they are frauds and do not deserve the success they have achieved. Proof of success is dismissed as luck, timing, or as a result of deceiving others into thinking they are more intelligent and competent than they believe themselves to be.”⁴

- † They created training programs, pushed the famous *Testing on the Toilet* newsletter (which they posted in the bathrooms), created the Test Certified roadmap and certification program, and led multiple “fix-it” days (i.e., improvement blitzes), which helped teams improve their automated testing processes so they could replicate the amazing outcomes that the GWS team was able to achieve.
- ‡ In Development, *continuous integration* often refers to the continuous integration of multiple code branches into trunk and ensuring that it passes unit tests. However, in the context of continuous delivery and DevOps, continuous integration also mandates running on production-like environments and passing acceptance and integration tests. Jez Humble and David Farley disambiguate these by calling the latter CI+. In this book, *continuous integration* will always refer to CI+ practices.¹³
- § If we create containers in our deployment pipeline and have an architecture such as microservices, we can enable each developer to build immutable artifacts where developers assemble and run all the service components in an environment identical to production on their workstation. This enables developers to build and run more tests on their workstation instead of on testing servers, giving us even faster feedback on their work.
- ¶ We may even require that these tools are run before changes are accepted into version control (e.g., get pre-commit hooks). We may also run these tools within the developer integrated development environment (IDE; where the developer edits, compiles, and runs code), which creates an even faster feedback loop.
- ** We can also use containers, such as Docker, as the packaging mechanism. Containers enable the capability to write once, run anywhere. These containers are created as part of our build process and can be quickly deployed and run in any environment. Because the same container is run in every environment, we help enforce the consistency of all our build artifacts.
- †† It is exactly this problem that led to the development of continuous integration practices.
- ‡‡ There is a broad category of architectural and testing techniques used to handle the problems of tests requiring input from external integration points, including “stubs,” “mocks,” “service virtualization,” and so forth. This becomes even more important for acceptance and integration testing, which place far more reliance on external states.

- §§ We should do this only when our teams already value automated testing—this type of metric is easily gamed by developers and managers.
- ¶¶ Nachi Nagappan, E. Michael Maximilien, and Laurie Williams (from Microsoft Research, IBM Almaden Labs, and North Carolina State University, respectively) conducted a study that showed teams using TDD produced code 60%–90% better in terms of defect density than non-TDD teams, while taking only 15%–35% longer.²⁰
- *** If the process for rolling back the code is not well known, a potential countermeasure is to schedule a *pair programmed rollback* so that it can be better documented.
- ††† This is sometimes called the *water-Scrum-fall anti-pattern*, which refers to when an organization claims to be using Agile-like practices, but, in reality, all testing and defect fixing are performed at the end of the project.²⁶

11

ENABLE AND PRACTICE CONTINUOUS INTEGRATION

In the previous chapter, we created automated testing practices to ensure that developers get fast feedback on the quality of their work. This becomes even more important as we increase the number of developers and the number of branches they work on in version control.

The ability to “branch” in version control systems was created primarily to enable developers to work on different parts of the software system in parallel without the risk of individual developers checking in changes that could destabilize or introduce errors into trunk (sometimes also called master or mainline).*

However, the longer developers are allowed to work in their branches in isolation, the more difficult it becomes to integrate and merge everyone’s changes back into trunk. In fact, integrating those changes becomes exponentially more difficult as we increase the number of branches and the number of changes in each code branch.

Integration problems result in a significant amount of rework to get back into a deployable state, including conflicting changes that must be manually merged or merges that break our automated or manual tests, usually requiring multiple developers to successfully resolve. And because integration has traditionally been done at the end of the project, when it takes far longer than planned, we are often forced to cut corners to make the release date.

This causes another downward spiral: when merging code is painful, we tend to do it less often, making future merges even worse. Continuous integration was designed to solve this problem by making merging into trunk a part of everyone's daily work.

HP's LaserJet Firmware (2014)

The surprising breadth of problems that continuous integration solves, as well as the solutions themselves, are exemplified in Gary Gruver's experience as the director of engineering for HP's LaserJet Firmware division, which builds the firmware that runs all their scanners, printers, and multifunction devices.¹

The team consisted of four hundred developers distributed across the US, Brazil, and India. Despite the size of their team, they were moving far too slowly. For years, they were unable to deliver new features as quickly as the business needed.

Gruver described the problem thus, "Marketing would come to us with a million ideas to dazzle our customer, and we'd just tell them, 'Out of your list, pick the two things you'd like to get in the next six to twelve months.'"²

They were only completing two firmware releases per year, with the majority of their time spent porting code to support new products. Gruver estimated that only 5% of their time was spent creating new features—the rest of the time was spent on non-productive work associated with their technical debt, such as managing multiple code branches and manual testing, as shown below:³

- 20% on detailed planning (their poor throughput and high lead times were misattributed to faulty

estimation, and so, hoping to get a better answer, they were asked to estimate the work in greater detail)

- 25% spent porting code, all maintained on separate code branches
- 10% spent integrating their code between developer branches
- 15% spent completing manual testing

Gruver and his team created a goal of increasing the time spent on innovation and new functionality by a factor of ten. The team hoped this goal could be achieved through:⁴

- continuous integration and trunk-based development
- significant investment in test automation
- creation of a hardware simulator so tests could be run on a virtual platform
- the reproduction of test failures on developer workstations
- a new architecture to support running all printers off a common build and release

Before this, each product line would require a new code branch, with each model having a unique firmware build with capabilities defined at compile time.[†] The new architecture would have all developers working in a common code base, with a single firmware release supporting all LaserJet models built off of trunk, with printer capabilities being established at runtime in an XML configuration file.

Four years later, they had one codebase supporting all twenty-four HP LaserJet product lines being developed on trunk. Gruver admits trunk-based development requires a big mindset shift.⁶ Engineers thought it would never work, but

once they started they couldn't imagine ever going back. Over the years, several engineers left HP, and they would call to tell Gruver about how backward development was in their new companies, pointing out how difficult it is to be effective and release good code when there is no feedback that continuous integration gives them.⁷

However, trunk-based development required them to build more effective automated testing. Gruver observed, "Without automated testing, continuous integration is the fastest way to get a big pile of junk that never compiles or runs correctly."⁸ In the beginning, a full manual testing cycle required six weeks.

In order to have all firmware builds automatically tested, they invested heavily in their printer simulators and created a testing farm in six weeks—within a few years, two thousand printer simulators ran on six racks of servers that would load the firmware builds from their deployment pipeline. Their continuous integration (CI) system ran their entire set of automated unit, acceptance, and integration tests on builds from trunk, just as described in the previous chapter. Furthermore, they created a culture that halted all work anytime a developer broke the deployment pipeline, ensuring that developers quickly brought the system back into a green state.⁹

Automated testing created fast feedback that enabled developers to quickly confirm that their committed code actually worked. Unit tests would run on their workstations in minutes, and three levels of automated testing would run on every commit as well as every two and four hours. The final full regression testing would run every twenty-four hours. During this process, they:¹⁰

- reduced the build to one build per day, eventually doing ten to fifteen builds per day
- went from around twenty commits per day performed by a “build boss” to over one hundred commits per day performed by individual developers
- enabled developers to change or add between seventy-five thousand and one-hundred thousand lines of code each day
- reduced regression test times from six weeks to one day

This level of productivity could never have been supported prior to adopting continuous integration, when merely creating a green build required days of heroics. The resulting business benefits were astonishing.¹¹

- Time spent on driving innovation and writing new features increased from 5% of developer time to 40%.
- Overall development costs were reduced by approximately 40%.
- Programs under development were increased by about 140%.
- Development costs per program were decreased by 78%.

What Gruver’s experience shows is that after comprehensive use of version control, continuous integration is one of the most critical practices that enable the fast flow of work in our value stream, enabling many development teams to independently develop, test, and deliver value. Nevertheless, continuous integration remains a controversial practice.

The remainder of this chapter describes the practices required to implement continuous integration, as well as how to overcome common objections.

Small Batch Development and What Happens When We Commit Code to Trunk Infrequently

As we've described, whenever changes are introduced into version control that cause our deployment pipeline to fail, we quickly swarm the problem to fix it, bringing our deployment pipeline back into a green state.

However, significant problems result when developers work in long-lived private branches (also known as “feature branches”), only merging back into trunk sporadically, resulting in a large batch-size of changes. As described in the HP LaserJet example, what results is significant chaos and rework in order to get their code into a releasable state.

Jeff Atwood, founder of the Stack Overflow site and author of the *Coding Horror* blog, observes that while there are many branching strategies, they can all be put on the following spectrum:¹²

- **Optimize for individual productivity:** Every single person on the project works in their own private branch. Everyone works independently, and nobody can disrupt anyone else's work; however, merging becomes a nightmare. Collaboration becomes almost comically difficult—every person's work has to be painstakingly merged with everyone else's work to see even the smallest part of the complete system.
- **Optimize for team productivity:** Everyone works in the same common area. There are no branches, just a long, unbroken straight line of development. Commits

are simple, but each commit can break the entire project and bring all progress to a screeching halt.

Atwood's observation is absolutely correct—stated more precisely, the required effort to successfully merge branches back together increases exponentially as the number of branches increases. The problem lies not only in the rework this “merge hell” creates, but also in the delayed feedback we receive from our deployment pipeline. For instance, instead of performance testing against a fully integrated system happening continuously, it will likely happen only at the end of our process.

Furthermore, as we increase the rate of code production and add more developers, we increase the probability that any given change will impact someone else and increase the number of developers who will be impacted when someone breaks the deployment pipeline.

Here is one last troubling side effect of large batch size merges: when merging is difficult, we become less able and motivated to improve and refactor our code, because refactorings are more likely to cause rework for everyone else. When this happens, we are more reluctant to modify code that has dependencies throughout the codebase, which is (tragically) where we may have the highest payoffs.

This is how Ward Cunningham, developer of the first wiki, originally described technical debt: “when we do not aggressively refactor our codebase, it becomes more difficult to make changes and to maintain over time, slowing down the rate at which we can add new features.”¹³

Solving this problem was one of the primary reasons behind the creation of continuous integration and trunk-based development practices, to optimize for team productivity over individual productivity. We'll go into a little

more detail about adopting trunk-based development practices in the next section of this book.

Adopt Trunk-Based Development Practices

Our countermeasure to large batch size merges is to institute continuous integration and trunk-based development practices, where all developers check their code into trunk at least once per day. Checking in code this frequently reduces our batch size to the work performed by our entire developer team in a single day. The more frequently developers check their code into trunk, the smaller the batch size and the closer we are to the theoretical ideal of single-piece flow.

Frequent code commits to trunk mean we can run all automated tests on our software system as a whole and receive alerts when a change breaks some other part of the application or interferes with the work of another developer. And because we can detect merge problems when they are small, we can correct them faster.

We may even configure our deployment pipeline to reject any commits (e.g., code or environment changes) that take us out of a deployable state. This method is called *gated commits*, where the deployment pipeline first confirms that the submitted change will successfully merge, build as expected, and pass all the automated tests before actually being merged into trunk. If not, the developer will be notified, allowing corrections to be made without impacting anyone else in the value stream.

The discipline of daily code commits also forces us to break our work down into smaller chunks while still keeping trunk in a working, releasable state. Version control becomes an integral mechanism of how the team communicates with each other—everyone has a better shared understanding of the system, is aware of the state of the deployment pipeline,

and can help each other when it breaks. As a result, we achieve higher quality and faster deployment lead times.

Having these practices in place, we can now again modify our definition of “done” (modification is in bold text): “At the end of each development interval, we must have integrated, tested, working, and potentially shippable code, demonstrated in a production-like environment, **created from trunk using a one-click process, and validated with automated tests.**”

Adhering to this revised definition of done helps us further ensure the ongoing testability and deployability of the code we’re producing. By keeping our code in a deployable state, we are able to eliminate the common practice of having a separate test and stabilization phase at the end of the project.

CASE STUDY

Continuous Integration at Bazaarvoice (2012)

Ernest Mueller, who helped engineer the DevOps transformation at National Instruments, later helped transform the development and release processes at Bazaarvoice in 2012.¹⁴ Bazaarvoice supplies customer-generated content (e.g., reviews, ratings) for thousands of retailers, such as Best Buy, Nike, and Walmart.

At that time, Bazaarvoice had \$120 million in revenue and was preparing for an IPO.[‡] The business was primarily driven by the Bazaarvoice Conversations application, a monolithic Java application comprising nearly five million lines of code dating back to 2006, spanning fifteen thousand files. The service ran on 1,200

servers across four data centers and multiple cloud service providers.¹⁵

Partially as a result of switching to an Agile development process and two-week development intervals, there was a tremendous desire to increase release frequency from their current ten-week production release schedule. They had also started to decouple parts of their monolithic application, breaking it down into microservices.

Their first attempt at a two-week release schedule was in January of 2012. Mueller observed, “It didn’t go well. It caused massive chaos, with forty-four production incidents filed by our customers. The major reaction from management was basically ‘Let’s not ever do that again.’”¹⁶

Mueller took over the release processes shortly afterward, with the goal of doing biweekly releases without causing customer downtime. The business objectives for releasing more frequently included enabling faster A/B testing (described in upcoming chapters) and increasing the flow of features into production. Mueller identified three core problems:¹⁷

- Lack of test automation made any level of testing during the two-week intervals inadequate to prevent large-scale failures.
- The version control branching strategy allowed developers to check in new code right up to the production release.
- The teams running microservices were also performing independent releases, which were

often causing issues during the monolith release.

Mueller concluded that the monolithic Conversations application deployment process needed to be stabilized, which required continuous integration. In the six weeks that followed, developers stopped doing feature work to focus instead on writing automated testing suites, including unit tests in JUnit, regression tests in Selenium, and getting a deployment pipeline running in TeamCity. “By running these tests all the time, we felt like we could make changes with some level of safety. And most importantly, we could immediately find when someone broke something, as opposed to discovering it only after it’s in production.”¹⁸

They also changed to a trunk/branch release model, where every two weeks they created a new dedicated release branch, with no new commits allowed to that branch unless there was an emergency—all changes would be worked through a sign-off process, either per-ticket or per-team through their internal wiki. That branch would go through a QA process, which would then be promoted into production. The improvements to predictability and quality of the releases were startling:¹⁹

- **January 2012 release:** forty-four customer incidents (continuous integration effort begins)
- **March 6, 2012 release:** five days late, five customer incidents
- **March 22, 2012 release:** on time, one customer incident

- **April 5, 2012 release:** on time, zero customer incidents

Mueller further described how successful this effort was:

We had such success with releases every two weeks, we went to weekly releases, which required almost no changes from the engineering teams. Because releases became so routine, it was as simple as doubling the number of releases on the calendar and releasing when the calendar told us to.

Seriously, it was almost a non-event. The majority of changes required were in our customer service and marketing teams, who had to change their processes, such as changing the schedule of their weekly customer emails to make sure customers knew that feature changes were coming.

After that, we started working toward our next goals, which eventually led to speeding up our testing times from three plus hours to less than an hour, reducing the number of environments from four to three (Dev, Test, Production, eliminating Staging), and moving to a full continuous delivery model where we enable fast, one-click deployments.²⁰

By systematically identifying and addressing three core problems, this case study illustrates the power of practices like feature freezes (in this case to work on automated testing) and using trunk-based development to enable small batch sizes and accelerate release cycles.

CONTINUOUS LEARNING

Continuous integration makes it easy for teams to get fast feedback, and it contributes to continuous delivery and elite performance. Research shows this capability is important, with the 2014–2019 *State of DevOps Reports* backing up the stories shared in this chapter with data.

Trunk-based development is likely the most controversial practice discussed in this book. However, data from DORA's 2016 and 2017 *State of DevOps Reports* is clear: trunk-based development predicts higher throughput, better stability, and better availability if they follow these practices.²¹

- have three or fewer active branches in the application's code repository
- merge branches to trunk at least daily
- don't have code freezes or integration phases

The benefits of continuous integration and trunk-based development extend beyond our ability to deliver software. DORA's research shows that it contributes to higher job satisfaction and lower rates of burnout.²²

Conclusion

In this chapter, we discussed the automation capabilities and behavioral practices that will allow us to ship our “done” code quickly and often. We have created the cultural norm of

developing on trunk and checking in code at least once a day. These practices and norms allow us to scale, accepting code from several or hundreds of developers. We will be able to ship code at any time, without painful code freezes or integration phases.

While convincing developers may be difficult at first, once they see the extraordinary benefits, they will likely become lifetime converts, as the HP LaserJet and Bazaarvoice examples illustrate. Continuous integration practices set the stage for the next step, which is automating the deployment process and enabling low-risk releases.

* Branching in version control has been used in many ways, but is typically used to divide work between team members by release, promotion, task, component, technology platforms, and so forth.

† Compile flags (`#define` and `#ifdef`) were used to enable/disable code execution for presence of copiers, paper size supported, and so on.⁵

‡ The production release was delayed due to their (successful) IPO.

12

AUTOMATE AND ENABLE LOW-RISK RELEASES

Chuck Rossi was the director of release engineering at Facebook for a decade. One of his responsibilities was overseeing the daily code push. In 2012, Rossi described their process as follows:

Starting around 1 PM, I switch over to “operations mode” and work with my team to get ready to launch the changes that are going out to [Facebook.com](https://www.facebook.com) that day. This is the more stressful part of the job and really relies heavily on my team’s judgment and past experience. We work to make sure that everyone who has changes going out is accounted for and is actively testing and supporting their changes.¹

Just prior to the production push, all developers with changes going out must be present and check in on their IRC chat channel—any developers not present have their changes automatically removed from the deployment package.² Rossi continued, “If everything looks good and our test dashboards and canary tests^{*} are green, we push the big red button and the entire [Facebook.com](https://www.facebook.com) server fleet gets the new code delivered. Within twenty minutes, thousands and thousands of machines are up on new code with no visible impact to the people using the site.”^{3†}

Later that year, Rossi doubled their software release frequency to twice daily.⁴ He explained that the second code push gave engineers not on the US West Coast the ability to

“move and ship as quickly as any other engineer in the company,” and also gave everyone a second opportunity each day to ship code and launch features.⁵




Figure 12.1: Number of Developers Deploying per Week at Facebook

Source: Chuck Rossi, “Ship early and ship twice as often.”

Kent Beck, the creator of the Extreme Programming (XP) methodology, one of the leading proponents of Test-Driven Development and a technical coach at Facebook, further comments on their code release strategy in an article posted on his Facebook page:

Chuck Rossi made the observation that there seem to be a fixed number of changes Facebook can handle in one deployment. If we want more changes, we need more deployments. This has led to a steady increase in deployment pace over the past five years, from weekly to daily to thrice daily deployments of our PHP code and from six to four to two week cycles for deploying our mobile apps. This improvement has been driven primarily by the release engineering team.⁶

By using continuous integration and making code deployment a low-risk process, Facebook has enabled code deployment to be a part of everyone's daily work and sustain developer productivity. This requires that code deployment be automated, repeatable, and predictable. In the practices described in the book so far, even though our code and environments have been tested together, most likely we are not deploying to production very often because deployments are manual, time consuming, painful, tedious, and error prone, and they often involve an inconvenient and unreliable handoff between Development and Operations.

And because it is painful, we tend to do it less and less frequently, resulting in another self-reinforcing downward spiral. By deferring production deployments, we accumulate ever-larger differences between the code to be deployed and what's running in production, increasing the deployment batch size. As deployment batch size grows, so does the risk of unexpected outcomes associated with the change, as well as the difficulty fixing them.

In this chapter, we reduce the friction associated with production deployments, ensuring that they can be performed frequently and easily, either by Operations or Development. We do this by extending our deployment pipeline.

Instead of merely continually integrating our code in a production-like environment, we will enable the promotion into production of any build that passes our automated test and validation process, either on demand (i.e., at the push of a button) or automatically (i.e., any build that passes all the tests is automatically deployed).

Because of the number of practices presented, extensive footnotes are provided with numerous examples and

additional information without interrupting the presentation of concepts in the chapter.

Automate Our Deployment Process

Achieving outcomes like those at Facebook requires that we have an automated mechanism that deploys our code into production. Especially if we have a deployment process that has existed for years, we need to fully document the steps in the deployment process, such as in a value stream mapping exercise, which we can assemble in a workshop or document incrementally (e.g., in a wiki).

Once we have the process documented, our goal is to simplify and automate as many of the manual steps as possible, such as:

- packaging code in ways suitable for deployment
- creating preconfigured virtual machine images or containers
- automating the deployment and configuration of middleware
- copying packages or files onto production servers
- restarting servers, applications, or services
- generating configuration files from templates
- running automated smoke tests to make sure the system is working and correctly configured
- running testing procedures
- scripting and automating database migrations

Where possible, we will re-architect to remove steps, particularly those that take a long time to complete. We also want to not only reduce our lead times but also the number of

handoffs as much as possible in order to reduce errors and loss of knowledge.

Having developers focus on automating and optimizing the deployment process can lead to significant improvements in deployment flow, such as ensuring that small application configuration changes no longer need new deployments or new environments. However, this requires that Development works closely with Operations to ensure that all the tools and processes we co-create can be used downstream, as opposed to alienating Operations or reinventing the wheel.

Many tools that provide continuous integration and testing also support the ability to extend the deployment pipeline so that validated builds can be promoted into production, typically after the production acceptance tests are performed (e.g., CircleCI, the Jenkins Build Pipeline plugin, Go.cd, Microsoft Visual Studio Team Services, and Pivotal Concourse).

The requirements for our deployment pipeline include:

- **Deploying the same way to every environment:** By using the same deployment mechanism for every environment (e.g., development, test, and production), our production deployments are likely to be far more successful, since we know that they have been successfully performed many times already earlier in the pipeline.
- **Smoke testing our deployments:** During the deployment process, we should test that we can connect to any supporting systems (e.g., databases, message buses, external services) and run a single test transaction through the system to ensure that our system is performing as designed. If any of these tests fail, we should fail the deployment.

- **Ensure we maintain consistent environments:** In previous steps, we created a single-step environment build process so that the development, test, and production environments had a common build mechanism. We must continually ensure that these environments remain synchronized.

Of course, when any problems occur during deployment, we pull the Andon cord and swarm the problem until the problem is resolved, just as we do when our deployment pipeline fails in any of the earlier steps.

CASE STUDY

Daily Deployments at CSG International (2013)

CSG International is North America's largest SaaS-based customer care and billing provider, with over sixty-five million subscribers and a tech stack that covers everything from Java to mainframe.⁷ Scott Prugh, Chief Architect and VP of Development, led an effort to improve the predictability and reliability of their software releases. To achieve this, they doubled their release frequency from two per year to four per year (halving their deployment interval from twenty-eight weeks to fourteen weeks).⁸

Although the Development teams were using continuous integration to deploy their code into test environments daily, the production releases were being performed by the Operations team. Prugh observed,

It was as if we had a “practice team” that practiced daily (or even more frequently) in low-risk test environments, perfecting their processes and tools. But our production

“game team” got very few attempts to practice, only twice per year. Worse, they were practicing in the high-risk production environments, which were often very different than the pre-production environments with different constraints—the development environments were missing many production assets such as security, firewalls, load balancers, and a SAN [storage area network].⁹

To solve this problem, they created a Shared Operations Team (SOT) that was responsible for managing all the environments (development, test, production) performing daily deployments into those development and test environments, as well as doing production deployments and releases every fourteen weeks. Because the SOT was doing deployments every day, any problems they encountered that were left unfixed would simply occur again the next day. This created tremendous motivation to automate tedious or error-prone manual steps and to fix any issues that could potentially happen again. Because the deployments were performed nearly one hundred times before the production release, most problems were found and fixed long before then.¹⁰

Doing this revealed problems that were previously only experienced by the Ops team, which were then problems for the entire value stream to solve. The daily deployments enabled daily feedback on which practices worked and which didn't.¹¹

They also focused on making all their environments look as similar as possible, including the restricted

security access rights and load balancers. Prugh writes, “We made non-production environments as similar to production as possible, and we sought to emulate production constraints in as many ways as possible. Early exposure to production-class environments altered the designs of the architecture to make them friendlier in these constrained or different environments. Everyone gets smarter from this approach.”¹²

Prugh also observes:

We have experienced many cases where changes to database schemas are either 1) handed off to a DBA team for them to “go and figure it out” or 2) automated tests that run on unrealistically small data sets (i.e., “100’s of MB vs. 100’s of GBs”), which led to production failures. In our old way of working, this would become a late-night blame game between teams trying to unwind the mess.

We created a development and deployment process that removed the need for handoffs to DBAs by cross-training developers, automating schema changes, and executing them daily. We created realistic load testing against sanitized customer data, ideally running migrations every day. By doing this, we run our service hundreds of times with realistic scenarios before seeing actual production traffic.^{13†}

Their results were astonishing. By doing daily deployments and doubling the frequency of production releases, the number of production incidents went

down by 91%, MTTR went down by 80%, and the deployment lead time required for the service to run in production in a “fully hands-off state” went from fourteen days to one day.¹⁵

Prugh reported that deployments became so routine that the Ops team was playing video games by the end of the first day. In addition to deployments going more smoothly for Dev and Ops, 50% of the time the customer received the value in half the time.¹⁶




Figure 12.2: Daily Deployments at CSG International

Daily deployments and increasing release frequency resulted in a decrease in the number of production incidents and MTTR.

Source: “DOES15 - Scott Prugh & Erica Morrison—Conway & Taylor Meet the Strangler (v2.0),” YouTube video, 29:39, posted by DevOps Enterprise Summit, November 5, 2015, <https://www.youtube.com/watch?v=tKdIHCL0DUg>.

This case study underscores how more frequent deployments are good for Development, QA, Operations, and the customer. Frequent deployments allowed problems to be identified earlier, motivated teams to fix errors sooner and led to the delivery of cleaner code faster.

Enable Automated Self-Service Deployments

Consider the following quote from Tim Tischler, Director of Operations Automation at Nike, Inc., who describes the common experience of a generation of developers: “As a developer, there has never been a more satisfying point in my career than when I wrote the code, when I pushed the button to deploy it, when I could see the production metrics confirm that it actually worked in production, and when I could fix it myself if it didn’t.”¹⁷

Developers’ ability to self-deploy code into production, to quickly see happy customers when their feature works, and to quickly fix any issues without having to open up a ticket with Operations has diminished over the last decade—in part as a result of a need for control and oversight, perhaps driven by security and compliance requirements.

The resulting common practice is for Operations to perform code deployments because this is a common implementation of separation of duties, a widely accepted practice to reduce the risk of production outages and fraud, which requires that no one person has end-to-end control over a process, such as software delivery. However, DORA’s research shows that we can substantially improve software delivery performance through implementing separation of duties through code review. In this implementation, we require another developer to review and approve every code change.¹⁸

This implementation is substantially more effective when supplemented with an automated test suite that must pass before changes can be submitted. All deployments are then self-serviced through an automated system that meets the requirements described in the next section.

When there are shared goals that span Development and Operations and there is transparency, responsibility, and accountability for deployment outcomes, it doesn't matter who performs the deployment. In fact, we may even have other roles, such as testers or project managers, able to deploy to certain environments so they can get their own work done quickly, such as setting up demonstrations of specific features in test or UAT environments.

To better enable fast flow, we want a code promotion process that can be performed by either Development or Operations, ideally without any manual steps or handoffs. This affects the following steps:

- **Build:** Our deployment pipeline must create packages from version control that can be deployed to any environment, including production.
- **Test:** Anyone should be able to run any or all of our automated test suite on their workstation or on our test systems.
- **Deploy:** Anybody should be able to deploy these packages to any environment where they have access, executed by running scripts that are also checked into version control.

Integrate Code Deployment into the Deployment Pipeline

Once the code deployment process is automated, we can make it part of the deployment pipeline. Consequently, our deployment automation must provide the following capabilities:

- ensure that packages created during the continuous integration process are suitable for deployment into production

- show the readiness of production environments at a glance
- provide a push-button, self-service method for any suitable version of the packaged code to be deployed into production
- record automatically, for auditing and compliance purposes, which commands were run on which machines when, who authorized it, and what the output was; record also the hash of all binaries deployed and the source control version from which all configuration information and scripts were taken
- run a smoke test to ensure the system is operating correctly and the configuration settings, including items such as database connection strings, are correct
- provide fast feedback for the deployer so they can quickly determine whether their deployment was successful (e.g., did the deployment succeed, is the application performing as expected in production, etc.)

Our goal is to ensure that deployments are fast—we don't want to have to wait hours to determine whether our code deployment succeeded or failed and then need hours to deploy any needed code fixes. Now that we have technologies such as containers, it is possible to complete even the most complex deployments in seconds or minutes.

In DORA's 2019 *State of DevOps Report*, data showed that elite performers deploy on demand, with deployment lead times measured in minutes or hours, while the lowest performers had deployment lead times measured in months. Over the years, many deployment statistics improved, with Puppet Labs' 2014 *State of DevOps Report* showing high performers typically deploying code in one hour to one day

and the lowest performers having lead times of six months or more, as was presented in the first edition of this book.¹⁹

By building this capability, we now have a “deploy code” button that allows us to safely and quickly promote changes to our code and our environments into production through our deployment pipeline.

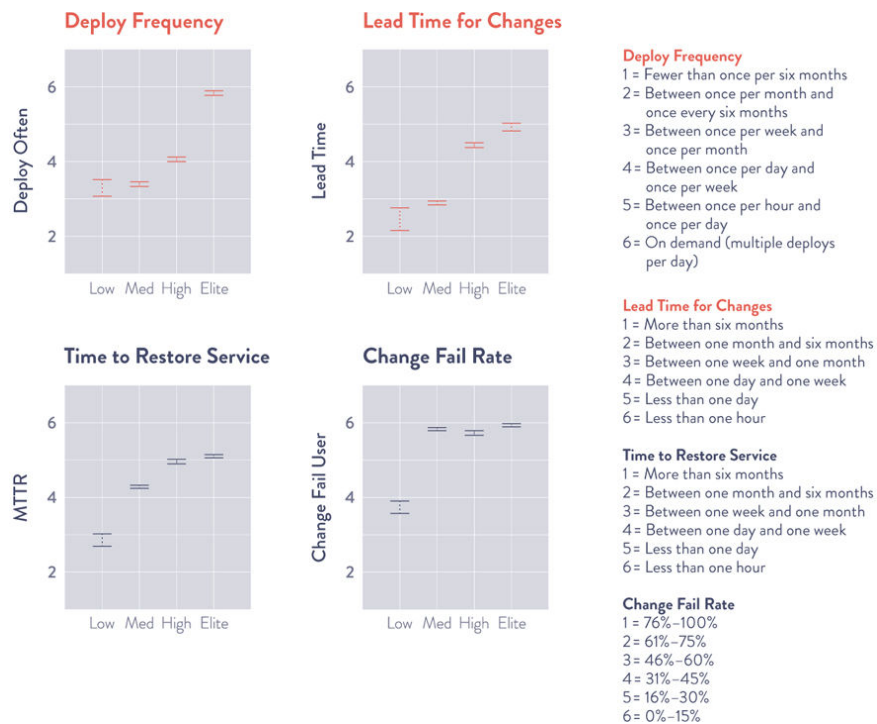


Figure 12.3: Elite and High Performers Achieve Faster Deployment Lead Times and MTTR (2019)[§]

Source: Forsgren, et. al., *Accelerate: State of DevOps Report*, 2019.

CASE STUDY

Etsy—Self-Service Developer Deployment: An Example of Continuous Deployment (2014)

Unlike at Facebook, where deployments are managed by release engineers, at Etsy deployments are performed by anyone who wants to perform a deployment, such as Development, Operations, or Infosec. The deployment

process at Etsy became so safe and routine that new engineers could perform a production deployment on their first day at work. Etsy board members and even dogs have deployed to production!²⁰

As Noah Sussman, a test architect at Etsy, wrote, “By the time 8am rolls around on a normal business day, 15 or so people and dogs are starting to queue up, all of them expecting to collectively deploy up to 25 changesets before the day is done.”²¹

Engineers who want to deploy their code first go to a chat room, where engineers add themselves to the deploy queue, see the deployment activity in progress, see who else is in the queue, broadcast their activities, and get help from other engineers when they need it. When it’s an engineer’s turn to deploy, they are notified in the chat room.²²

The goal at Etsy has been to make it easy and safe to deploy into production with the fewest number of steps and the least amount of ceremony. Likely before the developer even checks in code, they will run on their workstation all 4,500 unit tests, which takes less than one minute. All calls to external systems, such as databases, have been stubbed out.²³

After they check their changes into trunk in version control, over seven thousand automated trunk tests are instantly run on their continuous integration (CI) servers. Sussman writes, “Through trial-and-error, we’ve settled on about 11 minutes as the longest that the automated tests can run during a push. That leaves time to re-run the tests once during a deployment [if someone breaks something and needs to fix it], without going too far past the 20 minute time limit.”²⁴

If all the tests were run sequentially, Sussman states that “the 7,000 trunk tests would take about half an hour to execute. So we split these tests up into subsets, and distribute those onto the 10 machines in our Jenkins [CI] cluster... . Splitting up our test suite and running many tests in parallel, gives us the desired 11 minute runtime.”²⁵

The next tests to run are the *smoke tests*, which are system-level tests that run cURL to execute PHPUnit test cases. Following these tests, the functional tests are run, which execute end-to-end GUI-driven tests on a live server—this server is either their QA environment or staging environment (nicknamed “Princess”), which is actually a production server that has been taken out of rotation, ensuring that it exactly matches the production environment.

Once it is an engineer’s turn to deploy, Erik Kastner writes, “you go to Deployinator [an internally developed tool, see [Figure 12.4](#)] and push the button to get it on QA. From there it visits Princess... . Then, when it’s ready to go live, you hit the ‘Prod’ button and soon your code is live, and everyone in IRC [chat channel] knows who pushed what code, complete with a link to the diff. For anyone not on IRC, there’s the email that everyone gets with the same information.”²⁶

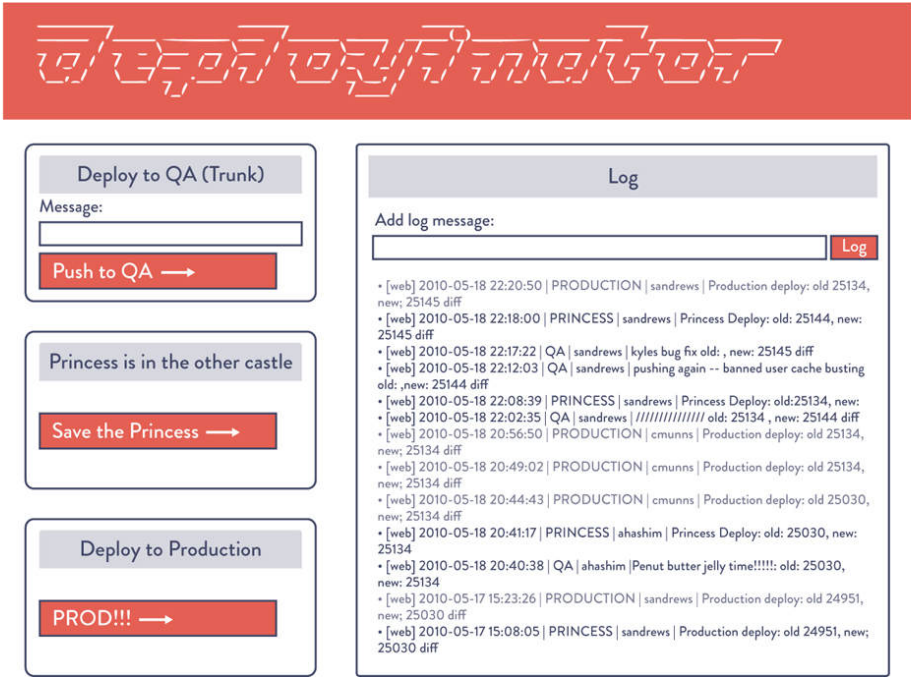


Figure 12.4: The Deployinator Console at Etsy

Source: Erik Kastner, “Quantum of Deployment,” [CodeasCraft.com](https://codeascraft.com/2010/05/20/quantum-of-deployment/), May 20, 2010, <https://codeascraft.com/2010/05/20/quantum-of-deployment/>.

In 2009, the deployment process at Etsy was a cause of stress and fear. By 2011, it had become a routine operation, happening twenty-five to fifty times per day, helping engineers get their code quickly into production, delivering value to their customers.

Automating the majority of the deployment process and including extensive automated testing created an easy and efficient self-service deployment pipeline, and teams experienced less stress and gained more confidence.

Decouple Deployments from Releases

In the traditional launch of a software project, releases are driven by our marketing launch date. On the prior evening, we deploy our completed software (or as close to complete as we could get) into production. The next morning, we

announce our new capabilities to the world, start taking orders, deliver the new functionality to customers, etc.

However, all too often things don't go according to plan. We may experience production loads that we never tested or designed for, causing our service to fail spectacularly, both for our customers and our organization. Worse, restoring service may require a painful rollback process or an equally risky *fix forward* operation, where we make changes directly in production. This can all be a truly miserable experience for workers. When everything is finally working, everyone breathes a sigh of relief, grateful that production deployments and releases don't happen more often.

Of course, we know that we need to be deploying more frequently to achieve our desired outcome of smooth and fast flow, not less frequently. To enable this, we need to decouple our production deployments from our feature releases. In practice, the terms *deployment* and *release* are often used interchangeably. However, they are two distinct actions that serve two very different purposes:

- Deployment is the installation of a specified version of software to a given environment (e.g., deploying code into an integration test environment or deploying code into production). Specifically, a deployment may or may not be associated with a release of a feature to customers.
- Release is when we make a feature (or set of features) available to all our customers or a segment of customers (e.g., we enable the feature to be used by 5% of our customer base). Our code and environments should be architected in such a way that the release of functionality does not require changing our application code.[¶]

In other words, when we conflate deployment and release, it makes it difficult to create accountability for successful outcomes—decoupling these two activities allows us to empower Development and Operations to be responsible for the success of fast and frequent deployments, while enabling product owners to be responsible for the successful business outcomes of the release (i.e., “Was building and launching the feature worth our time?”).

The practices described so far in this book ensure that we are doing fast and frequent production deployments throughout feature development, with the goal of reducing the risk and impact of deployment errors. The remaining risk is release risk, which is whether the features we put into production achieve the desired customer and business outcomes.

If we have extremely long deployment lead times, this dictates how frequently we can release new features to the marketplace. However, as we become able to deploy on demand, how quickly we expose new functionality to customers becomes a business and marketing decision, not a technical decision. There are two broad categories of release patterns we can use (and which can be combined):

- **Environment-based release patterns:** This is where we have two or more environments that we deploy into, but only one environment is receiving live customer traffic (e.g., by configuring our load balancers). New code is deployed into a non-live environment, and the release is performed moving traffic to this environment. These are extremely powerful patterns because they typically require little or no change to our applications. These patterns include *blue-green deployments*, *canary releases*, and

cluster immune systems, all of which will be discussed shortly.

- **Application-based release patterns:** This is where we modify our application so that we can selectively release and expose specific application functionality by small configuration changes. For instance, we can implement feature flags that progressively expose new functionality in production to the development team, all internal employees, 1% of our customers, or, when we are confident that the release will operate as designed, our entire customer base. As discussed earlier, this enables a technique called dark launching, where we stage all the functionality to be launched in production and test it with production traffic before our release. For instance, we may invisibly test our new functionality with production traffic for weeks before our launch in order to expose problems so that they can be fixed before our actual launch.

Environment-Based Release Patterns

Decoupling deployments from our releases dramatically changes how we work. We no longer have to perform deployments in the middle of the night or on weekends to lower the risk of negatively impacting customers. Instead, we can do deployments during typical business hours, enabling Ops to finally have normal working hours, just like everyone else.

This section focuses on environment-based release patterns, which require no changes to application code. We do this by having multiple environments to deploy into, but only one of them receives live customer traffic. By doing this, we can significantly decrease the risk associated with production releases and reduce the deployment lead time.

The Blue-Green Deployment Pattern

The simplest of the three patterns is blue-green deployment. In this pattern, we have two production environments: blue and green. At any time, only one of these is serving customer traffic (see [Figure 12.5](#)).




Figure 12.5: Blue-Green Deployment Patterns

Source: Humble and North, *Continuous Delivery*, 261.

To release a new version of our service, we deploy to the inactive environment, where we can perform our testing without interrupting the user experience. When we are confident that everything is functioning as designed, we execute our release by directing traffic to the blue environment. Thus, blue becomes live and green becomes staging. Rollback is performed by sending customer traffic back to the green environment.**

The blue-green deployment pattern is simple, and it is extremely easy to retrofit onto existing systems. It also has incredible benefits, such as enabling the team to perform deployments during normal business hours and conduct simple changeovers (e.g., changing a router setting, changing a symlink) during off-peak times. This alone can dramatically improve the work conditions for the team performing the deployment.

Dealing with Database Changes

Having two versions of our application in production creates problems when they depend upon a common database—when the deployment requires database schema changes or

adding, modifying, or deleting tables or columns, the database cannot support both versions of our application. There are two general approaches to solving this problem:

- **Create two databases (i.e., a blue and green database):** Each version—blue (old) and green (new)—of the application has its own database. During the release, we put the blue database into read-only mode, perform a backup of it, restore onto the green database, and finally switch traffic to the green environment. The problem with this pattern is that if we need to roll back to the blue version, we can potentially lose transactions if we don't manually migrate them from the green version first.
- **Decouple database changes from application changes:** Instead of supporting two databases, we decouple the release of database changes from the release of application changes by doing two things: First, we make only additive changes to our database; we never mutate existing database objects. Second, we make no assumptions in our application about which database version will be in production. This is very different from how we've been traditionally trained to think about databases, where we avoid duplicating data. The process of decoupling database changes from application changes was used by IMVU (among others) around 2009, enabling them to do fifty deployments per day, some of which required database changes.^{27††}

CASE STUDY

Dixons Retail—Blue-Green Deployment for Point-Of-Sale System (2008)

Dan Terhorst-North, technology and organizational change consultant, and Dave Farley, co-author of *Continuous Delivery*, were working on a project for Dixons Retail, a large British retailer, involving thousands of point-of-sale (POS) systems that resided in hundreds of retail stores and operated under a number of different customer brands. Although blue-green deployments are mostly associated with online web services, North and Farley used this pattern to significantly reduce the risk and changeover times for POS upgrades.²⁹

Traditionally, upgrading POS systems is a big-bang, waterfall project: the POS clients and the centralized server are upgraded at the same time, which requires extensive downtime (often an entire weekend), as well as significant network bandwidth to push out the new client software to all the retail stores. When things don't go entirely according to plan, it can be incredibly disruptive to store operations.

For this upgrade, there was not enough network bandwidth to upgrade all the POS systems simultaneously, which made the traditional strategy impossible. To solve this problem, they used the blue-green strategy and created two production versions of the centralized server software, enabling them to simultaneously support the old and new versions of the POS clients.

After they did this, weeks before the planned POS upgrade, they started sending out new versions of client POS software installers to the retail stores over the slow network links, deploying the new software onto the POS systems in an inactive state. Meanwhile, the old version kept running as normal.

When all the POS clients had everything staged for the upgrade (the upgraded client and server had tested together successfully, and new client software had been deployed to all clients), the store managers were empowered to decide when to release the new version.

Depending on their business needs, some managers wanted to use the new features immediately and released right away, while others wanted to wait. In either case, whether releasing features immediately or waiting, it was significantly better for the managers than having the centralized IT department choose for them when the release would occur.

This case study demonstrates how DevOps patterns can be universally applied to different technologies, often in very surprising ways but with the same fantastic outcomes.

The Canary and Cluster Immune System Release Patterns

The blue-green release pattern is easy to implement and can dramatically increase the safety of software releases. There are variants of this pattern that can further improve safety and deployment lead times using automation, but with the potential trade-off of additional complexity.

The *canary release pattern* automates the release process of promoting to successively larger and more critical environments as we confirm that the code is operating as designed. The term *canary release* comes from the tradition of coal miners bringing caged canaries into mines to provide early detection of toxic levels of carbon monoxide. If there was too much gas in the cave, it would kill the canaries before it killed the miners, alerting them to evacuate.

In this pattern, when we perform a release, we monitor how the software in each environment is performing. When something appears to be going wrong, we roll back; otherwise, we deploy to the next environment.†† Figure 12.6 shows the groups of environments Facebook created to support this release pattern:

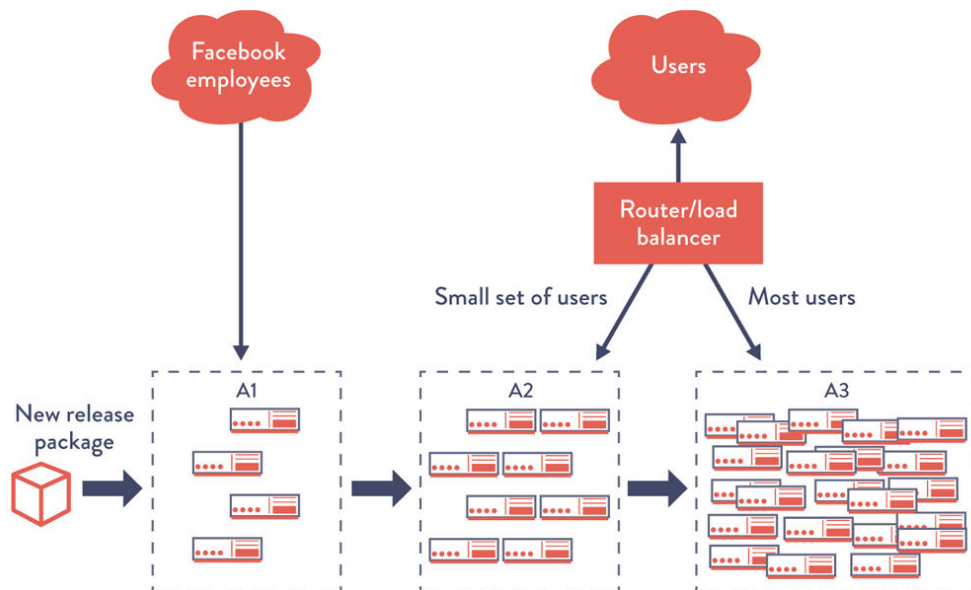


Figure 12.6: The Canary Release Pattern

Source: Humble and Farley, *Continuous Delivery*, 263.

- **A1 group:** production servers that only serve internal employees.
- **A2 group:** production servers that only serve a small percentage of customers and are deployed when certain acceptance criteria have been met (either automated or manual).
- **A3 group:** the rest of the production servers, which are deployed after the software running in the A2 cluster meets certain acceptance criteria.

The *cluster immune system* expands upon the canary release pattern by linking our production monitoring system with our release process and by automating the rollback of code when the user-facing performance of the production

system deviates outside of a predefined expected range, such as when the conversion rates for new users drops below our historical norms of 15%–20%.

There are two significant benefits to this type of safeguard. First, we protect against defects that are hard to find through automated tests, such as a web page change that renders some critical page element invisible (e.g., CSS change). Second, we reduce the time required to detect and respond to the degraded performance created by our change.^{SS}

Application-Based Patterns to Enable Safer Releases

In the previous section, we created environment-based patterns that allowed us to decouple our deployments from our releases by using multiple environments and by switching between which environment was live, which can be entirely implemented at the infrastructure level.

In this section, we describe application-based release patterns that we can implement in our code, allowing even greater flexibility in how we safely release new features to our customer, often on a per-feature basis. Because application-based release patterns are implemented in the application, these require involvement from Development.

Implement Feature Toggles

The primary way we enable application-based release patterns is by implementing feature toggles (also called feature flags), which provide us with the mechanism to selectively enable and disable features without requiring a production code deployment. Feature toggles can also control which features are visible and available to specific user segments (e.g., internal employees, segments of customers).

Feature toggles are usually implemented by wrapping application logic or UI elements with a conditional statement, where the feature is enabled or disabled based on a configuration setting stored somewhere. This can be as simple as an application configuration file (e.g., configuration files in JSON, XML), or it might be through a directory service or even a web service specifically designed to manage feature toggling.¶¶

Feature toggles also enable us to do the following:

- **Roll back easily:** Features that create problems or interruptions in production can be quickly and safely disabled by merely changing the feature toggle setting. This is especially valuable when deployments are infrequent—switching off one particular stakeholder’s features is usually much easier than rolling back an entire release.
- **Gracefully degrade performance:** When our service experiences extremely high loads that would normally require us to increase capacity or, worse, risk having our service fail in production, we can use feature toggles to reduce the quality of service. In other words, we can increase the number of users we serve by reducing the level of functionality delivered (e.g., reduce the number of customers who can access a certain feature, disable CPU-intensive features such as recommendations, etc.).
- **Increase our resilience through a service-oriented architecture:** If we have a feature that relies on another service that isn’t complete yet, we can still deploy our feature into production but hide it behind a feature toggle. When that service finally becomes available, we can toggle the feature on. Similarly,

when a service we rely upon fails, we can turn off the feature to prevent calls to the downstream service while keeping the rest of the application running.

- **Perform A/B testing:** Modern feature toggle frameworks, such as LaunchDarkly, Split, and Optimizely, also enable product teams to run experiments to test new features and see their impact on business metrics. In this way, we can demonstrate a causal relationship between new features and the outcomes we care about. This is an incredibly powerful tool that enables a scientific, hypothesis-driven approach to product development (a technique described later in this book).

To ensure that we find errors in features wrapped in feature toggles, our automated acceptance tests should run with all feature toggles on. (We should also test that our feature toggling functionality works correctly too!)

Perform Dark Launches

Feature toggles allow us to deploy features into production without making them accessible to users, enabling a technique known as *dark launching*. This is where we deploy all the functionality into production and then perform testing of new functionality while it is still invisible to customers. For large or risky changes, we often do this for weeks before the production launch, enabling us to safely test with the anticipated production-like loads.

For instance, suppose we dark launch a new feature that poses significant release risk, such as new search features, account creation processes, or new database queries. After all the code is in production, keeping the new feature disabled, we may modify user session code to make calls to new

functions—instead of displaying the results to the user, we simply log or discard the results.

For example, we may have 1% of our online users make invisible calls to a new feature scheduled to be launched to see how our new feature behaves under load. After we find and fix any problems, we progressively increase the simulated load by increasing the frequency and number of users exercising the new functionality. By doing this, we are able to safely simulate production-like loads, giving us confidence that our service will perform as it needs to.

Furthermore, when we launch a feature, we can progressively roll out the feature to small segments of customers, halting the release if any problems are found. That way, we minimize the number of customers who are given a feature only to have it taken away because we find a defect or are unable to maintain the required performance.

In 2009, when John Allspaw was VP of Operations at Flickr, he wrote to the Yahoo! executive management team about their dark launch process, saying it:

increases *everyone's* confidence almost to the point of apathy, as far as fear of load-related issues are concerned. I have no idea how many code deploys there were made to production on any given day in the past 5 years ... because for the most part I don't care, because those changes made in production have such a low chance of causing issues. When they have caused issues, everyone on the Flickr staff can find on a webpage *when* the change was made, *who* made the change, and exactly (line-by-line) *what* the change was.^{32**}

Later, when we have built adequate production telemetry in our application and environments, we can also enable faster feedback cycles to validate our business assumptions

and outcomes immediately after we deploy the feature into production.

By doing this, we no longer wait until a big-bang release to test whether customers want to use the functionality we build. Instead, by the time we announce and release our big feature, we have already tested our business hypotheses and run countless experiments to continually refine our product with real customers, which helps us validate that the features will achieve the desired customer outcomes.

CASE STUDY

Dark Launch of Facebook Chat (2008)

For nearly a decade, Facebook has been one of the most widely visited Internet sites, as measured by pages viewed and unique site users. In 2008, it had over seventy million daily active users, which created a challenge for the team that was developing the new Facebook Chat functionality.³⁴⁺⁺⁺

Eugene Letuchy, an engineer on the Chat team, wrote about how the number of concurrent users presented a huge software engineering challenge, “The most resource-intensive operation performed in a chat system is not sending messages. It is rather keeping each online user aware of the online-idle-offline states of their friends, so that conversations can begin.”³⁶

Implementing this computationally intensive feature was one of the largest technical undertakings ever at Facebook and took almost a year to complete.⁺⁺⁺ Part of the complexity of the project was due to the wide variety of technologies needed to achieve the desired performance, including C++, JavaScript, and PHP, as well

as their first use of Erlang in their back-end infrastructure.³⁷

Throughout the course of the year-long endeavor, the Chat team checked their code into version control, where it would be deployed into production at least once per day. At first, the Chat functionality was visible only to the Chat team. Later, it was made visible to all internal employees, but it was completely hidden from external Facebook users through Gatekeeper, the Facebook feature toggling service.

As part of their dark launch process, every Facebook user session, which runs JavaScript in the user browser, had a test harness loaded into it—the chat UI elements were hidden, but the browser client would send invisible test chat messages to the back-end chat service that was already in production, enabling them to simulate production-like loads throughout the entire project, allowing them to find and fix performance problems long before the customer release.

The Chat release and launch required only two steps: modifying the Gatekeeper configuration setting to make the Chat feature visible to some portion of external users, and having Facebook users load new JavaScript code that rendered the Chat UI and disabled the invisible test harness. If something went wrong, the two steps would be reversed.

When the launch day of Facebook Chat arrived, it was surprisingly successful and uneventful, seeming to scale effortlessly from zero to seventy million users overnight. During the release, they incrementally enabled the chat functionality to ever-larger segments of the customer population—first to all internal

Facebook employees, then to 1% of the customer population, then to 5%, and so forth. As Letuchy wrote, “The secret for going from zero to seventy million users overnight is to avoid doing it all in one fell swoop.”³⁸

In this case study, every Facebook user was part of a massive load testing program, which enabled the team to gain confidence that their systems could handle realistic production-like loads.

Survey of Continuous Delivery and Continuous Deployment in Practice

In *Continuous Delivery*, Jez Humble and David Farley define the term *continuous delivery*. The term *continuous deployment* was first mentioned by Tim Fitz in his blog post “Continuous Deployment at IMVU: Doing the impossible fifty times a day.” However, in 2015, during the development of the first edition of *The DevOps Handbook*, Jez Humble commented,

In the last five years, there has been confusion around the terms continuous delivery versus continuous deployment—and, indeed, my own thinking and definitions have changed since we wrote the book. Every organization should create their variations, based on what they need. The key thing we should care about is not the form, but the outcomes: deployments should be low-risk, push-button events we can perform on demand.³⁹

His updated definitions of continuous delivery and continuous deployment are as follows:

When all developers are working in small batches on trunk, or everyone is working off trunk in short-lived feature branches that get merged to trunk regularly, and

when trunk is always kept in a releasable state, and when we can release on demand at the push of a button during normal business hours, we are doing continuous delivery. Developers get fast feedback when they introduce any regression errors, which include defects, performance issues, security issues, usability issues, etc. When these issues are found, they are fixed immediately so that trunk is always deployable.

In addition to the above, when we are deploying good builds into production on a regular basis through self-service (being deployed by Dev or by Ops)—which typically means that we are deploying to production at least once per day per developer, or perhaps even automatically deploying every change a developer commits—this is when we are engaging in continuous deployment.⁴⁰

Defined this way, continuous delivery is the prerequisite for continuous deployment just as continuous integration is a prerequisite for continuous delivery. Continuous deployment is likely applicable in the context of web services that are delivered online. However, continuous delivery is applicable in almost every context where we desire deployments and releases that have high quality, fast lead times, and highly predictable, low-risk outcomes, including for embedded systems, COTS products, and mobile apps.

At Amazon and Google, most teams practice continuous delivery, although some perform continuous deployment—thus, there is considerable variation between teams in how frequently they deploy code and how deployments are performed. Teams are empowered to choose how to deploy based on the risks they are managing.

Similarly, most of the case studies presented in this book are also continuous delivery, such as the embedded software running on HP LaserJet printers, the CSG bill printing operations running on twenty technology platforms including a COBOL mainframe application, Facebook, and Etsy. These same patterns can be used for software that runs on mobile phones, ground control stations that control satellites, and so forth.

CONTINUOUS LEARNING

DORA's 2018 and 2019 *State of DevOps Reports* show that continuous delivery is a key predictor of elite performance. The research found this included both technical and cultural components:⁴¹

- Teams can deploy on demand to production or to end users throughout the software delivery life cycle.
- Fast feedback on the quality and deployability of the system is available to everyone on the team.
- Team members prioritize keeping the system in a deployable state.

→ CASE STUDY: NEW TO SECOND EDITION

Creating a Win-Win for Dev & Ops at CSG (2016)

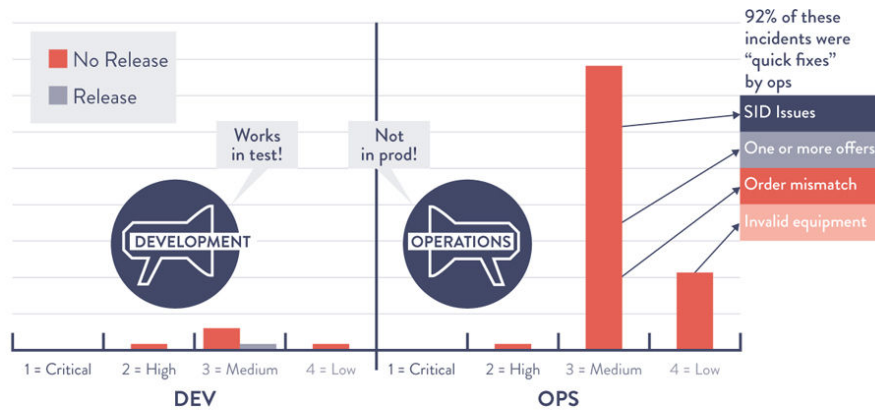
After successfully improving releases between 2012 and 2015, CSG further evolved their organizational structure to improve their day-to-day operational stance. At the DevOps Enterprise Summit in 2016, Scott Prugh, Chief

Architect and VP Software Engineering at the time, spoke about a dramatic organizational transformation that combined disparate development and operations teams into cross-functional build/run teams.

Prugh described the start of this journey:

We had made dramatic improvements in our release processes and release quality but continued to have escalations and conflicts with our operations team members. The development teams felt confident about their code quality and continued to push releases faster and more frequently.

On the other hand, our ops teams complained about production outages and the rapid changes breaking the environment. To battle these opposing forces, our change and program management teams ramped up their processes to improve coordination and to attempt to control the chaos. Sadly, this did little to improve production quality, our operational teams' experience, or the relationship between the development and operations teams.⁴²



Problem management and multi-tiered support (L1→L2→L3) won't save you.

Figure 12.7: How Structure Influences Behavior and Quality

Image courtesy of Scott Prugh.

To understand more of what was going on, the team dug into the incident data, which revealed some surprising and alarming trends:⁴³

- Release impact and incidents had improved almost 90% (from 201 incidents to 24).
- Release incidents represented 2% of the occurring incidents (98% were in production).
- And, 92% of these production incidents were quick restorations that were fixed by operations.

Prugh further observed, “We had basically improved our development stance significantly but had done little to improve the production operations environment. We got the exact result we had optimized for: great code quality and poor operations quality.”⁴⁴

In order to find a solution, Prugh asked the following questions:

- Were different organizational goals working against system goals?

- Did development's lack of operations understanding result in hard-to-run software?
- Did a lack of shared mission create a lack of empathy across teams?
- Did our handoffs contribute to elongated lead time?
- Did a lack of engineering skills in operations prevent improvements and encourage duct-tape engineering?

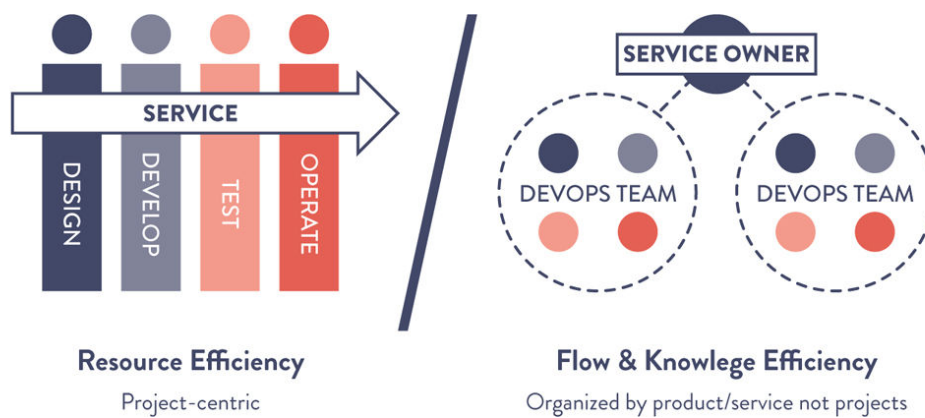


Figure 12.8: From Siloed Approach to Cross-Functional Teams

Image courtesy of Scott Prugh.

At the same time Prugh was making these discoveries, the customer escalations of production issues had repeatedly been raised to the executive leadership. CSG's customers were irate, and the executives asked what could be done to improve CSG's operational stance.

After several passes at the problem, Prugh suggested creating "Service Delivery Teams" that build and run the software. Basically, he suggested bringing together Dev and Ops onto one team.⁴⁵

At first, the proposal was viewed as polarizing. But after representing previous successes with shared operations teams, Prugh further argued that bringing Dev and Ops together would create a win-win for both teams by:

- improving understanding so the team could improve the entire delivery chain (development to operations)
- improving flow and knowledge efficiency and creating unified accountability for design, build, test, and operations
- making operations an engineering problem^{§§§}
- bring other benefits, like improving communication, reducing meetings, creating shared planning, improving collaboration, creating shared work visibility, and a shared leadership vision

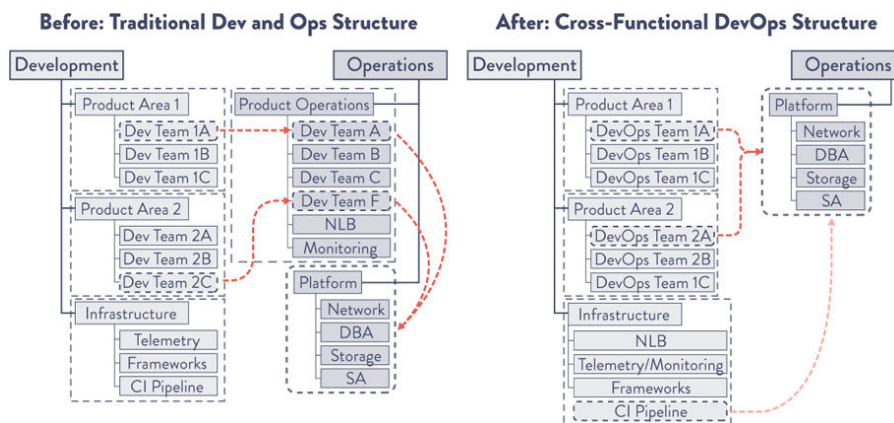


Figure 12.9: Conventional vs. Cross-Functional Structure

Image courtesy of Scott Prugh.

The next steps involved re-creating a new team structure of combined development and operations teams and leaders. Managers and leaders of the new

teams were selected from the current pool, and team members were re-recruited onto the new cross-functional teams. After the changes, development managers and leaders got real experience in running the software they had created.⁴⁶

It was a shocking experience. The leaders realized that creating build/run teams was only the first step in a very long journey. Erica Morrison, VP Software Engineering, recalls:

As I got more involved with the Network Load Balancer team, I quickly started to feel like I was in *The Phoenix Project*. While I had seen many parallels to the book in previous work experiences, it was nothing like this. There was invisible work/work in multiple systems: one system for stories, another for incidents, another for CRQs, another for new requests. And TONS of email. And some stuff wasn't in any system. My brain was exploding, trying to track it all.

The cognitive load from managing all the work was huge. It was also impossible to follow up with teams and stakeholders. Basically whoever screamed the loudest went to the top of the queue. Almost every item was a priority #1 due to the lack of a coordinated system to track and prioritize work.

We also realized that a ton of technical debt had accumulated, which prevented many critical vendor upgrades, leaving us on outdated hardware, software, and OS. There

was also a lack of standards. When we did have them, they were not universally applied or rolled out to production.

Critical people bottlenecks were also prolific, creating an unsustainable work environment.

Finally, all changes went through a traditional CAB [change advisory board] process, which created a massive bottleneck to get things approved. Additionally, there was little automation supporting the change process, making every change manual, not traceable, and very high risk.”⁴⁷

To address these issues, the CSG team took a multi-pronged approach. First, they created a bias for action and culture change by applying the learnings from John Shook’s *Model of Change: “Change Behavior to Change Culture.”* The leadership team understood that in order to change the culture they had to change behavior, which would then affect values and attitudes and result in eventual culture change.

Next, the team brought in developers to supplement the operational engineers and demonstrate what would be possible with great automation and engineering applied to key operational problems. Automation was added to traffic reporting and device reporting. Jenkins was used to orchestrate and automate basic jobs that were being done by hand. Telemetry and monitoring to CSG’s common platform (StatHub) were added. And finally, deployments were automated to remove errors and support rollback.

The team then invested in getting all the config in code and version control. This included CI practices as well as lower environments that could test and practice deployments to devices that would not impact production. The new processes and tools allowed easy peer review since all code went through a pipeline on the way to production.

Finally, the team invested in bringing all work into a single backlog. This included automation to pull tickets from many systems into one common system that the team could collaborate in and thus prioritize the work.

Erica Morrison recalls her final learnings:

We worked really hard in this journey to bring some of the best practices we know from Development to the Ops world. There have been many things in this journey that went well, but there were a lot of misses and a lot of surprises. Probably our biggest surprise is how hard Ops really is. It's one thing to read about it but quite another to experience it first-hand.

Also, change management is scary and not visible to developers. As a development team, we had no details about the change process and changes going in. We now deal with change on a daily basis. Change can be overwhelming and consume much of what a team does each day.

We also reaffirmed that change is one of the key intersection points of opposing goals between Dev and Ops. Developers want their changes to make it to production as soon as

possible. But when Ops is responsible for putting that change in and dealing with the fallout, it creates a natural reaction to want to go slower. We now understand that getting Dev and Ops working together to both design and implement the change creates a win-win that improves both speed and stability.⁴⁸

This updated case study from CSG shows the importance of the development and operations teams working together to design and implement changes, creating low-risk, low-drama releases.

Conclusion

As the Facebook, Etsy, and CSG examples have shown, releases and deployments do not have to be high-risk, high-drama affairs that require tens or hundreds of engineers to work around the clock to complete. Instead, they can be made entirely routine and a part of everyone's daily work. By doing this, we can reduce our deployment lead times from months to minutes, allowing our organizations to quickly deliver value to our customers without causing chaos and disruption. Furthermore, by having Dev and Ops work together, we can finally make Operations work humane.

-
- * A canary release test is when software is deployed to a small group of production servers to make sure nothing terrible happens to them with live customer traffic.
 - † The Facebook front-end codebase is primarily written in PHP. In 2010, to increase site performance, the PHP code was converted into C++ by their internally developed HipHop compiler, which was then compiled into a 1.5 GB executable. This file was then copied onto production servers using BitTorrent, enabling the copy operation to be completed in fifteen minutes.
 - ‡ In their experiments, they found that SOT teams were successful regardless of whether they were managed by Development or Operations, as long as the

teams were staffed with the right people and were dedicated to SOT success.¹⁴

- § We have replaced this figure with an updated metric for the second edition. This new figure more accurately portrays experiences over the last five years.
- ¶ Operation Desert Shield may serve as an effective metaphor. Starting on August 7, 1990, thousands of men and materials were safely deployed over four months into the production theater, culminating in a single, multidisciplinary, highly coordinated release.
- ** Other ways that we can implement the blue-green pattern include setting up multiple Apache/NGINX web servers to listen on different physical or virtual interfaces; employing multiple virtual roots on Windows IIS servers bound to different ports; using different directories for every version of the system, with a symbolic link determining which one is live (e.g., as Capistrano does for Ruby on Rails); running multiple versions of services or middleware concurrently, with each listening on different ports; using two different data centers and switching traffic between the data centers, instead of using them merely as hot- or warm-spares for disaster recovery purposes (incidentally, by routinely using both environments, we are continually ensuring that our disaster recovery process works as designed); or using different availability zones in the cloud.
- †† This pattern is also commonly referred to as the expand/contract pattern, which Timothy Fitz described when he said, “We do not change (mutate) database objects, such as columns or tables. Instead, we first expand, by adding new objects, then, later, contract by removing the old ones.”²⁸ Furthermore, increasingly, there are technologies that enable virtualization, versioning, labeling, and rolling back databases, such as Redgate, Delphix, DBMaestro, and Datical, as well as open source tools, such as DBDeploy, that make database changes dramatically safer and faster.
- ‡‡ Note that canary releases require having multiple versions of our software running in production simultaneously. However, because each additional version we have in production creates additional complexity to manage, we should keep the number of versions to a minimum. This may require the use of the expand/contract database pattern described earlier.
- §§ The cluster immune system was first documented by Eric Ries while working at IMVU. This functionality is also supported by Etsy in their Feature API library, as well as by Netflix.³⁰
- ¶¶ One sophisticated example of such a service is Facebook’s Gatekeeper, an internally developed service that dynamically selects which features are visible to specific users based on demographic information such as location, browser type, and user profile data (age, gender, etc.). For instance, a particular feature could be configured so that it is only accessible by internal employees, 10% of

their user base, or only users between the ages of twenty-five and thirty-five. Other examples include the Etsy Feature API and the Netflix Archaius library.³¹

*** Similarly, as Chuck Rossi, Director of Release Engineering at Facebook, described, “All the code supporting every feature we’re planning to launch over the next six months has already been deployed onto our production servers. All we need to do is turn it on.”³³

††† By 2015, Facebook had over one billion active users, growing 17% over the previous year.³⁵

‡‡‡ This problem has a worst-case computational characteristic of $O(n^3)$. In other words, the compute time increases exponentially as the function of the number of online users, the size of their friend lists, and the frequency of online/offline state change.

§§§ The team could now inject and cross-train engineering principles/skills into operations, and they could evolve operations from being solely a process activity to an engineering activity with true continuous improvement. This would shift operations left and stop the duct-tape engineering activities.

13

ARCHITECT FOR LOW-RISK RELEASES

Almost every well-known DevOps exemplar has had near-death experiences due to architectural problems, such as in the stories presented about LinkedIn, Google, eBay, Amazon, and Etsy. In each case, they were able to successfully migrate to a more suitable architecture that addressed their current problems and organizational needs.

This is the principle of evolutionary architecture—Jez Humble observes that architecture of “any successful product or organization will necessarily evolve over its life cycle.”¹ Before his tenure at Google, Randy Shoup served as Chief Engineer and Distinguished Architect at eBay from 2004 to 2011. He observes that “both eBay and Google are each on their fifth entire rewrite of their architecture from top to bottom.”²

He reflects, “Looking back with 20/20 hindsight, some technology [and architectural choices] look prescient and others look shortsighted. Each decision most likely best served the organizational goals at the time. If we had tried to implement the 1995 equivalent of micro-services out of the gate, we would have likely failed, collapsing under our own weight and probably taking the entire company with us.”^{3*}

The challenge is how to keep migrating from the architecture we have to the architecture we need. In the case of eBay, when they needed to re-architect, they would first do a small pilot project to prove to themselves that they understood the problem well enough to even undertake the effort. For instance, when Shoup’s team was planning on

moving certain portions of the site to full-stack Java in 2006, they looked for the area that would get them the biggest bang for their buck by sorting the site pages by revenue produced. They chose the highest revenue areas, stopping when there was not enough of a business return to justify the effort.⁵

What Shoup's team did at eBay is a textbook example of evolutionary design, using a technique called the *strangler fig application pattern*—instead of “ripping out and replacing” old services with architectures that no longer support our organizational goals, we put the existing functionality behind an API and avoid making further changes to it. All new functionality is then implemented in the new services that use the new desired architecture, making calls to the old system when necessary.

The strangler fig application pattern is especially useful for helping migrate portions of a monolithic application or tightly coupled services to one that is more loosely coupled. All too often, we find ourselves working within an architecture that has become too tightly coupled and too interconnected, often having been created years (or decades) ago.

The consequences of overly tight architectures are easy to spot: every time we attempt to commit code into trunk or release code into production, we risk creating global failures (e.g., we break everyone else's tests and functionality or the entire site goes down). To avoid this, every small change requires enormous amounts of communication and coordination over days or weeks, as well as approvals from any group that could potentially be affected.

Deployments become problematic as well—the number of changes that are batched together for each deployment grows, further complicating the integration and test effort,

and increasing the already high likelihood of something going wrong.

Even deploying small changes may require coordinating with hundreds (or even thousands) of other developers, with any one of them able to create a catastrophic failure, potentially requiring weeks to find and fix the problem. (This results in another symptom: “My developers spend only 15% of their time coding—the rest of their time is spent in meetings.”)

These all contribute to an extremely unsafe system of work, where small changes have seemingly unknowable and catastrophic consequences. It also often contributes to a fear of integrating and deploying our code, and the self-reinforcing downward spiral of deploying less frequently.

From an enterprise architecture perspective, this downward spiral is the consequence of the Second Law of Architectural Thermodynamics, especially in large, complex organizations. Charles Betz, author of *Architecture and Patterns for IT: Service Management, Resource Planning, and Governance: Making Shoes for the Cobbler’s Children*, observes, “[IT project owners] are not held accountable for their contributions to overall system entropy.”⁶ In other words, reducing our overall complexity and increasing the productivity of all our development teams is rarely the goal of an individual project.

In this chapter, we will describe steps we can take to reverse the downward spiral, review the major architectural archetypes, examine the attributes of architectures that enable developer productivity, testability, deployability, and safety, as well as evaluate strategies that allow us to safely migrate from whatever current architecture we have to one

that better enables the achievement of our organizational goals.

An Architecture that Enables Productivity, Testability, and Safety

In contrast to a tightly coupled architecture that can impede everyone's productivity and ability to safely make changes, a loosely coupled architecture with well-defined interfaces that enforce how modules connect with each other promotes productivity and safety. It enables small, productive teams that are able to make small changes that can be safely and independently deployed. And because each service also has a well-defined API, it enables easier testing of services and the creation of contracts and SLAs between teams.

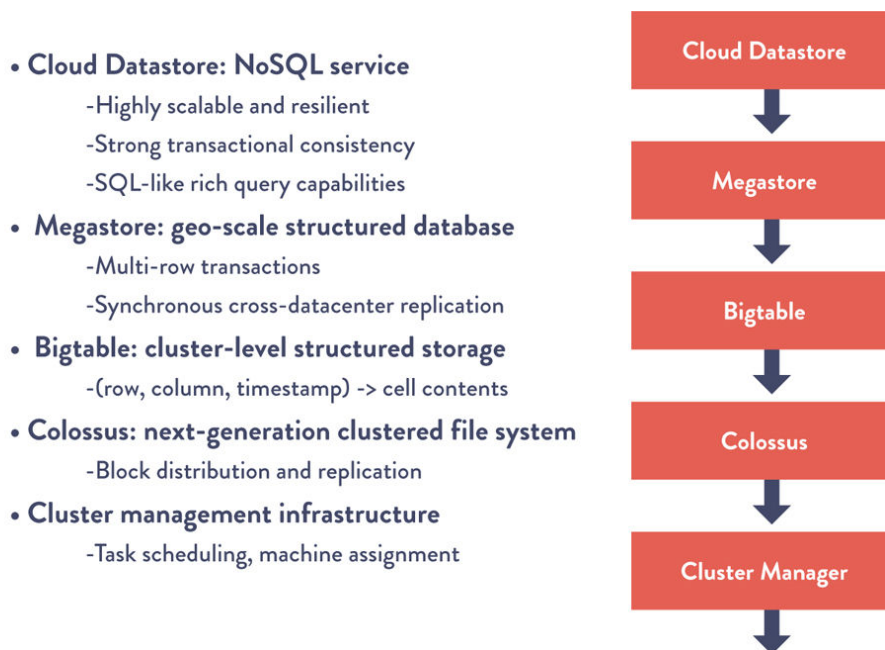


Figure 13.1: Google Cloud Datastore

Source: Shoup, "From the Monolith to Micro-services."

As Randy Shoup describes,

This type of architecture has served Google extremely well—for a service like Gmail, there's five or six other layers of services underneath it, each very focused on a

very specific function. Each service is supported by a small team, who builds it and runs their functionality, with each group potentially making different technology choices. Another example is the Google Cloud Datastore service, which is one of the largest NoSQL services in the world—and yet it is supported by a team of only about eight people, largely because it is based on layers upon layers of dependable services built upon each other.⁷

This kind of service-oriented architecture allows small teams to work on smaller and simpler units of development that each team can deploy independently, quickly, and safely. Shoup notes, “Organizations with these types of architectures, such as Google and Amazon, show how it can impact organizational structures, [creating] flexibility and scalability. These are both organizations with tens of thousands of developers, where small teams can still be incredibly productive.”⁸

Architectural Archetypes: Monoliths vs. Microservices

At some point in their history, most DevOps organizations were hobbled by tightly coupled, monolithic architectures that—while extremely successful at helping them achieve product/market fit—put them at risk of organizational failure once they had to operate at scale (e.g., eBay’s monolithic C++ application in 2001, Amazon’s monolithic Obidos application in 2001, Twitter’s monolithic Rails front end in 2009, and LinkedIn’s monolithic Leo application in 2011). In each of these cases, the organization was able to re-architect their systems and set the stage not only to survive but also to thrive and win in the marketplace.

Monolithic architectures are not inherently bad—in fact, they are often the best choice for an organization early in a

product life cycle. As Randy Shoup observes,

There is no one perfect architecture for all products and all scales. Any architecture meets a particular set of goals or range of requirements and constraints, such as time to market, ease of developing functionality, scaling, etc. The functionality of any product or service will almost certainly evolve over time—it should not be surprising that our architectural needs will change as well. What works at scale 1x rarely works at scale 10x or 100x.⁹

The major architectural archetypes are shown in [Table 13.1](#). Each row indicates a different evolutionary need for an organization, with each column giving the pros and cons of each of the different archetypes. As the table shows, a monolithic architecture that supports a startup (e.g., rapid prototyping of new features and potential pivots or large changes in strategies) is very different from an architecture that needs hundreds of teams of developers, each of whom must be able to independently deliver value to the customer. By supporting evolutionary architectures, we can ensure that our architecture always serves the current needs of the organization.

Table 13.1: Architectural Archetypes

	Pros	Cons
Monolithic v1 (all functionality in one application)	<ul style="list-style-type: none"> • Simple at first • Low interprocess latencies • Single code base; one deployment unit • Resource efficient at small scales 	<ul style="list-style-type: none"> • Coordination overhead increases as team grows • Poor enforcement of modularity • Poor scaling • All-or-nothing deploy (downtime, failures) • Long build times
Monolithic v2 (sets of monolithic tiers: “front end presentation,” “application server,” “database layer”)	<ul style="list-style-type: none"> • Simple at first • Join queries are easy • Single schema, deployment • Resource efficient at small scales 	<ul style="list-style-type: none"> • Tendency for increased coupling over time • Poor scaling and redundancy (all-or-nothing, vertical only) • Difficult to tune properly • All-or-nothing schema management
Microservice (modular, independent, graph relationship vs. tiers, isolated persistence)	<ul style="list-style-type: none"> • Each unit is simple • Independent scaling and performance • Independent testing and deployment • Can optimally tune performance 	<ul style="list-style-type: none"> • Many cooperating units • Many small repos • Requires more sophisticated tooling and dependency management • Network latencies

	(caching, replication, etc.)	
--	---------------------------------	--

Source: Shoup, "From the Monolith to Micro-services."

CASE STUDY

Evolutionary Architecture at Amazon (2002)

One of the most studied architecture transformations occurred at Amazon. In an interview with ACM Turing Award-winner and Microsoft Technical Fellow Jim Gray, Amazon CTO Werner Vogels explains that [Amazon.com](https://www.amazon.com) started in 1996 as a "monolithic application, running on a web server, talking to a database on the back end. This application, dubbed Obidos, evolved to hold all the business logic, all the display logic, and all the functionality that Amazon eventually became famous for: similarities, recommendations, Listmania, reviews, etc."¹⁰

As time went by, Obidos grew too tangled, with complex sharing relationships meaning individual pieces could not be scaled as needed. Vogels tells Gray that this meant "many things that you would like to see happening in a good software environment couldn't be done anymore; there were many complex pieces of software combined into a single system. It couldn't evolve anymore."¹¹

Describing the thought process behind the new desired architecture, he tells Gray, "We went through a period of serious introspection and concluded that a service-oriented architecture would give us the level of isolation that would allow us to build many software components rapidly and independently."¹²

Vogels notes, “The big architectural change that Amazon went through in the past five years [from 2001–2005] was to move from a two-tier monolith to a fully distributed, decentralized services platform serving many different applications. A lot of innovation was necessary to make this happen, as we were one of the first to take this approach.”¹³ The lessons from Vogel’s experience at Amazon that are important to our understanding of architecture shifts include the following:

- **Lesson 1:** When applied rigorously, strict service orientation is an excellent technique to achieve isolation; you achieve a level of ownership and control that was not seen before.
- **Lesson 2:** Prohibiting direct database access by clients makes performing scaling and reliability improvements to your service state possible without involving your clients.
- **Lesson 3:** Development and operational process greatly benefits from switching to service orientation. The services model has been a key enabler in creating teams that can innovate quickly with a strong customer focus. Each service has a team associated with it, and that team is completely responsible for the service—from scoping out the functionality to architecting, building, and operating it.

The extent to which applying these lessons enhances developer productivity and reliability is breathtaking. In 2011, Amazon was performing approximately fifteen thousands deployment per day.¹⁴

By 2015, they were performing nearly 136,000 deployments per day.¹⁵

This case study illustrates how evolving from a monolithic structure to one of microservices helped to decouple the architecture, allowing it to better serve the needs of the organization.

Use the Strangler Fig Application Pattern to Safely Evolve Our Enterprise Architecture

The term *strangler fig application* was coined by Martin Fowler in 2004 after he was inspired by seeing massive strangler vines during a trip to Australia, writing, “They seed in the upper branches of a fig tree and gradually work their way down the tree until they root in the soil. Over many years they grow into fantastic and beautiful shapes, meanwhile strangling and killing the tree that was their host.”¹⁶

If we have determined that our current architecture is too tightly coupled, we can start safely decoupling parts of the functionality from our existing architecture. By doing this, we enable teams supporting the decoupled functionality to independently develop, test, and deploy their code into production with autonomy and safety, and reduce architectural entropy.

As described earlier, the strangler fig application pattern involves placing existing functionality behind an API, where it remains unchanged, and implementing new functionality using our desired architecture, making calls to the old system when necessary. When we implement the strangler fig applications, we seek to access all services through versioned APIs, also called *versioned services* or *immutable services*.¹⁷

Versioned APIs enable us to modify the service without impacting the callers, which allows the system to be more loosely coupled—if we need to modify the arguments, we create a new API version and migrate teams who depend on our service to the new version. After all, we are not achieving our re-architecting goals if we allow our new strangler fig application to get tightly coupled into other services (e.g., connecting directly to another service’s database).

If the services we call do not have cleanly defined APIs, we should build them or at least hide the complexity of communicating with such systems within a client library that has a cleanly defined API.

By repeatedly decoupling functionality from our existing tightly coupled system, we move our work into a safe and vibrant ecosystem where developers can be far more productive, resulting in the legacy application shrinking in functionality. It might even disappear entirely as all the needed functionality migrates to our new architecture.

By creating strangler fig applications, we avoid merely reproducing existing functionality in some new architecture or technology—often, our business processes are far more complex than necessary due to the idiosyncrasies of the existing systems, which we will end up replicating. (By researching the user, we can often re-engineer the process so that we can design a far simpler and more streamlined means to achieving the business goal.)[†]

An observation from Martin Fowler underscores this risk.

Much of my career has involved rewrites of critical systems. You would think such a thing is easy—just make the new one do what the old one did. Yet they are always much more complex than they seem, and overflowing with risk. The big cut-over date looms, and

the pressure is on. While new features (there are always new features) are liked, old stuff has to remain. Even old bugs often need to be added to the rewritten system.¹⁹

As with any transformation, we seek to create quick wins and deliver early incremental value before continuing to iterate. Up-front analysis helps us identify the smallest possible piece of work that will usefully achieve a business outcome using the new architecture.

CASE STUDY

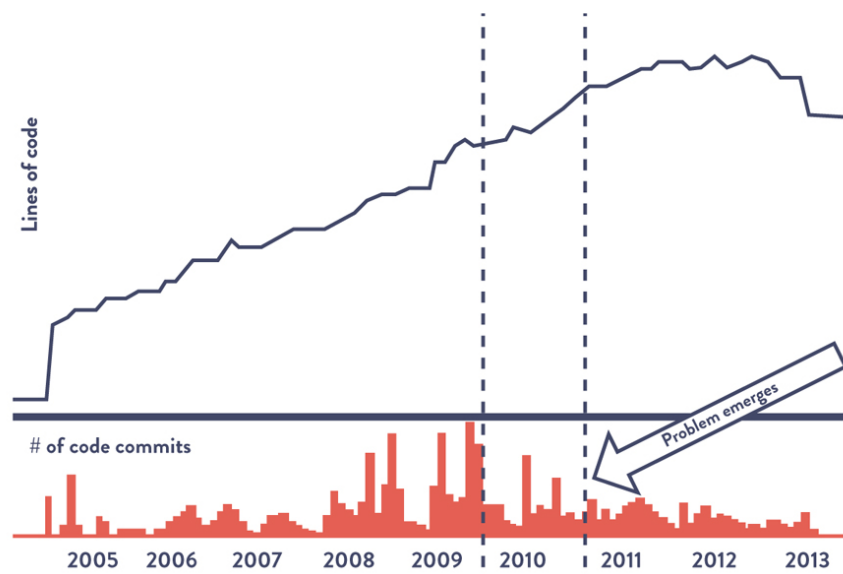
Strangler Fig Pattern at Blackboard Learn (2011)

Blackboard Inc. is one of the pioneers of providing technology for educational institutions, with annual revenue of approximately \$650 million in 2011. At that time, the development team for their flagship Learn product, packaged software that was installed and run on premises at their customer sites, was living with the daily consequences of a legacy J2EE codebase that went back to 1997.²⁰ As David Ashman, their chief architect, observed, “we still have fragments of Perl code still embedded throughout our codebase.”²¹

In 2010, Ashman was focused on the complexity and growing lead times associated with the old system, observing that “our build, integration, and testing processes kept getting more and more complex and error prone. And the larger the product got, the longer our lead times and the worse the outcomes for our customers. To even get feedback from our integration process would require twenty-four to thirty-six hours.”²²

How this started to impact developer productivity was made visible to Ashman in graphs generated from their source code repository going all the way back to 2005.

In [Figure 13.2](#), the top graph represents the number of lines of code in the monolithic Blackboard Learn code repository; the bottom graph represents the number of code commits. The problem that became evident to Ashman was that the number of code commits started to decrease, objectively showing the increasing difficulty of introducing code changes, while the number of lines of code continued to increase. Ashman noted, “To me, it said we needed to do something, otherwise the problems would keep getting worse, with no end in sight.”²³



**Figure 13.2: Blackboard Learn Code Repository:
Before Building Blocks**

Source: “DOES14—David Ashman—Blackboard Learn—Keep Your Head in the Clouds,” YouTube video, 30:43, posted by DevOps Enterprise Summit 2014, October 28, 2014, <https://www.youtube.com/watch?v=SSmixnMpsI4>.

As a result, in 2012 Ashman focused on implementing a code re-architecting project that used the strangler fig pattern. The team accomplished this by creating what they internally called *Building Blocks*, which allowed developers to work in separate modules that were decoupled from the monolithic codebase and accessed through fixed APIs. This enabled them to work with far more autonomy, without having to constantly communicate and coordinate with other development teams.

When Building Blocks were made available to developers, the size of the monolith source code repository began to decrease (as measured by number of lines of code). Ashman explained that this was because developers were moving their code into the Building Block modules source code repository. “In fact,” Ashman reported, “every developer given a choice would work in the Building Block codebase, where they could work with more autonomy and freedom and safety.”²⁴

Figure 13.3 shows the connection between the exponential growth in the number of lines of code and the exponential growth of the number of code commits for the Building Blocks code repositories. The new Building Blocks codebase allowed developers to be more productive, and they made the work safer because mistakes resulted in small, local failures instead of major catastrophes that impacted the global system.

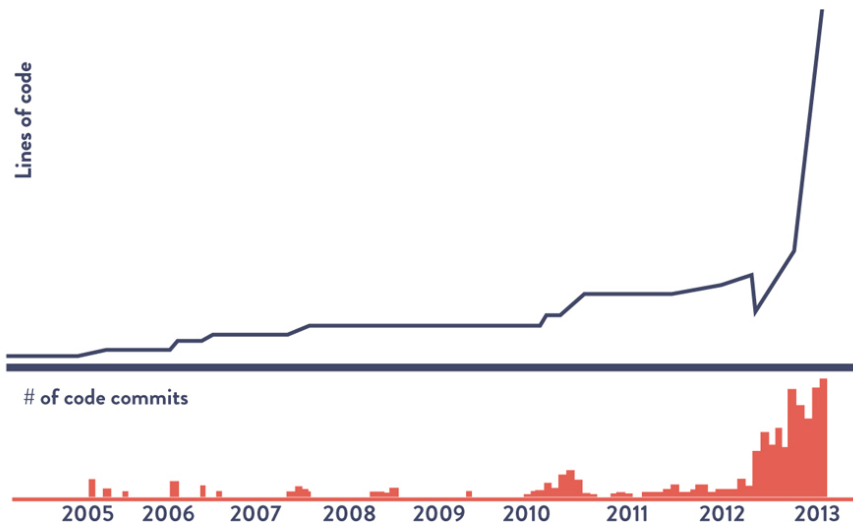


Figure 13.3: Blackboard Learn Code Repository: After Building Blocks

Source: “DOES14—David Ashman—Blackboard Learn—Keep Your Head in the Clouds,” YouTube video, 30:43, posted by DevOps Enterprise Summit 2014, October 28, 2014, <https://www.youtube.com/watch?v=SSmixnMpsI4>.

Ashman concluded, “Having developers work in the Building Blocks architecture made for impressive improvements in code modularity, allowing them to work with more independence and freedom. In combination with the updates to our build process, they also got faster, better feedback on their work, which meant better quality.”²⁵

Utilizing the Strangler Fig Application Pattern and creating a modular codebase, teams at Blackboard were able to work with more autonomy and tackle individual problems safer and faster.

CONTINUOUS LEARNING

Data supports the importance of architecture and its role in driving elite performance, with the DORA and Puppet 2017 *State of DevOps Report* finding that architecture was the largest contributor to continuous delivery.

The analysis found that teams who scored highest on architectural capabilities could complete their work independently of other teams, and change their systems without dependencies.²⁶

These findings were echoed in DORA's 2018 and 2019 *State of DevOps Reports*, showing the continued importance of having a loosely coupled architecture in teams' abilities to deploy and release quickly and with low friction.²⁷

Conclusion

To a large extent, the architecture that our services operate within dictates how we test and deploy our code. Because we are often stuck with architectures that were optimized for a different set of organizational goals, or for an era long passed, we must be able to safely migrate from one architecture to another. The case studies presented in this chapter, as well as the Amazon case study previously presented, describe techniques like the strangler fig pattern that can help us migrate between architectures incrementally, enabling us to adapt to the needs of the organization.

* eBay's architecture went through the following phases: Perl and files (v1, 1995),

C++ and Oracle (v2, 1997), XSL and Java (v3, 2002), full-stack Java (v4, 2007), Polyglot microservices (2013+).⁴

† The strangler fig application pattern involves incrementally replacing a whole system, usually a legacy system, with a completely new one. Conversely, *branching by abstraction*, a term coined by Paul Hammant, is a technique where we create an abstraction layer between the areas that we are changing. This enables evolutionary design of the application architecture while allowing everybody to work off trunk/master and practice continuous integration.¹⁸

PART III CONCLUSION

Within the previous chapters of Part III, we have implemented the architecture and technical practices that enable the fast flow of work from Dev to Ops, so that value can be quickly and safely delivered to customers.

In [Part IV: The Second Way: The Technical Practices of Feedback](#), we will create the architecture and mechanisms to enable the reciprocal fast flow of feedback from right to left, to find and fix problems faster, radiate feedback, and ensure better outcomes from our work. This enables our organization to further increase the rate at which it can adapt.

Additional Resources

The Unicorn Project, the companion novel to *The Phoenix Project*, illustrates first hand the experiences of a developer looking to improve, among other things, the flow of work (itrevolution.com/the-unicorn-project/).

Accelerate: The Science of Lean Software and DevOps: Building and Scaling High Performing Technology Organizations succinctly brings all the research from four years of *State of DevOps Reports* together and effectively shows what metrics lead to better quality software (itrevolution.com/accelerate-book)

Jez Humble and David Farley's *Continuous Delivery: Reliable Software Releases Through Build, Test, and Deployment Automation* remains a must read for anyone in Development or Operations.

Elisabeth Hendrickson's *Explore It!: Reduce Risk and Increase Confidence with Exploratory Testing* is an excellent source to help you figure out how to build effective tests.

And Martin Fowler's blog explaining his Strangler Fig Application Pattern is still an essential read. (martinfowler.com/bliki/StranglerFigApplication.html).



Part IV

THE SECOND WAY: THE TECHNICAL PRACTICES
OF FEEDBACK

PART IV: INTRODUCTION

In [Part III](#), we described the architecture and technical practices required to create fast flow from Development into Operations. Now in Part IV, we describe how to implement the technical practices of the Second Way, which are required to create fast and continuous feedback from Operations to Development.

By doing this, we shorten and amplify feedback loops so that we can see problems as they occur and radiate this information to everyone in the value stream. This allows us to quickly find and fix problems earlier in the software development life cycle, ideally long before they cause a catastrophic failure.

Furthermore, we will create a system of work where knowledge acquired downstream in Operations is integrated into the upstream work of Development and Product Management. This allows us to quickly create improvements and learnings, whether it's from a production issue, a deployment issue, early indicators of problems, or our customer usage patterns.

Additionally, we will create a process that allows everyone to get feedback on their work, makes information visible to enable learning, and enables us to rapidly test product hypotheses, helping us determine if the features we are building are helping us achieve our organizational goals.

We will also demonstrate how to create telemetry from our build, test, and deploy processes, as well as from user behavior, production issues and outages, audit issues, and security breaches. By amplifying signals as part of our daily work, we make it possible to see and solve problems as they occur, and we grow safe systems of work that allow us to

confidently make changes and run product experiments, knowing we can quickly detect and remediate failures. We will do all of this by exploring the following:

- creating telemetry to enable seeing and solving problems
- using our telemetry to better anticipate problems and achieve goals
- integrating user research and feedback into the work of product teams
- enabling feedback so Dev and Ops can safely perform deployments
- enabling feedback to increase the quality of our work through peer reviews and pair programming

The patterns in this chapter help reinforce the common goals of Product Management, Development, QA, Operations, and Infosec, and encourage them to share in the responsibility of ensuring that services run smoothly in production and collaborate on the improvement of the system as a whole. Where possible, we want to link cause to effect. The more assumptions we can invalidate, the faster we can discover and fix problems, but also the more capable we are at learning and innovating.

Throughout the following chapters, we will implement feedback loops, enabling everyone to work together toward shared goals, to see problems as they occur, to enable quick detection and recovery, and to ensure that features not only operate as designed in production but also achieve organizational goals and support organizational learning.

14

CREATE TELEMETRY TO ENABLE SEEING AND SOLVING PROBLEMS

A fact of life in Operations is that things go wrong—small changes may result in many unexpected outcomes, including outages and global failures that impact all our customers. This is the reality of operating complex systems; no single person can see the whole system and understand how all the pieces fit together.

When production outages and other problems occur in our daily work, we often don't have the information we need to solve the problem. For example, during an outage we may not be able to determine whether the issue is due to a failure in our application (e.g., defect in the code), in our environment (e.g., a networking problem, server configuration problem), or something entirely external to us (e.g., a massive denial of service attack).

In Operations, we may deal with this problem with the following rule of thumb: When something goes wrong in production, we just reboot the server. If that doesn't work, reboot the server next to it. If that doesn't work, reboot all the servers. If that doesn't work, blame the developers; they're always causing outages.¹

In contrast, the Microsoft Operations Framework (MOF) study in 2001 found that organizations with the highest service levels rebooted their servers twenty times less frequently than average and had five times fewer “blue screens of death.”² In other words, they found that the best-performing organizations were much better at diagnosing

and fixing service incidents, in what Kevin Behr, Gene Kim, and George Spafford called a “culture of causality” in *The Visible Ops Handbook*. High performers used a disciplined approach to solving problems, using production telemetry to understand possible contributing factors to focus their problem solving, as opposed to lower performers who would blindly reboot servers.³

To enable this disciplined problem-solving behavior, we need to design our systems so that they are continually creating *telemetry*, widely defined as “an automated communications process by which measurements and other data are collected at remote points and are subsequently transmitted to receiving equipment for monitoring.”⁴ Our goal is to create telemetry within our applications and environments, both in our production and pre-production environments as well as in our deployment pipeline.

DevOps Transformation at Etsy (2012)

Michael Rembetsy and Patrick McDonnell described how production monitoring was a critical part of Etsy’s DevOps transformation that started in 2009. This was because they were standardizing and transitioning their entire technology stack to the LAMP stack (Linux, Apache, MySQL, and PHP), abandoning myriad different technologies being used in production that were increasingly difficult to support.

At the 2012 Velocity Conference, McDonnell described how much risk this created:

We were changing some of our most critical infrastructure, which, ideally, customers would never notice. However, they’d definitely notice if we screwed something up. We needed more metrics to give us confidence that we weren’t actually breaking things

while we were doing these big changes, both for our engineering teams and for team members in the nontechnical areas, such as marketing.⁵

McDonnell explained further,

We started collecting all our server information in a tool called Ganglia, displaying all the information into Graphite, an open-source tool we invested heavily into. We started aggregating metrics together, everything from business metrics to deployments. This is when we modified Graphite with what we called “our unparalleled and unmatched vertical line technology” that overlaid onto every metric graph when deployments happened. By doing this, we could more quickly see any unintended deployment side effects. We even started putting TV screens all around the office so that everyone could see how our services were performing.⁶

By enabling developers to add telemetry to their features as part of their daily work, Etsy created enough telemetry to help make deployments safe. By 2011, Etsy was tracking over two hundred thousand production metrics at every layer of the application stack (e.g., application features, application health, database, operating system, storage, networking, security, etc.) with the top thirty most important business metrics prominently displayed on their “deploy dashboard.” By 2014, they were tracking over eight hundred thousand metrics, showing their relentless goal of instrumenting everything and making it easy for engineers to do so.⁷

As Ian Malpass, an engineer at Etsy, quipped,

If Engineering at Etsy has a religion, it’s the Church of Graphs. If it moves, we track it. Sometimes we’ll draw a graph of something that isn’t moving yet, just in case it

decides to make a run for it... . Tracking everything is key to moving fast, but the only way to do it is to make tracking anything easy... . We enable engineers to track what they need to track, at the drop of a hat, without requiring time-sucking configuration changes or complicated processes.⁸

In the 2015 *State of DevOps Report*, high performers resolved production incidents 168 times faster than their peers, with the median performer having an MTTR measured in minutes, while the low performer had an MTTR measured in days.⁹ In DORA's 2019 *State of DevOps Report*, elite performers resolved production incidents 2,604 times faster than their low-performing peers, with the median elite performer having a MTTR measured in minutes, while the median low performer had an MTTR measured in weeks.¹⁰

As was created at Etsy, our goal in this chapter is to ensure that we always have enough telemetry so that we can confirm that our services are correctly operating in production. And when problems do occur, our goal is to make it possible to quickly determine what is going wrong and make informed decisions on how best to fix it, ideally long before customers are impacted. Furthermore, telemetry helps us assemble our best understanding of reality and detect when our understanding of reality is incorrect.

Create Our Centralized Telemetry Infrastructure

Operational monitoring and logging is by no means new—multiple generations of Operations engineers have used and customized monitoring frameworks (e.g., HP OpenView, IBM Tivoli, and BMC Patrol/BladeLogic) to ensure the health of production systems. Data was typically collected through agents that ran on servers or through agent-less monitoring (e.g., SNMP traps or polling-based monitors). There was often

a graphical user interface (GUI) front end, and back-end reporting was often augmented through tools such as Crystal Reports.

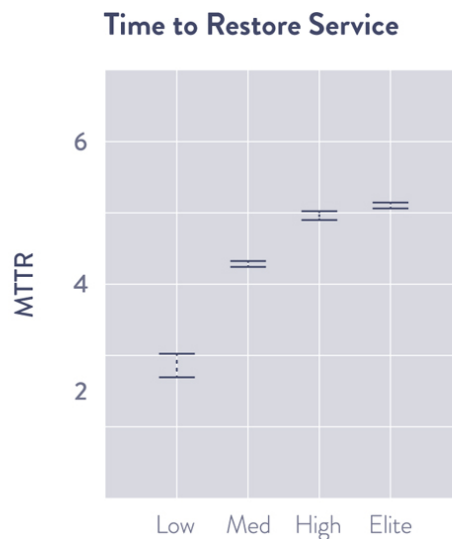


Figure 14.1: Incident Resolution Time for Elite, High, Medium, and Low Performers (2019)

Source: Forsgren et al., *Accelerate: State of DevOps* (2019).

Similarly, the practices of developing applications with effective logging and managing the resulting telemetry are not new—a variety of mature logging libraries exist for almost all programming languages.

However, for decades we have ended up with silos of information, where Development only creates logging events that are interesting to developers, and Operations only monitors whether the environments are up or down. As a result, when inopportune events occur, no one can determine why the entire system is not operating as designed or which specific component is failing, impeding our ability to bring our system back to a working state.

In order for us to see all problems as they occur, we must design and develop our applications and environments so that they generate sufficient telemetry, allowing us to understand how our system is behaving as a whole. When all

levels of our application stack have monitoring and logging, we enable other important capabilities, such as graphing and visualizing our metrics, anomaly detection, proactive alerting and escalation, etc.

In *The Art of Monitoring*, James Turnbull describes a modern monitoring architecture, which has been developed and used by Operations engineers at web-scale companies (e.g., Google, Amazon, Facebook). The architecture often consisted of open-source tools, such as Nagios and Zenoss, that were customized and deployed at a scale that was difficult to accomplish with licensed commercial software at the time.¹¹

This architecture has the following components:

- **Data collection at the business logic, application, and environments layer:** In each of these layers, we are creating telemetry in the form of events, logs, and metrics. Logs may be stored in application-specific files on each server (e.g., `/var/log/httpd-error.log`), but preferably we want all our logs sent to a common service that enables easy centralization, rotation, and deletion. This is provided by most operating systems, such as syslog for Linux, the Event Log for Windows, etc.

Furthermore, we gather metrics at all layers of the application stack to better understand how our system is behaving. At the operating system level, we can collect metrics such as CPU, memory, disk, or network usage over time using tools like collected, Ganglia, etc. The Cloud Native Computing Foundation has created an open standard for metrics and tracking data called OpenTelemetry, which is understood by many open-source and commercial tools. Other tools

that collect performance information include Apache Skywalking, AppDynamics, and New Relic.

- **An event router responsible for storing our events and metrics:** This capability potentially enables visualization, trending, alerting, anomaly detection, and so forth. By collecting, storing, and aggregating all our telemetry, we better enable further analysis and health checks. This is also where we store configurations related to our services (and their supporting applications and environments) and is likely where we do threshold-based alerting and health checks. Examples of tools in this space include Prometheus, Honeycomb, DataDog, and Sensu.

Once we have centralized our logs, we can transform them into metrics by counting them in the event router—for example, a log event such as “child pid 14024 exit signal Segmentation fault” can be counted and summarized as a single segfault metric across our entire production infrastructure.

By transforming logs into metrics, we can now perform statistical operations on them, such as using anomaly detection to find outliers and variances even earlier in the problem cycle. For instance, we might configure our alerting to notify us if we went from “ten segfaults last week” to “thousands of segfaults in the last hour,” prompting us to investigate further.

In addition to collecting telemetry from our production services and environments, we must also collect telemetry from our deployment pipeline when important events occur, such as when our automated tests pass or fail and when we perform deployments to any environment. We should also collect telemetry on how long it takes us to execute our builds

and tests. By doing this, we can detect conditions that could indicate problems, such as if the performance test or our build takes twice as long as normal, allowing us to find and fix errors before they go into production.

Furthermore, we should ensure that it is easy to enter and retrieve information from our telemetry infrastructure. Preferably, everything should be done through self-service APIs, as opposed to requiring people to open up tickets and wait to get reports.

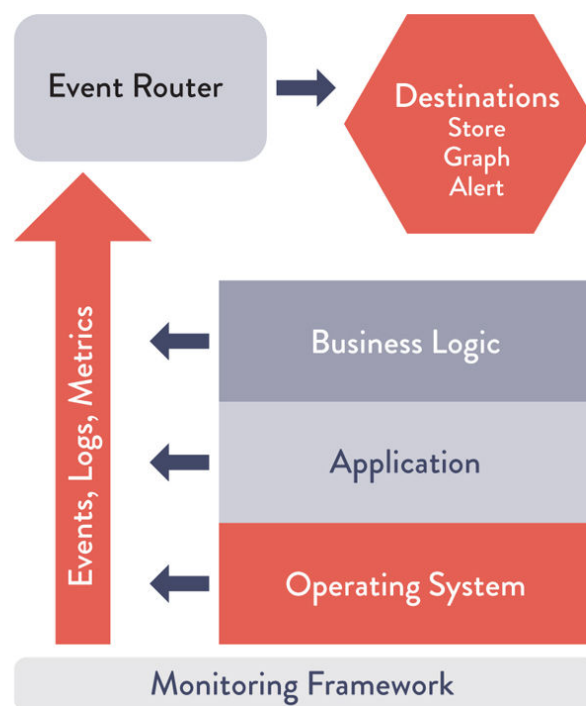


Figure 14.2: Monitoring Framework

Source: Turnbull, *The Art of Monitoring*, Kindle edition, chap. 2.

Ideally, we will create telemetry that tells us exactly when anything of interest happens, as well as where and how. Our telemetry should also be suitable for manual and automated analysis and should be able to be analyzed without having the application that produced the logs on hand.* As Adrian Cockcroft pointed out, “Monitoring is so important that our monitoring systems need to be more available and scalable than the systems being monitored.”¹²

Create Application Logging Telemetry That Helps Production

Now that we have a centralized telemetry infrastructure, we must ensure that the applications we build and operate are creating sufficient telemetry. We do this by having Dev and Ops engineers create production telemetry as part of their daily work, both for new and existing services.

Scott Prugh, CTO at CSG, said,

Every time NASA launches a rocket, it has millions of automated sensors reporting the status of every component of this valuable asset. And yet, we often don't take the same care with software—we found that creating application and infrastructure telemetry to be one of the highest return investments we've made. In 2014, we created over one billion telemetry events per day, with over one hundred thousand code locations instrumented.¹³

In the applications we create and operate, every feature should be instrumented. If it was important enough for an engineer to implement, then it is important enough to generate enough production telemetry to confirm that it is operating as designed and that the desired outcomes are being achieved.[†]

Every member of our value stream will use telemetry in a variety of ways. For example, developers may temporarily create more telemetry in their application to better diagnose problems on their workstation, while Ops engineers may use telemetry to diagnose a production problem. In addition, Infosec and auditors may review the telemetry to confirm the effectiveness of a required control, and a product manager

may use them to track business outcomes, feature usage, or conversion rates.

To support these various usage models, we have different logging levels, some of which may also trigger alerts, such as the following:¹⁴

- **DEBUG level:** Information at this level is about anything that happens in the program, most often used during debugging. Often, debug logs are disabled in production but temporarily enabled during troubleshooting.
- **INFO level:** Information at this level consists of actions that are user-driven or system specific (e.g., “beginning credit card transaction”).
- **WARN level:** Information at this level tells us of conditions that could potentially become an error (e.g., a database call taking longer than some predefined time). These will likely initiate an alert and troubleshooting, while other logging messages may help us better understand what led to this condition.
- **ERROR level:** Information at this level focuses on error conditions (e.g., API call failures, internal error conditions).
- **FATAL level:** Information at this level tells us when we must terminate (e.g., a network daemon can’t bind a network socket).

Choosing the right logging level is important. Dan North, a former ThoughtWorks consultant who was involved in several projects in which the core continuous delivery concepts took shape, observes, “When deciding whether a message should be ERROR or WARN, imagine being woken up at 4 AM. Low printer toner is not an ERROR.”¹⁵

To help ensure that we have information relevant to the reliable and secure operations of our service, we should ensure that all potentially significant application events generate logging entries, including those provided on this list assembled by Anton A. Chuvakin, a research VP at Gartner's GTP Security and Risk Management group:¹⁶

- authentication/authorization decisions (including logoff)
- system and data access
- system and application changes (especially privileged changes)
- data changes, such as adding, editing, or deleting data
- invalid input (possible malicious injection, threats, etc.)
- resources (RAM, disk, CPU, bandwidth, or any other resource that has hard or soft limits)
- health and availability
- startups and shutdowns
- faults and errors
- circuit breaker trips
- delays
- backup success/failure

To make it easier to interpret and give meaning to all these log entries, we should (ideally) create logging hierarchical categories, such as for non-functional attributes (e.g., performance, security) and for attributes related to features (e.g., search, ranking).

Use Telemetry to Guide Problem Solving

As described in the beginning of this chapter, high performers use a disciplined approach to solving problems. This is in contrast to the more common practice of using rumor and hearsay, which can lead to the unfortunate metric of *mean time until declared innocent*—how quickly can we convince everyone else that we didn't cause the outage.

When there is a culture of blame around outages and problems, groups may avoid documenting changes and displaying telemetry where everyone can see them to avoid being blamed for outages.

Other negative outcomes due to lack of public telemetry include a highly charged political atmosphere, the need to deflect accusations, and, worse, the inability to create institutional knowledge around how the incidents occurred and the learnings needed to prevent these errors from happening again in the future.[‡]

In contrast, telemetry enables us to use the scientific method to formulate hypotheses about what is causing a particular problem and what is required to solve it. Examples of questions we can answer during problem resolution include:

- What evidence do we have from our monitoring that a problem is actually occurring?
- What are the relevant events and changes in our applications and environments that could have contributed to the problem?
- What hypotheses can we formulate to confirm the link between the proposed causes and effects?
- How can we prove which of these hypotheses are correct and successfully affect a fix?

The value of fact-based problem-solving lies not only in significantly faster MTTR (and better customer outcomes), but also in its reinforcement of the perception of a win/win relationship between Development and Operations.

Enable Creation of Production Metrics as Part of Daily Work

To enable everyone to be able to find and fix problems in their daily work, we need to enable everyone to create metrics in their daily work that can be easily created, displayed, and analyzed.

To do this, we must create the infrastructure and libraries necessary to make it as easy as possible for anyone in Development or Operations to create telemetry for any functionality they build. In the ideal, it should be as easy as writing one line of code to create a new metric that shows up in a common dashboard where everyone in the value stream can see it.

This was the philosophy that guided the development of one of the most widely used metrics libraries, called StatsD, which was created and open-sourced at Etsy.¹⁸ As John Allspaw described, “We designed StatsD to prevent any developer from saying, ‘It’s too much of a hassle to instrument my code.’ Now they can do it with one line of code. It was important to us that for a developer, adding production telemetry didn’t feel as difficult as doing a database schema change.”¹⁹

StatsD can generate timers and counters with one line of code (in Ruby, Perl, Python, Java, and other languages) and is often used in conjunction with Graphite or Grafana, which render metric events into graphs and dashboards.

Figure 14.3 shows an example of how a single line of code creates a user login event (in this case, one line of PHP code: “StatsD::increment(“login.successes”)”). The resulting graph shows the number of successful and failed logins per minute, and overlaid on the graph are vertical lines that represent a production deployment.

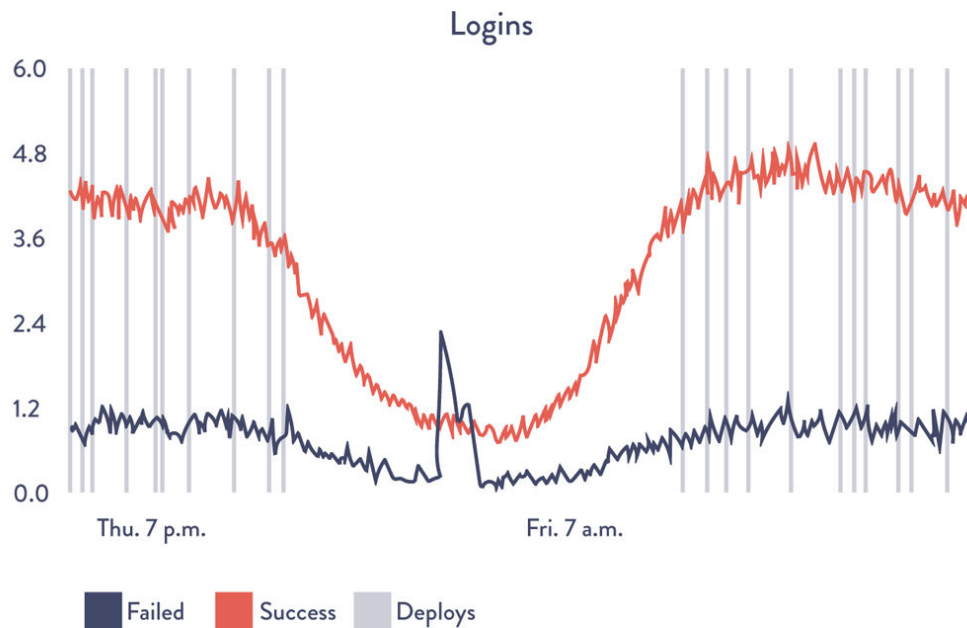


Figure 14.3: One Line of Code to Generate Telemetry using StatsD and Graphite at Etsy

Source: Ian Malpass, “Measure Anything, Measure Everything.”

When we generate graphs of our telemetry, we will also overlay onto them when production changes occur, because we know that the significant majority of production issues are caused by production changes, which include code deployments. This is part of what allows us to have a high rate of change while still preserving a safe system of work.

More recently, the emergence of the OpenTelemetry standard has provided a way for data collectors to communicate with metrics storage and processing systems. There are OpenTelemetry integrations with all major languages, frameworks, and libraries, and most popular metrics and observability tools accept OpenTelemetry data. §

By generating production telemetry as part of our daily work, we create an ever-improving capability to not only see problems as they occur but also to design our work so that problems in design and operations can be revealed, allowing an increasing number of metrics to be tracked, as we saw in the Etsy case study.

Create Self-Service Access to Telemetry and Information Radiators

In the previous steps, we enabled Development and Operations to create and improve production telemetry as part of their daily work. In this step, our goal is to radiate this information to the rest of the organization, ensuring that anyone who wants information about any of the services we are running can get it without needing production system access or privileged accounts, or having to open up a ticket and wait for days for someone to configure the graph for them.

By making telemetry fast, easy to get, and sufficiently centralized, everyone in the value stream can share a common view of reality. Typically, this means that production metrics will be radiated on web pages generated by a centralized server, such as Graphite or any of the other technologies described in the previous section.

We want our production telemetry to be highly visible, which means putting it in central areas where Development and Operations work, thus allowing everyone who is interested to see how our services are performing. At a minimum, this includes everyone in our value stream, such as Development, Operations, Product Management, and Infosec. This is often referred to as an *information radiator*, defined by the Agile Alliance as,

the generic term for any of a number of handwritten, drawn, printed, or electronic displays which a team places in a highly visible location, so that all team members as well as passers-by can see the latest information at a glance: count of automated tests, velocity, incident reports, continuous integration status, and so on. This idea originated as part of the Toyota Production System.²⁰

By putting information radiators in highly visible places, we promote responsibility among team members, actively demonstrating the following values:

- The team has nothing to hide from its visitors (customers, stakeholders, etc.).
- The team has nothing to hide from itself: it acknowledges and confronts problems.

Now that we possess the infrastructure to create and radiate production telemetry to the entire organization, we may also choose to broadcast this information to our internal customers and even to our external customers. For example, we might do this by creating publicly viewable service status pages so that customers can learn how the services they depend upon are performing.

Although there may be some resistance to providing this amount of transparency, Ernest Mueller describes the value of doing so:

One of the first actions I take when starting in an organization is to use information radiators to communicate issues and detail the changes we are making—this is usually extremely well-received by our business units, who were often left in the dark before. And for Development and Operations groups who must work together to deliver a service to others, we need

that constant communication, information, and feedback.²¹

We may even extend this transparency further—instead of trying to keep customer-impacting problems a secret, we can broadcast this information to our external customers. This demonstrates that we value transparency, thereby helping to build and earn customers’ trust.[¶] (See [Appendix 10](#).)

CASE STUDY

Creating Self-Service Metrics at LinkedIn (2011)

As described in [Part III](#), LinkedIn was created in 2003 to help users connect “to your network for better job opportunities.” By November 2015, LinkedIn had over 350 million members generating tens of thousands of requests per second, resulting in millions of queries per second on the LinkedIn back-end systems.

Prachi Gupta, Director of Engineering at LinkedIn, wrote in 2011 about the importance of production telemetry:

At LinkedIn, we emphasize making sure the site is up and our members have access to complete site functionality at all times. Fulfilling this commitment requires that we detect and respond to failures and bottlenecks as they start happening. That’s why we use these time-series graphs for site monitoring to detect and react to incidents within minutes... . This monitoring technique has proven to be a great tool for engineers. It

lets us move fast and buys us time to detect, triage, and fix problems.²²

However, in 2010, even though there was an incredibly large volume of telemetry being generated, it was extremely difficult for engineers to get access to the data, let alone analyze it. Thus began Eric Wong's summer intern project at LinkedIn, which turned into the production telemetry initiative that created InGraphs.

Wong wrote, "To get something as simple as CPU usage of all the hosts running a particular service, you would need to file a ticket and someone would spend 30 minutes putting [a report] together."²³

At the time, LinkedIn was using Zenoss to collect metrics, but as Wong explains, "Getting data from Zenoss required digging through a slow web interface, so I wrote some python scripts to help streamline the process. While there was still manual intervention in setting up metric collection, I was able to cut down the time spent navigating Zenoss' interface."²⁴

Over the course of the summer, he continued to add functionality to InGraphs so that engineers could see exactly what they wanted to see, adding the ability to make calculations across multiple datasets, view week-over-week trending to compare historical performance, and even define custom dashboards to pick exactly which metrics would be displayed on a single page.

In writing about the outcomes of adding functionality to InGraphs and the value of this capability, Gupta notes, "The effectiveness of our monitoring system was highlighted in an instant where our InGraphs monitoring functionality tied to a major web-

mail provider started trending downwards and the provider realized they had a problem in their system only after we reached out to them!”²⁵

What started off as a summer internship project is now one of the most visible parts of LinkedIn operations. InGraphs has been so successful that the real-time graphs are featured prominently in the company’s engineering offices where visitors can’t fail to see them.

Self-service metrics can empower problem-solving and decision-making at the individual and team level, and provide necessary transparency to build and earn customers’ trust.

Find and Fill Any Telemetry Gaps

We have now created the infrastructure necessary to quickly create production telemetry throughout our entire application stack and radiate it throughout our organization.

In this section, we will identify any gaps in our telemetry that impede our ability to quickly detect and resolve incidents—this is especially relevant if Dev and Ops currently have little (or no) telemetry. We will use this data later to better anticipate problems, as well as to enable everyone to gather the information they need to make better decisions to achieve organizational goals.

Achieving this requires that we create enough telemetry at all levels of the application stack for all our environments, as well as for the deployment pipelines that support them. We need metrics from the following levels:

- **Business level:** Examples include the number of sales transactions, revenue of sales transactions, user sign-ups, churn rate, A/B testing results, etc.

- **Application level:** Examples include transaction times, user response times, application faults, etc.
- **Infrastructure level (e.g., database, operating system, networking, storage):** Examples include web server traffic, CPU load, disk usage, etc.
- **Client software level (e.g., JavaScript on the client browser, mobile application):** Examples include application errors and crashes, user-measured transaction times, etc.
- **Deployment pipeline level:** Examples include build pipeline status (e.g., red or green for our various automated test suites), change deployment lead times, deployment frequencies, test environment promotions, and environment status.

By having telemetry coverage in all of these areas, we will be able to see the health of everything that our service relies upon, using data and facts instead of rumors, finger-pointing, blame, and so forth.

Further, we better enable detection of security-relevant events by monitoring any application and infrastructure faults (e.g., abnormal program terminations, application errors and exceptions, and server and storage errors). Not only does this telemetry better inform Development and Operations when our services are crashing, but these errors are often indicators that a security vulnerability is being actively exploited.

By detecting and correcting problems earlier, we can fix them while they are small and easy to fix with fewer customers impacted. Furthermore, after every production incident, we should identify any missing telemetry that could have enabled faster detection and recovery; or, better yet, we

can identify these gaps during feature development in our peer review process.

Application and Business Metrics

At the application level, our goal is to ensure that we are generating telemetry not only around application health (e.g., memory usage, transaction counts, etc.) but also to measure to what extent we are achieving our organizational goals (e.g., number of new users, user login events, user session lengths, percent of users active, how often certain features are being used, and so forth).

For example, if we have a service that is supporting e-commerce, we want to ensure that we have telemetry around all of the user events that lead up to a successful transaction that generates revenue. We can then instrument all the user actions that are required for our desired customer outcomes.

These metrics will vary according to different domains as well as organizational goals. For instance, for an e-commerce site, we may want to *maximize* the time spent on the site, increasing likelihood of a sale. However, for search engines, we may want to *reduce* the time spent on the site, since long sessions may indicate that users are having difficulty finding what they're looking for.

In general, business metrics will be part of a *customer acquisition funnel*, which is the theoretical steps a potential customer will take to make a purchase. For instance, in an e-commerce site, the measurable journey events include total time on site, product link clicks, shopping cart adds, and completed orders.

Ed Blankenship, Senior Product Manager for Microsoft Visual Studio Team Services, describes it like this: “Often, feature teams will define their goals in an acquisition funnel, with the goal of their feature being used in every customer’s

daily work. Sometimes they're informally described as 'tire kickers,' 'active users,' 'engaged users,' and 'deeply engaged users,' with telemetry supporting each stage."²⁶

Our goal is to have every business metric be *actionable*—these top metrics should help inform how to change our product and be amenable to experimentation and A/B testing. When metrics aren't actionable, they are likely vanity metrics that provide little useful information—these we want to store, but likely not display, let alone alert on.

Ideally, anyone viewing our information radiators will be able to make sense of the information we are showing in the context of desired organizational outcomes, such as goals around revenue, user attainment, conversion rates, etc. We should define and link each metric to a business outcome metric at the earliest stages of feature definition and development, and measure the outcomes after we deploy them in production. Furthermore, doing this helps product owners describe the business context of each feature for everyone in the value stream.

Further business context can be created by being aware of and visually displaying time periods relevant to high-level business planning and operations, such as high transaction periods associated with peak holiday selling seasons, end-of-quarter financial close periods, or scheduled compliance audits. This information may be used as a reminder to avoid scheduling risky changes when availability is critical or avoid certain activities when audits are in progress.

By radiating how customers interact with what we build in the context of our goals (see [Figure 14.4](#)), we enable fast feedback to feature teams so they can see whether the capabilities we are building are actually being used and to what extent they are achieving business goals. As a result, we

reinforce the cultural expectations that instrumenting and analyzing customer usage is also a part of our daily work, so we better understand how our work contributes to our organizational goals.

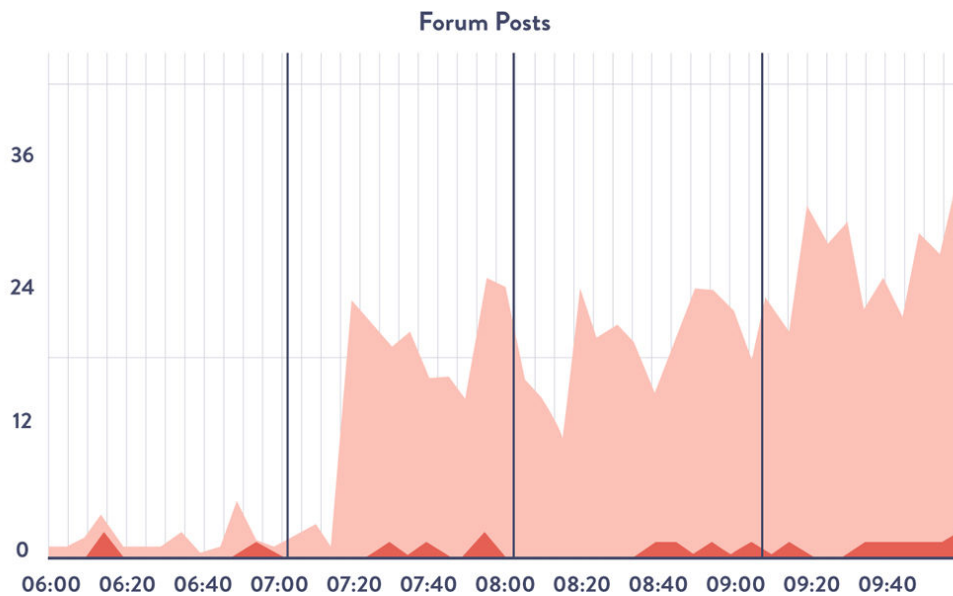


Figure 14.4: User Excitement of New Features in User Forum Posts after Deployments

Source: Mike Brittain, "Tracking Every Release," [CodeasCraft.com](https://codeascraft.com/2010/12/08/track-every-release/), December 8, 2010, <https://codeascraft.com/2010/12/08/track-every-release/>.

Infrastructure Metrics

Just as with application metrics, our goal for production and non-production infrastructure is to ensure that we are generating enough telemetry so that if a problem occurs in any environment, we can quickly determine whether infrastructure is a contributing cause of the problem. Furthermore, we must be able to pinpoint exactly what in the infrastructure is contributing to the problem (e.g., database, operating system, storage, networking, etc.).

We want to make as much infrastructure telemetry visible as possible, across all the technology stakeholders, ideally organized by service or application. In other words, when something goes wrong with something in our environment,

we need to know exactly what applications and services could be or are being affected.**

In decades past, creating links between a service and the production infrastructure it depended on was often a manual effort (such as ITIL CMDBs or creating configuration definitions inside alerting tools in tools such as Nagios). However, increasingly these links are now registered automatically within our services, which are then dynamically discovered and used in production through tools such as ZooKeeper, Etcd, Consul, Istio, etc.²⁷

These tools enable services to register themselves, storing information that other services need to interact with it (e.g., IP address, port numbers, URIs). This solves the manual nature of the ITIL CMDB and is absolutely necessary when services are made up of hundreds (or thousands, or even millions) of nodes, each with dynamically assigned IP addresses.^{††}

Regardless of how simple or complex our services are, graphing our business metrics alongside our application and infrastructure metrics allows us to detect when things go wrong. For instance, we may see that new customer sign-ups drop to 20% of daily norms, and then immediately also see that all our database queries are taking five times longer than normal, enabling us to focus our problem solving.

Furthermore, business metrics create context for our infrastructure metrics, enabling Development and Operations to better work together toward common goals. As Jody Mulkey, CTO of Ticketmaster/LiveNation, observes, “Instead of measuring Operations against the amount of downtime, I find it’s much better to measure both Dev and Ops against the real business consequences of downtime: how much revenue should we have attained, but didn’t.”^{29††}

CONTINUOUS LEARNING

DORA's 2019 *State of DevOps Report* found that infrastructure monitoring contributed to continuous delivery. This is because the visibility and fast feedback it provides to all stakeholders is key to help everyone see the outcomes of build, test, and deployment outcomes.³⁰

Note that in addition to monitoring our production services, we also need telemetry for those services in our pre-production environments (e.g., development, test, staging, etc.). Doing this enables us to find and fix issues before they go into production, such as detecting when we have ever-increasing database insert times due to a missing table index.

Overlaying Other Relevant Information Onto Our Metrics

Even after we have created our deployment pipeline that allows us to make small and frequent production changes, changes still inherently create risk. Operational side effects are not just outages but also significant disruptions and deviations from standard operations.

To make changes visible, we make work visible by overlaying all production deployment activities on our graphs. For instance, for a service that handles a large number of inbound transactions, production changes can result in a significant *settling period*, where performance degrades substantially as all cache lookups miss.

To better understand and preserve quality of service, we want to understand how quickly performance returns to normal and, if necessary, take steps to improve performance. Similarly, we want to overlay other useful operational

activities, such as when the service is under maintenance or being backed up, in places where we may want to display or suppress alerts.

Conclusion

The improvements enabled by production telemetry from Etsy and LinkedIn show us how critical it is to see problems as they occur so we can search out the cause and quickly remedy the situation. By having all elements of our service emitting telemetry that can be analyzed whether it is in our application, database, or environment, and making that telemetry widely available, we can find and fix problems long before they cause something catastrophic, ideally long before a customer even notices that something is wrong. The result is not only happier customers, but, by reducing the amount of firefighting and crises when things go wrong, we have a happier and more productive workplace with less stress and lower levels of burnout.

-
- * From here on, the term *telemetry* will be used interchangeably with *metrics*, which includes all event logging and metrics created by our services at all levels of our application stack and generated from all our production and pre-production environments, as well as from our deployment pipeline.
 - † A variety of application logging libraries exist that make it easy for developers to create useful telemetry, and we should choose one that allows us to send all our application logs to the centralized logging infrastructure that we created in the previous section. Popular examples include `rrd4j` and `log4j` for Java, and `log4r` and `ruby-cabin` for Ruby.
 - ‡ In 2004, Gene Kim, Kevin Behr, and George Spafford described this as a symptom of lacking a “culture of causality,” noting that high-performing organizations recognize that 80% of all outages are caused by change and 80% of MTTR is spent trying to determine what changed.¹⁷
 - § A whole other set of tools to aid in monitoring, aggregation, and collection include Splunk, Zabbix, Sumo Logic, DataDog, as well as Nagios, Cacti, Sensu, RRDTool, Netflix Atlas, Riemann, and others. Analysts often call this broad category of tools “application performance monitors.”

- ¶ Creating a simple dashboard should be part of creating any new product or service—automated tests should confirm that both the service and dashboard are working correctly, helping both our customers and our ability to safely deploy code.
- ** Exactly as an ITIL configuration management database (CMDB) would prescribe.
- †† Consul may be of specific interest, as it creates an abstraction layer that easily enables service mapping, monitoring, locks, and key-value configuration stores, as well as host clustering and failure detection.²⁸
- ‡‡ This could be the cost of production downtime or the costs associated with a late feature. In product development terms, the second metric is known as cost of delay and is key to making effective prioritization decisions.

15

ANALYZE TELEMETRY TO BETTER ANTICIPATE PROBLEMS AND ACHIEVE GOALS

As we saw in the previous chapter, we need sufficient production telemetry in our applications and infrastructure to see and solve problems as they occur. In this chapter, we will create tools that allow us to discover variances and ever-weaker failure signals hidden in our production telemetry so we can avert catastrophic failures. Numerous statistical techniques will be presented, along with case studies demonstrating their use.

Telemetry at Netflix (2012)

A great example of analyzing telemetry to proactively find and fix problems before customers are impacted can be seen at Netflix. Netflix had revenue of \$6.2 billion from 75 million subscribers in 2015, \$5.7 billion in revenue in March 2020, and 209 million subscribers as of July 2021.¹ One of their goals is to provide the best experience to those watching videos online around the world, which requires a robust, scalable, and resilient delivery infrastructure.

Roy Rapoport describes one of the challenges of managing the Netflix cloud-based video delivery service: “Given a herd of cattle that should all look and act the same, which cattle look different from the rest? Or more concretely, if we have a thousand-node stateless compute cluster, all running the same software and subject to the same approximate traffic load, our challenge is to find any nodes that don’t look like the rest of the nodes.”²

One of the statistical techniques that the team used at Netflix in 2012 was *outlier detection*, defined by Victoria J. Hodge and Jim Austin of the University of York as detecting “abnormal running conditions from which significant performance degradation may well result, such as an aircraft engine rotation defect or a flow problem in a pipeline.”³

Rapoport explains that Netflix “used outlier detection in a very simple way, which was to first compute what was the ‘current normal’ right now, given the population of nodes in a compute cluster. And then we identified which nodes didn’t fit that pattern, and removed those nodes from production.”⁴

Rapoport continues,

We can automatically flag misbehaving nodes without having to actually define what the ‘proper’ behavior is in any way. And since we’re engineered to run resiliently in the cloud, we don’t tell anyone in Operations to do something—instead, we just kill the sick or misbehaving compute node, and then log it or notify the engineers in whatever form they want.⁵

By implementing the server outlier detection process, Rapoport states, Netflix has “massively reduced the effort of finding sick servers, and, more importantly, massively reduced the time required to fix them, resulting in improved service quality. The benefit of using these techniques to preserve employee sanity, work/life balance, and service quality cannot be overstated.”^{6*}

Throughout this chapter, we will explore many statistical and visualization techniques (including outlier detection) that we can use to analyze our telemetry to better anticipate problems. This enables us to solve problems faster, cheaper, and earlier than ever, before our customer or anyone in our

organization is impacted. Furthermore, we will also create more context for our data to help us make better decisions and achieve our organizational goals.

Use Means and Standard Deviations to Detect Potential Problems

One of the simplest statistical techniques that we can use to analyze a production metric is computing its *mean* (or average) and *standard deviations*. By doing this, we can create a filter that detects when this metric is significantly different from its norm, and even configure our alerting so that we can take corrective action (e.g., notify on-call production staff at 2 AM to investigate when database queries are significantly slower than average).

When critical production services have problems, waking people at 2 AM may be the right thing to do. However, when we create alerts that are not actionable or are false-positives, we've unnecessarily woken up people in the middle of the night. As John Vincent, an early leader in the DevOps movement, observed, "Alert fatigue is the single biggest problem we have right now... . We need to be more intelligent about our alerts or we'll all go insane."⁷

We create better alerts by increasing the signal-to-noise ratio, focusing on the variances or outliers that matter. Suppose we are analyzing the number of unauthorized login attempts per day. Our collected data has a Gaussian distribution (i.e., normal or bell curve distribution) that matches the graph in the [Figure 15.1](#). The vertical line in the middle of the bell curve is the mean, and the first, second, and third standard deviations indicated by the other vertical lines contain 68%, 95%, and 99.7% of the data, respectively.




Figure 15.1: Standard Deviations (σ) & Mean (μ) with Gaussian Distribution

Source: *Wikipedia*, “Normal Distribution,”
https://en.wikipedia.org/wiki/Normal_distribution.

A common use of standard deviations is to periodically inspect the data set for a metric and alert if it has significantly varied from the mean. For instance, we may set an alert for when the number of unauthorized login attempts per day is three standard deviations greater than the mean. Provided that this data set has Gaussian distribution, we would expect that only 0.3% of the data points would trigger the alert.

Even this simple type of statistical analysis is valuable because no one had to define a static threshold value, something which is infeasible if we are tracking thousands or hundreds of thousands of production metrics.[†]

Instrument and Alert on Undesired Outcomes

Tom Limoncelli, co-author of *The Practice of Cloud System Administration: Designing and Operating Large Distributed Systems* and a former Site Reliability Engineer at Google, relates the following story on monitoring:

When people ask me for recommendations on what to monitor, I joke that in an ideal world, we would delete all the alerts we currently have in our monitoring system. Then, after each user-visible outage, we'd ask what indicators would have predicted that outage and then add those to our monitoring system, alerting as needed. Repeat. Now we only have alerts that prevent outages, as opposed to being bombarded by alerts after an outage already occurred.⁸

In this step, we will replicate the outcomes of such an exercise. One of the easiest ways to do this is to analyze our most severe incidents in the recent past (e.g., thirty days) and create a list of telemetry that could have enabled earlier and faster detection and diagnosis of the problem, as well as easier and faster confirmation that an effective fix had been implemented. For instance, if we had an issue where our NGINX web server stopped responding to requests, we would look at the leading indicators that could have warned us earlier that we were starting to deviate from standard operations, such as:

- **Application level:** increasing web page load times, etc.
- **OS level:** server free memory running low, disk space running low, etc.
- **Database level:** database transaction times taking longer than normal, etc.
- **Network level:** number of functioning servers behind the load balancer dropping, etc.

Each of these metrics is a potential precursor to a production incident. For each, we would configure our alerting systems to notify us when they deviate sufficiently from the mean so that we can take corrective action.

By repeating this process on ever-weaker failure signals, we find problems ever earlier in the life cycle, resulting in fewer customer-impacting incidents and near misses. In other words, we are preventing problems as well as enabling quicker detection and correction.

Problems That Arise When Our Telemetry Data Has Non-Gaussian Distribution

Using means and standard deviations to detect variance can be extremely useful. However, using these techniques on many of the telemetry data sets that we use in Operations will not generate the desired results. As Dr. Toufic Boubez observes, “Not only will we get wakeup calls at 2 AM, we’ll get them at 2:37 AM, 4:13 AM, 5:17 AM. This happens when the underlying data that we’re monitoring doesn’t have a Gaussian distribution.”⁹

In other words, when the distribution of the data set does not have the Gaussian shape described earlier, the properties associated with standard deviations do not apply. For example, consider the scenario in which we are monitoring the number of file downloads per minute from our website. We want to detect periods when we have unusually high numbers of downloads, such as when our download rate is greater than three standard deviations from our average, so that we can proactively add more capacity.

[Figure 15.2](#) shows our number of simultaneous downloads per minute over time, with a bar overlaid on top. When the bar is dark, the number of downloads within a given period (sometimes called a “sliding window”) is at least three standard deviations from the average. Otherwise, it is light.




Figure 15.2: Downloads per Minute: Over-Alerting when Using “Three Standard Deviations” Rule

Source: Dr. Toufic Boubez, “Simple math for anomaly detection.”

The obvious problem that the graph shows is that we are alerting almost all of the time. This is because in almost any given period of time, we have instances when the download count exceeds our three standard deviation threshold.

To confirm this, when we create a histogram (see [Figure 15.3](#)) that shows the frequency of downloads per minute, we can see that it does not have the classic symmetrical bell curve shape. Instead, it is obvious that the distribution is skewed toward the lower end, showing that the majority of the time we have very few downloads per minute but that download counts frequently spike three standard deviations higher.

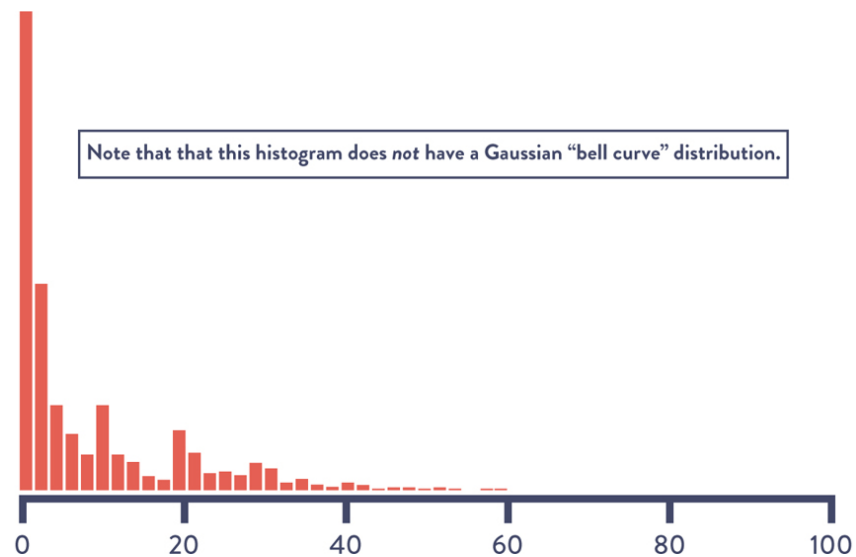


Figure 15.3: Downloads per Minute: Histogram of Data Showing Non-Gaussian Distribution

Source: Dr. Toufic Boubez, “Simple math for anomaly detection.”

Many production data sets are non-Gaussian distribution. Dr. Nicole Forsgren explains, “In Operations, many of our data sets have a ‘chi square’ distribution. Using standard deviations for this data not only results in over- or under-alerting, but it also results in nonsensical results.” She continues, “When you compute the number of simultaneous downloads that are three standard deviations below the mean, you end up with a negative number, which obviously doesn’t make sense.”¹⁰

Over-alerting causes Operations engineers to be woken up in the middle of the night for protracted periods of time, even when there are few actions that they can appropriately take. The problem associated with under-alerting is just as significant.

For instance, suppose we are monitoring the number of completed transactions, and the completed transaction count drops by 50% in the middle of the day due to a software component failure. If this is still within three standard deviations of the mean, no alert will be generated, meaning

that our customers will discover the problem before we do, at which point the problem may be much more difficult to solve. Fortunately, there are techniques we can use to detect anomalies even in non-Gaussian data sets, which are described next.

CASE STUDY

Auto-Scaling Capacity at Netflix (2012)

Another tool developed at Netflix to increase service quality, Scryer, addresses some of the shortcomings of Amazon Auto Scaling (AAS), which dynamically increases and decreases AWS compute server counts based on workload data. Scryer works by predicting what customer demands will be based on historical usage patterns and provisions the necessary capacity.¹¹

Scryer addressed three problems with AAS. The first was dealing with rapid spikes in demand. Because AWS instance startup times could be ten to forty-five minutes, additional compute capacity was often delivered too late to deal with spikes in demand.

The second problem was that after outages, the rapid decrease in customer demand led to AAS removing too much compute capacity to handle future incoming demand. The third problem was that AAS didn't factor in known usage traffic patterns when scheduling compute capacity.¹²

Netflix took advantage of the fact that their consumer viewing patterns were surprisingly consistent and predictable, despite not having Gaussian distributions. [Figure 15.4](#) is a chart reflecting customer requests per second throughout the work week,

showing regular and consistent customer viewing patterns Monday through Friday.¹³




Figure 15.4: Netflix Customer Viewing Demand for Five Days

Source: Jacobson, Yuan, and Joshi, “Scryer: Netflix’s Predictive Auto Scaling Engine,” The Netflix Tech Blog, November 5, 2013, <http://techblog.netflix.com/2013/11/scryer-netflixs-predictive-auto-scaling.html>.

Scryer uses a combination of outlier detections to throw out spurious data points and then uses techniques such as fast fourier transform (FFT) and linear regression to smooth the data while preserving legitimate traffic spikes that recur in their data. The result is that Netflix can forecast traffic demand with surprising accuracy. (See [Figure 15.5](#)).




Figure 15.5: Netflix Scryer Forecasting Customer Traffic and the Resulting AWS Schedule of Compute Resources

(Source: Jacobson, Yuan, Joshi, “Scryer: Netflix’s Predictive Auto Scaling Engine.”)

Only months after first using Scryer in production, Netflix significantly improved their customer viewing experience, improved service availability, and reduced Amazon EC2 costs.

The Netflix case study illustrates how the Scryer tool leveraged the power of non-Gaussian data sets to better understand customers, and then used their behavior to detect and predict issues.

Using Anomaly Detection Techniques

When our data does not have Gaussian distribution, we can still find noteworthy variances using a variety of methods. These techniques are broadly categorized as *anomaly detection*, often defined as “the search for items or events which do not conform to an expected pattern.”¹⁴ Some of these capabilities can be found inside our monitoring tools, while others may require help from people with statistical skills.

Tarun Reddy, VP of Development and Operations at Rally Software, actively advocates this active collaboration between Operations and statistics, observing:

to better enable service quality, we put all our production metrics into Tableau, a statistical analysis software package. We even have an Ops engineer trained in statistics who writes R code (another statistical package)—this engineer has her own backlog, filled with requests from other teams inside the company who want to find variance ever earlier, before it causes an even larger variance that could affect our customers.¹⁵

One of the statistical techniques we can use is called *smoothing*, which is especially suitable if our data is a time series, meaning each data point has a time stamp (e.g., download events, completed transaction events, etc.). Smoothing often involves using moving averages (or rolling averages), which transform our data by averaging each point with all the other data within our sliding window. This has the effect of smoothing out short-term fluctuations and highlighting longer-term trends or cycles.[‡]

An example of this smoothing effect is shown in [Figure 15.6](#). The light line represents the raw data, while the dark line indicates the thirty day moving average (i.e., the average of the trailing thirty days).[§]




Figure 15.6: Autodesk Share Price and Thirty-Day Moving Average Filter

Source: Jacobson, Yuan, Joshi, “Scriber: Netflix’s Predictive Auto Scaling Engine.”

More exotic filtering techniques exist, such as fast fourier transforms, which has been widely used in image processing, and the Kolmogorov-Smirnov test (found in Graphite and Grafana), which is often used to find similarities or differences in periodic/seasonal metric data.

We can expect that a large percentage of telemetry concerning user data will have periodic/seasonal similarities—web traffic, retail transactions, movie watching, and many other user behaviors have very regular and surprisingly predictable daily, weekly, and yearly patterns. This enables us to be able to detect situations that vary from historical norms, such as when our order transaction rate on a Tuesday afternoon drops to 50% of our weekly norms.

Because of the usefulness of these techniques in forecasting, we may be able to find people in the Marketing or Business Intelligence departments with the knowledge and skills necessary to analyze this data. We may want to seek these people out and explore working together to identify

shared problems and use improved anomaly detection and incident prediction to solve them.¹⁵

CASE STUDY

Advanced Anomaly Detection (2014)

At Monitorama in 2014, Dr. Toufic Boubez described the power of using anomaly detection techniques, specifically highlighting the effectiveness of the Komogorov-Smirnov test, a technique that is often used in statistics to determine whether two data sets differ significantly and is found in the popular Graphite and Grafana tools.¹⁶

Figure 15.7 shows the number of transactions per minute at an e-commerce site. Note the weekly periodicity of the graph, with transaction volume dropping on the weekends. By visual inspection, we can see that something peculiar seems to happen on the fourth week when normal transaction volume doesn't return to normal levels on Monday. This suggests an event we should investigate.




Figure 15.7: Transaction Volume: Under-Alerting Using “Three Standard Deviations” Rule

(Source: Dr. Toufic Boubez, “Simple math for anomaly detection.”)

Using the three standard deviations rule would only alert us twice, missing the critical Monday dropoff in transaction volume. Ideally, we would also want to be alerted that the data has drifted from our expected Monday pattern.

“Even saying ‘Kolmogorov-Smirnov’ is a great way to impress everyone,” Dr. Boubez jokes.¹⁷

But what Ops engineers should tell statisticians is that these types of non-parametric techniques are great for Operations data, because it makes no assumptions about normality or any other probability distribution, which is crucial for us to understand what’s going on in our very complex systems. These techniques compare two probability distributions, allowing us to compare periodic or seasonal data, which

helps us find variances in data that varies from day to day or week to week.¹⁸

Figure 15.8 shows the same data set with the K-S filter applied, with the third area highlighting the anomalous Monday where transaction volume didn't return to normal levels. This would have alerted us of a problem in our system that would have been virtually impossible to detect using visual inspection or using standard deviations. In this scenario, this early detection could prevent a customer impacting event, as well as better enable us to achieve our organizational goals.

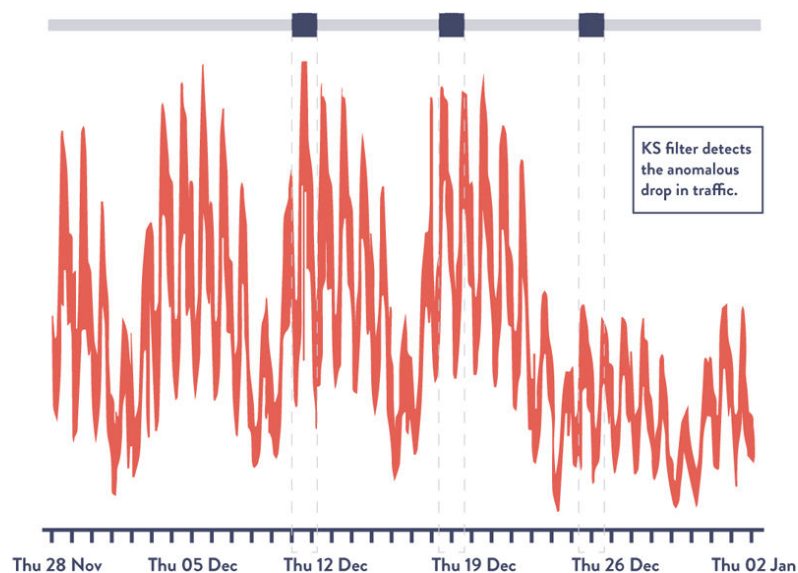


Figure 15.8: Transaction Volume: Using Kolmogorov-Smirnov Test to Alert on Anomalies

Source: Dr. Toufic Boubez, "Simple math for anomaly detection."

This case study demonstrates how we can still find valuable variances in data, even without Guassian distribution. It illustrates how we can use these techniques in our work, as well as how they're likely being used in our organizations in completely different applications.

Conclusion

In this chapter, we explored several different statistical techniques that can be used to analyze our production telemetry so we can find and fix problems earlier than ever, often when they are still small and long before they cause catastrophic outcomes. This enables us to find ever-weaker failure signals that we can then act upon, creating an ever-safer system of work, as well as increasing our ability to achieve our goals.

Specific case studies were presented, including how Netflix used these techniques to proactively remove servers from production and auto-scale their compute infrastructure. We also discussed how to use a moving average and the Kolmogorov-Smirnov filter, both of which can be found in popular telemetry graphing tools.

In the next chapter, we will describe how to integrate production telemetry into the daily work of Development in order to make deployments safer and improve the system as a whole.

* The work done at Netflix highlights one very specific way we can use telemetry to mitigate problems before they impact our customers and before they snowball into significant problems for our teams.

† For the remainder of this book, we will use the terms telemetry, metric, and data sets interchangeably—in other words, a metric (e.g., “page load times”) will map to a data set (e.g., 2 ms, 8 ms, 11 ms, etc.), the term used by statisticians to describe a matrix of data points where each column represents a variable of which statistical operations are performed.

‡ Smoothing and other statistical techniques are also used to manipulate graphic and audio files—for instance, image smoothing (or blurring) as each pixel is replaced by the average of all its neighbors.

§ Other examples of smoothing filters include weighted moving averages or exponential smoothing (which linearly or exponentially weight more recent data points over older data points, respectively), and so forth.

¶ Tools we can use to solve these types of problems include Microsoft Excel (which remains one of the easiest and fastest ways to manipulate data for one-time

purposes), as well as statistical packages such as SPSS, SAS, and the open source R project, now one of the most widely used statistical packages. Many other tools have been created, including several that Etsy has open-sourced, such as Oculus, which finds graphs with similar shapes that may indicate correlation; Opsweekly, which tracks alert volumes and frequencies; and Skyline, which attempts to identify anomalous behavior in system and application graphs.

16

ENABLE FEEDBACK SO DEVELOPMENT AND OPERATIONS CAN SAFELY DEPLOY CODE

In 2006, Nick Galbreath was VP of Engineering at Right Media and responsible for both the Development and Operations departments for an online advertising platform that displayed and served over ten billion impressions daily.¹

Galbreath described the competitive landscape they operated in:

In our business, ad inventory levels were extremely dynamic, so we needed to respond to market conditions within minutes. This meant that Development had to be able to quickly make code changes and get them into production as soon as possible, otherwise we would lose to faster competitors. We found that having a separate group for testing, and even deployment, was simply too slow. We had to integrate all these functions into one group, with shared responsibilities and goals. Believe it or not, our biggest challenge was getting developers to overcome their fear of deploying their own code!²

There is an interesting irony here: Dev often complains about Ops being afraid to deploy code. But in this case, when given the power to deploy their own code, developers became just as afraid to perform code deployments.

The fear of deploying code that was shared by both Dev and Ops at Right Media is not unusual. However, Galbreath observed that providing faster and more frequent feedback to engineers performing deployments (whether Dev or Ops), as

well as reducing the batch size of their work, created safety and then confidence.³

After observing many teams go through this transformation, Galbreath describes their progression as follows:

We start with no one in Dev or Ops being willing to push the “deploy code” button that we’ve built that automates the entire code deployment process, because of the paralyzing fear of being the first person to potentially bring all of the production systems down. Eventually, when someone is brave enough to volunteer to push their code into production, inevitably, due to incorrect assumptions or production subtleties that weren’t fully appreciated, the first production deployment doesn’t go smoothly—and because we don’t have enough production telemetry, we only find out about the problems when customers tell us.⁴

To fix the problem, teams should urgently fix the code and push it into production, but this time with more production telemetry added to applications and environment. This way, we can actually confirm that our fix restored service correctly, and we’ll be able to detect this type of problem before a customer tells us next time.

Later, more developers start to push their own code into production. And because we’re working in a complex system, we’ll still probably break something in production, but this time we’ll be able to quickly see what functionality broke, and quickly decide whether to roll back or fix forward, resolving the problem. This is a huge victory for the entire team and everyone celebrates—now we’re on a roll.

However, the team wants to improve the outcomes of their deployments, so developers proactively get more peer

reviews of their code changes (described in [Chapter 18](#)), and everyone helps each other write better automated tests so we can find errors before deployment. And because everyone now knows that the smaller our production changes, the fewer problems we will have, developers start checking ever-smaller increments of code more frequently into the deployment pipeline, ensuring that their change is working successfully in production before moving to their next change.

We are now deploying code more frequently than ever, and service stability is better than ever too. We have rediscovered that the secret to smooth and continuous flow is making small, frequent changes that anyone can inspect and easily understand.

Galbreath observes that the above progression benefits everyone, including Development, Operations, and Infosec.

As the person who is also responsible for security, it's reassuring to know that we can deploy fixes into production quickly, because changes are going into production throughout the entire day. Furthermore, it always amazes me how interested every engineer becomes in security when you find problems in their code that they are responsible for and that they can quickly fix themselves.⁵

The Right Media story shows that it is not enough to merely automate the deployment process—we must also integrate the monitoring of production telemetry into our deployment work, as well as establish the cultural norm that everyone is equally responsible for the health of the entire value stream.

In this chapter, we create the feedback mechanisms that enable us to improve the health of the value stream at every

stage of the service life cycle, from product design through development and deployment and into operation and eventually retirement. By doing this, we ensure that our services are “production ready,” even at the earliest stages of the project, as well as integrate the learnings from each release and production problem into our future work, resulting in better safety and productivity for everyone.

Use Telemetry to Make Deployments Safer

In this step, we ensure that we are actively monitoring our production telemetry when anyone performs a production deployment, as was illustrated in the Right Media story. This allows whoever is doing the deployment, be it Dev or Ops, to quickly determine whether features are operating as designed after the new release is running in production. After all, we should never consider our code deployment or production change to be done until it is operating as designed in the production environment.

We do this by actively monitoring the metrics associated with our feature during our deployment to ensure we haven’t inadvertently broken our service—or worse, that we broke another service. If our change breaks or impairs any functionality, we quickly work to restore service, bringing in whoever else is required to diagnose and fix the issue.*

As described in [Part III](#), our goal is to catch errors in our deployment pipeline before they get into production. However, there will still be errors that we don’t detect, and we rely on production telemetry to quickly restore service. We may choose to turn off broken features with feature toggles (which is often the easiest and least risky option since it involves no deployments to production), or *fix forward* (i.e., make code changes to fix the defect, which are then pushed into production through the deployment pipeline), or *roll back*

(e.g., switch back to the previous release by using feature toggles or by taking broken servers out of rotation using the blue-green or canary release patterns, etc.).

Although fixing forward can often be dangerous, it can be extremely safe when we have automated testing, fast deployment processes, and sufficient telemetry that allows us to quickly confirm whether everything is functioning correctly in production.

Figure 16.1 shows a deployment of PHP code change at Etsy that generated a spike in PHP runtime warnings—in this case, the developer quickly noticed the problem within minutes and generated a fix and deployed it into production, resolving the issue in less than ten minutes.




Figure 16.1: Deployment to [Etsy.com](#) Causes PHP Run-Time Warnings and Is Quickly Fixed

Source: Mike Brittain, “Tracking Every Release.”

Because production deployments are one of the top causes of production issues, each deployment and change event is overlaid onto our metric graphs to ensure that everyone in the value stream is aware of relevant activity,

enabling better communication and coordination, as well as faster detection and recovery.

Dev Shares Pager Rotation Duties with Ops

Even when our production deployments and releases go flawlessly, in any complex service we will still have unexpected problems, such as incidents and outages that happen at inopportune times (every night at 2 AM). Left unfixed, these can cause recurring problems and suffering for Ops engineers downstream, especially when these problems are not made visible to the upstream engineers responsible for creating the problem.

Even if the problem results in a defect being assigned to the feature team, it may be prioritized below the delivery of new features. The problem may keep recurring for weeks, months, or even years, causing continual chaos and disruption in Operations. This is an example of how upstream work centers can locally optimize for themselves but actually degrade performance for the entire value stream.

To prevent this from happening, we will have everyone in the value stream share the downstream responsibilities of handling operational incidents. We can do this by putting developers, development managers, and architects on pager rotation, just as Pedro Canahuati, Facebook Director of Production Engineering, did in 2009.⁶ This ensures everyone in the value stream gets visceral feedback on any upstream architectural and coding decisions they make.

By doing this, Operations doesn't struggle, isolated and alone, with code-related production issues; instead, everyone is helping find the proper balance between fixing production defects and developing new functionality, regardless of where they reside in the value stream. As Patrick Lightbody, SVP of Product Management at New Relic, observed in 2011, "We

found that when we woke up developers at 2 AM, defects were fixed faster than ever.”⁷

One side effect of this practice is that it helps Development management see that business goals are not achieved simply because features have been marked as “done.” Instead, the feature is only done when it is performing as designed in production, without causing excessive escalations or unplanned work for either Development or Operations.[†]

This practice is equally applicable for market-oriented teams responsible for both developing the feature and running it in production, and for functionally oriented teams. As Arup Chakrabarti, Operations Engineering Manager at PagerDuty, observed during a 2014 presentation, “It’s becoming less and less common for companies to have dedicated on-call teams; instead, everyone who touches production code and environments is expected to be reachable in the event of downtime.”⁸

Regardless of how we’ve organized our teams, the underlying principles remain the same: when developers get feedback on how their applications perform in production, which includes fixing it when it breaks, they become closer to the customer. This creates a buy-in that everyone in the value stream benefits from.

Have Developers Follow Work Downstream

One of the most powerful techniques in interaction and user experience design (UX) is contextual inquiry. This is when the product team watches a customer use the application in their natural environment, often working at their desk. Doing so often uncovers startling ways that customers struggle with the application, such as requiring scores of clicks to perform

simple tasks in their daily work, cutting and pasting text from multiple screens, or writing down notes on paper. All of these are examples of compensatory behaviors and workarounds for usability issues.

The most common reaction for developers after participating in a customer observation is dismay, often stating, “How awful it was seeing the many ways we have been inflicting pain on our customers.” These customer observations almost always result in significant learning and a fervent desire to improve the situation for the customer.

Our goal is to use this same technique to observe how our work affects our internal customers. Developers should follow their work downstream, so they can see how downstream work centers must interact with their product to get it running in production.[‡]

Developers want to follow their work downstream—by seeing customer difficulties firsthand, they make better and more informed decisions in their daily work. By doing this, we create feedback on the non-functional aspects of our code—all the elements that are not related to the customer-facing feature—and identify ways that we can improve deployability, manageability, operability, and so on.

UX observation often has a powerful impact on the observers. When describing his first customer observation, Gene Kim, the founder and CTO at Tripwire for thirteen years and co-author of this book, says:

One of the worst moments of my professional career was in 2006 when I spent an entire morning watching one of our customers use our product. I was watching him perform an operation that we expected customers to do weekly, and, to our extreme horror, we discovered that it required sixty-three clicks. This person kept

apologizing, saying things like, “Sorry, there’s probably a better way to do this.”

Unfortunately, there wasn’t a better way to do that operation. Another customer described how initial product setup took 1,300 steps. Suddenly, I understood why the job of managing our product was always assigned to the newest engineer on the team—no one wanted the job of running our product. That was one of the reasons I helped create the UX practice at my company, to help atone for the pain we were inflicting on our customers.

UX design enables the creation of quality at the source and results in far greater empathy for fellow team members in the value stream. Ideally, UX design helps us as we create codified non-functional requirements to add to our shared backlog of work, eventually allowing us to proactively integrate them into every service we build, which is an important part of creating a DevOps work culture.[§]

Have Developers Initially Self-Manage Their Production Service

Even when Developers are writing and running their code in production-like environments in their daily work, Operations may still experience disastrous production releases because it is the first time we actually see how our code behaves during a release and under true production conditions. This result occurs because operational learnings often happen too late in the software life cycle.

Left unaddressed, the result is often production software that is difficult to operate. As an anonymous Ops engineer once said

In our group, most system administrators lasted only six months. Things were always breaking in production, the hours were insane, and application deployments were painful beyond belief—the worst part was pairing the application server clusters, which would take us six hours. During each moment, we all felt like the developers personally hated us.¹⁰

This can be an outcome of not having enough Ops engineers to support all the product teams and the services we already have in production, which can happen in both functionally and market-oriented teams.

One potential countermeasure is to do what Google does, which is have Development groups self-manage their services in production to prove they are stable before they become eligible for an SRE (site reliability engineering) team to manage. By having developers be responsible for deployment and production support, we are far more likely to have a smooth transition to Operations.[¶]

To prevent the possibility of problematic, self-managed services going into production and creating organizational risk, we may define launch requirements that must be met in order for services to interact with real customers and be exposed to real production traffic. Furthermore, to help the product teams, Ops engineers should act as consultants to help them make their services production-ready.

By creating launch guidance, we help ensure that every product team benefits from the cumulative and collective experience of the entire organization, especially Operations. Launch guidance and requirements will likely include the following:¹¹

- **Defect counts and severity:** Does the application actually perform as designed?
- **Type/frequency of pager alerts:** Is the application generating an unsupportable number of alerts in production?
- **Monitoring coverage:** Is the coverage of monitoring sufficient to restore service when things go wrong?
- **System architecture:** Is the service loosely coupled enough to support a high rate of changes and deployments in production?
- **Deployment process:** Is there a predictable, deterministic, and sufficiently automated process to deploy code into production?
- **Production hygiene:** Is there evidence of enough good production habits that would allow production support to be managed by anyone else?

Superficially, these requirements may appear similar to traditional production checklists we have used in the past. However, the key differences are that we require effective monitoring to be in place, deployments to be reliable and deterministic, and architecture to be built that supports fast and frequent deployments.

If any deficiencies are found during the review, the assigned Ops engineer should help the feature team resolve the issues or even help re-engineer the service if necessary, so that it can be easily deployed and managed in production.

At this time, we may also want to learn whether this service is subject to any regulatory compliance objectives or if it is likely to be in the future:

- Does the service generate a significant amount of revenue? (For example, if it is more than 5% of total

revenue of a publicly held US corporation, it is a “significant account” and in scope for compliance with Section 404 of the Sarbanes-Oxley Act of 2002 [SOX].)

- Does the service have high user traffic or have high outage/impairment costs? (i.e., do operational issues risk creating availability or reputational risk?)
- Does the service store payment cardholder information (such as credit card numbers) or personally identifiable information (such as Social Security numbers or patient care records)? Are there other security issues that could create regulatory, contractual obligation, privacy, or reputation risk?
- Does the service have any other regulatory or contractual compliance requirements associated with it, such as US export regulations, PCI-DSS, HIPAA, and so forth?

This information helps ensure that we effectively manage not only the technical risks associated with this service, but also any potential security and compliance risks. It also provides essential input into the design of the production control environment.

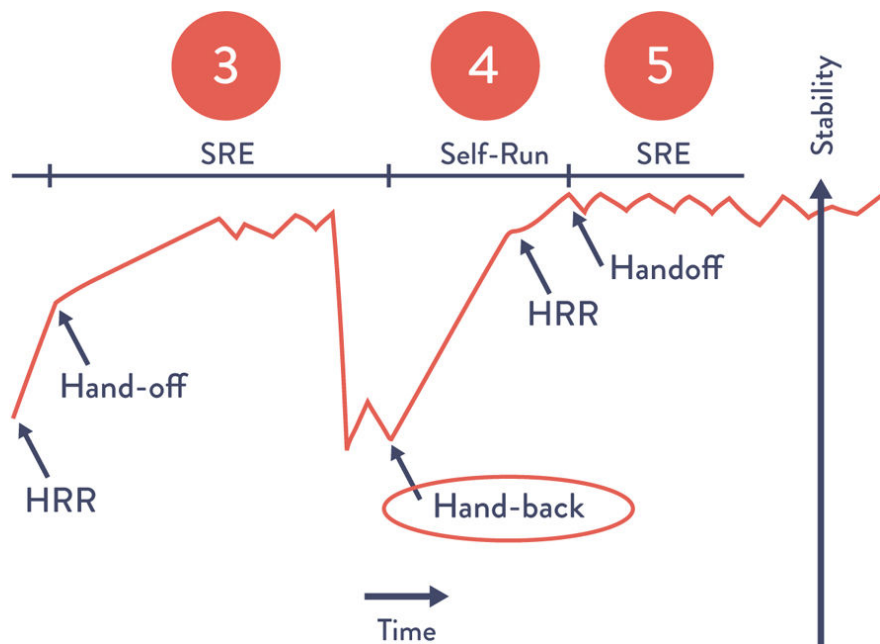


Figure 16.2: The “Service Handback” at Google

Source: “SRE@Google: Thousands of DevOps Since 2004,” YouTube video, 45:57, posted by USENIX, January 12, 2012, <https://www.youtube.com/watch?v=iIuTnhdTzK0>.

By integrating operability requirements into the earliest stages of the development process and having Development initially self-manage their own applications and services, the process of transitioning new services into production becomes smoother, easier, and more predictable to complete. However, for services already in production, we need a different mechanism to ensure that Operations is never stuck with an unsupportable service in production. This is especially relevant for functionally oriented Operations organizations.

In this step, we may create a *service handback mechanism*—in other words, when a production service becomes sufficiently fragile, Operations has the ability to return production support responsibility back to Development.

When a service goes back into a developer-managed state, the role of Operations shifts from production support to

consultation, helping the team make the service production ready.

This mechanism serves as our pressure escape valve, ensuring that we never put Operations in a situation where they are trapped into managing a fragile service while an ever-increasing amount of technical debt buries them and amplifies a local problem into a global problem. This mechanism also helps ensure that Operations has enough capacity to work on improvement work and preventive projects.

The handback remains a long-standing practice at Google and is perhaps one of the best demonstrations of the mutual respect between Dev and Ops engineers. By doing this, Development is able to quickly generate new services, with Ops engineers joining the team when the services become strategically important to the company and, in rare cases, handing them back when they become too troublesome to manage in production.** The following case study of site reliability engineering at Google describes how the handoff readiness review and launch readiness review processes evolved, and the benefits that resulted.

CASE STUDY

The Launch and HandOff Readiness Review at Google (2010)

One of the many surprising facts about Google is that they have a functional orientation for their Ops engineers, who are referred to as “site reliability engineers” (SRE), a term coined by Ben Treynor Sloss in 2004. That year, Treynor Sloss started off with a staff of seven SREs that grew to over 1,200 SREs by 2014.†† As

Trey Sloss said, “If Google ever goes down, it’s my fault.” Trey Sloss has resisted creating a single sentence definition of what SREs are, but he once described SREs as “what happens when a software engineer is tasked with what used to be called operations.”¹²

Every SRE reports to Trey Sloss’s organization to help ensure consistency of quality of staffing and hiring, and they are embedded into product teams across Google (which also provide their funding). However, SREs are still so scarce they are assigned only to the product teams that have the highest importance to the company or those that must comply with regulatory requirements. Furthermore, those services must have low operational burden. Products that don’t meet the necessary criteria remain in a developer-managed state.

Even when new products become important enough to the company to warrant being assigned an SRE, developers still must have self-managed their service in production for at least six months before it becomes eligible to have an SRE assigned to the team.¹³

To help ensure that these self-managed product teams can still benefit from the collective experience of the SRE organization, Google created two sets of safety checks for two critical stages of releasing new services called the *launch readiness review* and the *handoff readiness review* (LRR and HRR, respectively).

The LRR must be performed and approved before any new Google service is made publicly available to customers and receives live production traffic, while the HRR is performed when the service is transitioned to an Ops-managed state, usually months after the LRR. The

LRR and HRR checklists are similar, but the HRR is far more stringent and has higher acceptance standards, while the LRR is self-reported by the product teams.

Any product team going through an LRR or HRR has an SRE assigned to them to help them understand the requirements and to help them achieve those requirements. The LRR and HRR launch checklists have evolved over time so every team can benefit from the collective experiences of all previous launches, whether successful or unsuccessful. Tom Limoncelli noted during his “SRE@Google: Thousands of DevOps Since 2004” presentation in 2012, “Every time we do a launch, we learn something. There will always be some people who are less experienced than others doing releases and launches. The LRR and HRR checklists are a way to create that organizational memory.”¹⁴

Requiring product teams to self-manage their own services in production forces Development to walk in the shoes of Ops, but guided by the LRR and HRR, which not only makes service transition easier and more predictable but also helps create empathy between upstream and downstream work centers.

Limoncelli noted, “In the best case, product teams have been using the LRR checklist as a guideline, working on fulfilling it in parallel with developing their service, and reaching out to SREs to get help when they need it.”¹⁵

Furthermore, Limoncelli observed,

The teams that have the fastest HRR production approval are the ones that worked with SREs earliest, from the early design stages up until launch. And the great thing is,

it's always easy to get an SRE to volunteer to help with your project. Every SRE sees value in giving advice to project teams early, and will likely volunteer a few hours or days to do just that.¹⁶

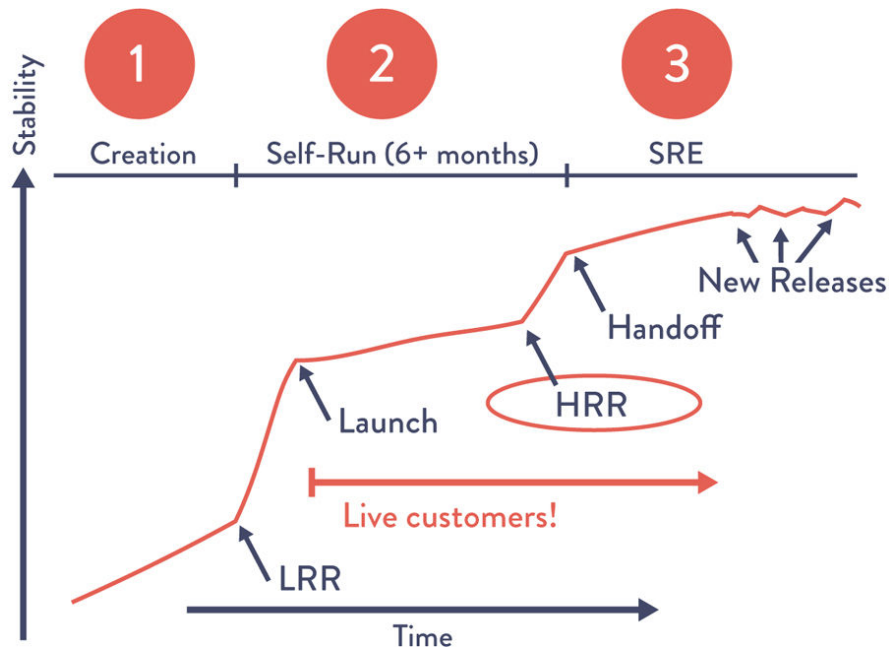


Figure 16.3: The LRR and HRR at Google

Source: “SRE@Google: Thousands of DevOps Since 2004,” YouTube video, 45:57, posted by USENIX, January 12, 2012, <https://www.youtube.com/watch?v=iFuTnhdTzK0>.

The practice of SREs helping product teams early is an important cultural norm that is continually reinforced at Google. Limoncelli explained, “Helping product teams is a long-term investment that will pay off many months later when it comes time to launch. It is a form of ‘good citizenship’ and ‘community service’ that is valued, it is routinely considered when evaluating engineers for SRE promotions.”¹⁷

At Google, having product teams self manage their own services brings learning to the front, giving them valuable insight into how their code behaves in true production conditions. This

practice also strengthens the relationship and understanding between Dev and Ops, creating a cultural feedback loop.

Conclusion

In this chapter, we discussed the feedback mechanisms that enable us to improve our service at every stage of our daily work, whether it is deploying changes into production, fixing code when things go wrong and paging engineers, having developers follow their work downstream, creating non-functional requirements that help development teams write more production-ready code, or even handing problematic services back to be self-managed by Development.

By creating these feedback loops, we make production deployments safer, increase the production readiness of code created by Development, and help create a better working relationship between Development and Operations by reinforcing shared goals, responsibilities, and empathy.

In the next chapter, we explore how telemetry can enable hypothesis-driven development and A/B testing to perform experiments that help us achieve our organizational goals and win in the marketplace.

* By doing this, and by using the required architecture, we “optimize for MTTR, instead of MTBF,” a popular DevOps maxim to describe our desire to optimize for recovering from failures quickly, as opposed to attempting to prevent failures.

† ITIL defines *warranty* as when a service can run in production reliably without intervention for a predefined period of time (e.g., two weeks). This definition of warranty should ideally be integrated into our collective definition of “done.”

‡ By following work downstream, we may uncover ways to help improve flow, such as automating complex, manual steps (e.g., pairing application server clusters that require six hours to successfully complete); performing packaging of code once instead of creating it multiple times at different stages of QA and Production deployment; working with testers to automate manual test suites, thus removing a common bottleneck for more frequent deployment; and creating

more useful documentation instead of having someone decipher developer application notes to build packaged installers.

- § More recently, Jeff Sussna attempted to further codify how to better achieve UX goals in what he calls “digital conversations,” which are intended to help organizations understand the customer journey as a complex system, broadening the context of quality. The key concepts include designing for service, not software; minimizing latency and maximizing strength of feedback; designing for failure and operating to learn; using Operations as an input to design; and seeking empathy.⁹
- ¶ We further increase the likelihood of production problems being fixed by ensuring that the Development teams remain intact and not disbanded after the project is complete.
- ** In organizations with project-based funding, there may be no developers to hand the service back to, as the team has already been disbanded or may not have the budget or time to take on service responsibility. Potential countermeasures include holding an improvement blitz to improve the service, temporarily funding or staffing improvement efforts, or retiring the service.
- †† In this book, we use the term “Ops engineer,” but the two terms, “Ops engineer” and “site reliability engineer,” are intended to be interchangeable.

INTEGRATE HYPOTHESIS-DRIVEN DEVELOPMENT AND A/B TESTING INTO OUR DAILY WORK

All too often in software projects, developers work on features for months or years spanning multiple releases without ever confirming whether the desired business outcomes are being met, such as whether a particular feature is achieving the desired results or even being used at all.

Worse, even when we discover that a given feature isn't achieving the desired results, making corrections to the feature may be out-prioritized by other new features, ensuring that the under-performing feature will never achieve its intended business goal. In general, Jez Humble, co-author of this book, observes, "the most inefficient way to test a business model or product idea is to build the complete product to see whether the predicted demand actually exists."¹

Before we build a feature, we should rigorously ask ourselves, "Should we build it and why?" We should then perform the cheapest and fastest experiments possible to validate through user research whether the intended feature will actually achieve the desired outcomes. We can use techniques such as hypothesis-driven development, customer acquisition funnels, and A/B testing, concepts we explore throughout this chapter. Intuit, Inc. provides a dramatic example of how organizations use these techniques to create products that customers love, to promote organizational learning, and to win in the marketplace.

Hypothesis-Driven Development at Intuit, Inc. (2012)

Intuit is focused on creating business and financial management solutions to simplify life for small businesses, consumers, and accounting professionals. In 2012, they had \$4.5 billion in revenue and 8,500 employees, with flagship products that include QuickBooks, TurboTax, Mint, and, until recently, Quicken.^{2*}

Scott Cook, the founder of Intuit, has long advocated building a culture of innovation, encouraging teams to take an experimental approach to product development, and exhorting leadership to support them. As he said, “Instead of focusing on the boss’s vote ... the emphasis is on getting real people to really behave in real experiments and basing your decisions on that.”³ This is the epitome of a scientific approach to product development.

Cook explained that what is needed is “a system where every employee can do rapid, high-velocity experiments... . Dan Maurer runs our consumer division... . [which] runs the TurboTax website. When he took over, we did about seven experiments a year.”⁴

He continued, “By installing a rampant innovation culture [in 2010], they now do 165 experiments in the three months of the [US] tax season. Business result? [The] conversion rate of the website is up 50 percent... . The folks [team members] just love it, because now their ideas can make it to market.”⁵

Aside from the effect on the website conversion rate, one of the most surprising elements of this story is that TurboTax performed production experiments during their peak traffic seasons. For decades, especially in retailing, the risk of revenue-impacting outages during the holiday season were so

high that they would often put into place a change freeze from mid-October to mid-January.

However, by making software deployments and releases fast and safe, the TurboTax team made online user experimentation and any required production changes a low-risk activity that could be performed during the highest traffic and revenue-generating periods.

This highlights the notion that the period when experimentation has the highest value is during peak traffic seasons. Had the TurboTax team waited until April 16th, the day after the US tax filing deadline, to implement these changes, the company could have lost many of its prospective customers, and even some of its existing customers, to the competition.

The faster we can experiment, iterate, and integrate feedback into our product or service, the faster we can learn and out-experiment the competition. And how quickly we can integrate our feedback depends on our ability to deploy and release software.

The Intuit example shows that the Intuit TurboTax team was able to make this situation work for them and won in the marketplace as a result.

A Brief History of A/B Testing

As the Intuit TurboTax story highlights, an extremely powerful user research technique is defining the customer acquisition funnel and performing A/B testing. A/B testing techniques were pioneered in *direct response marketing*, which is one of the two major categories of marketing strategies. The other is called *mass marketing* or *brand marketing* and often relies on placing as many ad impressions in front of people as possible to influence buying decisions.

In previous eras, before email and social media, direct response marketing meant sending thousands of postcards or flyers via postal mail and asking prospects to accept an offer by calling a telephone number, returning a postcard, or placing an order.⁶

In these campaigns, experiments were performed to determine which offer had the highest conversion rates. They experimented with modifying and adapting the offer, rewording the offer, varying the copywriting styles, design and typography, packaging, and so forth to determine which was most effective in generating the desired action (e.g., calling a phone number, ordering a product).

Each experiment often required doing another design and print run, mailing out thousands of offers, and waiting weeks for responses to come back. Each experiment typically cost tens of thousands of dollars per trial and required weeks or months to complete. However, despite the expense, iterative testing easily paid off if it significantly increased conversion rates (e.g., the percentage of respondents ordering a product going from 3–12%).

Well-documented cases of A/B testing include campaign fundraising, internet marketing, and the Lean Startup methodology. Interestingly, it has also been used by the British government to determine which letters were most effective in collecting overdue tax revenue from delinquent citizens.^{7†}

Integrating A/B Testing into Our Feature Testing

The most commonly used A/B technique in modern UX practice involves a website where visitors are randomly selected to be shown one of two versions of a page, either a control (the “A”) or a treatment (the “B”). Based on statistical

analysis of the subsequent behavior of these two cohorts of users, we demonstrate whether there is a significant difference in the outcomes of the two, establishing a *causal* link between the treatment (e.g., a change in a feature, design element, background color) and the outcome (e.g., conversion rate, average order size).

For example, we may conduct an experiment to see whether modifying the text or color on a “buy” button increases revenue or whether slowing down the response time of a website (by introducing an artificial delay as the treatment) reduces revenue. This type of A/B testing allows us to establish a dollar value on performance improvements. Sometimes, A/B tests are also known as online controlled experiments and split tests. It’s also possible to run experiments with more than one variable. This allows us to see how the variables interact, a technique known as multivariate testing.

The outcomes of A/B tests are often startling. Dr. Ronny Kohavi, previous Distinguished Engineer and General Manager of the Analysis and Experimentation group at Microsoft, observed that after “evaluating well-designed and executed experiments that were designed to improve a key metric, only about one-third were successful at improving the key metric!”⁸ In other words, two-thirds of features either have a negligible impact or actually make things worse. Kohavi goes on to note that all these features were originally thought to be reasonable, good ideas, further elevating the need for user testing over intuition and expert opinions.⁹

CONTINUOUS LEARNING

If you're interested in more detail about experiment design and A/B testing, check out *Trustworthy Online Controlled Experiments: A Practical Guide to A/B Testing*, from authors Dr. Diane Tang, Dr. Ron Kohavi, and Dr. Ya Xu. The authors share stories from several companies that used online experimentation and A/B testing to improve their products. The book also includes tips and tricks so anyone can improve their products and services using trusted experiments and not poorly informed data that may lead them astray.

The implications of Dr. Kohavi's data are staggering. If we are not performing user research, the odds are that two-thirds of the features we are building deliver zero or negative value to our organization, even as they make our codebase ever more complex, thus increasing our maintenance costs over time and making our software more difficult to change.

Furthermore, the effort to build these features is often made at the expense of delivering features that would deliver value (i.e., opportunity cost). Jez Humble joked, "Taken to an extreme, the organization and customers would have been better off giving the entire team a vacation, instead of building one of these non-value-adding features."¹⁰

Our countermeasure is to integrate A/B testing into the way we design, implement, test, and deploy our features. Performing meaningful user research and experiments ensures that our efforts help achieve our customer and organizational goals, and help us win in the marketplace.

Integrate A/B Testing into Our Release

Fast and iterative A/B testing is made possible by being able to quickly and easily do production deployments on demand, using feature toggles, and potentially delivering multiple versions of our code simultaneously to customer segments. Doing this requires useful production telemetry at all levels of the application stack.

By hooking into our feature toggles, we can control which percentage of users see the treatment version of an experiment. For example, we may have one-half of our customers be our treatment group and one-half get shown the following: “Similar items link on unavailable items in the cart.” As part of our experiment, we compare the behavior of the control group (no offer made) against the treatment group (offer made), perhaps measuring the number of purchases made in that session.

Etsy open-sourced their experimentation framework Feature API (formerly known as the Etsy A/B API), which not only supports A/B testing but also online ramp-ups, enabling throttling exposure to experiments. Other A/B testing products include Optimizely, Google Analytics, etc. In a 2014 interview with Kendrick Wang of Apptimize, Lacy Rhoades at Etsy described their journey:

Experimentation at Etsy comes from a desire to make informed decisions, and ensure that when we launch features for our millions of members, they work. Too often, we had features that took a lot of time and had to be maintained without any proof of their success or any popularity among users. A/B testing allows us to ... say a feature is worth working on as soon as it's underway.¹¹

Integrating A/B Testing into Our Feature Planning

Once we have the infrastructure to support A/B feature release and testing, we must ensure that product owners think about each feature as a hypothesis and use our production releases as experiments with real users to prove or disprove that hypothesis. Constructing experiments should be designed in the context of the overall customer acquisition funnel. Barry O'Reilly, co-author of *Lean Enterprise: How High Performance Organizations Innovate at Scale*, described how we can frame hypotheses in feature development in the following form:¹²

- **We believe** increasing the size of hotel images on the booking page
- **Will result** in improved customer engagement and conversion
- **We will have confidence to proceed when** we see a 5% increase in customers who review hotel images who then proceed to book in forty-eight hours.

Adopting an experimental approach to product development requires us to not only break down work into small units (stories or requirements) but also validate whether each unit of work is delivering the expected outcomes. If it does not, we modify our road map of work with alternative paths that will actually achieve those outcomes.

CASE STUDY

Doubling Revenue Growth through Fast Release Cycle Experimentation at Yahoo! Answers (2010)

In 2009, Jim Stoneham was General Manager of the Yahoo! Communities group that included Flickr and Answers. Previously, he had been primarily responsible for Yahoo! Answers, competing against other Q&A companies such as Quora, Aardvark, and Stack Exchange.¹³

At that time, Answers had approximately 140 million monthly visitors, with over twenty million active users answering questions in over twenty different languages. However, user growth and revenue had flattened, and user engagement scores were declining.¹⁴

Stoneham observes that,

Yahoo Answers was and continues to be one of the biggest social games on the Internet; tens of millions of people are actively trying to “level up” by providing quality answers to questions faster than the next member of the community. There were many opportunities to tweak the game mechanic, viral loops, and other community interactions. When you’re dealing with these human behaviors, you’ve got to be able to do quick iterations and testing to see what clicks with people.¹⁵

He continues,

These [experiments] are the things that Twitter, Facebook, and Zynga did so well. Those organizations were doing experiments at least twice per week—they were even reviewing the changes they made before their deployments, to make sure they were still on track. So here I am, running [the] largest Q&A site in the market, wanting to do rapid

iterative feature testing, but we can't release any faster than once every 4 weeks. In contrast, the other people in the market had a feedback loop 10x faster than us.¹⁶

Stoneham observed that as much as product owners and developers talk about being metrics-driven, if experiments are not performed frequently (daily or weekly), the focus of daily work is merely on the feature they're working on, as opposed to customer outcomes.¹⁷

As the Yahoo! Answers team was able to move to weekly deployments, and later multiple deployments per week, their ability to experiment with new features increased dramatically. Their astounding achievements and learnings over the next twelve months of experimentation included increased monthly visits by 72%, increased user engagement by threefold, and doubled the team's revenue.¹⁸

To continue their success, the team focused on optimizing the following top metrics:

- **Time to first answer:** How quickly was an answer posted to a user question?
- **Time to best answer:** How quickly did the user community award a best answer?
- **Upvotes per answer:** How many times was an answer upvoted by the user community?
- **Answers per week per person:** How many answers were users creating?
- **Second search rate:** How often did visitors have to search again to get an answer? (Lower is better.)

Stoneham concluded, “This was exactly the learning that we needed to win in the marketplace—and it changed more than our feature velocity. We transformed from a team of employees to a team of owners. When you move at that speed, and are looking at the numbers and the results daily, your investment level radically changes.”¹⁹

This story from Yahoo! Answers demonstrates how dramatically outcomes can be affected by faster cycle times. The faster we can iterate and integrate feedback into the product or service we are delivering to customers, the faster we can learn and the bigger the impact we can create.

Conclusion

Success requires us to not only deploy and release software quickly but also to out-experiment our competition. Techniques such as hypothesis-driven development, defining and measuring our customer acquisition funnel, and A/B testing allow us to perform user experiments safely and easily, enabling us to unleash creativity and innovation and create organizational learning. And, while succeeding is important, the organizational learning that comes from experimentation also gives employees ownership of business objectives and customer satisfaction. In the next chapter, we examine and create review and coordination processes as a way to increase the quality of our current work.

* In 2016, Intuit sold the Quicken business to the private equity firm HIG Capital.

† There are many other ways to conduct user research before embarking on development. Among the most inexpensive methods include performing surveys, creating prototypes (either mock-ups using tools such as Balsamiq or interactive versions written in code), and performing usability testing. Alberto Savoia, Director of Engineering at Google, coined the term *pretotyping* for the practice of

using prototypes to validate whether we are building the right thing. User research is so inexpensive and easy relative to the effort and cost of building a useless feature in code that, in almost every case, we shouldn't prioritize a feature without some form of validation.

18

CREATE REVIEW AND COORDINATION PROCESSES TO INCREASE QUALITY OF OUR CURRENT WORK

In the previous chapters, we created the telemetry necessary to see and solve problems in production and at all stages of our deployment pipeline and created fast feedback loops from customers to help enhance organizational learning—learning that encourages ownership and responsibility for customer satisfaction and feature performance, which help us succeed.

Our goal in this chapter is to enable Development and Operations to reduce the risk of production changes before they are made. Traditionally, when we review changes for deployment, we tend to rely heavily on reviews, inspections, and approvals just prior to deployment. Frequently, those approvals are given by external teams who are often too far removed from the work to make informed decisions on whether a change is risky or not, and the time required to get all the necessary approvals also lengthens our change lead times.

Peer Review at GitHub (2011)

The peer review process at GitHub is a striking example of how inspection can increase quality, make deployments safe, and be integrated into the flow of everyone's daily work. GitHub pioneered the process called *pull request*, one of the most popular forms of peer review that span Dev and Ops.

Scott Chacon, CIO and co-founder of GitHub, wrote on his website that pull requests are the mechanism that lets

engineers tell others about changes they have pushed to a repository on GitHub. Once a pull request is sent, interested parties can review the set of changes, discuss potential modifications, and even push follow-up commits if necessary. Engineers submitting a pull request will often request a “+1,” “+2,” or so forth, depending on how many reviews they need, or “@mention” engineers that they’d like to get reviews from.¹

At GitHub, pull requests are also the mechanism used to deploy code into production through a collective set of practices they call “GitHub Flow”—it’s how engineers request code reviews, gather and integrate feedback, and announce that code will be deployed to production (i.e., “master” branch).

The screenshot shows a GitHub pull request interface. On the left, a diff view shows line numbers 27 through 38. The code changes are as follows:

```
+      opts[:options] [:stripnl] || = false
35      28
36      29
37      30
38      31
```

Two comments are visible:

- brianmario** (repo collab):
So what are the defaults here if no encoding or lexer is passed?
Also there's at least one other place where the API is expected to take an :encoding key (not nested under an :options key/hash) - <https://github.com/github/github/blob/master/app/models/gist.rb#L114>
Only reason I did it that way was to sorta abstract the fact that we're using pygments for coloring currently (not that we have plans to change that anytime soon...)
- Josh** (repo collab):
Alright, I'll push that down to colorize.

At the bottom, there is a button labeled "Add a line note".

Figure 18.1: Comments and Suggestions on a GitHub Pull Request

Source: Scott Chacon, “GitHub Flow,” [ScottChacon.com](http://scottchacon.com), August 31, 2011, <http://scottchacon.com/2011/08/31/github-flow.html>.

GitHub Flow is composed of five steps:

1. To work on something new, the engineer creates a descriptively named branch off of master (e.g., “new-

oauth2-scopes”).

2. The engineer commits to that branch locally, regularly pushing their work to the same named branch on the server.
3. When they need feedback or help, or when they think the branch is ready for merging, they open a pull request.
4. When they get their desired reviews and get any necessary approvals of the feature, the engineer can then merge it into master.
5. Once the code changes are merged and pushed to master, the engineer deploys them into production.

These practices, which integrate review and coordination into daily work, have allowed GitHub to quickly and reliably deliver features to market with high quality and security. For example, in 2012 they performed an amazing 12,602 deployments. In particular, on August 23rd, after a company-wide summit where many exciting ideas were brainstormed and discussed, the company had their busiest deployment day of the year, with 563 builds and 175 successful deployments into production, all made possible through the pull request process.²

Throughout this chapter, we will integrate practices, such as those used at GitHub, to shift our reliance away from periodic inspections and approvals and move to integrated peer review performed continually as a part of our daily work. Our goal is to ensure that Development, Operations, and Infosec are continuously collaborating so that changes we make to our systems will operate reliably, securely, safely, and as designed.

The Dangers of Change Approval Processes

The Knight Capital failure is one of the most prominent software deployment errors in recent memory. A fifteen-minute deployment error resulted in a \$440 million trading loss, during which the engineering teams were unable to disable the production services. The financial losses jeopardized the firm's operations and forced the company to be sold over the weekend so they could continue operating without jeopardizing the entire financial system.³

John Allspaw observed that when high-profile incidents occur, such as the Knight Capital deployment accident, there are typically two *counterfactual* narratives for why the accident occurred.^{4*}

The first narrative is that the accident was due to a *change control failure*, which seems valid because we can imagine a situation where better change control practices could have detected the risk earlier and prevented the change from going into production. And if we couldn't prevent it, we might have taken steps to enable faster detection and recovery.

The second narrative is that the accident was due to a *testing failure*. This also seems valid: with better testing practices we could have identified the risk earlier and canceled the risky deployment or we could have at least taken steps to enable faster detection and recovery.

The surprising reality is that in environments that have low-trust and command-and-control cultures the outcomes of these types of change control and testing countermeasures often result in an increased likelihood that problems will occur again, potentially with even worse outcomes.

Gene Kim (co-author of this book) describes his realization that change and testing controls can potentially have the opposite effect than intended as "one of the most important moments of my professional career. This 'aha'

moment was the result of a conversation in 2013 with John Allspaw and Jez Humble about the Knight Capital accident, making me question some of my core beliefs that I've formed over the last ten years, especially having been trained as an auditor.”

He continues, “However upsetting it was, it was also a very formative moment for me. Not only did they convince me that they were correct, we tested these beliefs in the 2014 *State of DevOps Report*, which led to some astonishing findings that reinforce that building high-trust cultures is likely the largest management challenge of this decade.”

Potential Dangers of “Overly Controlling Changes”

Traditional change controls can lead to unintended outcomes, such as contributing to long lead times and reducing the strength and immediacy of feedback from the deployment process. In order to understand how these unintended outcomes happen, let us examine the controls we often put in place when change control failures occur:

- Adding more questions to the change request form.
- Requiring more authorizations, such as one more level of management approval (e.g., instead of merely the VP of Operations approving, we now require that the CIO also approve) or more stakeholders (e.g., network engineering, architecture review boards, etc.).
- Requiring more lead time for change approvals so that change requests can be properly evaluated.

These controls often add more friction to the deployment process by multiplying the number of steps and approvals, as well as increasing batch sizes and deployment lead times, which we know reduces the likelihood of successful

production outcomes for both Dev and Ops. These controls also reduce how quickly we get feedback from our work.

One of the core beliefs in the Toyota Production System is that “people closest to a problem typically know the most about it.”⁵ This becomes more pronounced as the work being performed and the system the work occurs in become more complex and dynamic, as is typical in DevOps value streams. In these cases, creating approval steps from people who are located further and further away from the work may actually reduce the likelihood of success. As has been proven time and again, the further the distance between the person doing the work (i.e., the change implementer) and the person deciding to do the work (i.e., the change authorizer), the worse the outcome.

CONTINUOUS LEARNING

Research shows that elite organizations do change approvals quickly, with clarity, and without friction—and that this contributes to better software delivery performance. DORA's 2019 *State of DevOps Report* found that a clear, lightweight change process, where developers are confident they can go from “submitted” to “accepted” for all typical approvals, contributes to high performance.⁶ In contrast, heavyweight change processes, where an external review board or senior manager is required for significant changes, has a negative impact on performance.⁷ This echoes findings from the 2014 Puppet *State of DevOps Report*, which found that high-performing teams relied more on peer review and less on external approval of changes.⁸

In many organizations, change advisory boards serve an important role in coordinating and governing the delivery process, but their job should not be to manually evaluate every change, nor does ITIL mandate such a practice. To understand why this is the case, consider the predicament of being on a change advisory board, reviewing a complex change composed of hundreds of thousands of lines of code changes and created by hundreds of engineers.

At one extreme, we cannot reliably predict whether a change will be successful either by reading a hundred-word description of the change or by merely validating that a checklist has been completed. At the other extreme, painfully scrutinizing thousands of lines of code changes is unlikely to

reveal any new insights. Part of this is the nature of making changes inside of a complex system. Even the engineers who work inside the codebase as part of their daily work are often surprised by the side effects of what should be low-risk changes.

For all these reasons, we need to create effective control practices that more closely resemble peer review, reducing our reliance on external bodies to authorize our changes. We also need to coordinate and schedule changes effectively. We explore both of these in the next two sections.

→ CASE STUDY: NEW TO THE SECOND EDITION

From Six-Eye Principle to Release at Scale at Adidas (2020)

In November of 2020, after a tiring year dealing with the impacts of the Covid-19 pandemic and just before their peak sales period, Adidas found themselves in crisis mode. After experiencing five really bad outages, it was clear they had reached their limit. They no longer had things under control.⁹

In the year before the crisis, Adidas had experienced 10x growth, increasing digital business revenue by almost 50%. This growth also meant that they experienced two to three times more visitors to their platforms and tracked exponentially more data from them. This created 10x more technical traffic and load. Along with this growth, the number of technical teams and capabilities was constantly growing, as well as the dependencies among them per Brooke's Law.¹⁰

On a peak day, order rates were reaching three thousand orders per minute, and Adidas was sending

out 11 billion touchpoints every day. But their strategy was to double that rate. For hype drop products (special product releases), they were shooting for 1.5 million hits per second.¹¹

Going back to November 2020, after years of growth and freedom, Adidas had found themselves in the nightmarish situation where they had to bring three VPs into a room to approve every single change or release during the two-month peak sales period.¹²

Fernando Cornago, Vice President Digital Tech, was one of those VPs in the room. “I can tell you the reality of how clueless at the end we were on some of the details.”¹³

At that time, Adidas had more than 550 million lines of code in their ecosystem and almost two thousand engineers. By the end of the crisis, it was clear they had to change.¹⁴ Thus began a new wave of handling operations and release management for readiness.

They began by asking three questions on stability:

- How do we detect interruptions as fast as possible?
- How do we fix interruptions fast?
- And, how do we make sure interruptions don't even land in the production environment?

They brought in ITIL and SRE practices to help them answer these questions. And what they began to see was that interruptions were not affecting only one product. “And that’s where we realized that everything is connected,” explains Vikalp Yadav, Senior Director—Head of Digital SRE Operations.¹⁵

They realized they had to look at the value stream as a thought process. By doing so, they developed the “%revenue bleed versus net sales” KPI, which measured the larger impact of an outage. Finally, to ensure they achieved observability, resilience, and release excellence, they adopted what they called: release fitness.¹⁶

Adidas’s environment had become huge, and the complexity had grown as well. “Now, it’s not only about how my system is but also how is the whole ecosystem?” Andreia Otto, SRE Champion WnM Services, explains.¹⁷ Release processes needed to be standardized.

They worked with the product teams and service management team to find a set of KPIs or factors to check before any release. It began as an Excel spreadsheet. Each team had to fill out the spreadsheet before each release to check to see if they had a go or no-go.

Obviously, this wasn’t a sustainable or well-received option. Manually filling out a spreadsheet for every release was tedious and time-consuming. They knew they needed to find a way to automate the assessment, which is exactly what they did.

They developed a dashboard that looked at every release from three different angles: system level (how my product is), value stream (upstream/downstream dependencies, etc.), and environment (platform, events, etc.). Checking across these three angles, the dashboard gives a clear go/no-go for release.¹⁸

Before anything is released into production, this automated check is completed. If everything is fine, it goes into production. If not, the release is stopped and

the team responsible checks the dashboard to see what the problem is. There might be an event, like Adidas' hype drop days where no changes are allowed, or it might be the team doesn't have any more error budget, etc.

With their release fitness program, Adidas was able to build a system that is able to self-adjust and self-regulate. On the one hand, they have strict release guidelines. But on the other hand, they have automated checks and error budgets telling any of their developers if they can deploy or not. They no longer need three VPs in a room to approve every change. Last but not least, in a growing ecosystem where Adidas onboards approximately one hundred new engineers on a monthly basis, the release fitness automation massively reduces their onboarding time

Adidas automated change reviews, ensuring the quality of the code across numerous dependencies, without the need for costly and slow review boards.

Enable Coordination and Scheduling of Changes

Whenever we have multiple groups working on systems that share dependencies, our changes will likely need to be coordinated to ensure that they don't interfere with each other (e.g., marshaling, batching, and sequencing the changes). In general, the more loosely coupled our architecture, the less we need to communicate and coordinate with other component teams—when the architecture is truly service oriented, teams can make changes with a high degree of autonomy, where local changes are unlikely to create global disruptions.

However, even in a loosely coupled architecture, when many teams are doing hundreds of independent deployments per day, there may be a risk of changes interfering with each other (e.g., simultaneous A/B tests). To mitigate these risks, we may use chat rooms to announce changes and proactively find collisions that may exist.

For more complex organizations and organizations with more tightly coupled architectures, we may need to deliberately schedule our changes, where representatives from the teams get together, not to authorize changes, but to schedule and sequence their changes in order to minimize accidents.

However, certain areas, such as global infrastructure changes (e.g., core network switch changes) will always have a higher risk associated with them. These changes will always require technical countermeasures, such as redundancy, failover, comprehensive testing, and (ideally) simulation.

Enable Peer Review of Changes

Instead of requiring approval from an external body prior to deployment, we may require engineers to get peer reviews of their changes. In Development, this practice has been called code review, but it is equally applicable to any change we make to our applications or environments, including servers, networking, and databases. The goal is to find errors by having fellow engineers close to the work scrutinize our changes. This review improves the quality of our changes, which also creates the benefits of cross-training, peer learning, and skill improvement.

A logical place to require reviews is prior to committing code to trunk in source control, where changes could potentially have a team-wide or global impact. At a minimum, fellow engineers should review our change, but for higher risk

areas, such as database changes or business-critical components with poor automated test coverage, we may require further review from a subject matter expert (e.g., information security engineer, database engineer) or multiple reviews (e.g., “+2” instead of merely “+1”).[†]

The principle of small batch sizes also applies to code reviews. The larger the size of the change that needs to be reviewed, the longer it takes to understand and the larger the burden on the reviewing engineer. As Randy Shoup observed, “There is a non-linear relationship between the size of the change and the potential risk of integrating that change—when you go from a ten line code change to a one hundred line code, the risk of something going wrong is more than ten times higher, and so forth.”¹⁹ This is why it’s so essential for developers to work in small, incremental steps rather than on long-lived feature branches.

Furthermore, our ability to meaningfully critique code changes goes down as the change size goes up. As Giray Özil tweeted, “Ask a programmer to review ten lines of code, he’ll find ten issues. Ask him to do five hundred lines, and he’ll say it looks good.”²⁰

Guidelines for code reviews include:

- Everyone must have someone to review their changes (e.g., to the code, environment, etc.) before committing to trunk.
- Everyone should monitor the commit stream of their fellow team members so that potential conflicts can be identified and reviewed.
- Define which changes qualify as high risk and may require review from a designated subject matter

expert (e.g., database changes, security-sensitive modules such as authentication, etc.).[‡]

- If someone submits a change that is too large to reason about easily—in other words, you can't understand its impact after reading through it a couple of times, or you need to ask the submitter for clarification—it should be split up into multiple, smaller changes that can be understood at a glance.

To ensure that we are not merely rubber-stamping reviews, we may also want to inspect the code review statistics to determine the number of proposed changes approved versus not approved, and perhaps sample and inspect specific code reviews.

Code reviews come in various forms:

- **Pair programming:** Programmers work in pairs (see section below).
- **“Over the shoulder:”** One developer looks over the author's shoulder as the latter walks through the code.
- **Email pass-around:** A source code management system emails code to reviewers automatically after the code is checked in.
- **Tool-assisted code review:** Authors and reviewers use specialized tools designed for peer code review (e.g., Gerrit, GitHub pull requests, etc.) or facilities provided by the source code repositories (e.g., GitHub, Mercurial, Subversion, as well as other platforms such as Gerrit, Atlassian Stash, and Atlassian Crucible).

Close scrutiny of changes in many forms is effective in locating errors previously overlooked. Code reviews can facilitate increased code commits and production deployments, and support trunk-based deployment and

continuous delivery at scale, as we will see in the following case study.

CASE STUDY

Code Reviews at Google (2010)

As noted earlier in this book, Eran Messeri described that in 2013 the processes at Google enabled over thirteen thousand developers to work off of trunk on a single source code tree, performing over fifty-five hundred code commits per week, resulting in hundreds of production deployments per week.²¹

In 2016, Google's twenty-five thousand developers worldwide were committing sixteen thousand changes into trunk on a typical workday, with a further twenty-four thousand changes per day committed by automated systems.²²

This requires considerable discipline from Google team members and mandatory code reviews, which cover the following areas:²³

- code readability for languages (enforces style guide)
- ownership assignments for code sub-trees to maintain consistency and correctness
- code transparency and code contributions across teams

Figure 18.2 shows how code review lead times are affected by the size of the change. On the x-axis is the size of the change and on the y-axis is the lead time required for the code review process. In general, the

larger the change submitted for code reviews, the longer the lead time required to get the necessary sign-offs. And the data points in the upper-left corner represent the more complex and potentially risky changes that required more deliberation and discussion.

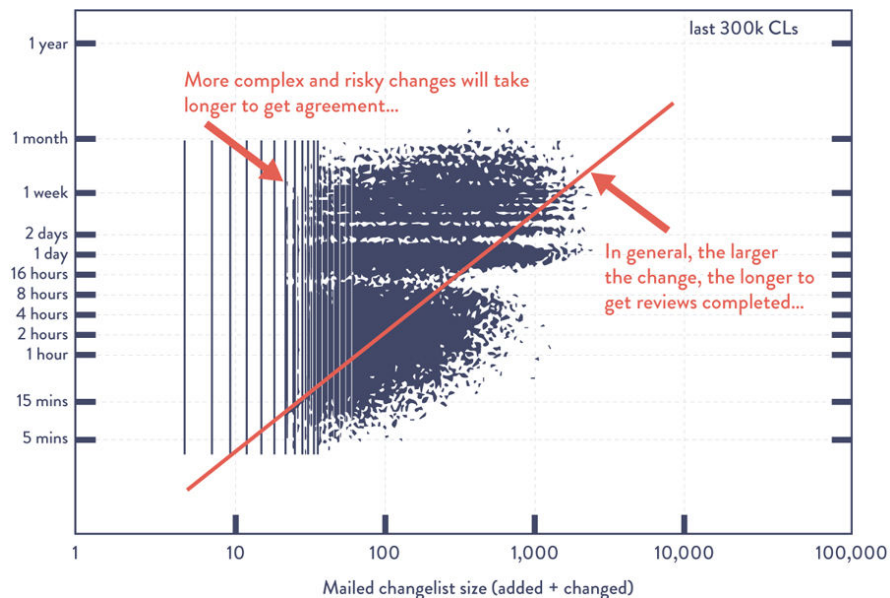


Figure 18.2: Size of Change vs. Lead Time for Reviews at Google

Source: Ashish Kumar, “Development at the Speed and Scale of Google,” presentation at QCon, San Francisco, CA, 2010, qconsf.com/sf2010/dl/qcon-sanfran-2010/slides/AshishKumar_DevelopingProductsattheSpeedandScaleofGoogle.pdf.

While he was working as a Google engineering director, Randy Shoup started a personal project to solve a technical problem that the organization was facing. He said,

I worked on that project for weeks and finally got around to asking a subject matter expert to review my code. It was nearly three thousand lines of code, which took the reviewer days of work to go through. He told me, “Please don’t do that to me again.” I was grateful that this engineer took the time to

do that. That was also when I learned how to make code reviews a part of my daily work.²⁴

Google is an excellent example of a company that employs trunk-based development and continuous delivery at scale through the use of code reviews.

Potential Dangers of Doing More Manual Testing and Change Freezes

Now that we have created peer reviews that reduce our risk, shortened lead times associated with change approval processes, and enabled continuous delivery at scale, such as we saw in the Google case study, let us examine the effects of how testing countermeasures can sometimes backfire. When testing failures occur, our typical reaction is to do more testing. However, if we are merely performing more testing at the end of the project, we may worsen our outcomes.

This is especially true if we are doing manual testing, because manual testing is naturally slower and more tedious than automated testing and performing “additional testing” often has the consequence of taking significantly longer to test, which means we are deploying less frequently, thus increasing our deployment batch size. And we know from both theory and practice that when we increase our deployment batch size, our change success rates go down and our incident counts and MTTR go up—the opposite of the outcomes we want.

Instead of performing testing on large batches of changes that are scheduled around change freeze periods, we want to fully integrate testing into our daily work as part of the smooth and continual flow into production and increase our deployment frequency. By doing this, we build in quality,

which allows us to test, deploy, and release in ever smaller batch sizes.

Enable Pair Programming to Improve All Our Changes

Pair programming is when two engineers work together at the same workstation, a method popularized by Extreme Programming and Agile in the early 2000s. As with code reviews, this practice started in Development but is equally applicable to the work that any engineer does in our value stream.⁵

In one common pattern of pairing, one engineer fills the role of the *driver*, the person who actually writes the code, while the other engineer acts as the *navigator*, *observer*, or *pointer*, the person who reviews the work as it is being performed. While reviewing, the observer may also consider the strategic direction of the work, coming up with ideas for improvements and likely future problems to address. This frees the driver to focus all of their attention on the tactical aspects of completing the task, using the observer as a safety net and guide. When the two have differing specialties, skills are transferred as an automatic side effect, whether it's through ad hoc training or by sharing techniques and workarounds.

Another pair programming pattern reinforces test-driven development (TDD) by having one engineer write the automated test and the other engineer implement the code.

Jeff Atwood, one of the founders of Stack Exchange, wrote, "I can't help wondering if pair programming is nothing more than code review on steroids... . The advantage of pair programming is its gripping immediacy: it is impossible to ignore the reviewer when he or she is sitting right next to you."²⁵

He continued, “Most people will passively opt out [of reviewing code] if given the choice. With pair programming, that’s not possible. Each half of the pair *has* to understand the code, right then and there, as it’s being written. Pairing may be invasive, but it can also force a level of communication that you’d otherwise never achieve.”²⁶

Dr. Laurie Williams performed a study in 2001 that showed

paired programmers are 15% slower than two independent individual programmers, while “error-free” code increased from 70% to 85%. Since testing and debugging are often many times more costly than initial programming, this is an impressive result. Pairs typically consider more design alternatives than programmers working alone and arrive at simpler, more maintainable designs; they also catch design defects early.²⁷

Dr. Williams also reported that 96% of her respondents stated that they enjoyed their work more when they programmed in pairs than when they programmed alone.^{28¶}

Pair programming has the additional benefit of spreading knowledge throughout the organization and increasing information flow within the team. Having more experienced engineers review while the less experienced engineer codes is also an effective way to teach and be taught.

CASE STUDY

Pair Programming Replacing Broken Code Review Processes at Pivotal Labs (2011)

Elisabeth Hendrickson, VP of Engineering at Pivotal Software, Inc. and author of *Explore It!: Reduce Risk and Increase Confidence with Exploratory Testing*, has spoken extensively about making every team responsible for their own quality, as opposed to making separate departments responsible. She argues that doing so not only increases quality, but significantly increases the flow of work into production.²⁹

In her 2015 DevOps Enterprise Summit presentation, she described how in 2011, there were two accepted methods of code review at Pivotal: pair programming (which ensured that every line of code was inspected by two people) or a code review process that was managed by Gerrit (which ensured that every code commit had two designated people “+1” the change before it was allowed into trunk).³⁰

The problem Hendrickson observed with the Gerrit code review process was that it would often take an entire week for developers to receive their required reviews. Worse, skilled developers were experiencing the “frustrating and soul-crushing experience of not being able to get simple changes into the codebase because we had inadvertently created intolerable bottlenecks.”³¹

Hendrickson lamented that

the only people who had the ability to “+1” the changes were senior engineers, who had many other responsibilities and often didn’t care as much about the fixes the more junior developers were working on or their productivity. It created a terrible situation—while you were waiting for your changes to

get reviewed, other developers were checking in their changes. So for a week, you would have to merge all their code changes onto your laptop, re-run all the tests to ensure that everything still worked, and (sometimes) you'd have to resubmit your changes for review again!³²

To fix the problem and eliminate all of these delays, they ended up dismantling the entire Gerrit code review process, instead requiring pair programming to implement code changes into the system. By doing this, they reduced the amount of time required to get code reviewed from weeks to hours.

Hendrickson is quick to note that code reviews work fine in many organizations, but it requires a culture that values reviewing code as highly as it values writing the code in the first place.³³

This case study illustrates just one form of code review. Particularly when the culture is not yet in place, pair programming can serve as a valuable interim practice.

Evaluating the Effectiveness of Pull Request Processes

Because the peer review process is an important part of our control environment, we need to be able to determine whether it is working effectively or not. One method is to look at production outages and examine the peer review process for any relevant changes.

Another method comes from Ryan Tomayko, CIO and co-founder of GitHub and one of the inventors of the pull request process. When asked to describe the difference between a bad pull request and a good pull request, he said it

has little to do with the production outcome. Instead, a bad pull request is one that doesn't have enough context for the reader, having little or no documentation of what the change is intended to do—for example, a pull request that merely has the following text: “Fixing issue #3616 and #3841.”^{34**}

That was an actual internal GitHub pull request, which Tomayko critiqued,

This was probably written by a new engineer here. First off, no specific engineers were specifically @mentioned—at a minimum, the engineer should have mentioned their mentor or a subject matter expert in the area that they're modifying to ensure that someone appropriate reviews their change. Worse, there isn't any explanation of what the changes actually are, why it's important, or exposing any of the implementer's thinking.³⁵

On the other hand, when asked to describe a great pull request that indicates an effective review process, Tomayko quickly listed off the essential elements: there must be sufficient detail on why the change is being made, how the change was made, as well as any identified risks and resulting countermeasures.³⁶

Tomayko also looks for good discussion of the change, enabled by all the context that the pull request provided—often, there will be additional risks pointed out, ideas on better ways to implement the desired change, ideas on how to better mitigate the risk, and so forth. And if something bad or unexpected happens upon deployment, it is added to the pull request, with a link to the corresponding issue. All discussion happens without placing blame; instead, there is a candid conversation on how to prevent the problem from recurring in the future.

As an example, Tomayko produced another internal GitHub pull request for a database migration. It was many pages long, with lengthy discussions about the potential risks, leading up to the following statement by the pull request author: “I am pushing this now. Builds are now failing for the branch, because of a missing column in the CI servers. (Link to Post-Mortem: MySQL outage).”³⁷

The change submitter then apologized for the outage, describing what conditions and mistaken assumptions led to the accident, as well as a list of proposed countermeasures to prevent it from happening again. This was followed by pages and pages of discussion. Reading through the pull request, Tomayko smiled, “Now *that* is a great pull request.”³⁸

As described above, we can evaluate the effectiveness of our peer review process by sampling and examining pull requests, either from the entire population of pull requests or those that are relevant to production incidents.

Fearlessly Cut Bureaucratic Processes

So far, we have discussed peer review and pair programming processes that enable us to increase the quality of our work without relying on external approvals for changes. However, many companies still have long-standing processes for approval that require months to navigate. These approval processes can significantly increase lead times, not only preventing us from delivering value quickly to customers, but potentially increasing the risk to our organizational objectives. When this happens, we must re-engineer our processes so that we can achieve our goals more quickly and safely.

As Adrian Cockcroft observed, “A great metric to publish widely is how many meetings and work tickets are mandatory

to perform a release—the goal is to relentlessly reduce the effort required for engineers to perform work and deliver it to the customer.”³⁹

Similarly, Dr. Tapabrata Pal, technical fellow at Capital One, described a program at Capital One called Got Goo?, which involves a dedicated team removing obstacles—including tools, processes, and approvals—that impede work completion.⁴⁰

Jason Cox, Senior Director of Systems Engineering at Disney, described in his presentation at the DevOps Enterprise Summit in 2015 a program called Join the Rebellion that aimed to remove toil and obstacles from daily work.⁴¹

At Target in 2012, a combination of the Technology Enterprise Adoption Process and Lead Architecture Review Board (TEAP-LARB process) resulted in complicated, long approval times for anyone attempting to bring in new technology. The TEAP form needed to be filled out by anyone wanting to propose new technologies to be adopted, such as a new database or monitoring technologies. These proposals were evaluated, and those deemed appropriate were put onto the monthly LARB meeting agenda.⁴²

Heather Mickman and Ross Clanton, Director of Development and Director of Operations at Target, Inc., respectively, were helping to lead the DevOps movement at Target. During their DevOps initiative, Mickman had identified a technology needed to enable an initiative from the lines of business (in this case, Tomcat and Cassandra). The decision from the LARB was that Operations could not support it at the time. However, because Mickman was so convinced that this technology was essential, she proposed that her Development team be responsible for service

support as well as integration, availability, and security, instead of relying on the Operations team.⁴³ Mickman observed,

As we went through the process, I wanted to better understand why the TEAP-LARB process took so long to get through, and I used the technique of “the five why’s”... . Which eventually led to the question of why TEAP-LARB existed in the first place. The surprising thing was that no one knew, outside of a vague notion that we needed some sort of governance process. Many knew that there had been some sort of disaster that could never happen again years ago, but no one could remember exactly what that disaster was, either.⁴⁴

Mickman concluded that this process was not necessary for her group if they were responsible for the operational responsibilities of the technology she was introducing. She added, “I let everyone know that any future technologies that we would support wouldn’t have to go through the TEAP-LARB process, either.”⁴⁵

The outcome was that Cassandra was successfully introduced inside Target and eventually widely adopted. Furthermore, the TEAP-LARB process was eventually dismantled. Out of appreciation, her team awarded Mickman the Lifetime Achievement Award for removing barriers to get technology work done within Target.⁴⁶

Conclusion

In this chapter, we discussed how to integrate practices into our daily work that increase the quality of our changes and reduce the risk of poor deployment outcomes, reducing our reliance on approval processes. Case studies from GitHub and Target show that these practices not only improve our

outcomes but also significantly reduce lead times and increase developer productivity. To do this kind of work requires a high-trust culture.

Consider a story that John Allspaw told about a newly hired junior engineer: The engineer asked if it was okay to deploy a small HTML change, and Allspaw responded, “I don’t know, is it?” He then asked “Did you have someone review your change? Do you know who the best person to ask is for changes of this type? Did you do everything you absolutely could to assure yourself that this change operates in production as designed? If you did, then don’t ask me—just make the change!”⁴⁷

By responding this way, Allspaw reminded the engineer that she was solely responsible for the quality of her change—if she did everything she felt she could to give herself confidence that the change would work, then she didn’t need to ask anyone for approval, she should make the change.

Creating the conditions that enable change implementers to fully own the quality of their changes is an essential part of the high-trust, generative culture we are striving to build. Furthermore, these conditions enable us to create an ever-safer system of work, where we are all helping each other achieve our goals, spanning whatever boundaries necessary to get there.

* *Counterfactual thinking* is a term used in psychology that involves the human tendency to create possible alternatives to life events that have already occurred. In reliability engineering, it often involves narratives of the “system as imagined” as opposed to the “system in reality.”

† In this book, the terms *code review* and *change review* will be used interchangeably.

‡ Incidentally, a list of high-risk areas of code and environments has likely already been created by the change advisory board.

- § In this book, we will use the term *pairing* and *pair programming* interchangeably to indicate that the practice is not just for developers.
- ¶ Some organizations may require pair programming, while in others, engineers find someone to pair program with when working in areas where they want more scrutiny (such as before checking in) or for challenging tasks. Another common practice is to set *pairing hours* for a subset of the working day, perhaps four hours from mid-morning to mid-afternoon.
- ** Gene Kim is grateful to Shawn Davenport, James Fryman, Will Farr, and Ryan Tomayko at GitHub for discussing the differences between good and bad pull requests.

PART IV CONCLUSION

Part IV has shown us that by implementing feedback loops we can enable everyone to work together toward shared goals, see problems as they occur, and, with quick detection and recovery, ensure that features not only operate as designed in production, but also achieve organizational goals and drive organizational learning. We have also examined how to enable shared goals spanning Dev and Ops so that they can improve the health of the entire value stream.

We are now ready to enter [Part V: The Third Way, The Technical Practices of Learning](#), so we can create opportunities for learning that happen earlier and ever more quickly and cheaply, and so that we can unleash a culture of innovation and experimentation that enables everyone to do meaningful work that helps our organization succeed.

Additional Resources

Two white papers that can help you learn more about feedback and metrics are *Measuring Software Quality and Measure Efficiency, Effectiveness, and Culture to Optimize DevOps Transformation: Metrics for DevOps Initiatives* (itrevolution.com).

Mark Schwartz, former CIO of the US Center for Immigration Services, has expert advice on how to cut bloat and make bureaucracy lean, learning, and enabling in *The (Delicate) Art of Bureaucracy: Digital Transformation with the Monkey, the Razor, and the Sumo Wrestler*.

Elisabeth Hendrickson's 2015 DevOps Enterprise Summit presentation "It's All About Feedback" has a wealth of information and is always worth a watch (videolibrary.doesvirtual.com/?video=524439999).

You can also learn even more from Elisabeth Hendrickson on feedback in her interviews with Gene Kim on *The Idealcast* podcast. (itrevolution.com/the-idealcast-podcast/.)

Rachel Potvin and Josh Levenberg write more extensively on Google's practices in the paper "Why Google Stores Billions of Lines of Code in a Single Repository" (<https://cacm.acm.org/magazines/2016/7/204032-why-google-stores-billions-of-lines-of-code-in-a-single-repository/fulltext>).



Part V

THE THIRD WAY: THE TECHNICAL
PRACTICES OF CONTINUAL
LEARNING AND EXPERIMENTATION

PART V: INTRODUCTION

In [Part III: The First Way: The Technical Practices of Flow](#), we discussed implementing the practices required to create fast flow in our value stream. In [Part IV: The Second Way: The Technical Practices of Feedback](#), our goal was to create as much feedback as possible, from as many areas in our system as possible—sooner, faster, and cheaper.

In [Part V: The Third Way: The Technical Practices of Continual Learning and Experimentation](#), we present the practices that create opportunities for learning as quickly, frequently, cheaply, and soon as possible. This includes creating learnings from accidents and failures, which are inevitable when we work within complex systems, as well as organizing and designing our systems of work so that we are constantly experimenting and learning, continually making our systems safer. The results include higher resilience and an ever-growing collective knowledge of how our system actually works, so that we are better able to achieve our goals.

In the following chapters, we will institutionalize rituals that increase safety, continuous improvement, and learning by doing the following:

- establish a just culture to make safety possible
- inject production failures to create resilience
- convert local discoveries into global improvements
- reserve time to create organizational improvements and learning

We will also create mechanisms so that any new learnings generated in one area of the organization can be rapidly used across the entire organization, turning local improvements into global advancements. In this way, we not only learn

faster than our competition, helping us win in the marketplace, but we also create a safer, more resilient work culture that people are excited to be a part of and that helps them achieve their highest potential.

19

ENABLE AND INJECT LEARNING INTO DAILY WORK

When we work within a complex system, it is impossible for us to predict all the outcomes for the actions we take. This contributes to unexpected and sometimes catastrophic accidents, even when we use static precautionary tools, such as checklists and runbooks, which codify our current understanding of the system.

To enable us to safely work within complex systems, our organizations must become ever-better at self-diagnostics and self-improvement and must be skilled at detecting problems, solving them, and multiplying the effects by making the solutions available throughout the organization. This creates a dynamic system of learning that allows us to understand our mistakes and translate that understanding into actions that helps to prevent those mistakes from recurring in the future.

The result is what Dr. Steven Spear describes as resilient organizations who are “skilled at detecting problems, solving them, and multiplying the effect by making the solutions available throughout the organization.”¹ These organizations can heal themselves. “For such an organization, responding to crises is not idiosyncratic work. It is something that is done all the time. It is this responsiveness that is their source of reliability.”²

AWS US-East and Netflix (2011)

A striking example of the incredible resilience that can result from these principles and practices was seen on April 21,

2011, when an entire availability zone in Amazon Web Services' (AWS) US-East region went down, taking down virtually all of the organizations who depended on it, including Reddit and Quora.^{3*} However, Netflix was a surprising exception, seemingly unaffected by this massive AWS outage.

Following the event, there was considerable speculation about how Netflix kept their services running. A popular theory ran that since Netflix was one of the largest customers of Amazon Web Services, it was given some special treatment that allowed them to keep running. However, a *Netflix Engineering* blog post explained that their architectural design decisions in 2009 enabled their exceptional resilience.

Back in 2008, Netflix's online video delivery service ran on a monolithic J2EE application hosted in one of their data centers. However, starting in 2009, they began re-architecting this system to be what they called *cloud native*—it was designed to run entirely in the Amazon public cloud and to be resilient enough to survive significant failures.⁵

One of their specific design objectives was to ensure Netflix services kept running, even if an entire AWS availability zone went down, such as what happened with US-East. To do this required that their system be loosely coupled, with each component having aggressive timeouts and circuit breakers[†] to ensure that failing components didn't bring the entire system down.

Instead, each feature and component was designed to gracefully degrade. For example, during traffic surges that created CPU-usage spikes, instead of showing a list of movies personalized to the user, they would show static content, such as cached or un-personalized results, which required less computation.⁶

Furthermore, as the blog post explained, in addition to implementing these architectural patterns, they also built and had been running a surprising and audacious service called *Chaos Monkey*, which simulated AWS failures by constantly and randomly killing production servers. They did so because they wanted all “engineering teams to be used to a constant level of failure in the cloud” so that services could “automatically recover without any manual intervention.”⁷

In other words, the Netflix team ran Chaos Monkey to gain assurance that they had achieved their operational resilience objectives, constantly injecting failures into their pre-production and production environments.

As one might expect, when they first ran Chaos Monkey in their production environments, services failed in ways they never could have predicted or imagined. By constantly finding and fixing these issues during normal working hours, Netflix engineers quickly and iteratively created a more resilient service, while simultaneously creating organizational learnings (during normal working hours!) that enabled them to evolve their systems far beyond their competition.

Chaos Monkey is just one example of how learning can be integrated into daily work. The story also shows how learning organizations think about failures, accidents, and mistakes—as an opportunity for learning and not something to be punished. This chapter explores how to create a system of learning and how to establish a *just culture*, as well as how to routinely rehearse and deliberately create failures to accelerate learning.

Establish a Just, Learning Culture

One of the prerequisites for a learning culture is that when accidents occur (which they undoubtedly will), the response to those accidents is seen as “just.” Dr. Sidney Dekker, who

helped codify some of the key elements of safety culture and coined the term *just culture*, writes, “When responses to incidents and accidents are seen as unjust, it can impede safety investigations, promoting fear rather than mindfulness in people who do safety-critical work, making organizations more bureaucratic rather than more careful, and cultivating professional secrecy, evasion, and self-protection.”⁸

This notion of punishment is present, either subtly or prominently, in the way many managers have operated during the last century. The thinking goes, in order to achieve the goals of the organization, leaders must command, control, establish procedures to eliminate errors, and enforce compliance of those procedures.

Dr. Dekker calls this notion of eliminating error by eliminating the people who caused the errors the *bad apple theory*. He asserts that this is invalid, because “human error is not our cause of troubles; instead, human error is a consequence of the design of the tools that we gave them.”⁹

If accidents are not caused by “bad apples” but rather are due to inevitable design problems in the complex system that we created, then instead of “naming, blaming, and shaming” the person who caused the failure, our goal should always be to maximize opportunities for organizational learning, continually reinforcing that we value actions that expose and share more widely the problems in our daily work. This is what enables us to improve the quality and safety of the system we operate within and reinforce the relationships between everyone who operates within that system.

By turning information into knowledge and building the results of the learning into our systems, we start to achieve the goals of a just culture, balancing the needs for safety and

accountability. As John Allspaw, CTO of Etsy, states, “Our goal at Etsy is to view mistakes, errors, slips, lapses, and so forth with a perspective of learning.”¹⁰

When engineers make mistakes and feel safe when giving details about it, they are not only willing to be held accountable, but also enthusiastic in helping the rest of the company avoid the same error in the future. This is what creates organizational learning. On the other hand, if we punish that engineer, everyone is dis-incentivized to provide the necessary details to get an understanding of the mechanism, pathology, and operation of the failure, which guarantees that the failure will occur again.

Two effective practices that help create a just, learning-based culture are blameless post-mortems (also called retrospectives or learning reviews) and the controlled introduction of failures into production to create opportunities to practice for the inevitable problems that arise within complex systems. We will first look at a retrospective and follow that with an exploration of why failure can be a good thing.

Schedule Retrospective Meetings after Accidents Occur

To help enable a just culture, when accidents and significant incidents occur (e.g., failed deployment, production issue that affected customers), we should conduct a retrospective after the incident has been resolved. † Retrospectives help us examine “mistakes in a way that focuses on the situational aspects of a failure’s mechanism and the decision-making process of individuals proximate to the failure.”¹¹

To do this, we conduct the retrospective as soon as possible after the accident occurs and before memories and

the links between cause and effect fade or circumstances change. (Of course, we wait until after the problem has been resolved so as not to distract the people who are still actively working on the issue.)

In the meeting, we will do the following:

- Construct a timeline and gather details from multiple perspectives on failures, ensuring we don't punish people for making mistakes.
- Empower all engineers to improve safety by allowing them to give detailed accounts of their contributions to failures.
- Enable and encourage people who do make mistakes to be the experts who educate the rest of the organization on how not to make them in the future.
- Accept that there is always a discretionary space where humans can decide to take action or not, and that the judgment of those decisions lies in hindsight.
- Propose countermeasures to prevent a similar accident from happening in the future and ensure these countermeasures are recorded with a target date and an owner for follow-up.

To enable us to gain this understanding, the following stakeholders need to be present at the meeting:

- the people involved in decisions that may have contributed to the problem
- the people who identified the problem
- the people who responded to the problem
- the people who diagnosed the problem
- the people who were affected by the problem

- anyone else who is interested in attending the meeting

Our first task in the retrospective is to record our best understanding of the timeline of relevant events as they occurred. This includes all actions we took and at what time (ideally supported by chat logs, such as IRC or Slack), what effects we observed (ideally in the form of the specific metrics from our production telemetry, as opposed to merely subjective narratives), all investigation paths we followed, and what resolutions were considered.

To enable these outcomes, we must be rigorous about recording details and reinforcing a culture within which information can be shared without fear of punishment or retribution. Because of this, especially for our first few retrospectives, it may be helpful to have the meeting led by a trained facilitator who wasn't involved in the accident.

During the meeting and the subsequent resolution, we should explicitly disallow the phrases “would have” or “could have,” as they are *counterfactual* statements that result from our human tendency to create possible alternatives to events that have already occurred.

Counterfactual statements, such as “I could have ... ” or “If I had known about that, I should have ... ,” frame the problem in terms of the *system as imagined* instead of in terms of the *system that actually exists*, which is the context we need to restrict ourselves to. (See [Appendix 8](#).)

One of the potentially surprising outcomes of these meetings is that people will often blame themselves for things outside of their control or question their own abilities. Ian Malpass, an engineer at Etsy observes,

In that moment when we do something that causes the entire site to go down, we get this “ice-water down the

spine” feeling, and likely the first thought through our head is, “I suck and I have no idea what I’m doing.” We need to stop ourselves from doing that, as it is route to madness, despair, and feelings of being an imposter, which is something that we can’t let happen to good engineers. The better question to focus on is, “Why did it make sense to me when I took that action?”¹²

In the meeting, we must reserve enough time for brainstorming and deciding which countermeasures to implement. Once the countermeasures have been identified, they must be prioritized and given an owner and timeline for implementation. Doing this further demonstrates that we value improvement of our daily work more than daily work itself.

Dan Milstein, one of the principal engineers at Hubspot, writes that he begins all retrospectives by saying, “We’re trying to prepare for a future where we’re as stupid as we are today.”¹³ In other words, it is not acceptable to have a countermeasure to merely “be more careful” or “be less stupid”—instead, we must design real countermeasures to prevent these errors from happening again.

Examples of such countermeasures include new automated tests to detect dangerous conditions in our deployment pipeline, adding further production telemetry, identifying categories of changes that require additional peer review, and conducting rehearsals of this category of failure as part of regularly scheduled game day exercises.

Publish Our Retrospective Reviews as Widely as Possible

After we conduct a retrospective, we should widely announce the availability of the meeting notes and any associated

artifacts (e.g., timelines, IRC chat logs, external communications). This information should (ideally) be placed in a centralized location where our entire organization can access it and learn from the incident. Conducting retrospectives is so important that we may even prohibit production incidents from being closed until the retrospective has been completed.

Doing this helps us translate local learnings and improvements into global learnings and improvements. Randy Shoup, former engineering director for Google App Engine, describes how documentation of retrospectives can have tremendous value to others in the organization: “As you can imagine at Google, everything is searchable. All the retrospective documents are in places where other Googlers can see them. And trust me, when any group has an incident that sounds similar to something that happened before, these retrospective documents are among the first documents being read and studied.”^{14§}

Widely publishing retrospectives and encouraging others in the organization to read them increases organizational learning. It is also becoming increasingly commonplace for online service companies to publish retrospectives for customer-impacting outages. This often significantly increases the transparency we have with our internal and external customers, which in turn increases their trust in us.

This desire to conduct as many retrospective meetings as necessary at Etsy led to some problems—over the course of four years, Etsy accumulated a large number of post-mortem meeting notes (retrospectives) in wiki pages, which became increasingly difficult to search, save, and collaborate from. To help with this issue, they developed a tool called Morgue to easily record aspects of each accident, such as the incident MTTR and severity, better address time zones (which became

relevant as more Etsy employees were working remotely), and other data, such as rich text in Markdown format, embedded images, tags, and history.¹⁶

Morgue was designed to make it easy for the team to record:

- whether the problem was due to a scheduled or unscheduled incident
- the retrospective owner
- relevant IRC chat logs (especially important for 3 AM issues when accurate note-taking may not happen)
- relevant JIRA tickets for corrective actions and their due dates (information particularly important to management)
- links to customer forum posts (where customers complain about issues)

After developing and using Morgue, the number of recorded retrospectives at Etsy increased significantly compared to when they used wiki pages, especially for P2, P3, and P4 incidents (i.e., lower severity problems).¹⁷ This result reinforced the hypothesis that if they made it easier to document retrospectives through tools such as Morgue, more people would record and detail the outcomes of their retrospective, enabling more organizational learning.

CONTINUOUS LEARNING

Conducting retrospectives does more than help us learn from failure. DORA's *2018 State of DevOps Report* found that they contribute to our culture, helping teams feel better about sharing information, taking smart risks, and understanding the value of learning. In addition, the research found that elite performers were 1.5 times more likely to consistently hold retrospectives and use them to improve their work—so these elite teams continue to reap the benefits.¹⁸

Dr. Amy C. Edmondson, Novartis Professor of Leadership and Management at Harvard Business School and co-author of *Building the Future: Big Teaming for Audacious Innovation*, writes:

Again, the remedy—which does not necessarily involve much time and expense—is to reduce the stigma of failure. Eli Lilly has done this since the early 1990s by holding “failure parties” to honor intelligent, high-quality scientific experiments that fail to achieve the desired results. The parties don't cost much, and redeploying valuable resources—particularly scientists—to new projects earlier rather than later can save hundreds of thousands of dollars, not to mention kickstart potential new discoveries.¹⁹

Decrease Incident Tolerances to Find Ever-Weaker Failure Signals

Inevitably, as organizations learn how to see and solve problems efficiently, they need to decrease the threshold of

what constitutes a problem in order to keep learning. To do this, we seek to amplify weak failure signals. For example, as described in [Chapter 4](#), when Alcoa was able to reduce the rate of workplace accidents so that they were no longer commonplace, Paul O’Neill, CEO of Alcoa, started to be notified of accident near-misses in addition to actual workplace accidents.²⁰

Dr. Spear summarizes O’Neill’s accomplishments at Alcoa when he writes, “Though it started by focusing on problems related to workplace safety, it soon found that safety problems reflected process ignorance and that this ignorance would also manifest itself in other problems such as quality, timeliness, and yield versus scrap.”²¹

When we work within complex systems, this need to amplify weak failure signals is critical to averting catastrophic failures. The way NASA handled failure signals during the space shuttle era serves as an illustrative example. In 2003, sixteen days into the *Columbia* space shuttle mission, it exploded as it re-entered the earth’s atmosphere. We now know that a piece of insulating foam had broken off the external fuel tank during takeoff.

However, prior to *Columbia’s* re-entry, a handful of mid-level NASA engineers had reported this incident, but their voices had gone unheard. They observed the foam strike on video monitors during a post-launch review session and immediately notified NASA’s managers, but they were told that the foam issue was nothing new. Foam dislodgement had damaged shuttles in previous launches but had never resulted in an accident. It was considered a maintenance problem and not acted upon until it was too late.²²

Michael Roberto, Richard M. J. Bohmer, and Amy C. Edmondson wrote in a 2006 article for *Harvard Business*

Review how NASA culture contributed to this problem. They described how organizations are typically structured in one of two models: a *standardized model*, where routine and systems govern everything, including strict compliance with timelines and budgets, or an *experimental model*, where every day, every exercise and every piece of new information is evaluated and debated in a culture that resembles a research and design (R&D) laboratory.²³

They observe, “Firms get into trouble when they apply the wrong mind-set to an organization [which dictates how they respond to *ambiguous threats* or, in the terminology of this book, *weak failure signals*]... . By the 1970s, NASA had created a culture of rigid standardization, promoting to Congress the space shuttle as a cheap and reusable spacecraft.”²⁴

NASA favored strict process compliance instead of an experimental model where every piece of information needed to be evaluated as it occurred without bias. The absence of continuous learning and experimentation had dire consequences. The authors conclude that it is culture and mindset that matters, not just “being careful”—as they write, “vigilance alone will not prevent ambiguous threats [*weak failure signals*] from turning into costly (and sometimes tragic) failures.”²⁵

Our work in the technology value stream, like space travel, should be approached as a fundamentally experimental endeavor and managed that way. All work we do is a potentially important hypothesis and a source of data, rather than a routine application and validation of past practice. Instead of treating technology work as entirely standardized, where we strive for process compliance, we must continually seek to find ever-weaker failure signals so

that we can better understand and manage the system we operate in.

Redefine Failure and Encourage Calculated Risk-Taking

Leaders of an organization, whether deliberately or inadvertently, reinforce the organizational culture and values through their actions. Audit, accounting, and ethics experts have long observed that the “tone at the top” predicts the likelihood of fraud and other unethical practices. To reinforce our culture of learning and calculated risk-taking, we need leaders to continually reinforce that everyone should feel both comfortable with and responsible for surfacing and learning from failures.

On failures, Roy Rapoport from Netflix observes,

What the *2014 State of DevOps Report* proved to me is that high performing DevOps organizations will fail and make mistakes more often. Not only is this okay, it's what organizations need! You can even see it in the data: if high performers are performing thirty times more frequently but with only half the change failure rate, they're obviously having more failures...²⁶

I was talking with a co-worker about a massive outage we just had at Netflix—it was caused by, frankly, a dumb mistake. In fact, it was caused by an engineer who had taken down Netflix twice in the last eighteen months. But, of course, this is a person we'd never fire. In that same eighteen months, this engineer moved the state of our operations and automation forward not by miles but by light-years. That work has enabled us to do deployments safely on a daily basis, and has personally performed huge numbers of production deployments.²⁷

He concludes, “DevOps must allow this sort of innovation and the resulting risks of people making mistakes. Yes, you’ll have more failures in production. But that’s a good thing and should not be punished.”²⁸

Inject Production Failures to Enable Resilience and Learning

As we saw in the chapter introduction, injecting faults into the production environment (such as using Chaos Monkey) is one way we can increase our resilience. In this section, we describe the processes involved in rehearsing and injecting failures into our system to confirm that we have designed and architected our systems properly, so that failures happen in specific and controlled ways. We do this by regularly (or even continuously) performing tests to make certain that our systems fail gracefully.

As Michael Nygard, author of *Release It! Design and Deploy Production-Ready Software*, comments, “Like building crumple zones into cars to absorb impacts and keep passengers safe, you can decide what features of the system are indispensable and build in failure modes that keep cracks away from those features. If you do not design your failure modes, then you will get whatever unpredictable—and usually dangerous—ones happen to emerge.”²⁹

Resilience requires that we first define our failure modes and then perform testing to ensure that these failure modes operate as designed. One way we do this is by injecting faults into our production environment and rehearsing large-scale failures so we are confident we can recover from accidents when they occur, ideally without even impacting our customers.

The 2012 story about Netflix and the Amazon AWS US-East outage presented at the beginning of this chapter is just one example. An even more interesting example of resilience at Netflix was during the “Great Amazon Reboot of 2014,” when nearly 10% of the entire Amazon EC2 server fleet had to be rebooted to apply an emergency Xen security patch.³⁰

As Christos Kalantzis of Netflix Cloud Database Engineering recalled, “When we got the news about the emergency EC2 reboots, our jaws dropped. When we got the list of how many Cassandra nodes would be affected, I felt ill.”³¹ But, Kalantzis continues, “Then I remembered all the Chaos Monkey exercises we’ve gone through. My reaction was, ‘Bring it on!’”³²

Once again, the outcomes were astonishing. Of the 2,700+ Cassandra nodes used in production, 218 were rebooted, and twenty-two didn’t reboot successfully. As Kalantzis and Bruce Wong from Netflix Chaos Engineering wrote, “Netflix experienced 0 downtime that weekend. Repeatedly and regularly exercising failure, even in the persistence [database] layer, should be part of every company’s resilience planning. If it wasn’t for Cassandra’s participation in Chaos Monkey, this story would have ended much differently.”³³

Even more surprising, not only was no one at Netflix working active incidents due to failed Cassandra nodes, no one was even in the office—they were in Hollywood at a party celebrating an acquisition milestone.³⁴ This is another example demonstrating that proactively focusing on resilience often means that a firm can handle events that may cause crises for most organizations in a manner that is routine and mundane.[¶] (See [Appendix 9](#).)

Institute Game Days to Rehearse Failures

In this section, we describe specific disaster recovery rehearsals called game days, a term popularized by Jesse Robbins, one of the founders of the Velocity Conference community and co-founder of Chef. While at Amazon, Robbins was responsible for creating programs to ensure site availability and was widely known internally as the “Master of Disaster.”³⁶ The concept of game days comes from the discipline of resilience engineering. Robbins defines *resilience engineering* as “an exercise designed to increase resilience through large-scale fault injection across critical systems.”³⁷

Robbins observes that “whenever you set out to engineer a system at scale, the best you can hope for is to build a reliable software platform on top of components that are completely unreliable. That puts you in an environment where complex failures are both inevitable and unpredictable.”³⁸

Consequently, we must ensure that services continue to operate when failures occur, potentially throughout our system, ideally without crisis or even manual intervention. As Robbins quips, “a service is not really tested until we break it in production.”³⁹

Our goal for game day is to help teams simulate and rehearse accidents to give them the ability to practice. First, we schedule a catastrophic event, such as the simulated destruction of an entire data center, to happen at some point in the future. We then give teams time to prepare, to eliminate all the single points of failure, and to create the necessary monitoring procedures, failover procedures, etc.

Our game day team defines and executes drills, such as conducting database failovers (i.e., simulating a database

failure and ensuring that the secondary database works) or turning off an important network connection to expose problems in the defined processes. Any problems or difficulties that are encountered are identified, addressed, and tested again.

At the scheduled time, we then execute the outage. As Robbins describes, at Amazon they “would literally power off a facility—without notice—and then let the systems fail naturally and [allow] the people to follow their processes wherever they led.”⁴⁰

By doing this, we start to expose the *latent defects* in our system, which are the problems that appear only because of having injected faults into the system. Robbins explains, “You might discover that certain monitoring or management systems crucial to the recovery process end up getting turned off as part of the failure you’ve orchestrated. You would find some single points of failure you didn’t know about that way.”⁴¹ These exercises are then conducted in an increasingly intense and complex way with the goal of making them feel like just another part of an average day.

By executing game days, we progressively create a more resilient service and a higher degree of assurance that we can resume operations when inopportune events occur, as well as create more learnings and a more resilient organization.

An excellent example of simulating disaster is Google’s Disaster Recovery Program (DiRT). Kripa Krishnan, Technical Program Director at Google at the time, had led the program for over seven years. During that time, they had simulated an earthquake in Silicon Valley, which resulted in the entire Mountain View campus being disconnected from Google; major data centers having complete loss of power; and even aliens attacking cities where engineers resided.⁴²

As Krishnan wrote, “An often-overlooked area of testing is business process and communications. Systems and processes are highly intertwined, and separating testing of systems from testing of business processes isn’t realistic: a failure of a business system will affect the business process, and conversely a working system is not very useful without the right personnel.”⁴³

Some of the learnings gained during these disasters included:⁴⁴

- When connectivity was lost, the failover to the engineer workstations didn’t work.
- Engineers didn’t know how to access a conference call bridge or the bridge only had capacity for fifty people or they needed a new conference call provider who would allow them to kick off engineers who had subjected the entire conference to hold music.
- When the data centers ran out of diesel for the backup generators, no one knew the procedures for making emergency purchases through the supplier, resulting in someone using a personal credit card to purchase \$50,000 worth of diesel.

By creating failure in a controlled situation, we can practice and create the playbooks we need. One of the other outputs of game days is that people actually know who to call and know who to talk to—by doing this, they develop relationships with people in other departments so they can work together during an incident, turning conscious actions into unconscious actions that are able to become routine.

➔ **CASE STUDY: NEW TO
SECOND EDITION**

Turning an Outage into a Powerful Learning Opportunity at CSG (2021)

CSG is North America's largest SaaS-based customer care and billing provider, with over sixty-five million subscribers and a tech stack that covers everything from Java to mainframe. At the DevOps Enterprise Summit in 2020, Erica Morrison, Vice President of Software Engineering, shared the story of CSG's worst outage—the result of a complex system failure that pushed CSG beyond the limits of its response systems, processes, and culture.⁴⁵

But in the face of that adversity, they were able to find opportunity and use the lessons they learned to improve how they understand incidents, respond to them, and prevent them in the first place.

The 2/4 Outage, as it later came to be known, lasted thirteen hours. It started abruptly and large portions of CSG's product were unavailable. On the initial calls as the outage began, the team was troubleshooting blind as they had trouble accessing the tools they normally use, including their health monitoring system and server access. With the number of vendors and customers involved, the initial calls were particularly chaotic.

In the end, it would take several days to figure out what had actually happened by reproducing the outage in their lab. The issue started with routine server maintenance on an OS that was different than most of the servers they ran. When that server rebooted, it put an LLDP packet out on the network. Due to a bug, CSG's network software picked it up and interpreted it as a spanning tree. It broadcast it out to the network and it was then picked up by their load balancer. Due to a

misconfiguration in the load balancer, this got rebroadcast to the network, creating a network loop, taking the network down.

The aftermath was severe. The extent of angry customers required leadership to pivot their focus from their planned work (strategic initiatives, etc.) to just this outage. Throughout the company, there was also a large sense of loss and heartbreak over having failed their customers so severely. Morale was extremely low. Hurtful things were said, like “DevOps doesn’t work.”

CSG knew they wanted to respond to this failure differently. They needed to maximize learnings while also reducing the likelihood of an incident like this happening again. Their first step was incident analysis.

Their standard incident analysis was a structured process to help them understand what happened and identify opportunities for improvement. They did this by understanding the timeline of the incident; asking, “What happened? How can we detect it sooner? How can we recover sooner? What went well?” Understanding system behavior. And maintaining a blameless culture and avoiding finger pointing.

Because of this incident they also knew they needed to up their game. They reached out to Dr. Richard Cook and John Allspaw of Adaptive Capacity Labs to analyze the incident. Through two weeks of intense interviews and research, they gained a more thorough understanding of the events, and, in particular, learned the different perspectives of the people who were working on the outage.

From this intensive retrospective, they created an operation improvement program based on the Incident

Command System. They broke the program into four categories: incident response, tool reliability, data center/platform resiliency, and application reliability.

Even before the whole organization had been through the training for their new incident management process, people started to see observable improvements in how outage calls were run: clutter on the calls had been removed; status reports had a known, steady cadence; and having an LNO (liaison officer) helped avoid interruptions on the incident calls.

The next biggest improvement was a sense of control over chaos. The simple act of having predictable cadences and patterns to follow helped everyone feel more confident and in control. It also allowed activities to run in parallel until the set time for a status update, allowing that activity to run without interruption.

Additionally, decision-making was updated from the old system by giving the incident commander clear command and authority so there's no question about who can make decisions.

Now, CSG has a stronger organizational ability to perform incident management. They've reinforced and broadened culture norms around safety and, most impactfully, implemented the incident management system that changed how they run outage calls.

In this case study, a blameless post-mortem (retrospective) led CSG to completely revamp the way they handle incidents. They directly applied their learnings regarding how they conduct their work, changing their culture and not blaming an individual or team.

Conclusion

To create a just culture that enables organizational learning, we have to re-contextualize so-called failures. When treated properly, errors that are inherent in complex systems can create a dynamic learning environment where all of the shareholders feel safe enough to come forward with ideas and observations, and where groups rebound more readily from projects that don't perform as expected.

Both retrospectives and injecting production failures reinforce a culture that everyone should feel both comfortable with and responsible for surfacing and learning from failures. In fact, when we sufficiently reduce the number of accidents, we decrease our tolerance so that we can keep learning. As Peter Senge is known to say, “The only sustainable competitive advantage is an organization's ability to learn faster than the competition.”⁴⁶

* In January 2013 at re:Invent, James Hamilton, VP and Distinguished Engineer for Amazon Web Services, said that the US-East region had more than ten data centers all by itself, and added that a typical data center has between fifty thousand and eighty thousand servers. By this math, the 2011 EC2 outage affected customers on more than half a million servers.⁴

† See Martin Fowler's article on circuit breakers for more on this: <https://martinfowler.com/bliki/CircuitBreaker.html>.

‡ This practice has also been called *blameless post-incident reviews* as well as *post-event retrospectives*. There is also a noteworthy similarity to the routine retrospectives that are a part of many iterative and Agile development practices.

§ We may also choose to extend the philosophies of *transparent uptime* to our post-mortem reports and, in addition to making a service dashboard available to the public, we may choose to publish (maybe sanitized) post-mortem meetings to the public. Some of the most widely admired public post-mortems include those posted by the Google App Engine team after a significant 2010 outage, as well as the post-mortem of the 2015 Amazon DynamoDB outage. Interestingly, Chef publishes their post-mortem meeting notes on their blog, as well as recorded videos of the actual post-mortem meetings.¹⁵

¶ Specific architectural patterns that they implemented included fail fasts (setting aggressive timeouts such that failing components don't make the entire system crawl to a halt), fallbacks (designing each feature to degrade or fall back to a lower quality representation), and feature removal (removing non-critical features when they run slowly from any given page to prevent them from impacting the member experience). Another astonishing example of the resilience that the Netflix team created beyond preserving business continuity during the AWS outage was that Netflix went over six hours into the AWS outage before declaring a Sev 1 incident, assuming that AWS service would eventually be restored (i.e., "AWS will come back ... it usually does, right?"). Only after six hours into the outage did they activate any business continuity procedures.³⁵

20

CONVERT LOCAL DISCOVERIES INTO GLOBAL IMPROVEMENTS

In the previous chapter, we discussed developing a safe learning culture by encouraging everyone to talk about mistakes and accidents through blameless post-mortems. We also explored finding and fixing ever-weaker failure signals, as well as reinforcing and rewarding experimentation and risk-taking. Furthermore, we helped make our system of work more resilient by proactively scheduling and testing failure scenarios, making our systems safer by finding latent defects and fixing them.

In this chapter, we will create mechanisms that make it possible for new learnings and improvements discovered locally to be captured and shared globally throughout the entire organization, multiplying the effect of global knowledge and improvement. By doing this, we elevate the state of the practice of the entire organization so that everyone doing work benefits from the cumulative experience of the organization.

Use Chat Rooms and Chat Bots to Automate and Capture Organizational Knowledge

Many organizations have created chat rooms to facilitate fast communication within teams. Chat rooms can also be used to trigger automation.

This technique was pioneered in the ChatOps journey at GitHub. The goal was to put automation tools into the middle of the conversation in their chat rooms, helping create transparency and documentation of their work. As Jesse

Newland, a systems engineer at GitHub, describes, “Even when you’re new to the team, you can look in the chat logs and see how everything is done. It’s as if you were pair-programming with them all the time.”¹

They created *Hubot*, a software application that interacted with the Ops team in their chat rooms, where it could be instructed to perform actions merely by sending it a command (e.g., “@hubot deploy owl to production”). The results would also be sent back into the chat room.²

Having this work performed by automation in the chat room (as opposed to running automated scripts via command line) had numerous benefits, including:

- Everyone saw everything that was happening.
- Engineers, on their first day of work, could see what daily work looked like and how it was performed.
- People were more apt to ask for help when they saw others helping each other.
- Rapid organizational learning was enabled and accumulated.

Furthermore, beyond the above tested benefits, chat rooms inherently record and make all communications public; in contrast, emails are private by default, and the information in them cannot easily be discovered or propagated within an organization.

Integrating our automation into chat rooms helps document and share our observations and problem-solving as an inherent part of performing our work. This reinforces a culture of transparency and collaboration in everything we do.

Hubot at GitHub

This is also an extremely effective way of converting local learning to global knowledge. At GitHub, all the Operations staff worked remotely—in fact, no two engineers worked in the same city. As Mark Imbriaco, former VP of Operations at GitHub, recalls, “There was no physical watercooler at GitHub. The chat room was the water cooler.”³

GitHub enabled Hubot to trigger their automation technologies, including Puppet, Capistrano, Jenkins, resque (a Redis-backed library for creating background jobs), and graphme (which generates graphs from Graphite).⁴

Actions performed through Hubot included checking the health of services, doing Puppet pushes or code deployments into production, and muting alerts as services went into maintenance mode. Actions that were performed multiple times, such as pulling up the smoke test logs when a deployment failed, taking production servers out of rotation, reverting to master for production front-end services, or even apologizing to the engineers who were on call, also became Hubot actions.^{5*}

Similarly, commits to the source code repository and the commands that trigger production deployments both emit messages to the chat room. Additionally, as changes move through the deployment pipeline, their status is posted in the chat room.

A typical, quick chat room exchange might look like:

*@sr: @jnewland, how do you get that list of big repos?
disk_hogs or something?*

@jnewland: /disk-hogs

Newland observes that certain questions that were previously asked during the course of a project are rarely

asked now.⁶ For example, engineers may ask each other, “How is that deploy going?” or “Are you deploying that or should I?” or “How does the load look?”

Among all the benefits that Newland describes, which include faster onboarding of newer engineers and making all engineers more productive, the result that he felt was most important was that Ops work became more humane as Ops engineers were enabled to discover problems and help each other quickly and easily.⁷

GitHub created an environment for collaborative local learning that could be transformed into learnings across the organization. Throughout the rest of this chapter, we will explore ways to create and accelerate the spread of new organizational learnings.

Automate Standardized Processes in Software for Reuse

All too often, we codify our standards and processes for architecture, testing, deployment, and infrastructure management in prose, storing them in Word documents that are uploaded somewhere. The problem is that engineers who are building new applications or environments often don't know that these documents exist, or they don't have the time to implement the documented standards. The result is they create their own tools and processes, with all the disappointing outcomes we'd expect: fragile, insecure, and unmaintainable applications and environments that are expensive to run, maintain, and evolve.

Instead of putting our expertise into Word documents, we need to transform these documented standards and processes, which encompass the sum of our organizational learnings and knowledge, into an executable form that makes

them easier to reuse.⁸ One of the best ways we can make this knowledge reusable is by putting it into a centralized source code repository, making the tool available for everyone to search and use.

Justin Arbuckle was chief architect at GE Capital in 2013 when he said, “We needed to create a mechanism that would allow teams to easily comply with policy—national, regional, and industry regulations across dozens of regulatory frameworks, spanning thousands of applications running on tens of thousands of servers in tens of data centers.”⁹

The mechanism they created was called ArchOps, which “enabled our engineers to be builders, not bricklayers. By putting our design standards into automated blueprints that were able to be used easily by anyone, we achieved consistency as a byproduct.”¹⁰

By encoding our manual processes into code that is automated and executed, we enable the process to be widely adopted, providing value to anyone who uses them. Arbuckle concluded that “the actual compliance of an organization is in direct proportion to the degree to which its policies are expressed as code.”¹¹

By making this automated process the easiest means to achieve the goal, we allow practices to be widely adopted—we may even consider turning them into shared services supported by the organization.

Create a Single, Shared Source Code Repository for Our Entire Organization

A firm-wide, shared source code repository is one of the most powerful mechanisms used to integrate local discoveries across the entire organization. When we update anything in

the source code repository (e.g., a shared library), it rapidly and automatically propagates to every other service that uses that library, and it is integrated through each team's deployment pipeline.

Google is one of the largest examples of using an organization-wide shared source code repository. By 2015, Google had a single shared source code repository with over one billion files and over two billion lines of code. This repository is used by every one of their twenty-five thousand engineers and spans every Google property, including Google Search, Google Maps, Google Docs, Google Calendar, Gmail, and YouTube.^{12†}

One of the valuable results of this is that engineers can leverage the diverse expertise of everyone in the organization. Rachel Potvin, a Google engineering manager overseeing the Developer Infrastructure group, told *Wired* that every Google engineer can access “a wealth of libraries” because “almost everything has already been done.”¹⁴

Furthermore, as Eran Messeri, an engineer in the Google Developer Infrastructure group, explains, one of the advantages of using a single repository is that it allows users to easily access all of the code in its most up-to-date form, without the need for coordination.¹⁵

We put into our shared source code repository not only source code but also other artifacts that encode knowledge and learning, including:

- configuration standards for our libraries, infrastructure, and environments (Chef, Puppet, or Ansible scripts)
- deployment tools
- testing standards and tools, including security

- deployment pipeline tools
- monitoring and analysis tools
- tutorials and standards

Encoding knowledge and sharing it through this repository is one of the most powerful mechanisms we have for propagating knowledge. As Randy Shoup describes,

the most powerful mechanism for preventing failures at Google is the single code repository. Whenever someone checks in anything into the repo, it results in a new build, which always uses the latest version of everything. Everything is built from source rather than dynamically linked at runtime—there is always a single version of a library that is the current one in use, which is what gets statically linked during the build process.¹⁶

Tom Limoncelli is the co-author of *The Practice of Cloud System Administration: Designing and Operating Large Distributed Systems* and a former site reliability engineer at Google. In his book, he states that the value of having a single repository for an entire organization is so powerful it is difficult to even explain.

You can write a tool exactly once and have it be usable for all projects. You have 100% accurate knowledge of who depends on a library; therefore, you can refactor it and be 100% sure of who will be affected and who needs to test for breakage. I could probably list one hundred more examples. I can't express in words how much of a competitive advantage this is for Google.¹⁷

CONTINUOUS LEARNING

Research shows that good code-related practices contribute to elite performance. Based on her expertise building systems and leading development teams at Google, Rachel Potvin was an advisor on DORA's 2019 *State of DevOps Report*, which identified **code maintainability** as a key construct in helping teams to do continuous delivery successfully. This new construct, based on the benefits that Potvin saw from the infrastructure available at Google, helps teams think about structuring their work and their code.¹⁸

According to the report,

Teams that manage code maintainability well have systems and tools that make it easy for developers to change code maintained by other teams, find examples in the codebase, reuse other people's code, as well as add, upgrade, and migrate to new versions of dependencies without breaking their code. Having these systems and tools in place not only contributes to CD, but also helps decrease technical debt, which in turn improves productivity.¹⁹

At Google, every library (e.g., libc, OpenSSL, as well as internally developed libraries, such as Java threading libraries) has an owner who is responsible for ensuring that the library not only compiles but also successfully passes the tests for all projects that depend upon it, much like a real-

world librarian. That owner is also responsible for migrating each project from one version to the next.

Consider the real-life example of an organization that runs eighty-one different versions of the Java Struts framework library in production—all but one of those versions have critical security vulnerabilities, and maintaining all those versions, each with its own quirks and idiosyncrasies, creates significant operational burden and stress. Furthermore, all this variance makes upgrading versions risky and unsafe, which in turn discourages developers from upgrading. And the cycle continues. The single source repository solves much of this problem, as well as having automated tests that allow teams to migrate to new versions safely and confidently.

If we are not able to build everything off a single source tree, then we must find another means to maintain known good versions of the libraries and their dependencies. For instance, we may have an organization-wide repository such as Nexus, Artifactory, or a Debian or RPM repository, which we must then update when there are known vulnerabilities, both in these repositories and in production systems.

It is essential to ensure that dependencies are drawn only from within the organization's source control repository or package repository in order to prevent attacks through this "software supply chain" from compromising an organization's systems.

Spread Knowledge by Using Automated Tests as Documentation and Communities of Practice

When we have shared libraries being used across the organization, we should enable rapid propagation of expertise and improvements. Ensuring that each of these libraries has significant amounts of automated testing included means

these libraries become self-documenting and show other engineers how to use them.

This benefit will be nearly automatic if we have test-driven development (TDD) practices in place, where automated tests are written before we write the code. This discipline turns our test suites into a living, up-to-date specification of the system. Any engineer wishing to understand how to use the system can look at the test suite to find working examples of how to use the system's API.

Ideally, each library will have a single owner or a single team supporting it, representing where knowledge and expertise for the library resides. Furthermore, we should (ideally) only allow one version to be used in production, ensuring that whatever is in production leverages the best collective knowledge of the organization.

In this model, the library owner is also responsible for safely migrating each group using the repository from one version to the next. This in turn requires quick detection of regression errors through comprehensive automated testing and continuous integration for all systems that use the library.

In order to more rapidly propagate knowledge, we can also create discussion groups or chat rooms for each library or service, so anyone who has questions can get responses from other users, who are often faster to respond than the developers.

By using this type of communication tool instead of having isolated pockets of expertise spread throughout the organization, we facilitate an exchange of knowledge and experience, ensuring that workers are able to help each other with problems and new patterns.

Design for Operations through Codified Non-Functional Requirements

When Development follows their work downstream and participates in production incident resolution activities, the application becomes increasingly better designed for Operations. Furthermore, as we start to deliberately design our code and application so that it can accommodate fast flow and deployability, we will likely identify a set of non-functional requirements that we will want to integrate into all of our production services.

Implementing these non-functional requirements will enable our services to be easy to deploy and keep running in production, where we can quickly detect and correct problems and ensure they degrade gracefully when components fail. Examples of non-functional requirements include ensuring that we have:

- sufficient production telemetry in our applications and environments
- the ability to accurately track dependencies
- services that are resilient and degrade gracefully
- forward and backward compatibility between versions
- the ability to archive data to manage the size of the production data set
- the ability to easily search and understand log messages across services
- the ability to trace requests from users through multiple services
- simple, centralized runtime configuration using feature flags, etc.

By codifying these types of non-functional requirements, we make it easier for all of our new and existing services to leverage the collective knowledge and experience of the organization. These are all responsibilities of the team building the service.

Build Reusable Operations User Stories into Development

When there is Operations work that cannot be fully automated or made self-service, our goal is to make this recurring work as repeatable and deterministic as possible. We do this by standardizing the needed work, automating as much as possible, and documenting our work so that we can best enable product teams to better plan and resource this activity.

Instead of manually building servers and then putting them into production according to manual checklists, we should automate as much of this work as possible, including post-installation configuration management. Where certain steps cannot be automated (e.g., manually racking a server and having another team cable it), we should collectively define the handoffs as clearly as possible to reduce lead times and errors. This will also enable us to better plan and schedule these steps in the future.

For instance, we can use tools such as Terraform to automate provisioning and configuration management of cloud infrastructure. Ad hoc changes or project work can be captured in work ticket systems such as JIRA or ServiceNow, with changes to infrastructure configuration captured in version control and linked to work tickets and then applied to our system automatically (a paradigm known as infrastructure-as-cod or GitOps).

Ideally, for all our recurring Ops work, we will know the following: what work is required, who is needed to perform it, what the steps to complete it are, and so forth. For instance, “We know a high-availability rollout takes fourteen steps, requiring work from four different teams, and the last five times we performed this it took an average of three days.”

Just as we create user stories in Development that we put into the backlog and then pull into work, we can create well-defined “Ops user stories” that represent work activities that can be reused across all our projects (e.g., deployment, capacity, security, etc.). By creating these well-defined Ops user stories, we expose repeatable IT Operations work in a manner where it shows up alongside Development work, enabling better planning and more repeatable outcomes.

Ensure Technology Choices Help Achieve Organizational Goals

When one of our goals is to maximize developer productivity, and we have service-oriented architectures, small service teams can potentially build and run their service in whatever language or framework best serves their specific needs. In some cases, this is what best enables us to achieve our organizational goals.

However, there are scenarios when the opposite occurs, such as when expertise for a critical service resides only in one team, and only that team can make changes or fix problems, creating a bottleneck. In other words, we may have optimized for team productivity but inadvertently impeded the achievement of organizational goals.

This often happens when we have a functionally oriented Operations group responsible for any aspect of service support. In these scenarios, to ensure that we enable the deep skill sets in specific technologies, we want to make sure that

Operations can influence which components are used in production or give them the ability to not be responsible for unsupported platforms.

If we do not have a list of technologies that Operations will support, collectively generated by Development and Operations, we should systematically go through the production infrastructure and services, as well as all their dependencies that are currently supported, to find which ones are creating a disproportionate amount of failure demand and unplanned work.

Our goal is to identify the technologies that:

- impede or slow down the flow of work
- disproportionately create high levels of unplanned work
- disproportionately create large numbers of support requests
- are most inconsistent with our desired architectural outcomes (e.g., throughput, stability, security, reliability, business continuity)

By removing these problematic infrastructures and platforms from the technologies supported by Ops, we enable everyone to focus on infrastructure that best helps achieve the global goals of the organization.

CONTINUOUS LEARNING

The goal is to create infrastructure platforms where users (including development teams) can self-service the operations they need without having to raise tickets or send emails. This is a key capability enabled by modern cloud infrastructure—it is even one of the five essential characteristics of cloud computing defined by the US Federal Government's National Institute of Standards and Technology (NIST):²⁰

- **On-demand self service:** Consumers can automatically provision computing resources as needed, without human interaction from the provider.
- **Broad network access:** Capabilities can be accessed through heterogeneous platforms, such as mobile phones, tablets, laptops, and workstations.
- **Resource pooling:** Provider resources are pooled in a multi-tenant model, with physical and virtual resources dynamically assigned on demand. The customer may specify location at a higher level of abstraction, such as country, state, or data center.
- **Rapid elasticity:** Capabilities can be elastically provisioned and released to rapidly scale outward or inward on demand, appearing to be unlimited and

able to be appropriated in any quantity at any time.

- **Measured service:** cloud systems automatically control, optimize, and report resource use based on the type of service, such as storage, processing, bandwidth, and active user accounts.

It's possible to achieve success in building infrastructure platforms with private, public, and hybrid models—provided you also modernize your traditional data center operations' practices and processes so you can meet these five essential characteristics. If your technology platforms don't support these characteristics, it should be a priority to replace them with ones that do or modernize your existing platform to achieve these architectural outcomes as much as possible.

DORA's 2019 *State of DevOps Report* found that only 29% of respondents who said they were using cloud infrastructure agreed or strongly agreed that they met all five of the characteristics of essential cloud computing as defined by the NIST. And leveraging all five characteristics of cloud computing mattered; elite performers were twenty-four times more likely to have met all essential cloud characteristics when compared to low performers.²¹

This demonstrates two things: First, the apparent disconnect between teams who may say they are in the cloud but may not reap the

benefits—it takes executing on the characteristics described above to be successful. Second, the impact of technical and architectural capabilities on software delivery performance. By executing well, elite teams saw significant performance in speed and stability compared to their low-performing peers.

CASE STUDY

Standardizing a New Technology Stack at Etsy (2010)

In many organizations' adopting DevOps, a common story developers tell is, "Ops wouldn't provide us what we needed, so we just built and supported it ourselves." However, in the early stages of the Etsy transformation, technology leadership took the opposite approach, significantly reducing the number of supported technologies in production.

In 2010, after a nearly disastrous peak holiday season, the Etsy team decided to massively reduce the number of technologies used in production, choosing a few that the entire organization could fully support and eradicating the rest.[‡]

Their goal was to standardize and very deliberately reduce the supported infrastructure and configurations. One of the early decisions was to migrate Etsy's entire platform to PHP and MySQL. This was primarily a philosophical decision rather than a technological decision—they wanted both Dev and Ops to be able to understand the full technology stack so that everyone could contribute to a single platform, as well as enable

everyone to be able to read, rewrite, and fix each other's code.

Over the next several years, as Michael Rembetsy, who was Etsy's Director of Operations at the time, recalls, "We retired some great technologies, taking them entirely out of production," including lighttpd, Postgres, MongoDB, Scala, CoffeeScript, Python, and many others.²³

Similarly, Dan McKinley, a developer on the feature team that introduced MongoDB into Etsy in 2010, writes on his blog that all the benefits of having a schema-less database were negated by all the operational problems the team had to solve. These included problems concerning logging, graphing, monitoring, production telemetry, and backups and restoration, as well as numerous other issues that developers typically do not need to concern themselves with. The result was to abandon MongoDB, porting the new service to use the already supported MySQL database infrastructure.²⁴

This case study from Etsy shows that, by removing problematic infrastructure and platforms, an organization can shift its focus to architectures that best align with and help achieve their goals.

→ CASE STUDY: NEW TO THE SECOND EDITION

Crowdsourcing Technology Governance at Target (2018)

One of the key findings from the *State of DevOps Reports* has been that teams move faster when we don't control how they work and operate or what technologies they

use.²⁵ In the past, technology selection was an enforcement mechanism to limit variability in the enterprise. This led to perceived compliance in conforming to architecture, security, and business architecture needs. The tollgates, centralized approvals, and silos resulted in less automation, limited automation, and continued the “process & tool first” ideology ahead of results and outcomes.

But in 2015, Target began a new program: `recommend_tech`, which uses crowdsourcing to make technology choices. It started with a basic template on a single page layout for all technologies by domain, providing a scope (local versus enterprise) and a half-life that Target’s internal experts felt was applicable to a specific disposition.²⁶

During their 2018 DevOps Enterprise Summit presentation, Dan Cundiff, Principal Engineer, Levi Geinert, Director of Engineering, and Lucas Rettif, Principal Product Owner, explained how they wanted to move faster by shifting from governance to guidance when it came to technologies: libraries, frameworks, tools, etc. This guidance would provide guardrails that teams could feel comfortable operating within while also removing the friction of a strict governance process.²⁷

They found that the key to providing guidance versus governance was that it needed to be accessible (everyone can contribute), transparent (everyone should be able to see), flexible (easy to change), and cultural (community-driven) in the simplest way possible. Ultimately, guidance should be there to empower engineers, not constrain them.²⁸

Previously Target had in place what they called an architectural review board (ARB),⁵ a centralized group that met on a regular basis to make tool decisions for all product teams. This was neither efficient nor effective.

To improve, Dan Cundiff and his colleague Jason Walker created a repo in GitHub that featured a simple list of technology choices: collaboration tools, application frameworks, caching, datastores, etc. They named it `recommended_tech`. Each technology is listed as either *recommended*, *limited use*, or *do not use*. Each file shows why that technology is recommended or not, how you can use it, etc. Plus, the full history showing how the decision was made is available in the repo.²⁹ The context of the decisions—and most importantly, the tracking of discussions—provided more “why” answers to the engineering community. As mentioned above, the “half-life” of a disposition was a directional waypoint for teams to understand the possibility of a shift within a domain.

This list isn’t handed down to engineers without any input. Anybody at Target can open a pull request to any of the technology categories and suggest a change, a new technology, etc. Everyone can comment and discuss the various benefits or risks of that technology. And when it gets merged, that’s it. The technology choice is strongly recommended and loosely held until the next pattern emerges.³⁰

As changes are localized and easy for teams to adjust, the reversibility is easy and flexibility is high. For example, switching from Python to Golang for a given product’s API is considered highly flexible and easy to reverse. Changing a cloud provider or decommissioning

a data center, on the other hand, is rigid and has an extremely large blast radius.

For changes with a “steep cost,” the CIO is brought into the process. Any engineer can then pitch their idea directly to the CIO and a group of senior leaders. Ultimately, the recommended_tech approach is about empowering engineers at any level to be invested in their work in the simplest way possible.³¹

This simple solution shows how removing impediments and bottlenecks can empower teams while still making sure they operate within approved guardrails.

Conclusion

The techniques described in this chapter enable every new learning to be incorporated into the collective knowledge of the organization, multiplying its effect. We do this by actively and widely communicating new knowledge, such as through chat rooms and through technology such as architecture as code, shared source code repositories, technology standardization, and so forth. By doing this, we elevate the state of the practice of not only Dev and Ops but also the entire organization, so everyone who performs work does so with the cumulative experience of the entire organization.

* Hubot often performed tasks by calling shell scripts, which could then be executed from the chat room anywhere, including from an engineer’s phone.

† The Chrome and Android projects reside in a separate source code repository, and certain algorithms that are kept secret, such as PageRank, are available only to certain teams.¹³

‡ At that time, Etsy used PHP, lighttpd, Postgres, MongoDB, Scala, CoffeeScript, Python, as well as many other platforms and languages.²²

§ The TEP-LARB described in *The Unicorn Project* is partially based on the ARB at Target.

21

RESERVE TIME TO CREATE ORGANIZATIONAL LEARNING AND IMPROVEMENT

One of the practices that forms part of the Toyota Production System is called the *improvement blitz* (or sometimes a *kaizen blitz*), defined as a dedicated and concentrated period of time to address a particular issue, often over the course of several days.¹ Dr. Spear explains, “blitzes often take this form: A group is gathered to focus intently on a process with problems... . The blitz lasts a few days, the objective is process improvement, and the means are the concentrated use of people from outside the process to advise those normally inside the process.”²

Spear observes that the output of the improvement blitz team will often be a new approach to solving a problem, such as new layouts of equipment, new means of conveying material and information, a more organized workspace, or standardized work. They may also leave behind a to-do list of changes to be made down the road.³

Thirty-Day Challenge at Target (2015)

An example of a DevOps improvement blitz is the monthly challenge program at the Target DevOps Dojo. Ross Clanton, former Director of Operations at Target, was responsible for accelerating the adoption of DevOps. One of his primary mechanisms for this was the Technology Innovation Center, more popularly known as the DevOps Dojo.

Target's DevOps Dojo occupies about eighteen thousand square feet of open office space, where DevOps coaches help teams from across the Target technology organization elevate the state of their practice. The most intensive format is what they call "thirty-day challenges," where internal development teams come in for a month and work together with dedicated Dojo coaches and engineers. The team brings their work with them, with the goal of solving an internal problem they have been struggling with and creating a breakthrough in thirty days.

Throughout the thirty days, they work intensively with the Dojo coaches on the problem—planning, working, and doing demos in two-day sprints. When the thirty-day challenge is complete, the internal teams return to their lines of business, not only having solved a significant problem, but bringing their new learnings back to their teams.

Clanton described, "We currently have capacity to have eight teams doing 30-Day Challenges concurrently, so we are focused on the most strategic projects of the organization. So far, we've had some of our most critical capabilities come through the Dojo, including teams from Point Of Sale (POS), Inventory, Pricing, and Promotion."⁴

By having full-time assigned Dojo staff and being focused on only one objective, teams going through a thirty-day challenge make incredible improvements. Ravi Pandey, a Target development manager who went through this program, explains, "In the old days, we would have to wait six weeks to get a test environment. Now, we get it in minutes, and we're working side by side with Ops engineers who are helping us increase our productivity and building tooling for us to help us achieve our goals."⁵

Clanton expands on this idea: “It is not uncommon for teams to achieve in days what would usually take them three to six months. So far, two hundred learners have come through the Dojo, having completed fourteen challenges.”⁶

The Dojo also supports less intensive engagement models, including flash builds, where teams come together for one- to three-day events, with the goal of shipping a minimal viable product (MVP) or a capability by the end of the event. They also host open labs every two weeks, where anyone can visit the Dojo to talk to the Dojo coaches, attend demos, or receive training.

In this chapter, we will describe this and other ways of reserving time for organizational learning and improvement, further institutionalizing the practice of dedicating time for improving daily work.

Institutionalize Rituals to Pay Down Technical Debt

In this section, we schedule rituals that help enforce the practice of reserving Dev and Ops time for improvement work, such as non-functional requirements, automation, etc. One of the easiest ways to do this is to schedule and conduct day- or week-long improvement blitzes, where everyone on a team (or in the entire organization) self-organizes to fix problems they care about—no feature work is allowed. It could be a problematic area of the code, environment, architecture, tooling, and so forth. These teams span the entire value stream, often combining Development, Operations, and Infosec engineers. Teams that typically don’t work together combine their skills and effort to improve a chosen area and then demonstrate their improvement to the rest of the company.

In addition to the Lean-oriented terms kaizen blitz and improvement blitz, the technique of dedicated rituals for

improvement work has also been called *spring* or *fall cleanings* and *ticket queue inversion weeks*.⁷ Other terms have also been used, such as *hack days*, *hackathons*, and *20% innovation time*. Unfortunately, these specific rituals sometimes focus on product innovation and prototyping new market ideas rather than on improvement work, and, worse, they are often restricted to developers—which is considerably different than the goals of an improvement blitz.*

Our goal during these blitzes is not to simply experiment and innovate for the sake of testing out new technologies but to improve our daily work, such as solving our daily workarounds. While experiments can also lead to improvements, improvement blitzes are very focused on solving specific problems we encounter in our daily work.

We may schedule week-long improvement blitzes that prioritize Dev and Ops working together toward improvement goals. These improvement blitzes are simple to administer: One week is selected where everyone in the technology organization works on an improvement activity at the same time. At the end of the period, each team makes a presentation to their peers that discusses the problem they were tackling and what they built. This practice reinforces a culture in which engineers work across the entire value stream to solve problems. Furthermore, it reinforces fixing problems as part of our daily work and demonstrates that we value paying down technical debt.

What makes improvement blitzes so powerful is that we are empowering those closest to the work to continually identify and solve their own problems. Consider for a moment that our complex system is like a spider web, with intertwining strands that are constantly weakening and

breaking. If the right combination of strands breaks, the entire web collapses.

There is no amount of command-and-control management that can direct workers to fix each strand one by one. Instead, we must create the organizational culture and norms that lead to everyone continually finding and fixing broken strands as part of our daily work. As Dr. Spear observes, “No wonder then that spiders repair rips and tears in the web as they occur, not waiting for the failures to accumulate.”⁸

A great example of the success of the improvement blitz concept is described by Mark Zuckerberg, Facebook CEO. In an interview with Jessica Stillman of *Inc.*, he says,

Every few months we have a hackathon, where everyone builds prototypes for new ideas they have. At the end, the whole team gets together and looks at everything that has been built. Many of our most successful products came out of hackathons, including Timeline, chat, video, our mobile development framework and some of our most important infrastructure like the HipHop compiler.⁹

Of particular interest is the HipHop PHP compiler. In 2008, Facebook was facing significant capacity problems, with over one hundred million active users and rapidly growing, creating tremendous problems for the entire engineering team.¹⁰ During a hack day, Haiping Zhao, Senior Server Engineer at Facebook, started experimenting with converting PHP code to compilable C++ code, with the hope of significantly increasing the capacity of their existing infrastructure. Over the next two years, a small team was assembled to build what became known as the HipHop compiler, converting all Facebook production services from

interpreted PHP to compiled C++ binaries. HipHop enabled Facebook's platform to handle production loads six times higher than the native PHP.¹¹

In an interview with Cade Metz of *Wired*, Drew Paroski, one of the engineers who worked on the project, noted, "There was a moment where, if HipHop hadn't been there, we would have been in hot water. We would probably have needed more machines to serve the site than we could have gotten in time. It was a Hail Mary pass that worked out."¹²

Later, Paroski and fellow engineers Keith Adams and Jason Evans decided that they could beat the performance of the HipHop compiler effort and reduce some of its limitations that reduced developer productivity. The resulting project was the HipHop virtual machine project ("HHVM"), taking a just-in-time compilation approach. By 2012, HHVM had completely replaced the HipHop compiler in production, with nearly twenty engineers contributing to the project.¹³

By performing regularly scheduled improvement blitzes and hack weeks, we enable everyone in the value stream to take pride and ownership in the innovations they create, and we continually integrate improvements into our system, further enabling safety, reliability, and learning.

Enable Everyone to Teach and Learn

A dynamic culture of learning creates conditions so that everyone can not only learn but also teach, whether through traditional didactic methods (e.g., people taking classes, attending training) or more experiential or open methods (e.g., conferences, workshops, mentoring). One way that we can foster this teaching and learning is to dedicate organizational time to it.

Steve Farley, VP of Information Technology at Nationwide Insurance, said,

We have five thousand technology professionals, who we call “associates.” Since 2011, we have been committed to create a culture of learning—part of that is something we call Teaching Thursday, where each week we create time for our associates to learn. For two hours, each associate is expected to teach or learn. The topics are whatever our associates want to learn about—some of them are on technology, on new software development or process improvement techniques, or even on how to better manage their career. The most valuable thing any associate can do is mentor or learn from other associates.¹⁴

As has been made evident throughout this book, certain skills are becoming increasingly needed by all engineers, not just by developers. For instance, it is becoming more important for all Operations and Test engineers to be familiar with Development techniques, rituals, and skills, such as version control, automated testing, deployment pipelines, configuration management, and creating automation. Familiarity with Development techniques helps Operations engineers remain relevant as more technology value streams adopt DevOps principles and patterns.

Although the prospect of learning something new may be intimidating or cause a sense of embarrassment or shame, it shouldn't. After all, we are all lifelong learners, and one of the best ways to learn is from our peers. Karthik Gaekwad, who was part of the National Instruments DevOps transformation, said, “For Operations people who are trying to learn automation, it shouldn't be scary—just ask a friendly developer, because they would love to help.”¹⁵

We can further help teach skills through our daily work by jointly performing code reviews that include both parties so that we learn by doing, as well as by having Development and Operations work together to solve small problems. For instance, we might have Development show Operations how to authenticate an application and log in and run automated tests against various parts of the application to ensure that critical components are working correctly (e.g., key application functionality, database transactions, message queues). We would then integrate this new automated test into our deployment pipeline and run it periodically, sending the results to our monitoring and alerting systems so that we get earlier detection when critical components fail.

As Glenn O'Donnell from Forrester Research quipped in his 2014 DevOps Enterprise Summit presentation, "For all technology professionals who love innovating, love change, there is a wonderful and vibrant future ahead of us."¹⁶

CONTINUOUS LEARNING

ASREDS Learning Loop

Humans have an innate need to belong to groups. But this instinct to belong can also be an instinct to exclude. When it comes to shared learning, this group mentality can result in learning bubbles (learning silos) or the permanent loss of institutional knowledge when people leave the group (company).

When learnings become trapped in bubbles, knowledge is hidden and different teams unnecessarily struggle with similar issues, run similar experiments, develop the same antipatterns, and fail to use each other's learnings. In *Sooner Safer Happier*, the authors use the ASREDS learning loop to pop these learning bubbles.¹⁷

The loop asks teams to first Align on a goal, then Sense the context, Respond by designing one or more Experiments, Distill the results into insights and metrics, and then Share the results by publishing the learnings so others can pick them up again at Sense.

Practices like ASREDS help pop disconnected learning bubbles, and when combined with awards and Centers of Practice (see *Sooner Safer Happier* for more), these patterns eventually promote the creation of a learning ecosystem.

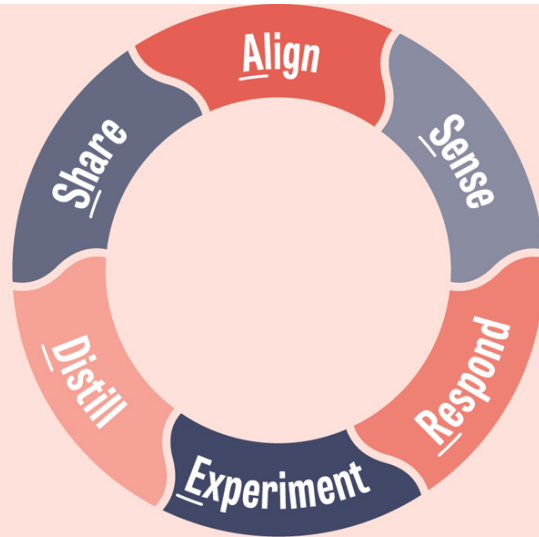


Figure 21.1: The ASREDS Learning Loop

Source: Smart et al., *Sooner Safer Happier: Antipatterns and Patterns for Business Agility* (Portland, OR: IT Revolution, 2020).

Share your Experiences from DevOps Conferences

In many cost-focused organizations, engineers are often discouraged from attending conferences and learning from their peers. To help build a learning organization, we should encourage our engineers (both from Development and Operations) to attend conferences, give talks at them, and, when necessary, create and organize internal or external conferences themselves.

DevOpsDays remains one of the most vibrant self-organized conference series today. Many DevOps practices have been shared and promulgated at these events. It has remained free or nearly free, supported by a vibrant community of practitioner communities and vendors.

The DevOps Enterprise Summit was created in 2014 for technology leaders to share their experiences adopting DevOps principles and practices in large, complex organizations. The program is organized primarily around experience reports from technology leaders on the DevOps journey, as well as subject matter experts on topics selected

by the community. By 2021, DevOps Enterprise Summit has held fourteen conferences, with nearly one thousand talks from technology experts from almost every industry vertical.

CASE STUDY

Internal Technology Conferences at Nationwide Insurance, Capital One, and Target (2014)

Along with attending external conferences, many companies have internal conferences for their technology employees. Nationwide Insurance is a leading provider of insurance and financial services and operates in heavily regulated industries. Their many offerings include auto and homeowner's insurance, and they are the top provider of public-sector retirement plans and pet insurance. As of 2014, they had \$195 billion in assets, with \$24 billion in revenue.¹⁸

Since 2005, Nationwide has been adopting Agile and Lean principles to elevate the state of practice for their five thousand technology professionals, enabling grassroots innovation.

Steve Farley, VP of Information Technology, remembers,

Exciting technology conferences were starting to appear around that time, such as the Agile national conference. In 2011, the technology leadership at Nationwide agreed that we should create a technology conference, called TechCon. By holding this event, we wanted to create a better way to teach ourselves, as well as ensure that everything had a Nationwide context, as

opposed to sending everyone to an external conference.¹⁹

Capital One, one of the largest banks in the US with over \$298 billion in assets and \$24 billion in revenue in 2015, held their first internal software engineering conference in 2015 as part of their goal to create a world-class technology organization.²⁰ The mission was to promote a culture of sharing and collaboration, and to build relationships between the technology professionals and enable learning. The conference had thirteen learning tracks and fifty-two sessions, and over 1,200 internal employees attended.²¹

Dr. Tapabrata Pal, a technical fellow at Capital One and one of the organizers of the event, describes, “We even had an expo hall, where we had twenty-eight booths, where internal Capital One teams were showing off all the amazing capabilities they were working on. We even decided very deliberately that there would be no vendors there, because we wanted to keep the focus on Capital One goals.”²²

Target is the sixth-largest retailer in the US, with \$72 billion in revenue in 2014 and 1,799 retail stores and 347,000 employees worldwide.²³ Heather Mickman, a director of Development, and Ross Clanton have held six internal DevOpsDays events since 2014 and have over 975 followers inside their internal technology community, modeled after the DevOpsDays held at ING in Amsterdam in 2013.^{24†}

After Mickman and Clanton attended the DevOps Enterprise Summit in 2014, they held their own internal conference, inviting many of the speakers from outside

firms so that they could re-create their experience for their senior leadership. Clanton describes, “2015 was the year when we got executive attention and when we built up momentum. After that event, tons of people came up to us, asking how they could get involved and how they could help.”²⁶

Institutions can enable their technologists by providing a dynamic culture of learning and teaching, not only through attending and presenting at external conferences but also by building and holding internal conferences. This can foster greater team and organizational trust, increase communication and innovation, and improve daily work.

CONTINUOUS LEARNING

The DORA 2019 *State of DevOps Report* investigated how organizations spread DevOps and Agile practices, asking them to choose from a variety of common methods, such as training centers, centers of excellence, various kinds of proof of concept, big bang, and communities of practice.

Analysis showed that:

high performers favor strategies that create community structures at both low and high levels in the organization, likely making them more sustainable and resilient to re-organizations and product changes. The top two strategies employed are communities of practice and grassroots, followed by proof of concept as a template (a pattern where the proof of concept gets reproduced elsewhere in the organization), and proof of concept as a seed.²⁷

Create Community Structures to Spread Practices

Earlier in the book, we began the story of how the Testing Grouplet built a world-class automated testing culture at Google starting in 2005. Their story continues here, as they try to improve the state of automated testing across all of Google by using dedicated improvement blitzes, internal coaches, and even an internal certification program.

Mike Bland said, at that time, there was a 20% innovation time policy at Google, enabling developers to spend roughly

one day per week on a Google-related project outside of their primary area of responsibility. Some engineers chose to form communities of practice they called *grouplets*, ad hoc teams of like-minded engineers who wanted to pool their 20% time, allowing them to do focused improvement blitzes.²⁸

A testing community of practice was formed by Bharat Mediratta and Nick Lesiecki, with the mission of driving the adoption of automated testing across Google. Even though they had no budget or formal authority, as Mike Bland described, “There were no explicit constraints put upon us, either. And we took advantage of that.”²⁹

They used several mechanisms to drive adoption, but one of the most famous was *Testing on the Toilet* (or TotT), their weekly testing periodical. Each week, they published a newsletter in nearly every bathroom in nearly every Google office worldwide. Bland said, “The goal was to raise the degree of testing knowledge and sophistication throughout the company. It’s doubtful an online-only publication would’ve involved people to the same degree.”³⁰

Bland continues, “One of the most significant TotT episodes was the one titled, ‘Test Certified: Lousy Name, Great Results,’ because it outlined two initiatives that had significant success in advancing the use of automated testing.”³¹

Test Certified (TC) provided a road map to improve the state of automated testing. As Bland describes, “It was intended to hack the measurement-focused priorities of Google culture ... and to overcome the first, scary obstacle of not knowing where or how to start. Level 1 was to quickly establish a baseline metric, Level 2 was setting a policy and

reaching an automated test coverage goal, and Level 3 was striving towards a long-term coverage goal.”³²

The second capability was providing test certified mentors to any team who wanted advice or help, and test mercenaries (i.e., a full-time team of internal coaches and consultants) to work hands-on with teams to improve their testing practices and code quality. The mercenaries did so by applying the testing grouplet’s knowledge, tools, and techniques to a team’s own code, using TC as both a guide and a goal.

Bland was eventually a leader of the testing grouplet from 2006 to 2007, and a member of the test mercenaries from 2007 to 2009.³³

Bland continues,

It was our goal to get every team to TC Level 3, whether they were enrolled in our program or not. We also collaborated closely with the internal testing tools teams, providing feedback as we tackled testing challenges with the product teams. We were boots on the ground, applying the tools we built, and eventually, we were able to remove ‘I don’t have time to test’ as a legitimate excuse...³⁴

The TC levels exploited the Google metrics-driven culture—the three levels of testing were something that people could discuss and brag about at performance review time. The Testing Grouplet eventually got funding for the Test Mercenaries, a staffed team of full-time internal consultants. This was an important step, because now management was fully onboard, not with edicts, but by actual funding.³⁵

Another important construct was leveraging company-wide “fixit” improvement blitzes. Bland describes fixits as

“when ordinary engineers with an idea and a sense of mission recruit all of Google engineering for one-day, intensive sprints of code reform and tool adoption.”³⁶

He organized four company-wide fixits, two pure testing fixits and two that were more tools-related, the last involving more than one hundred volunteers in over twenty offices in thirteen countries. He also led the Fixit Grouplet from 2007 to 2008.³⁷

These fixits, as Bland describes, mean that we should provide focused missions at critical points in time to generate excitement and energy, which helps advance the state of the art. This will help the long-term culture change mission reach a new plateau with every big, visible effort.³⁸

The results of the testing culture are self-evident in the amazing results Google has achieved, as presented throughout the book.

Conclusion

This chapter described how we can institute rituals that help reinforce the culture that we are all lifelong learners and that we value the improvement of daily work over daily work itself. We do this by reserving time to pay down technical debt, and creating community structures that allow everyone to learn from and teach each other, both inside our organization and outside it. By having everyone help each other learn in our daily work, we out-learn the competition, helping us win in the marketplace. But also, we help each other achieve our full potential as human beings.

* From here on, the terms “hack week” and “hackathon” are used interchangeably with “improvement blitz,” and not in the context of “you can work on whatever you want.”

† Incidentally, the first Target internal DevOpsDays event was modeled after the

first ING DevOpsDays that was organized by Ingrid Algra, Jan-Joost Bouwman, Evelijn Van Leeuwen, and Kris Buytaert in 2013, after some of the ING team attended the 2013 Paris DevOpsDays.²⁵

CONCLUSION TO PART V

Throughout Part V, we have explored the practices that help create a culture of learning and experimentation in your organization. Learning from incidents, creating shared repositories, and sharing learnings is essential when we work in complex systems, helping to make our work culture more just and our systems safer and more resilient.

In [Part VI](#), we'll explore how to extend flow, feedback, and learning and experimentation by using them to simultaneously help us achieve our Information Security goals.

Additional Resources

Amy Edmondson's *The Fearless Organization: Creating Psychological Safety in the Workplace for Learning, Innovation, and Growth* is a go-to book on building psychological safety in the workplace.

Team of Teams: New Rules of Engagement for a Complex World by General Stanley McChrystal showcases the art of leadership in the US military.

The Five Dysfunctions of a Team: A Leadership Fable from Patrick Lencioni brings a leadership fable that reveals the five dysfunctions that go to the very heart of why teams—even the best ones—often struggle.

Release It! Design and Deploy Production-Ready Software from Michael Nygard helps you avoid the pitfalls that cost companies millions of dollars in downtime and reputation.

Just Culture by Sidney Dekker can help you see how to build a culture of trust, learning, and accountability.

To learn more about incident command models, check out the 2018 DevOps Enterprise Summit presentation, “Mastering

Outages with Incident Command for DevOps,” from Brent Chapman, Principal at Great Circle Associates (<https://videolibrary.doesvirtual.com/?video=524038081>).



Part VI

THE TECHNOLOGICAL PRACTICES
OF INTEGRATING INFORMATION
SECURITY, CHANGE MANAGEMENT,
AND COMPLIANCE

PART VI: INTRODUCTION

In the previous chapters, we discussed enabling the fast flow of work from check-in to release, as well as creating the reciprocal fast flow of feedback. We explored the cultural rituals that reinforce the acceleration of organizational learning and the amplification of weak failure signals that help us to create an ever-safer system of work.

In Part VI, we further extend these activities so that we not only achieve Development and Operations goals but also simultaneously achieve Information Security goals, helping us to create a high degree of assurance around the confidentiality, integrity, and availability of our services and data.

Instead of injecting security into our product at the end of the process, we will create and integrate security controls into the daily work of Development and Operations, so that security is part of everyone's job every day. Ideally, much of this work will be automated and put into our deployment pipeline. Furthermore, we will augment our manual practices, acceptances, and approval processes with automated controls, relying less on controls such as separation of duties and change approval processes.

By automating these activities, we can generate evidence on demand to demonstrate that our controls are operating effectively, whether to auditors, assessors, or anyone else working in our value stream.

In the end, we will not only improve security but also create processes that are easier to audit and that attest to the effectiveness of controls, in support of compliance with regulatory and contractual obligations. We do this by:

- making security a part of everyone's job

- integrating preventative controls into our shared source code repository
- integrating security with our deployment pipeline
- integrating security with our telemetry to better enable detection and recovery
- protecting our deployment pipeline
- integrating our deployment activities with our change approval processes
- reducing reliance on separation of duty

When we integrate security work into everyone's daily work, making it everyone's responsibility, we help the organization have better security. Better security means that we are defensible and sensible with our data. It means that we are reliable and have business continuity by being more available and more capable of easily recovering from issues. We are also able to overcome security problems before they cause catastrophic results, and we can increase the predictability of our systems. And, perhaps most importantly, we can secure our systems and data better than ever.

INFORMATION SECURITY IS EVERYONE'S JOB EVERY DAY

One of the top objections to implementing DevOps principles and patterns has been: “Information Security and Compliance won’t let us.” And yet, DevOps may be one of the best ways to better integrate information security into the daily work of everyone in the technology value stream.

When Infosec is organized as a silo outside of Development and Operations, many problems arise. James Wickett, one of the creators of the Gauntlt security tool and organizer of DevOpsDays Austin and the Lonestar Application Security conference, observed:

One interpretation of DevOps is that it came from the need to enable developer’s productivity, because as the number of developers grew, there weren’t enough Ops people to handle all the resulting deployment work. This shortage is even worse in Infosec—the ratio of engineers in Development, Operations, and Infosec in a typical technology organization is 100:10:1. When Infosec is that outnumbered, without automation and integrating information security into the daily work of Dev and Ops, Infosec can only do compliance checking, which is the opposite of security engineering—and besides, it also makes everyone hate us.¹

James Wickett and Josh Corman, former CTO of Sonatype and respected information security researcher, have written about incorporating information security objectives

into DevOps, a set of practices and principles termed *Rugged DevOps*.^{2*}

Throughout *The DevOps Handbook*, we have explored how to fully integrate QA and Operations objectives throughout our entire technology value stream. In this chapter, we describe how to similarly integrate Infosec objectives into our daily work, where we can increase our safety and security while maintaining developer and operational productivity.

Integrate Security into Development Iteration Demonstrations

One of our goals is to have feature teams engaged with Infosec as early as possible, as opposed to primarily engaging at the end of the project. One way we can do this is by inviting Infosec to the product demonstrations at the end of each development interval so that they can better understand the team goals in the context of organizational goals, observe their implementations as they are being built, and provide guidance and feedback at the earliest stages of the project, when there is the most amount of time and freedom to make corrections.

Justin Arbuckle, former Chief Architect at GE Capital, observes,

When it came to information security and compliance, we found that blockages at the end of the project as much more expensive than at the beginning—and Infosec blockages were among the worst. ‘Compliance by demonstration’ became one of the rituals we used to shift all this complexity earlier in the process...⁴

By having Infosec involved throughout the creation of any new capability, we were able to reduce our use of static checklists dramatically and rely more on using

their expertise throughout the entire software development process.⁵

This helped the organization achieve its goals. Snehal Antani, former CIO of Enterprise Architecture at GE Capital Americas, described their top three key business measurements as “development velocity (i.e., speed of delivering features to market), failed customer interactions (i.e., outages, errors), and compliance response time (i.e., lead time from audit request to delivery of all quantitative and qualitative information required to fulfill the request).”⁶

When Infosec is an assigned part of the team, even if they are only being kept informed and observing the process, they gain the business context they need to make better risk-based decisions. Furthermore, Infosec is able to help feature teams learn what is required to meet security and compliance objectives.

Integrate Security into Defect Tracking and Post-Mortems

When possible, we want to track all open security issues in the same work tracking system that Development and Operations are using, ensuring the work is visible and can be prioritized against all other work. This is very different from how Infosec has traditionally worked, where all security vulnerabilities are stored in a GRC (governance, risk, and compliance) tool that only Infosec has access to. Instead, we will put any needed work in the systems that Dev and Ops use.

In a presentation at the 2012 Austin DevOpsDays, Nick Galbreath, who headed up Information Security at Etsy for many years, describes how his team treated security issues: “We put all security issues into JIRA, which all engineers use

in their daily work, and they were either ‘P1’ or ‘P2,’ meaning that they had to be fixed immediately or by the end of the week, even if the issue is only an internally facing application.”⁷

Furthermore, he states, “Any time we had a security issue, we would conduct a post-mortem, because it would result in better educating our engineers on how to prevent it from happening again in the future, as well as a fantastic mechanism for transferring security knowledge to our engineering teams.”⁸

Integrate Preventive Security Controls into Shared Source Code Repositories and Shared Services

In [Chapter 20](#), we created a shared source code repository that makes it easy for anyone to discover and reuse the collective knowledge of our organization—not only for our code but also for our toolchains, deployment pipeline, standards, etc. By doing this, anyone can benefit from the cumulative experience of everyone in the organization.

Now we will add to our shared source code repository any mechanisms or tools that help enable us to ensure our applications and environments are secure. We will add libraries that are pre-blessed by security to fulfill specific Infosec objectives, such as authentication and encryption libraries and services.

Because everyone in the DevOps value stream uses version control for anything they build or support, putting our information security toolchain and approved libraries there makes it much easier to influence the daily work of Dev and Ops, because anything we create is available, searchable, and reusable. Version control also serves as an omni-

directional communication mechanism to keep all parties aware of changes being made.

If we have a centralized shared services organization, we may also collaborate with them to create and operate shared security-relevant platforms, such as authentication, authorization, logging, and other security and auditing services that Dev and Ops require. When engineers use one of these predefined libraries or services, they won't need to schedule a separate security design review for that module; they'll be using the guidance we've created concerning configuration hardening, database security settings, key lengths, and so forth.

To further increase the likelihood that the services and libraries we provide will be used correctly, we can provide security training to Dev and Ops, as well as review what they've created to help ensure that security objectives are being implemented correctly, especially for teams using these tools for the first time.

Ultimately, our goal is to provide the security libraries or services that every modern application or environment requires, such as enabling user authentication, authorization, password management, data encryption, and so forth. Furthermore, we can provide Dev and Ops with effective security-specific configuration settings for the components they use in their application stacks, such as for logging, authentication, and encryption. We may include items such as:

- code libraries and their recommended configurations (e.g., 2FA [two-factor authentication library], bcrypt password hashing, logging)
- secret management (e.g., connection settings, encryption keys) using tools such as Vault, sneaker,

Keywhiz, credstash, Trousseau, Red October, etc.[†]

- OS packages and builds (e.g., NTP for time syncing, secure versions of OpenSSL with correct configurations, OSSEC or Tripwire for file integrity monitoring, syslog configuration to ensure logging of critical security into our centralized ELK stack)

By putting all these into our shared source code repository, we make it easy for any engineer to correctly create and use logging and encryption standards in their applications and environments, with no further work from us.

We should also collaborate with Ops teams to create a base cookbook or base image of our OS, databases, and other infrastructure (e.g., NGINX, Apache, Tomcat), showing they are in a known, secure, and risk-reduced state. Our shared repository not only becomes the place where we can get the latest versions but also becomes a place where we can collaborate with other engineers and monitor and alert on changes made to security-sensitive modules.

Now that Docker-based systems are ubiquitous, organizations should use a container registry to hold all base images. In order to secure the software supply chain, these source versions should be stored along with a secure hash of the image created. This hash must be validated whenever the image is used or deployed.

Integrate Security into Our Deployment Pipeline

In previous eras, in order to harden and secure our application, we would start our security review after development was completed. Often, the output of this review would be hundreds of pages of vulnerabilities in a PDF, which we'd give to Development and Operations, which would be

completely unaddressed due to project due date pressure or problems being found too late in the software life cycle to be easily corrected.

In this step, we will automate as many of our information security tests as possible, so that they run alongside all our other automated tests in our deployment pipeline, being performed (ideally) upon every code commit by Dev or Ops, and even in the earliest stages of a software project.

Our goal is to provide both Dev and Ops with fast feedback on their work so that they are notified whenever they commit changes that are potentially insecure. By doing this, we enable them to quickly detect and correct security problems as part of their daily work, which enables learning and prevents future errors.

Ideally, these automated security tests will be run in our deployment pipeline alongside the other static code analysis tools.

Tools such as Gauntlt have been designed to integrate into the deployment pipelines, which run automated security tests on our applications, our application dependencies, our environment, etc. Remarkably, Gauntlt even puts all its security tests in Gherkin syntax test scripts, which is widely used by developers for unit and functional testing. Doing this puts security testing in a framework they are likely already familiar with. This also allows security tests to easily run in a deployment pipeline on every committed change, such as static code analysis, checking for vulnerable dependencies, or dynamic testing.

By doing this, we provide everyone in the value stream with the fastest possible feedback about the security of what they are creating, enabling Dev and Ops engineers to find and fix issues quickly.

Jenkins					
S	W	Name	Last Success	Last Failure	Last Duration
●	☀️	Static analysis scan	7 days 1 hr - #2	N/A	6.3 sec
●	🕒	Check known vulnerabilities in dependencies	N/A	7 days 1 hr - #2	1.6 sec
●	☀️	Download and unit test	7 days 1 hr - #2	N/A	32 sec
●	☀️	Scan with OWASP ZAP	7 days 1 hr - #2	N/A	4 min 43 sec
●	☀️	Start	7 days 1 hr - #2	N/A	5 min 46 sec
●	☀️	Virus scanning	7 days 1 hr - #2	N/A	4.7 sec

Figure 22.1: Jenkins Running Automated Security Testing

Source: James Wicket and Gareth Rushgrove, “Battle-tested code without the battle,” Velocity 2014 conference presentation, posted to [Speakerdeck.com](https://speakerdeck.com/garethr/battle-tested-code-without-the-battle), June 24, 2014, <https://speakerdeck.com/garethr/battle-tested-code-without-the-battle>.

Ensure Security of the Application

Often, Development testing focuses on the correctness of functionality, looking at positive logic flows. This type of testing is often referred to as the *happy path*, which validates user journeys (and sometimes alternative paths) where everything goes as expected, with no exceptions or error conditions.

On the other hand, effective QA, Infosec, and fraud practitioners will often focus on the *sad paths*, which happen when things go wrong, especially in relation to security-related error conditions. (These types of security-specific conditions are often jokingly referred to as the *bad paths*.)

For instance, suppose we have an e-commerce site with a customer input form that accepts credit card numbers as part of generating a customer order. We want to define all the sad and bad paths required to ensure that invalid credit cards are properly rejected to prevent fraud and security exploits, such as SQL injections, buffer overruns, and other undesirable outcomes.

Instead of performing these tests manually, we would ideally generate them as part of our automated unit or functional tests so that they can be run continuously in our deployment pipeline.

As part of our testing, we will want to include the following:

Static analysis: This is testing that we perform in a non-runtime environment, ideally in the deployment pipeline. Typically, a static analysis tool will inspect program code for all possible runtime behaviors and seek out coding flaws, back doors, and potentially malicious code (this is sometimes known as “testing from the inside out”). Examples of tools include Brakeman, Code Climate, and searching for banned code functions (e.g., “exec”).

Dynamic analysis: As opposed to static testing, dynamic analysis consists of tests executed while a program is in operation. Dynamic tests monitor items such as system memory, functional behavior, response time, and overall performance of the system. This method (sometimes known as “testing from the outside in”) is similar to the manner in which a malicious third party might interact with an application. Examples include Arachni and OWASP ZAP (Zed Attack Proxy). † Some types of penetration testing can also be performed in an automated fashion and should be included as part of dynamic analysis using tools such as Nmap and Metasploit. Ideally, we should perform automated dynamic testing during the automated functional testing phase of our deployment pipeline, or even against our services

while they are in production. To ensure correct security handling, tools like OWASP ZAP can be configured to attack our services through a web browser proxy and inspect the network traffic within our test harness.

Dependency scanning: Another type of static testing we would normally perform at build time inside of our deployment pipeline involves inventorying all our dependencies for binaries and executables and ensuring that these dependencies, which we often don't have control over, are free of vulnerabilities or malicious binaries. Examples include Gemnasium and bundler audit for Ruby, Maven for Java, and the OWASP Dependency-Check.

Source code integrity and code signing: All developers should have their own PGP key, perhaps created and managed in a system such as keybase.io. All commits to version control should be signed—that is straightforward to configure using the open-source tools gpg and git. Furthermore, all packages created by the CI process should be signed, and their hash recorded in the centralized logging service for audit purposes.

Furthermore, we should define design patterns to help developers write code to prevent abuse, such as putting in rate limits for our services and graying out submit buttons after they have been pressed.

OWASP publishes a great deal of useful guidance such as the Cheat Sheet series, which includes:⁹

- how to store passwords
- how to handle forgotten passwords

- how to handle logging
- how to prevent cross-site scripting (XSS) vulnerabilities

CASE STUDY

Static Security Testing at Twitter (2009)

The “10 Deploys per Day: Dev and Ops Cooperation at Flickr” presentation by John Allspaw and Paul Hammond is famous for catalyzing the Dev and Ops community in 2009. The equivalent for the information security community is likely the presentation that Justin Collins, Alex Smolen, and Neil Matatall gave on their information security transformation work at Twitter at the AppSecUSA conference in 2012.

Twitter had many challenges due to hyper-growth. For years, the famous Fail Whale error page would be displayed when Twitter did not have sufficient capacity to keep up with user demand, showing a graphic of a whale being lifted by eight birds. The scale of user growth was breathtaking—between January and March 2009, the number of active Twitter users went from 2.5 million to 10 million.¹⁰

Twitter also had security problems during this period. In early 2009, two serious security breaches occurred. First, in January the @BarackObama Twitter account was hacked. Then in April, the Twitter administrative accounts were compromised through a brute-force dictionary attack. These events led the Federal Trade Commission to judge that Twitter was misleading its users into believing that their accounts were secure and issued an FTC consent order.¹¹

The consent order required that Twitter comply within sixty days by instituting a set of processes that were to be enforced for the following twenty years and would do the following:¹²

- Designate an employee or employees to be responsible for Twitter's information security plan.
- Identify reasonably foreseeable risks, both internal and external, that could lead to an intrusion incident and create and implement a plan to address these risks.⁸
- Maintain the privacy of user information, not just from outside sources but also internally, with an outline of possible sources of verification and testing of the security and correctness of these implementations.

The group of engineers assigned to solve this problem had to integrate security into the daily work of Dev and Ops and close the security holes that allowed the breaches to happen in the first place.

In their previously mentioned presentation, Collins, Smolen, and Matatall identified several problems they needed to address:¹³

- **Prevent security mistakes from being repeated:** They found that they were fixing the same defects and vulnerabilities over and over again. They needed to modify the system of work and automation tools to prevent the issues from happening again.
- **Integrate security objectives into existing developer tools:** They identified early on that

the major source of vulnerabilities were code issues. They couldn't run a tool that generated a huge PDF report and then email it to someone in Development or Operations. Instead, they needed to provide the developer who had created the vulnerability with the exact information needed to fix it.

- **Preserve trust of Development:** They needed to earn and maintain the trust of Development. That meant they needed to know when they sent Development false positives, so they could fix the error that prompted the false positive and avoid wasting Development's time.
- **Maintain fast flow through Infosec through automation:** Even when code vulnerability scanning was automated, Infosec still had to do lots of manual work and waiting. They had to wait for the scan to complete, get back the big stack of reports, interpret the reports, and then find the person responsible for fixing it. And when the code changed, it had to be done all over again. By automating the manual work, they removed "button-pushing" tasks, saving time and enabling them to use more creativity and judgment to solve new problems.
- **Make everything security related self-service, if possible:** They trusted that most people wanted to do the right thing, so it was necessary to provide them with all the context and information they needed to fix any issues.

- **Take a holistic approach to achieving Infosec objectives:** Their goal was to do analysis from all the angles: source code, the production environment, and even what their customers were seeing.

The first big breakthrough occurred during a company-wide hack week when they integrated static code analysis into the Twitter build process. The team used Brakeman, which scans Ruby on Rails applications for vulnerabilities. The goal was to integrate security scanning into the earliest stages of the Development process, not just when the code was committed into the source code repo.¹⁴

The results of integrating security testing into the development process were breathtaking. Over the years, by creating fast feedback for developers when they write insecure code and showing them how to fix the vulnerabilities, Brakeman has reduced the rate of vulnerabilities found by 60%,¹⁵ as shown in [Figure 22.2](#). (The spikes are usually associated with new releases of Brakeman.)

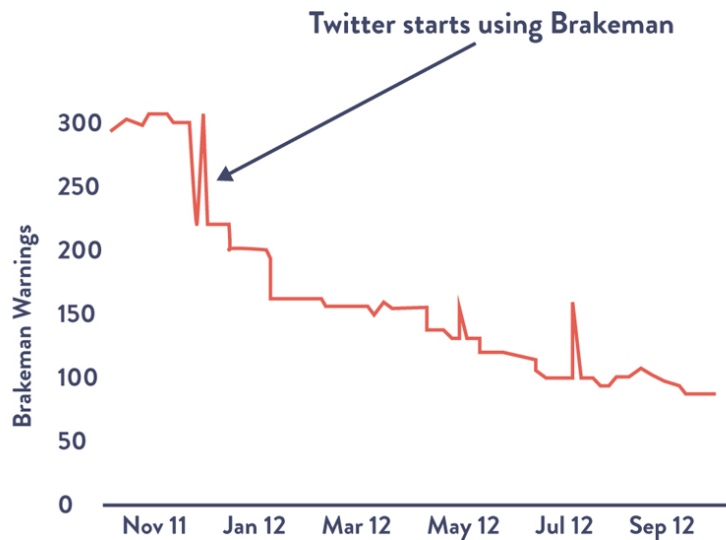


Figure 22.2: Number of Brakeman Security Vulnerabilities Detected

This case study illustrates just how necessary it is to integrate security into the daily work and tools of DevOps and how effectively it can work. Doing so mitigates security risk, reduces the probability of vulnerabilities in the system, and helps teach developers to write more secure code.

Ensure Security of Our Software Supply Chain

Josh Corman observed that as developers “we are no longer writing customized software—instead, we assemble what we need from open source parts, which has become the software supply chain that we are very much reliant upon.”¹⁶ In other words, when we use components or libraries—either commercial or open source—in our software, we not only inherit their functionality, but also any security vulnerabilities they contain.

CONTINUOUS LEARNING

In the *2020 State of the Octoverse* report, Dr. Nicole Forsgren and her team conducted a deep dive on open source and their dependencies in the Securing Software section of the report. They found the most frequent use of open-source dependencies in JavaScript (94%), Ruby (90%), and .NET (90%).¹⁷

Their research also found that when teams used automation to generate a pull request patch for detected vulnerabilities, it accelerated their supply chain security thirteen days sooner, or 1.4 times faster, than those who didn't. This shows the effectiveness of shifting left and integrating security into our development and operations workflows.¹⁸

When selecting software, we detect when our software projects are relying on components or libraries that have known vulnerabilities and help developers choose the components they use deliberately and with due care, selecting those components (e.g., open-source projects) that have a demonstrated history of quickly fixing software vulnerabilities. We also look for multiple versions of the same library being used across our production landscape, particularly the presence of older versions of libraries that contain known vulnerabilities.

Examining cardholder data breaches shows how important the security of open-source components we choose can be. Since 2008, the annual Verizon PCI Data Breach Investigation Report (DBIR) has been the most authoritative voice on data breaches where cardholder data was lost or

stolen. The 2014 report studied over eighty-five thousand breaches to better understand where attacks were coming from, how cardholder data was stolen, and factors leading to the breaches.

The 2014 DBIR found that ten vulnerabilities (i.e., CVEs) accounted for almost 97% of the exploits used in studied cardholder data breaches in 2014. Of these ten vulnerabilities, eight of them were over ten years old.¹⁹

CONTINUOUS LEARNING

In 2021, the authors of the DBIR analyzed the vulnerabilities on all the internet-facing assets for eighty-five organizations and found that most had vulnerabilities that were from 2010 or before. They wrote, “one might think that more recent vulnerabilities would be more common. However, as we saw last year, it is actually the older vulnerabilities that are leading the way.”²⁰

The 2019 *Sonatype State of the Software Supply Chain Report*, coauthored by Dr. Stephen Magill and Gene Kim, described the analysis of the Maven Central repository, which stores software components for the Java ecosystem (similar to what NPM does for JavaScript, PyPi for Python, or Gems for Ruby). In 2019, Maven Central contained over four million versions of 310,000 components, servicing over 146 billion download requests (68% year-over-year growth). In the study, the authors analyzed 4.2 million JAR artifacts (Java archive files) and 6,952 GitHub projects where they resided.²¹

The report included these startling findings:²²

- 9% of components had at least one vulnerability associated with them.
- When the component and all of its transitives are analyzed, 47% of

components have at least one vulnerability.

- The median time for software components to remediate software vulnerabilities was 326 days.




Figure 22.3: Time to Remediate vs. Time to Update Dependencies (TTU)

Source: Sonatype, 2019 *Software Supply Chain Report*.

The 2019 report also showed that when analyzing software components, the time required for those projects to remediate their security vulnerabilities (TTR) was correlated with the time required to update any of their dependencies (TTU).²³ In other words, projects that update more frequently tend to remediate their security vulnerabilities faster.

This fact motivates why Jeremy Long, founder of the OWASP Dependency Check project, suggests that the best security patching strategy is to remain current on all dependencies.²⁴ He speculates that “only 25% of organizations report vulnerabilities to users, and only 10% of vulnerabilities are reported as Common Vulnerabilities and

Exposures (CVE).”²⁵ Furthermore, the publication of a CVE is often for a vulnerability that was fixed in an earlier version of a component.

For example, the PrimeFaces CVE-2017-1000486 was published January 3, 2018, which cryptominers then started exploiting. However, the vulnerability was actually fixed in February 2016. Those who had already updated to a newer version were unaffected.²⁶

The 2019 study also found that the “popularity” of a software project (e.g., number of GitHub stars or forks or the number of Maven Central downloads) is not correlated with better security characteristics.²⁷ This is problematic because many engineers select open-source components based on the project’s popularity. However, project popularity has not been shown to correlate with their TTU (time to update their dependencies).²⁸

The 2019 Sonatype study found five behavioral clusters for open-source projects:²⁹

- **Small exemplar:** small development teams (1.6 devs), exemplary MTTU.
- **Large exemplar:** large development teams (8.9 devs), exemplary MTTU, very likely to be foundation supported, 11x more popular.
- **Laggards:** poor MTTU, high stale dependency count, more likely to be

commercially supported.

- **Features first:** frequent releases but poor TTU, still reasonably popular.
- **Cautious:** good TTU but seldom completely up to date.




Figure 22.4: Five Behavioral Clusters for Open-Source Projects

Source: Sonatype, 2019 *Software Supply Chain Report*.

The 2020 *State of the Software Supply Chain Report* surveyed developers to determine what practices help achieve developer productivity and security objectives. When comparing the high-performing versus low-performing clusters (as measured by developer productivity and security outcomes), they found that high performers had these properties:³⁰

Confidence of Changes:

- 15x more frequent deployments
- 4.9x less likely to have dependencies break application functionality

- 3.8x more likely to describe updating dependencies as easy (i.e., not painful)

Security of Components:

- 26x faster detection and remediation of vulnerable OSS components
- 33x more likely to be confident that OSS dependencies are secure (i.e., no known vulnerabilities)
- 4.6x more likely to be confident that OSS licenses of dependencies are compliant with internal requirements
- 2.1x more likely to have access to newer OSS component versions where prior defects have been fixed

Productivity:

- 5.7x less time required to developers to be productive when switching teams
- 26x less time to approve a new OSS dependency for use
- 1.5x more likely for employees to recommend their organizations as a great place to work

When comparing the practices between these clusters, performance differences can be explained by the degree to which governance objectives are automated and integrated into the processes that developers use in their daily work. High performers were.³¹

- 77% more likely to automate approval, management, and analysis of

dependencies

- 59% more likely to be using software composition analysis (SCA) tools
- 28% more likely to enforce governance policies in Continuous Integration (CI)
- 56% more likely to have centrally managed CI infrastructure (where information security governance policies can be enforced)
- 51% more likely to maintain a centralized record of all deployed artifacts, supporting the collection of a Software Bill of Materials (SBOM) for each application
- 96% more likely to be able to centrally scan all deployed artifacts for security and license compliance

Another study that confirms these statistics is by Dr. Dan Geer and Josh Corman, which showed that of the open-source projects with known vulnerabilities registered in the National Vulnerability Database, only 41% were ever fixed and required, on average, 390 days to publish a fix. For those vulnerabilities that were labeled at the highest severity (i.e., those scored as CVSS level 10), fixes required 224 days.³²⁴

CONTINUOUS LEARNING

The 2020 *State of the Octoverse* report showed open-source vulnerability timelines: Across GitHub, a vulnerability typically takes 218 weeks (just over four years) before being disclosed; then it takes about 4.4 weeks for the community to identify and release a fix. From there, it takes ten weeks to alert on the availability of a fix. For the repos that do apply the fix, it typically takes one week to resolve.³³

Two very prominent security breaches in recent years, SolarWinds and Codecov, involved attacks on the software supply chain. In spring 2020, a malicious payload was added to an update to the SolarWinds Orion network management software, which then affected over 18,000 of their customers. The payload used the privileged accounts to corporate network infrastructure to gain unauthorized access, ranging from reading emails to planting something far more destructive.³⁴

In April 2021, a “CI poisoning attack” was discovered in Codecov, a code coverage analysis tool. A malicious payload was added to the Codecov docker image and bash uploader, which stole the credentials from the CI environment. This likely impacted a significant number of their reported 29,000 customers.³⁵

Both of these attacks show how reliant organizations have become on automated

updates, how any CI/CD pipeline can be compromised to insert malicious payloads (this is addressed later in the book), and how new risks can emerge as new development practices are adopted. This serves as another example of how information security must continually examine threats presented by a thinking adversary.

Ensure Security of the Environment

In this step, we should do whatever is required to help ensure that the environments are in a hardened, risk-reduced state. Although we may have created known, good configurations already, we must put in monitoring controls to ensure that all production instances match these known good states.

We do this by generating automated tests to ensure that all appropriate settings have been correctly applied for configuration hardening, database security settings, key lengths, and so forth. Furthermore, we will use tests to scan our environments for known vulnerabilities.**

Another category of security verification is understanding actual environments (i.e., “as they actually are”). Examples of tools for this include Nmap, to ensure that only expected ports are open, and Metasploit, to ensure that we’ve adequately hardened our environments against known vulnerabilities, such as scanning with SQL injection attacks. The output of these tools should be put into our artifact repository and compared with the previous version as part of our functional testing process. Doing this will help us detect any undesirable changes as soon as they occur.

**CASE
STUDY**

18F Automating Compliance for the Federal Government with Compliance Masonry (2016)

US Federal Government agencies were projected to spend nearly \$80 billion on IT in 2016, supporting the mission of all the executive branch agencies.

Regardless of agency, to take any system from “dev complete” to “live in production” requires obtaining an authority to operate (ATO) from a designated approving authority (DAA).

The laws and policies that govern compliance in government are composed of tens of documents that together number over four thousand pages, littered with acronyms such as FISMA, FedRAMP, and FITARA. Even for systems that only require low levels of confidentiality, integrity, and availability, over one hundred controls must be implemented, documented, and tested. It typically takes between eight and fourteen months for an ATO to be granted following “dev complete.”³⁶

The 18F team in the federal government’s General Services Administration has taken a multipronged approach to solving this problem. Mike Bland explains, “18F was created within the General Services Administration to capitalize on the momentum generated by the [Healthcare.gov](#) recovery to reform how the government builds and buys software.”³⁷

One 18F effort is a platform as a service called [Cloud.gov](#), created from open-source components. [Cloud.gov](#) runs on AWS GovCloud in 2016. Not only does the platform handle many of the operational concerns delivery teams might otherwise have to take care of, such as logging, monitoring, alerting, and

service life cycle management, it also handles the bulk of compliance concerns.

By running on this platform, a large majority of the controls that government systems must implement can be taken care of at the infrastructure and platform level. Then, only the remaining controls that are in scope at the application layer have to be documented and tested, significantly reducing the compliance burden and the time it takes to receive an ATO.³⁸

AWS GovCloud has already been approved for use for federal government systems of all types, including those which require high levels of confidentiality, integrity, and availability. [Cloud.gov](#) has been approved for all systems that require moderate levels of confidentiality, integrity, and availability.^{††}

Furthermore, the [Cloud.gov](#) team has built framework to automate the creation of system security plans (SSPs), which are “comprehensive descriptions of the system’s architecture, implemented controls, and general security posture ... [which are] often incredibly complex, running several hundred pages in length.”³⁹ They developed a prototype tool called compliance masonry so that SSP data is stored in machine-readable YAML and then turned into GitBooks and PDFs automatically.

18F is dedicated to working in the open and publishes its work open source in the public domain. You can find compliance masonry and the components that make up [Cloud.gov](#) in 18F’s GitHub repositories—you can even stand up your own instance of [Cloud.gov](#). The work on open documentation for SSPs is being

done in close partnership with the OpenControl community.

This case study demonstrates how an organization—even one as monolithic as the federal government—can use a PaaS to generate automated tests and still meet compliance.

Integrate Information Security into Production Telemetry

Marcus Sachs observed in 2010, that

year after year, in the vast majority of cardholder data breaches, the organization detected the security breach months or quarters after the breach occurred. Worse, the way the breach was detected was not an internal monitoring control but was far more likely someone outside of the organization, usually a business partner or the customer who notices fraudulent transactions. One of the primary reasons for this is that no one in the organization was regularly reviewing the log files.⁴⁰

In other words, internal security controls are often ineffective in successfully detecting breaches in a timely manner, either because of blind spots in our monitoring or because no one in our organization is examining the relevant telemetry in their daily work.

In [Chapter 14](#), we discussed creating a culture in Dev and Ops where everyone in the value stream is creating production telemetry and metrics, making them visible in prominent public places so that everyone can see how our services are performing in production. Furthermore, we explored the necessity of relentlessly seeking ever-weaker failure signals so that we can find and fix problems before they result in a catastrophic failure.

Here, we deploy the monitoring, logging, and alerting required to fulfill our information security objectives throughout our applications and environments, as well as ensure that it is adequately centralized to facilitate easy and meaningful analysis and response.

We do this by integrating our security telemetry into the same tools that Development, QA, and Operations are using, giving everyone in the value stream visibility into how their application and environments are performing in a hostile threat environment where attackers are constantly attempting to exploit vulnerabilities, gain unauthorized access, plant backdoors, commit fraud, perform denials of service, and so forth.

By radiating how our services are being attacked in the production environment, we reinforce that everyone needs to be thinking about security risks and designing countermeasures in their daily work.

Creating Security Telemetry in Our Applications

In order to detect problematic user behavior that could be an indicator or enabler of fraud and unauthorized access, we must create the relevant telemetry in our applications. Examples may include:

- successful and unsuccessful user logins
- user password resets
- user email address resets
- user credit card changes

For instance, as an early indicator of brute-force login attempts to gain unauthorized access, we might display the ratio of unsuccessful login attempts to successful logins. And,

of course, we should create alerting around important events to ensure we can detect and correct issues quickly.

Creating Security Telemetry in Our Environment

In addition to instrumenting our application, we also need to create sufficient telemetry in our environments so that we can detect early indicators of unauthorized access, especially in the components that are running on infrastructure that we do not control (e.g., hosting environments, in the cloud).

We need to monitor and potentially alert on items, including the following:⁴¹

- OS changes (e.g., in production, in our build infrastructure)
- security group changes
- changes to all our production configurations (e.g., OSSEC, Puppet, Chef, Tripwire, Kubernetes, network infrastructure, middleware)
- cloud infrastructure changes (e.g., VPC, security groups, users and privileges)
- XSS attempts (i.e., cross-site scripting attacks)
- SQLi attempts (i.e., SQL injection attacks)
- web server errors (e.g., 4XX and 5XX errors)

We also want to confirm that we've correctly configured our logging so that all telemetry is being sent to the right place. When we detect attacks, in addition to logging that it happened, we may also choose to block access and store information about the source to aid us in choosing the best mitigation actions.

**CASE
STUDY**

Instrumenting the Environment at Etsy (2010)

In 2010, Nick Galbreath was director of engineering at Etsy and responsible for information security, fraud control, and privacy. Galbreath defined *fraud* as when “the system works incorrectly, allowing invalid or un-inspected input into the system, causing financial loss, data loss/theft, system downtime, vandalism, or an attack on another system.”⁴²

To achieve these goals, Galbreath did not create a separate fraud control or information security department; instead, he embedded those responsibilities throughout the DevOps value stream.

Galbreath created security-related telemetry that was displayed alongside all the other more Dev and Ops oriented metrics, which every Etsy engineer routinely saw:⁴³

- **Abnormal production program terminations (e.g., segmentation faults, core dumps, etc.):** “Of particular concern was why certain processes kept dumping core across our entire production environment, triggered from traffic coming from the one IP address, over and over again. Of equal concern were those HTTP ‘500 Internal Server Errors.’ These are indicators that a vulnerability was being exploited to gain unauthorized access to our systems, and that a patch needs to be urgently applied.”⁴⁴
- **Database syntax error:** “We were always looking for database syntax errors inside our code—these either enabled SQL injection attacks or were actual attacks in progress. For this reason, we had zero tolerance for database

syntax errors in our code, because it remains one of the leading attack vectors used to compromise systems.”⁴⁵

- **Indications of SQL injection attacks:** “This was a ridiculously simple test—we’d merely alert whenever ‘UNION ALL’ showed up in user-input fields, since it almost always indicates a SQL injection attack. We also added unit tests to make sure that this type of uncontrolled user input could never be allowed into our database queries.”⁴⁶

Figure 22.5 is an example of a graph that every developer would see, which shows the number of potential SQL injection attacks that were attempted in the production environment.




Figure 22.5: Developers See SQL Injection Attempts in Graphite at Etsy

Source: “DevOpsSec: Applying DevOps Principles to Security, DevOpsDays Austin 2012,” [SlideShare.net](http://www.slideshare.net/nickgsuperstar/devopssec-apply-devops-principles-to-security), posted by Nick Galbreath, April 12, 2012, <http://www.slideshare.net/nickgsuperstar/devopssec-apply-devops-principles-to-security>.

As Galbreath observed, “Nothing helps developers understand how hostile the operating environment is than seeing their code being attacked in real time.”⁴⁷ Galbreath continued, “One of the results of showing this

graph was that developers realized that they were being attacked all the time! And that was awesome, because it changed how developers thought about the security of their code as they were writing the code.”⁴⁸

Presenting security-related telemetry brought security into developers' daily work and made vulnerabilities more visible to everyone.

Protect Our Deployment Pipeline

The infrastructure that supports our continuous integration and continuous deployment processes also presents a new surface area vulnerable to attack. For instance, if someone compromises the servers running the deployment pipeline that has the credentials for our version control system, it could enable someone to steal source code. Worse, if the deployment pipeline has write access, an attacker could also inject malicious changes into our version control repository and, therefore, inject malicious changes into our application and services.

As Jonathan Claudius, former Senior Security Tester at TrustWave SpiderLabs, observed, “Continuous build and test servers are awesome, and I use them myself. But I started thinking about ways to use CI/CD as a way to inject malicious code. Which led to the question: Where would be a good place to hide malicious code? The answer was obvious: in the unit tests. No one actually looks at the unit tests, and they’re run every time someone commits code to the repo.”⁴⁹

This demonstrates that in order to adequately protect the integrity of our applications and environments, we must also mitigate the attack vectors on our deployment pipeline. Risks include developers introducing code that enables unauthorized access (which we’ve mitigated through controls

such as code testing, code reviews, and penetration testing) and unauthorized users gaining access to our code or environment (which we've mitigated through controls such as ensuring configurations match known, good states and through effective patching).

However, in order to protect our continuous build, integration, or deployment pipeline, our mitigation strategies may also include the following:

- hardening continuous build and integration servers and ensuring we can reproduce them in an automated manner, just as we would for infrastructure that supports customer-facing production services, to prevent our continuous build and integration servers from being compromised
- reviewing all changes introduced into version control, either through pair programming at commit time or by a code review process between commit and merge into trunk, to prevent continuous integration servers from running uncontrolled code (e.g., unit tests may contain malicious code that allows or enables unauthorized access)
- instrumenting our repository to detect when test code contains suspicious API calls (e.g., unit tests accessing the file system or network) is checked into the repository, perhaps quarantining it and triggering an immediate code review
- ensuring every CI process runs on its own isolated container or VM, and ensuring this is recreated from a known, good, verified base image at the start of every build
- ensuring the version control credentials used by the CI system are read-only

→ CASE STUDY: NEW TO THE SECOND EDITION

Shifting Security Left at Fannie Mae (2020)

Fannie Mae has a more than \$3 billion balance sheet and helps finance approximately one in four homes in the US as of 2020.⁵⁰ At Fannie Mae, safety and soundness is part of their mission.

They've experienced crises before. With a low risk tolerance, their challenge was ensuring that security strengthened everything they did. DevOps provided a solution to learn from chaos engineering to improve security, put security in the pipeline, and weave security transparently into the fabric of everything they did.

Chris Porter, Fannie Mae's CISO, and Kimberly Johnson, the Executive Vice President and COO, talked about their evolution at the 2020 DevOps Enterprise Summit. It boiled down to two key changes: changing culture and changing the way security communicated with Dev teams and how they integrated security tools.⁵¹

In the old way, Dev would hand off code that was ready for production. Security would conduct their own tests and send back a list of vulnerabilities that the Dev team would have to correct. It was inefficient and no one liked it. They needed to learn to shift security left.

They did this by relinquishing control over their security tools, making them more self-service, making them API-based, and integrating them with Jira and Jenkins. They trained developers to run the tools and to learn what the results meant, and they had to change their own nomenclature (instead of vulnerabilities, they talked about defects).⁵²

They also had to fully integrate all security tests within the CI/CD pipeline so that every time code was checked in they were running a test. Ultimately, this made it easier for developers to know what to do. They could see a test fail, understand why, and fix the problem.⁵³

“I call this the paved road. If you follow the paved road and you use the CI/CD pipeline, which has all the checks integrated into the pipeline, then it will be easier for you to deploy code,” says Chris Porter.⁵⁴

This was treated like an Andon cord. If the test didn’t pass, then it broke the line and had to be fixed before the line could continue. If you didn’t use the paved road, it was a much slower, bumpier journey.

According to Porter, a mindset change is needed from development and security. In the past, security’s mindset had been to protect developers from themselves. But in a DevOps model, the work has moved to “you build it, you own it.”⁵⁵ Everyone has shared accountability, and security is baked into the code instead of being inserted later.

As Kimberly Johnson put it:

In the old way, with Dev handing off production-ready code to Security for testing, we had a major bottleneck in the throughput of the Security team. For large organizations that operate at scale, it can be really hard to find enough Security talent to continually test everything that is developed. Building the security tests into the development pipeline unlocked a lot more productivity for us and reduced our

dependence on Security personnel for standard testing and routine deployments.

In addition to reducing our reliance on the Information Security team, shifting left and automating our testing has yielded better business results. Our deployment frequency has increased by 25% in the last year, and our deployment failure rate has fallen by about the same amount. We are getting critical business changes into production much faster, with fewer errors, using fewer resources, and generating less rework. Moving to DevSecOps has been a win-win-win for us.⁵⁶

By shifting security left, Fannie Mae was able to maintain security and soundness of their code without sacrificing speed, efficiency, and the happiness of their teams.

Conclusion

Throughout this chapter, we have described ways to integrate information security objectives into all stages of our daily work. We do this by integrating security controls into the mechanisms we've already created, ensuring that all on-demand environments are also in a hardened, risk-reduced state—by integrating security testing into the deployment pipeline and ensuring the creation of security telemetry in pre-production and production environments. By doing so, we enable developer and operational productivity to increase while simultaneously increasing our overall safety. Our next step is to protect the deployment pipeline.

* Rugged DevOps traces some of its history to *Visible Ops Security*, written by

Gene Kim, Paul Love, and George Spafford. Similar ideas were created by Dr. Tapabrata Pal, former Sr. Director and Platform Engineering Technical Fellow at Capital One, and the Capital One team; they describe their processes as DevOpsSec and described it as a set of processes where Infosec is integrated into all stages of the SDLC.³

- † Note that all major cloud providers now operate cloud-based secret management systems that provide a good alternative to running your own.
- ‡ The Open Web Application Security Project (OWASP) is a nonprofit organization focused on improving the security of software.
- § Strategies for managing these risks include providing employee training and management; rethinking the design of information systems, including network and software; and instituting processes designed to prevent, detect, and respond to attacks.
- ¶ Tools that can help ensure the integrity of our software dependencies include OWASP Dependency Check and Sonatype Nexus Lifecycle.
- ** Examples of tools that can help with security correctness testing (i.e., “as it should be”) include automated configuration management systems (e.g., Puppet, Chef, Ansible, Salt), as well as tools such as ServerSpec and the Netflix Simian Army (e.g., Conformity Monkey, Security Monkey, etc.).
- †† These approvals are known as FedRAMP JAB P-ATOs.

PROTECTING THE DEPLOYMENT PIPELINE

Throughout this chapter, we will look at how to protect our deployment pipeline, as well as how to achieve security and compliance objectives in our control environment, including change management and separation of duty.

Integrate Security and Compliance into Change Approval Processes

Almost any IT organization of any significant size will have existing change management processes, which are the primary controls to reduce operations and security risks. Compliance managers and security managers place reliance on change management processes for compliance requirements, and they typically require evidence that all changes have been appropriately authorized.

If we have constructed our deployment pipeline correctly so that deployments are low risk, the majority of our changes won't need to go through a manual change approval process, because we will have placed our reliance on controls such as automated testing and proactive production monitoring.

In this step, we will do what is required to ensure that we can successfully integrate security and compliance into any existing change management process. Effective change management policies will recognize that there are different risks associated with different types of changes and that those changes are all handled differently. These processes are defined in ITIL, which breaks changes down into three categories:

Standard changes: These are lower-risk changes that follow an established and approved process but can also be pre-approved. They include monthly updates of application tax tables or country codes, website content and styling changes, and certain types of application or operating system patches that have well-understood impacts. The change proposer does not require approval before deploying the change, and change deployments can be completely automated and should be logged so there is traceability.

Normal changes: These are higher-risk changes that require review or approval from the agreed-upon change authority. In many organizations, this responsibility is inappropriately placed on the change advisory board (CAB) or emergency change advisory board (ECAB), which may lack the required expertise to understand the full impact of the change, often leading to unacceptably long lead times. This problem is especially relevant for large code deployments, which may contain hundreds of thousands (or even millions) of lines of new code, submitted by hundreds of developers over the course of several months. In order for normal changes to be authorized, the CAB will almost certainly have a well-defined request for change (RFC) form that defines what information is required for the go/no-go decision. The RFC form usually includes the desired business outcomes, planned utility and warranty,^{*} a business case with risks and alternatives, and a proposed schedule.[†]

Urgent changes: These are emergency and, consequently, potentially high-risk changes that must be put into production immediately (e.g., urgent security patch, restore service). They often require senior management approval but allow documentation to be performed after the fact. A key goal of DevOps practices is to streamline our normal change process such that it is also suitable for emergency changes.

Recategorize Lower-Risk Changes as Standard Changes

Ideally, by having a reliable deployment pipeline in place, we will have already earned a reputation for fast, reliable, and undramatic deployments. At this point, we should seek to gain agreement from Operations and the relevant change authorities that our changes have been demonstrated to be low risk enough to be defined as standard changes, pre-approved by the CAB. This enables us to deploy into production without need for further approval, although the changes should still be properly recorded.

One way to support an assertion that our changes are low risk is to show a history of changes over a significant time period (e.g., months or quarters) and provide a complete list of production issues during that same period. If we can show high change success rates and low MTTR, we can assert that we have a control environment that is effectively preventing deployment errors, as well as prove that we can effectively and quickly detect and correct any resulting problems.

Even when our changes are categorized as standard changes, they still need to be visual and recorded in our change management systems (e.g., Remedy or ServiceNow). Ideally, deployments will be performed automatically by our configuration management and deployment pipeline tools

and the results will be automatically recorded. By doing this, everyone in our organization (DevOps or not) will have visibility into our changes in addition to all the other changes happening in the organization.

We may automatically link these change request records to specific items in our work planning tools (e.g., JIRA, Rally, LeanKit), allowing us to create more context for our changes, such as linking to feature defects, production incidents, or user stories. This can be accomplished in a lightweight way by including ticket † numbers from planning tools in the comments associated with version control check-ins. By doing this, we can trace a production deployment to the changes in version control and, from there, trace them further back to the planning tool tickets.

Creating this traceability and context should be easy and should not create an overly onerous or time-consuming burden for engineers. Linking to user stories, requirements, or defects is almost certainly sufficient—any further detail, such as opening a ticket for each commit to version control, is likely not useful, and thus unnecessary and undesired, as it will impose a significant level of friction on their daily work.

What to Do When Changes Are Categorized as Normal Changes

For those changes that we cannot get classified as standard changes, they will be considered *normal changes* and will require approval from at least a subset of the CAB before deployment. In this case, our goal is still to ensure that we can deploy quickly, even if it is not fully automated.

In this case, we must ensure that any submitted change requests are as complete and accurate as possible, giving the CAB everything they need to properly evaluate our change—

after all, if our change request is malformed or incomplete, it will be bounced back to us, increasing the time required for us to get into production and casting doubt on whether we actually understand the goals of the change management process.

We can almost certainly automate the creation of complete and accurate RFCs, populating the ticket with details of exactly what is to be changed. For instance, we could automatically create a ServiceNow change ticket with a link to the JIRA user story, along with the build manifests and test output from our deployment pipeline tool and links to the scripts that will be run and the dry run output of these commands.

Because our submitted changes will be manually evaluated by people, it is even more important that we describe the context of the change. This includes identifying why we are making the change (e.g., providing a link to the features, defects, or incidents), who is affected by the change, and what is going to be changed.

Our goal is to share the evidence and artifacts that give us confidence that the change will operate in production as designed. Although RFCs typically have free-form text fields, we should provide links to machine-readable data to enable others to integrate and process our data (e.g., links to JSON files).

In many toolchains, this can be done in a compliant and fully automated way by associating a ticket number with every commit in version control. When we release a new change, we can automatically collate the commits included in that change and then assemble an RFC by enumerating every ticket or bug that was completed or fixed as part of these changes.

Upon submission of our RFC, the relevant members of the CAB will review, process, and approve these changes as they would any other submitted change request. If all goes well, the change authorities will appreciate the thoroughness and detail of our submitted changes because we have allowed them to quickly validate the correctness of the information we've provided (e.g., viewing the links to artifacts from our deployment pipeline tools). However, our goal should be to continually show an exemplary track record of successful changes, so we can eventually gain their agreement that our automated changes can be safely classified as standard changes.

CASE STUDY

Automated Infrastructure Changes as Standard Changes at [Salesforce.com](https://www.salesforce.com) (2012)

Salesforce was founded in 2000 with the aim of making customer relationship management easily available and deliverable as a service. Salesforce's offerings were widely adopted by the marketplace, leading to a successful IPO in 2004.³ By 2007, the company had over fifty-nine thousand enterprise customers, processing hundreds of millions of transactions per day, with an annual revenue of \$497 million.⁴

However, around that same time, their ability to develop and release new functionality to their customers seemed to grind to a halt. In 2006, they had four major customer releases, but in 2007 they were only able to do one customer release despite having hired more engineers.⁵ The result was that the number of features delivered per team kept decreasing and the

days between major releases kept increasing. And because the batch size of each release kept getting larger, the deployment outcomes also kept getting worse.

Karthik Rajan, then VP of Infrastructure Engineering, reports in a 2013 presentation that 2007 marked “the last year when software was created and shipped using a waterfall process and when we made our shift to a more incremental delivery process.”⁶

At the 2014 DevOps Enterprise Summit, Dave Mangot and Reena Mathew described the resulting multiyear DevOps transformation that started in 2009. According to Mangot and Mathew, by implementing DevOps principles and practices, the company reduced their deployment lead times from six days to five minutes by 2013. As a result, they were able to scale capacity more easily, allowing them to process over one billion transactions per day.⁷

One of the main themes of the Salesforce transformation was to make quality engineering everyone’s job, regardless of whether they were part of Development, Operations, or Infosec. To do this, they integrated automated testing into all stages of the application and environment creation, as well as into the continuous integration and deployment process, and created the open-source tool Rouster to conduct functional testing of their Puppet modules.⁸

They also started to routinely perform *destructive testing*, a term used in manufacturing to refer to performing prolonged endurance testing under the most severe operating conditions until the component being tested is destroyed. The Salesforce team started

routinely testing their services under increasingly higher loads until the service broke, which helped them understand their failure modes and make appropriate corrections. Unsurprisingly, the result was significantly higher service quality with normal production loads.⁹

Information Security also worked with Quality Engineering at the earliest stages of their project, continually collaborating in critical phases such as architecture and test design, as well as properly integrating security tools into the automated testing process.¹⁰

For Mangot and Mathew, one of the key successes from all the repeatability and rigor they designed into the process was being told by their change management group that “infrastructure changes made through Puppet would now be treated as ‘standard changes,’ requiring far less or even no further approvals from the CAB.” However, they noted that “manual changes to infrastructure would still require approvals.”¹¹

Salesforce not only integrated their DevOps processes with the change management process but also created further motivation to automate the change process for more of their infrastructure.

Implement Separation of Duty through Code Review

For decades, we have used separation of duty as one of our primary controls to reduce the risk of fraud or mistakes in the software development process. It has been the accepted practice in most SDLCs to require developer changes to be submitted to a code librarian, who would review and approve

the change before IT Operations promoted the change into production.

There are plenty of other less contentious examples of separation of duty in Ops work, such as server administrators ideally being able to view logs but not delete or modify them, in order to prevent someone with privileged access from deleting evidence of fraud or other issues.

When we did production deployments less frequently (e.g., annually) and when our work was less complex, compartmentalizing our work and doing handoffs were tenable ways of conducting business. However, as complexity and deployment frequency increase, performing successful production deployments increasingly requires everyone in the value stream to quickly see the outcomes of their actions.

The traditional approach to implementing separation of duty can often impede this by slowing down and reducing the feedback engineers receive on their work. This prevents engineers from taking full responsibility for the quality of their work and reduces a firm's ability to create organizational learning.

Consequently, wherever possible, we should implement separation of duties as a control. Instead, we should choose controls such as pair programming, continuous inspection of code check-ins, and code review. These controls can give us the necessary reassurance about the quality of our work. Furthermore, by putting these controls in place, if separation of duties is required, we can show that we achieve equivalent outcomes with the controls we have created.

CASE STUDY

PCI Compliance and a Cautionary Tale of Separating Duties at Etsy (2014)[§]

Bill Massie is a development manager at Etsy and is responsible for the payment application called ICHT (an abbreviation for “I Can Haz Tokens”). ICHT takes customer credit orders through a set of internally developed payment processing applications that handle online order entry by taking customer-entered cardholder data, tokenizing it, communicating with the payment processor, and completing the order transaction.¹²

Because the scope of the Payment Card Industry Data Security Standards (PCI DSS) cardholder data environment (CDE) is “the people, processes and technology that store, process or transmit cardholder data or sensitive authentication data,” including any connected system components, the ICHT application is in scope for PCI DSS.¹³

To contain the PCI DSS scope, the ICHT application is physically and logically separated from the rest of the Etsy organization and is managed by a completely separate application team of developers, database engineers, networking engineers, and ops engineers. Each team member is issued two laptops: one for ICHT (which is configured differently to meet the DSS requirements, and is locked in a safe when not in use) and one for the rest of Etsy.

By doing this, they were able to decouple the CDE environment from the rest of the Etsy organization, limiting the scope of the PCI DSS regulations to one segregated area. The systems that form the CDE are separated (and managed differently) from the rest of

Etsy's environments at the physical, network, source code, and logical infrastructure levels. Furthermore, the CDE is built and operated by a cross-functional team that is solely responsible for the CDE.

The ICHT team had to modify their continuous delivery practices in order to accommodate the need for code approvals. According to Section 6.3.2 of the PCI DSS v3.1, teams should review all custom code prior to release to production or customers in order to identify any potential coding vulnerability (using either manual or automated processes) as follows:¹⁴

- Are code changes reviewed by individuals other than the originating code author, and by individuals knowledgeable about code-review techniques and secure coding practices?
- Do code reviews ensure code is developed according to secure coding guidelines?
- Are appropriate corrections implemented prior to release?
- Are code review results reviewed and approved by management prior to release?

To fulfill this requirement, the team initially decided to designate Massie as the change approver responsible for deploying any changes into production. Desired deployments would be flagged in JIRA, and Massie would mark them as reviewed and approved and manually deploy them into the ICHT production.¹⁵

This has enabled Etsy to meet their PCI DSS requirements and get their signed Report of Compliance from their assessors. However, with regard to the team, significant problems have resulted.

Massie observes that one troubling sideeffect “is a level of ‘compartmentalization’ that is happening in the ICHT team that no other group is having at Etsy. Ever since we implemented separation of duty and other controls required by the PCI DSS compliance, no one can be a full-stack engineer in this environment.”¹⁶

As a result, while the rest of the Development and Operations teams at Etsy work together closely and deploy changes smoothly and with confidence, Massie notes that:

within our PCI environment, there is fear and reluctance around deployment and maintenance because no one has visibility outside their portion of the software stack. The seemingly minor changes we made to the way we work seem to have created an impenetrable wall between developers and ops, and creates an undeniable tension that no one at Etsy has had since 2008. Even if you have confidence in your portion, it’s impossible to get confidence that someone else’s change isn’t going to break your part of the stack.¹⁷

This case study shows that compliance is possible in organizations using DevOps. However, the potentially cautionary tale here is that all the virtues that we associate with high-performing DevOps teams are fragile—even a team that has shared experiences with high trust and shared goals can begin to struggle when low-trust control mechanisms are put into place.

→ CASE STUDY: NEW TO THE SECOND EDITION

Biz and Tech Partnership toward Ten “No Fear Releases” Per Day at Capital One (2020)

Over the last seven years, Capital One has been undergoing an Agile/DevOps transformation. In that time, they’ve moved from waterfall to Agile, from outsource to insource and open-sources, from monolithic to microservices, from data centers to the cloud.

But they were still facing a big problem: an aging customer servicing platform. This platform serviced tens of millions of Capital One credit card customers and generated hundreds of millions of dollars in value to the business.¹⁸ It was a critical platform, but it was showing its age and was no longer meeting customer needs or the internal strategic needs of the company. They needed to not only solve the technology/cyber-risk problem of the aging platform but also increase the NPV (net present value) of the system.

“What we had was a mainframe-based vendor product that had been bandaged to the point where the systems and operational teams were as large as the product itself... . We needed a modern system to deliver on the business problem,” says Rakesh Goyal, Director, Technology Engineering at Capital One.¹⁹

They started with a set of principles to work from. First, they worked backwards from the customer’s needs. Second, they were determined to deliver value iteratively to maximize learnings and minimize risk. And third, they wanted to avoid anchoring bias. That is, they

wanted to make sure they were not just building a faster and stronger horse but actually solving a problem.²⁰

With these guiding principles in place, they set about making changes. First, they took a look at their platform and the set of customers. Then they divided them into segments based on what their needs were and what functionalities they needed. Importantly, they thought strategically about who their customers were, because it wasn't just credit card holders. Their customers were regulators, business analysts, internal employees who used the system, etc.

“We use very heavy human-centered design to ensure that we are actually meeting the needs [of our customers] and not just replicating what was there in the old system,” says Biswanath Bosu, Senior Business Director, Anti-Money Laundering-Machine Learning and Fraud at Captial One.²¹

Next they graded these segments on the sequence in which they would be deployed. Each segment represented a thin slice that they could experiment with, see what worked and what didn't, and then iterate from there.

“As much as we were looking for an MVP [minimum viable product], we were not looking for the least common denominator. We were looking for the minimum viable experience that we could give to our customers, not just any small product we could come up with. Once we test that piece out and it works, the next thing we will do is just essentially scale it up,” explained Bosu.²²

As part of the platform transformation, it was clear they would need to move to the cloud. They would also

need to invest in and evolve the tools in their toolbox, as well as invest in reskilling for their engineers to provide them with the appropriate tooling to be agile during this transformation.

They settled on building an API-driven microservice-based architecture system. The goal was to sustain and build it incrementally, slowly expanding into various business strategies.

“You can think of this as having a fleet of smart cars built for specific workloads rather than one futuristic car,” describes Goyal.²³

They began by leveraging proven enterprise tools. By standardizing, they could react faster to situations where engineers needed to contribute to other teams or move from one team to another.

Building out their CI/CD pipeline enabled incremental releases and empowered teams by reducing cycle time and risk. As a financial institution, they also had to address regulatory and compliance controls. Using the pipeline, they were able to block releases when certain controls were not met.

The pipeline also allowed teams to focus on product features, since the pipeline was a tool to leverage rather than a required investment from each team. At the height of their effort, they had twenty-five teams working and contributing simultaneously.

Focusing on customer needs and building a CI/CD pipeline helped Capital One not only meet business needs but move faster.

Ensure Documentation and Proof for Auditors and Compliance Officers

As technology organizations increasingly adopt DevOps patterns, there is more tension than ever between IT and Audit. These new DevOps patterns challenge traditional thinking about auditing, controls, and risk mitigation.

As Bill Shinn, a principal security solutions architect at Amazon Web Services, observes,

DevOps is all about bridging the gap between Dev and Ops. In some ways, the challenge of bridging the gap between DevOps and auditors and compliance officers is even larger. For instance, how many auditors can read code and how many developers have read NIST 800-37 or the Gramm-Leach-Bliley Act? That creates a gap of knowledge, and the DevOps community needs to help bridge that gap.²⁴

CASE STUDY

Proving Compliance in Regulated Environments (2015)

Helping large enterprise customers show that they can still comply with all relevant laws and regulations is among Bill Shinn's responsibilities as a principal security solutions architect at Amazon Web Services. Over the years, he has spent time with over one thousand enterprise customers, including Hearst Media, GE, Phillips, and Pacific Life, who have publicly referenced their use of public clouds in highly regulated environments.

Shinn notes, "One of the problems is that auditors have been trained in methods that aren't very suitable

for DevOps work patterns. For example, if an auditor saw an environment with ten thousand production servers, they have been traditionally trained to ask for a sample of one thousand servers, along with screenshot evidence of asset management, access control settings, agent installations, server logs, and so forth.”²⁵

“That was fine with physical environments,” Shinn continues. “But when infrastructure is code, and when auto-scaling makes servers appear and disappear all the time, how do you sample that? You run into the same problems when you have a deployment pipeline, which is very different than the traditional software development process, where one group writes the code and another group deploys that code into production.”²⁶

He explains, “In audit fieldwork, the most commonplace methods of gathering evidence are still screenshots and CSV files filled with configuration settings and logs. Our goal is to create alternative methods of presenting the data that clearly show auditors that our controls are operating and effective.”²⁷

To help bridge that gap, he has teams work with auditors in the control design process. They use an iterative approach, assigning a single control for each sprint to determine what is needed in terms of audit evidence. This has helped ensure that auditors get the information they need when the service is in production, entirely on demand.

Shinn states that the best way to accomplish this is to “send all data into our telemetry systems, such as Splunk or Kibana. This way auditors can get what they need, completely self-serviced. They don’t need to request a data sample—instead, they log into Kibana,

and then search for audit evidence they need for a given time range. Ideally, they'll see very quickly that there's evidence to support that our controls are working.”²⁸

Shinn continues, “With modern audit logging, chat rooms, and deployment pipelines, there's unprecedented visibility and transparency into what's happening in production, especially compared to how Operations used to be done, with far lower probability of errors and security flaws being introduced. So, the challenge is to turn all that evidence into something an auditor recognizes.”²⁹

That requires deriving the engineering requirements from the actual regulations. Shinn explains,

To discover what HIPAA requires from an information security perspective, you have to look into the forty-five CFR Part 160 legislation, go into Subparts A and C of Part 164. Even then, you need to keep reading until you get into ‘technical safeguards and audit controls.’ Only there will you see that what is required is that we need to determine activities that will be tracked and audited relevant to Patient Healthcare Information, document and implement those controls, select tools, and then finally review and capture the appropriate information.³⁰

Shinn continues, “How to fulfill that requirement is the discussion that needs to be happening between compliance and regulatory officers, and the security and DevOps teams, specifically around how to prevent, detect, and correct problems. Sometimes they can be

fulfilled in a configuration setting in version control. Other times, it's a monitoring control."³¹

Shinn gives an example: "We may choose to implement one of those controls using AWS CloudWatch, and we can test that the control is operating with one command line. Furthermore, we need to show where the logs are going—in the ideal, we push all this into our logging framework, where we can link the audit evidence with the actual control requirement."³²

To help solve this problem, the *DevOps Audit Defense Toolkit* describes the end-to-end narrative of the compliance and audit process for a fictitious organization (Parts Unlimited from *The Phoenix Project*). It starts by describing the entity's organizational goals, business processes, top risks, and resulting control environment, as well as how management could successfully prove that controls exist and are effective. A set of audit objections is also presented, as well as how to overcome them.³³

The toolkit describes how controls could be designed in a deployment pipeline to mitigate the stated risks, and provides examples of control attestations and control artifacts to demonstrate control effectiveness. It was intended to be general to all control objectives, including in support of accurate financial reporting, regulatory compliance (e.g., SEC SOX-404, HIPAA, FedRAMP, EU Model Contracts, and the proposed SEC Reg-SCI regulations), contractual obligations (e.g., PCI DSS, DOD DISA), and effective and efficient operations.

This case study shows how building documentation helps bridge the gap between Dev and Ops practices and auditor requirements, showing DevOps can comply with requirements and improve risk assessment and mitigation.

CASE STUDY

Relying on Production Telemetry for ATM Systems (2013)

Mary Smith (a pseudonym) heads up the DevOps initiative for the consumer banking property of a large US financial services organization. She made the observation that Information Security, auditors, and regulators often put too much reliance on code reviews to detect fraud. Instead, they should be relying on production monitoring controls in addition to using automated testing, code reviews, and approvals to effectively mitigate the risks associated with errors and fraud.³⁴

She observed:

Many years ago, we had a developer who planted a backdoor in the code that we deploy to our ATM cash machines. They were able to put the ATMs into maintenance mode at certain times, allowing them to take cash out of the machines. We were able to detect the fraud very quickly, and it wasn't through a code review. These types of backdoors are difficult, or even impossible, to detect when the perpetrators have sufficient means, motive, and opportunity.

However, we quickly detected the fraud during our regular operations review meeting when someone noticed that ATMs in a city were being put into maintenance mode at unscheduled times. We found the fraud even before the scheduled cash audit process, when they reconcile the amount of cash in the ATMs with authorized transactions.³⁵

In this case study, the fraud occurred despite separation of duties between Development and Operations and a change approval process, but it was quickly detected and corrected through effective production telemetry.

As this case study demonstrates, auditors' overreliance on code reviews and separation of duties between Dev and Ops can leave vulnerabilities. Telemetry helps provide the necessary visibility to detect and act upon errors and fraud, mitigating the perceived need to separate duties or create additional layers of change review boards.

Conclusion

Throughout this chapter, we have discussed practices that make information security everyone's job, where all of our information security objectives are integrated into the daily work of everyone in the value stream. By doing this, we significantly improve the effectiveness of our controls so that we can better prevent security breaches, as well as detect and recover from them faster. And we significantly reduce the work associated with preparing and passing compliance audits.

- * ITIL defines *utility* as “what the service does,” while *warranty* is defined as “how the service is delivered and can be used to determine whether a service is ‘fit for use.’”¹
- † To further manage risk changes, we may also have defined rules, such as certain changes can only be implemented by a certain group or individual (e.g., only DBAs can deploy database schema changes). Traditionally, the CAB meetings have been held weekly, where the change requests are approved and scheduled. From ITIL Version 3 onward, it is acceptable for changes to be approved electronically in a just-in-time fashion through a change management tool. It also specifically recommends that “standard changes should be identified early on when building the Change Management process to promote efficiency. Otherwise, a Change Management implementation can create unnecessarily high levels of administration and resistance to the Change Management process.”²
- ‡ The term *ticket* is used generically to indicate any uniquely identifiable work item.
- § The authors thank Bill Massie and John Allspaw for spending an entire day with Gene Kim, sharing their compliance experience.

PART VI CONCLUSION

Throughout the previous chapters, we explored how to take DevOps principles and apply them to Information Security, helping us achieve our goals and making sure security is a part of everyone's job, every day. Better security ensures that we are defensible and sensible with our data, that we can recover from security problems before they become catastrophic, and, most importantly, that we can make the security of our systems and data better than ever.

Additional Resources

You can dive deeper into considerations of DevOps and audit with this amazing auditor's panel for the 2019 DevOps Enterprise Summit. In it, representatives from each of the four big auditing firms took the time to talk about how DevOps and audit can work together (<https://videolibrary.doesvirtual.com/?video=485153001>).

Sooner Safer Happier: Antipatterns and Patterns for Business Agility has an excellent chapter on building intelligent control, outlining clear patterns and antipatterns for dealing with highly regulated industries. The authors all come from the banking industry and thus have a lot of hard-won experience.

Safety Differently: Human Factors for a New Era by Sidney Dekker addresses how to turn safety from bureaucratic accountability back into an ethical responsibility, and it embraces the human factor not as a problem to control but as a solution to harness.

You can also view a lecture from Dekker on the subject here: <https://www.youtube.com/watch?v=oMtLS0FNDZs>.

A CALL TO ACTION: CONCLUSION TO THE DEVOPS HANDBOOK

We have come to the end of a detailed exploration of both the principles and technical practices of DevOps. At a time when every technologist and technology leader is challenged with enabling security, reliability, and agility, and at a time when security breaches, time to market, and massive technology transformation is taking place, DevOps offers a solution. Hopefully, this book has provided an in-depth understanding of the problem and a road map to creating relevant solutions.

As we have explored throughout *The DevOps Handbook*, we know that, left unmanaged, an inherent conflict can exist between Development and Operations that creates ever-worsening problems, which result in slower time to market for new products and features, poor quality, increased outages and technical debt, reduced engineering productivity, as well as increased employee dissatisfaction and burnout.

DevOps principles and patterns enable us to break this core, chronic conflict. After reading this book, we hope you see how a DevOps transformation can enable the creation of dynamic learning organizations, achieving the amazing outcomes of fast flow and world-class reliability and security, as well as increased competitiveness and employee satisfaction.

DevOps requires potentially new cultural and management norms and changes in our technical practices and architecture. This requires a coalition that spans business leadership, Product Management, Development, QA, IT Operations, Information Security, and even Marketing, where many technology initiatives originate. When all these teams work together, we can create a safe system of work,

enabling small teams to quickly and independently develop and validate code that can be safely deployed to customers. This results in maximizing developer productivity, organizational learning, high employee satisfaction, and the ability to win in the marketplace.

Our goal in writing this book was to sufficiently codify DevOps principles and practices so that the amazing outcomes achieved within the DevOps community could be replicated by others. We hope to accelerate the adoption of DevOps initiatives and support their successful implementations while lowering the activation energy required for them to be completed.

We know the dangers of postponing improvements and settling for daily workarounds, as well as the difficulties of changing how we prioritize and perform our daily work. Furthermore, we understand the risks and effort required to get organizations to embrace a different way of working, as well as the perception that DevOps is another passing fad, soon to be replaced by the next buzzword.

We assert that DevOps is transformational to how we perform technology work, just as Lean forever transformed how manufacturing work was performed in the 1980s. Those that adopt DevOps will win in the marketplace at the expense of those that do not. They will create energized and continually learning organizations that outperform and out-innovate their competitors.

Because of this, DevOps is not just a technology imperative but an organizational imperative. The bottom line is, DevOps is applicable and relevant to any and all organizations that must increase flow of planned work through the technology organization, while maintaining quality, reliability, and security for our customers.

Our call to action is this: no matter what role you play in your organization, start finding people around you who want to change how work is performed. Show this book to others and create a coalition of like-minded thinkers to break out of the downward spiral. Ask organizational leaders to support these efforts or, better yet, sponsor and lead these efforts yourself.

Finally, since you've made it this far, we have a dirty secret to reveal. In many of our case studies, following the achievement of the breakthrough results presented, many of the change agents were promoted—but, in some cases, there was later a change of leadership which resulted in many of the people involved leaving, accompanied by a rolling back of the organizational changes they had created.

We believe it's important not to be cynical about this possibility. The people involved in these transformations knew up front that what they were doing had a high chance of failure, and they did it anyway. In doing so, perhaps most importantly, they inspired the rest of us by showing us what can be done. Innovation is impossible without risk-taking, and if you haven't managed to upset at least some people in management, you're probably not trying hard enough. Don't let your organization's immune system deter or distract you from your vision. As Jesse Robbins, previously “master of disaster” at Amazon, likes to say, “Don't fight stupid, make more awesome.”¹

DevOps benefits all of us in the technology value stream, whether we are Dev, Ops, QA, Infosec, Product Owners, or customers. It brings joy back to developing great products, with fewer death marches. It enables humane work conditions with fewer weekends worked and fewer missed holidays with our loved ones. It enables teams to work

together to survive, learn, thrive, delight our customers, and help our organization succeed.

We sincerely hope *The DevOps Handbook* helps you achieve these goals.

AFTERWORD TO THE SECOND EDITION

Nicole Forsgren

The questions I am asked most often—by leaders and developers alike—are about productivity and performance: How can we help our teams develop and deliver software more effectively? How can we improve developer productivity? Are performance improvements sustainable, or are we simply making tradeoffs? How should we measure and track these improvements?

Data and experience have reinforced the importance of using good automation, strategic process, and a culture that prioritizes trust and information flow to help teams achieve high software delivery performance. Even during the COVID-19 pandemic, teams and organizations with smart automation, flexible processes, and good communication were able to not only survive but grow and expand, with some pivoting in just days or weeks to serve new customers and markets.

GitHub's 2020 *State of the Octoverse* report found that developers spent more time at work, logging 4.2 to 4.7 more hours per day across the four time zones studied compared to the previous year.^{1*} And developers weren't just spreading out their work to accommodate housework or childcare; using push volume as a proxy for amount of work, the data showed that developers were also doing more. Developers pushed ten to seventeen more commits per day to their main branch during the workweek compared to a year before. Just as the data shows enterprise activity slowing on the weekend, open-source activity picks up, suggesting that developers are logging off of work and logging in to open source. Open-source project creation is up 25% since April 2020 year over year.




Figure AF.1: Average Development Window by Day of Week per User

Source: Forsgren et. al., 2020 *State of the Octoverse*.

While these statistics are impressive, and our ability to continue to innovate and deliver software in the face of a pandemic admirable, we must also step back to consider broader patterns. Simply pushing hard and delivering results when the conditions don't typically allow for it can mask underlying issues.

A recent study from Microsoft reports that “high productivity is masking an exhausted workforce.”² Those of us who have worked in technology for years may recognize the patterns and know they aren't sustainable; true improvement and transformation requires both improvement and balance. We need to ensure that the lessons we learned from our old ways of work—long days, brute force, adrenaline-fueled delivery schedules—aren't merely replicated now that we have better technology and methods.

The data and patterns above also highlight another important point: activity metrics like number of hours worked or number of commits don't tell the full story. Teams and organizations that only measure productivity using these shallow metrics risk missing the full picture, one that experienced and expert technology leaders already know: productivity is complex and must be measured holistically.

Based on decades of collected expertise and research, my colleagues and I recently published the SPACE framework to help developers, teams, and leaders think about and measure productivity.³ It includes five dimensions: satisfaction and well-being, performance, activity, communication and collaboration, and efficiency and flow. By including measures across at least three dimensions in the framework, teams and organizations can measure developer productivity in ways that more accurately capture productivity, will have a better understanding of how individuals and teams work, and will have superior information to enable better decisions.

For example, if you're already measuring commits (an activity measure), don't simply add the number of pull requests to your metrics dashboard, as this is another activity metric. To help capture productivity, add at least one metric from two different dimensions: perhaps satisfaction with the engineering system (a satisfaction measure and an important measure of developer experience) and pull request merge time (an efficiency and flow measure). By adding these two metrics, we can now see an indication of number of commits per individual or team and pull request merge time per individual or team—as well as how those balance with development time to ensure that reviews aren't interrupting coding time—and gain insight into how the engineering system is supporting the overall development and delivery pipeline.

We can see how this gives us much more insight than merely the number of commits and helps us make better decisions to support our development team. These metrics can also ensure more sustainable development and well-being for our developers by surfacing problems early and highlighting sustainability of our tools and inherent trade-offs our teams may be making.

As we look back over the past decade, it's rewarding to know that our improved processes, technology, and methods of working and communicating have enabled our teams to develop and deliver software at levels we probably never imagined were possible—even in the face of daunting and unforeseen change. At the same time, it's humbling to realize the continued (and perhaps even greater) responsibility we have to ensure our improvement journey continues. The opportunities are exciting, and we wish you well on the journey.

Gene Kim

I continue to be so inspired by how technology leaders are creating better ways of working, enabling the creation of value sooner, safer, and happier, as framed so wonderfully by Jon Smart. And I am so delighted by how the second edition of this book contains so many new case studies (and so many of them coming from the DevOps Enterprise community). The fact that these case studies come from so many different industry verticals is another testament to the universality of the problem DevOps helps to solve.

One of the most wonderful things I've observed is how more and more of these experience reports are given as joint presentations between the technology leader and their business-leader counterpart, clearly articulating how the achievement of their goals, dreams, and aspirations are made

possible through creating a world-class technology organization.

In one of the last pages in *The Phoenix Project*, the Yoda-like Erik character predicts that technology capabilities not only need to be a core competency in almost every organization but that they also need to be embedded throughout the organization, closest to where customer problems are being solved.

It's exhilarating to see evidence of these predictions coming true, and I look forward to technology helping every organization win, fully supported by the very highest levels in the organization.

Jez Humble

I think of DevOps as a movement of people who are working out how to build secure, rapidly changing, resilient distributed systems at scale. This movement was born from seeds planted by developers, testers, and sysadmins many years before, but really took off with the enormous growth of digital platforms. In the last five years, DevOps has become ubiquitous.

While I think we've learned a lot as a community over those years, I also see many of the same problems that have dogged the technology industry repeat themselves, which can basically be reduced to the fact that sustained process improvement, architectural evolution, culture change, and teamwork of the kind that produces a lasting impact is hard. It's easier to focus on tools and organizational structure—and those things are important, but they're not enough.

Since the book came out, I've used the practices we describe in the US Federal Government, in a four-person startup, and in Google. I've had discussions about them with people all over the world, thanks to the wonderful DevOps

community. I've been part of a team, led by Dr. Nicole Forsgren—who I'm delighted has contributed to this second edition—that's done world-leading research into how to build high-performing teams.

If I've learned anything, it's that high performance starts with organizations whose leadership focuses on building an environment where people from different backgrounds and with different identities, experiences, and perspectives can feel psychologically safe working together, and where teams are given the necessary resources, capacity, and encouragement to experiment and learn together in a safe and systematic way.

The world is changing and evolving constantly. While organizations and teams will come and go, we in our community have the responsibility of taking care of and supporting each other and sharing what we learn. That's the future of DevOps, and it's also the challenge of DevOps. My deepest gratitude goes to this community, especially those of you who do the critical work of creating psychologically safe environments and welcoming and encouraging new people of all backgrounds. I can't wait to see what you learn and share with us.

Patrick Debois

Initially, I saw DevOps as a way to improve the bottleneck between Development and Operations.

After running my own business, I understand there are many other groups in a company influencing this relationship. For example, when Marketing and Sales overpromise, this puts a lot of stress on the whole relationship. Or when HR keeps hiring the wrong people or bonuses are misaligned. This made me look at DevOps as a way of finding bottlenecks even at a higher level within a company.

Since the term was first coined, I've settled on my own definition of "DevOps": everything you do to overcome the friction between silos. All the rest is plain engineering.

This definition highlights that just building technology is not enough; you need to have the intent to overcome a friction point. And that friction point will move once you have improved the blocker. This continuous evaluation of the bottleneck is crucial.

Indeed, organizations keep optimizing the pipeline/automation but don't work on other friction points that are causing bottlenecks. This is a key challenge today. It's interesting how concepts such as FinOps, for example, have put more pressure on collaboration, or even improvements on a personal level, to better understand and articulate what people need and want. That broader view of improvement and thinking beyond the pipeline/automation is where most people/organizations get stuck.

As we move forward, I believe we'll keep seeing other bottlenecks being addressed under the umbrella of DevOps—DevSecOps is a good example of this, the bottleneck just moving somewhere else. I've seen people mention DesignOps, AIOps, FrontendOps, DataOps, NetworkOps ... all labels to balance the attention points of things to keep in mind.

At some point, it will not matter anymore if it's called DevOps or not. The idea that organizations must keep optimizing will just happen naturally. So I actually hope the future of DevOps is that nobody talks about the term anymore and just keeps improving the practices, mainly because we have reached peak inflation of the term.

John Willis

I met Gene about ten years ago, and he told me about a book he was writing based on Dr. Eliyahu Goldratt's, *The Goal*. That book was *The Phoenix Project*. At that time I knew very little about operations management, supply chain, and Lean. Gene had told me that he was working on another book as a follow up to *The Phoenix Project* that was going to be a more prescriptive work and that it was already in progress with my good friend Patrick Debois. I immediately begged to join the project, which ultimately became *The DevOps Handbook*. Early on we were basically focused on best-practice stories and less about the aforementioned meta concepts. Later on, Jez Humble joined the team and really added significant depth to the project in those areas.

However, to be honest, it took me over ten years to appreciate the real impact of operations management, supply chain, and Lean on what we call DevOps. The more I understand what happened in the Japanese manufacturing economy from 1950 to 1980, the more I realize the root impact on current knowledge economies. In fact, now there seems to be an interesting möbius loop between manufacturing and knowledge economy learnings. One example is autonomous vehicle production. Movements like Industrial DevOps are a great example of the loop between what we learned from manufacturing to knowledge and back to manufacturing economies.

One of the greatest challenges today is that most legacy organizations straddle two worlds. Their first-world habits are historical, systemic, and driven by calcified capital market forces. Their second-world emerging habits, like DevOps, are mostly counterintuitive to their first-world habits. Typically, these organizations struggle between their two worlds like two tectonic plates colliding, creating subduction between one and the other. These collisions tend to create temporal

successes with an organization, constantly bouncing between second-world success and second-world subduction.

The good news is, the evolution of movements like Lean, Agile, DevOps, and DevSecOps seem to trend favorably toward second-world habits. Over time, organizations that embrace these new habits as a true north tend to have more successes than subductions.

One of the bright spots over the past few years has been a reductionist approach to technology. Although technology is only one of the three main principles of high performance organization success (people, process, technology), it doesn't hurt when one of the three reduces enormous toil.

Over the past few years, we have seen less reliance on older legacy infrastructure. More than just the cloud, there is an increasing gravity to more atomic styles of computing. We see large institutions moving rapidly into cluster-based and function-style computing with an increasing emphasis on event-driven architectures (EDA). These reductionist-style technological processes leave far less room for aggregate toil across all three of the principles (people, process, technology). This, combined with the aforementioned evolutionary second-world habits, will accelerate more success than losses for larger legacy organizations.

* The time zones studied were UK, US Eastern, US Pacific, and Japan Standard.

APPENDICES

Appendix 1: The Convergence of DevOps

We believe that DevOps is benefiting from an incredible convergence of management movements, which are all mutually reinforcing and can help create a powerful coalition to transform how organizations develop and deliver IT products and services.

John Willis named this “the convergence of DevOps.” The various elements of this convergence are described below in approximate chronological order. (Note that these descriptions are not intended to be exhaustive, but merely enough to show the progression of thinking and the rather improbable connections that led to DevOps.)

The Lean Movement

The Lean Movement started in the 1980s as an attempt to codify the Toyota Production System with the popularization of techniques such as value stream mapping, kanban boards, and total productive maintenance.¹

Two major tenets of Lean were the deeply held belief that lead time (i.e., the time required to convert raw materials into finished goods) was the best predictor of quality, customer satisfaction, and employee happiness; and that one of the best predictors of short lead times was small batch sizes, with the theoretical ideal being “single piece flow” (i.e., “1x1” flow: inventory of 1, batch size of 1).

Lean principles focus on creating value for the customer —thinking systematically, creating constancy of purpose, embracing scientific thinking, creating flow and pull (versus push), assuring quality at the source, leading with humility, and respecting every individual.

The Agile Movement

Started in 2001, the Agile Manifesto was created by seventeen of the leading thinkers in software development, with the goal of turning lightweight methods such as DP and DSDM into a wider movement that could take on heavyweight software development processes such as waterfall development and methodologies such as the rational unified process.

A key principle was to “deliver working software frequently, from a couple of weeks to a couple of months, with a preference to the shorter timescale.”² Two other principles focus on the need for small, self-motivated teams, working in a high-trust management model and an emphasis on small batch sizes. Agile is also associated with a set of tools and practices such as Scrum, Standups, and so on.

The Velocity Conference Movement

Started in 2007, the Velocity Conference was created by Steve Souders, John Allspaw, and Jesse Robbins to provide a home for the IT Operations and Web Performance tribe. At the Velocity 2009 conference, John Allspaw and Paul Hammond gave the seminal “10 Deploys per Day: Dev and Ops Cooperation at Flickr,” presentation.

The Agile Infrastructure Movement

At the 2008 Agile Toronto conference, Patrick Debois and Andrew Shafer held a “birds of a feather” session on applying Agile principles to infrastructure as opposed to application code. They rapidly gained a following of like-minded thinkers, including John Willis. Later, Debois was so excited by Allspaw and Hammond’s “10 Deploys per Day: Dev and Ops Cooperation at Flickr” presentation that he created the first

DevOpsDays in Ghent, Belgium, in 2009, coining the word “DevOps.”

The Continuous Delivery Movement

Building upon the Development discipline of continuous build, test, and integration, Jez Humble and David Farley extended the concept of continuous delivery, which included a “deployment pipeline” to ensure that code and infrastructure are always in a deployable state and that all code checked into trunk is deployed into production.³ This idea was first presented at Agile 2006 and was also independently developed by Tim Fitz in a blog post titled “Continuous Deployment.”⁴

The Toyota Kata Movement

In 2009, Mike Rother wrote *Toyota Kata: Managing People for Improvement, Adaptiveness and Superior Results*, which described learnings over his twenty-year journey to understand and codify the causal mechanisms of the Toyota Production System. *Toyota Kata* describes the “unseen managerial routines and thinking that lie behind Toyota’s success with continuous improvement and adaptation ... and how other companies develop similar routines and thinking in their organizations.”⁵

His conclusion was that the Lean community missed the most important practice of all, which he described as the Improvement Kata. He explains that every organization has work routines, and the critical factor in Toyota was making improvement work habitual and building it into the daily work of everyone in the organization. The Toyota Kata institutes an iterative, incremental, scientific approach to problem-solving in the pursuit of a shared organizational true north.⁶

The Lean Startup Movement

In 2011, Eric Ries wrote *The Lean Startup: How Today's Entrepreneurs Use Continuous Innovation to Create Radically Successful Businesses*, codifying the lessons he learned at IMVU, a Silicon Valley startup, which built upon the work of Steve Blank in *The Four Steps to the Epiphany* as well as continuous deployment techniques. Eric Ries also codified related practices and terms including minimum viable product, the build-measure-learn cycle, and many continuous deployment technical patterns.⁷

The Lean UX Movement

In 2013, Jeff Gothelf wrote *Lean UX: Applying Lean Principles to Improve User Experience*, which codified how to improve the “fuzzy front end” and explained how product owners can frame business hypotheses, experiment, and gain confidence in those business hypotheses before investing time and resources in the resulting features. By adding Lean UX, we now have the tools to fully optimize the flow between business hypotheses, feature development, testing, deployment, and service delivery to the customer.

The Rugged Computing Movement

In 2011, Joshua Corman, David Rice, and Jeff Williams examined the apparent futility of securing applications and environments late in the life cycle. In response, they created a philosophy called “Rugged Computing,” which attempts to frame the non-functional requirements of stability, scalability, availability, survivability, sustainability, security, supportability, manageability, and defensibility.

Because of the potential for high release rates, DevOps can put incredible pressure on QA and Infosec, because when deploy rates go from monthly or quarterly to hundreds or

thousands daily, no longer are two-week turnaround times from Infosec or QA tenable. The Rugged Computing movement posited that the current approach to fighting the vulnerable industrial complex being employed by most information security programs is hopeless.

Appendix 2: Theory of Constraints and Core, Chronic Conflicts

The Theory of Constraints body of knowledge extensively discusses the use of creating core conflict clouds (often referred to as “C3”). Figure A.1 shows the conflict cloud for IT:

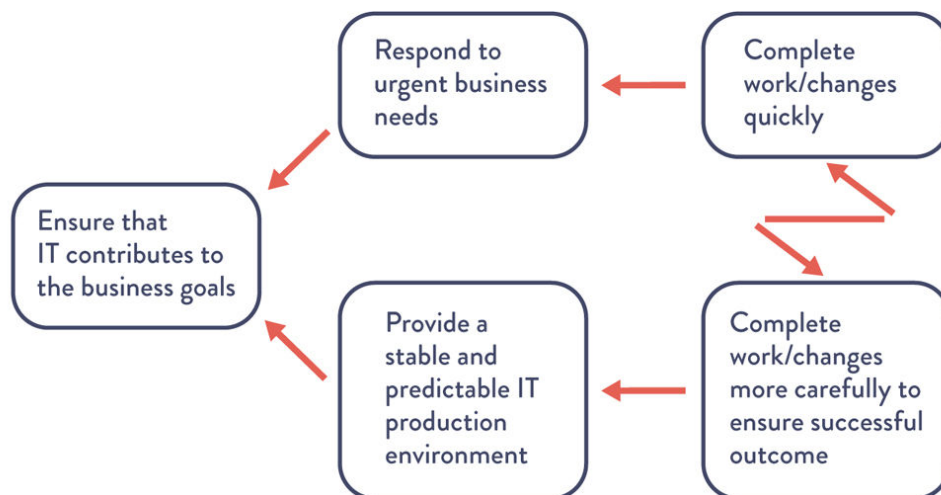


Figure A.1: The Core, Chronic Conflict Facing Every IT Organization

During the 1980s, there was a very well-known core, chronic conflict in manufacturing. Every plant manager had two valid business goals: protect sales and reduce costs. The problem was that in order to protect sales, sales management was incentivized to increase inventory to ensure that it was always possible to fulfill customer demand.

On the other hand, in order to reduce cost, production management was incentivized to decrease inventory to ensure that money was not tied up in work-in-progress that

wasn't immediately shippable to the customer in the form of fulfilled sales.

They were able to break the conflict by adopting Lean principles, such as reducing batch sizes, reducing work in process, and shortening and amplifying feedback loops. This resulted in dramatic increases in plant productivity, product quality, and customer satisfaction.

The principles behind DevOps work patterns are the same as those that transformed manufacturing, allowing us to optimize the IT value stream, converting business needs into capabilities and services that provide value for our customers.

Appendix 3: Tabular Form of Downward Spiral

The tabular form of the downward spiral depicted in *The Phoenix Project* is shown in [Table A.1](#).

Table A.1: The Downward Spiral

IT Operations sees ...	Development sees ...
Fragile applications are prone to failure	Fragile applications are prone to failure
Long time required to figure out which bit got flipped	More urgent, date-driven projects put into the queue
Detective control is a salesperson	Even more fragile code (less secure) put into production
Too much time required to restore service	More releases have increasingly turbulent installs
Too much firefighting and unplanned work	Release cycles lengthen to amortize cost of deployments
Urgent security rework and remediation	Failing bigger deployments more difficult to diagnose
Planned project work cannot be completed	Most senior and constrained IT ops resources have less time to fix underlying process problems
Frustrated customers leave	Ever increasing backlog of work that could help the business win
Market share goes down	Ever increasing amount of tension between IT Ops, Development, Design
Business misses Wall Street commitments	—

Business makes even larger promises to Wall Street

Appendix 4: The Dangers of Handoffs and Queues

The problem with high amounts of queue time is exacerbated when there are many handoffs, because that is where queues are created. [Figure A.2](#) shows wait time as a function of how busy a resource at a work center is. The asymptotic curve shows why a “simple thirty-minute change” often takes weeks to complete—specific engineers and work centers often become problematic bottlenecks when they operate at high utilization. As a work center approaches 100% utilization, any work required from it will languish in queues and won’t be worked on without someone expediting/escalating.

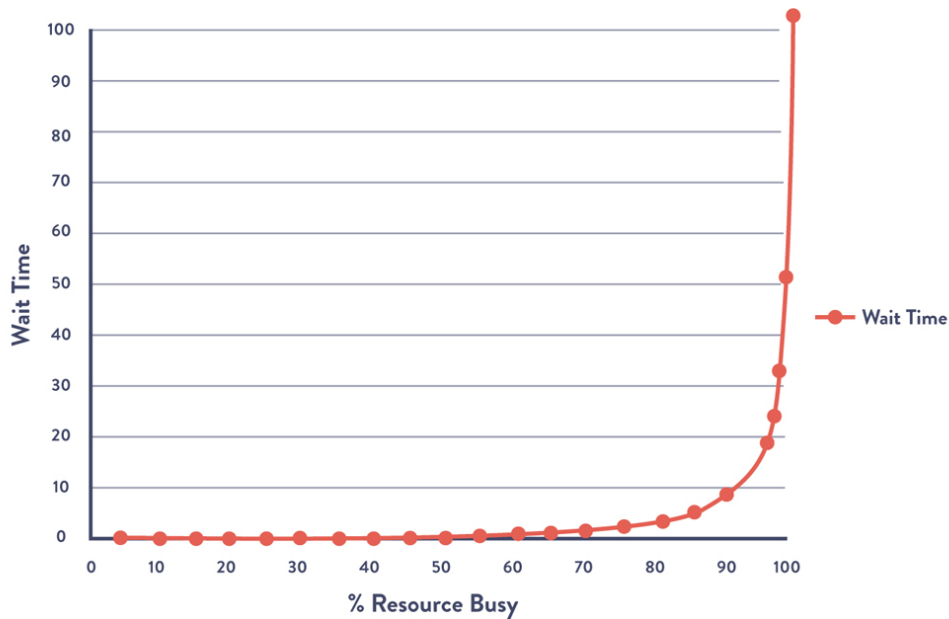


Figure A.2: Queue Size and Wait Times as a Function of Percent Utilization

Source: Kim, Behr, and Spafford, *The Phoenix Project*, ePub edition, 557.

In [Figure A.2](#), the x-axis is the percent busy for a given resource at a work center and the y-axis is the approximate wait time (or, more precisely stated, the queue length). What

the shape of the line shows is that as resource utilization goes past 80%, wait time goes through the roof.

In *The Phoenix Project*, Bill and his team realized the devastating consequences of this property on lead times for the commitments they were making to the project management office:⁸

I tell them about what Erik told me at MRP-8, about how wait times depend upon resource utilization. “The wait time is the ‘percentage of time busy’ divided by the ‘percentage of time idle.’ In other words, if a resource is fifty percent busy, then it’s fifty percent idle. The wait time is fifty percent divided by fifty percent, so one unit of time. Let’s call it one hour.

So, on average, our task would wait in the queue for one hour before it gets worked.

“On the other hand, if a resource is ninety percent busy, the wait time is ‘ninety percent divided by ten percent or nine hours’. In other words, our task would wait in queue nine times longer than if the resource were fifty percent idle.”

I conclude, “So ... For the Phoenix task, assuming we have seven handoffs, and that each of those resources is busy ninety percent of the time, the tasks would spend in queue a total of nine hours times the seven steps ...”

“What? Sixty-three hours, just in queue time?” Wes says, incredulously. “That’s impossible!”

Patty says with a smirk, “Oh, of course. Because it’s only thirty seconds of typing, right?”

Bill and team realize that their “simple thirty-minute task” actually requires seven handoffs (e.g., server team,

networking team, database team, virtualization team, and, of course, Brent, the “rockstar” engineer).

Assuming that all work centers were 90% busy, [Figure A.2](#) shows us that the average wait time at each work center is nine hours—and because the work had to go through seven work centers, the total wait time is seven times that: sixty-three hours.

In other words, the total % of *value added time* (sometimes known as process time) was only 0.16% of the total lead time (thirty minutes divided by sixty-three hours). That means that for 99.8% of our total lead time, the work was simply sitting in queue, waiting to be worked on.

Appendix 5: Myths of Industrial Safety

Decades of research into complex systems shows that countermeasures are based on several myths. In “Some Myths about Industrial Safety,” by Denis Besnard and Erik Hollnagel, they are summarized as such:

Myth 1: “Human error is the largest single cause of accidents and incidents.”⁹

Myth 2: “Systems will be safe if people comply with the procedures they have been given.”¹⁰

Myth 3: “Safety can be improved by barriers and protection; more layers of protection results in higher safety.”¹¹

Myth 4: “Accident analysis can identify the root cause (the ‘truth’) of why the accident happened.”¹²

Myth 5: “Accident investigation is the logical and rational identification of causes based on facts.”¹³

Myth 6: “Safety always has the highest priority and will never be compromised.”¹⁴

The differences between what is myth and what is true are shown in [Table A.2](#).

Table A.2: Two Stories

Myth	Reality
Human error is seen as the cause of failure.	Human error is seen as the effect of systemic vulnerabilities deeper inside the organization.
Saying what people should have done is a satisfying way to describe failure.	Saying what people should have done doesn't explain why it made sense for them to do what they did.
Telling people to be more careful will make the problem go away.	Only by constantly seeking out their vulnerabilities can organizations enhance safety.

Appendix 6: The Toyota Andon Cord

Many ask, how can any work be completed if the Andon cord is being pulled over five thousand times per day? To be precise, not every Andon cord pull results in stopping the entire assembly line. Rather, when the Andon cord is pulled, the team leader overseeing the specified work center has fifty seconds to resolve the problem. If the problem has not been resolved by the time the fifty seconds is up, the partially assembled vehicle will cross a physically drawn line on the floor, and the assembly line will be stopped.¹⁵

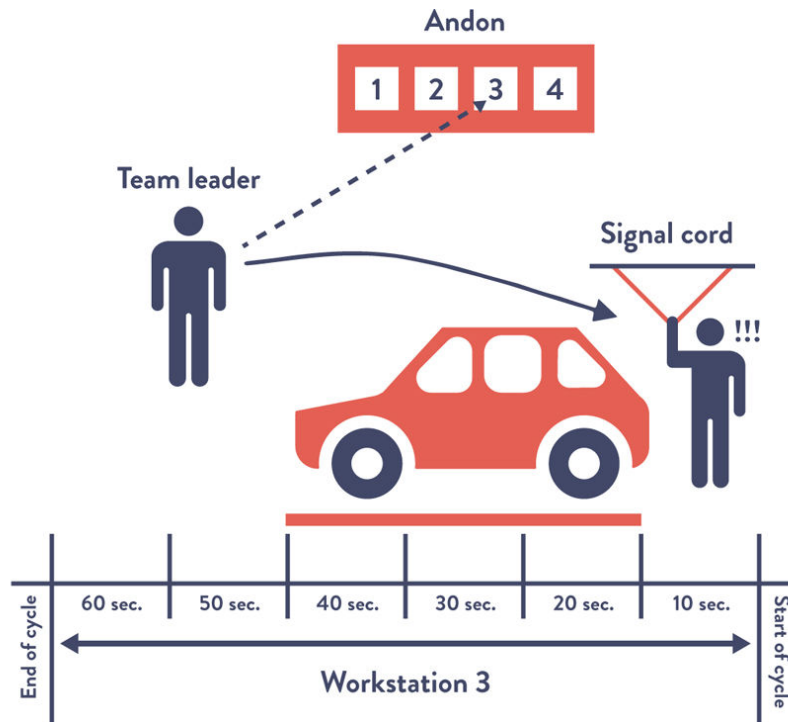


Figure A.3: The Toyota Andon Cord

Appendix 7: COTS Software

Currently, in order to get complex COTS (commercial off-the-shelf) software (e.g., SAP, IBM WebSphere, Oracle WebLogic) into version control, we may have to eliminate the use of graphical point-and-click vendor installer tools. To do that, we need to discover what the vendor installer is doing, and we may need to do an install on a clean server image, diff the file system, and put those added files into version control. Files that don't vary by environment are put into one place ("base install"), while environment-specific files are put into their own directory ("test" or "production"). By doing this, software install operations become merely a version control operation, enabling better visibility, repeatability, and speed.

We may also have to transform any application configuration settings so that they are in version control. For instance, we may transform application configurations that are stored in a database into XML files and vice versa.

Appendix 8: Post-Mortem Meetings (Retrospective)

A sample agenda of the post-mortem meeting is shown below:¹⁶

- An initial statement will be made by the meeting leader or facilitator to reinforce that this meeting is a blameless post-mortem and that we will not focus on past events or speculate on “would haves” or “could haves.” Facilitators might read the “Retrospective Prime Directive” from the website [Retrospective.com](https://www.retrospective.com).
- Furthermore, the facilitator will remind everyone that any countermeasures must be assigned to someone, and if the corrective action does not warrant being a top priority when the meeting is over, then it is not a corrective action. (This is to prevent the meeting from generating a list of good ideas that are never implemented.)
- Those at the meeting will reach an agreement on the complete timeline of the incident, including when and who detected the issue, how it was discovered (e.g., automated monitoring, manual detection, customer notified us), when service was satisfactorily restored, and so forth. We will also integrate into the timeline all external communications during the incident.
- When we use the word “timeline,” it may evoke the image of a linear set of steps of how we gained an understanding of the problem and eventually fixed it. In reality, especially in complex systems, there will likely be many events that contributed to the accident, and many troubleshooting paths and actions will have been taken in an effort to fix it. In this activity, we seek to chronicle all of these events and the

perspectives of the actors and establish hypotheses concerning cause and effect where possible.

- The team will create a list of all the factors that contributed to the incident, both human and technical. They may then sort them into categories, such as “design decision,” “remediation,” “discovering there was a problem,” and so forth. The team will use techniques such as brainstorming and the “infinite hows” to drill down on contributing factors they deem particularly important to discover deeper levels of contributing factors. All perspectives should be included and respected—nobody should be permitted to argue with or deny the reality of a contributing factor somebody else has identified. It’s important for the post-mortem facilitator to ensure that sufficient time is spent on this activity and that the team doesn’t try and engage in convergent behavior such as trying to identify one or more “root causes.”
- Those at the meeting will reach an agreement on the list of corrective actions that will be made top priorities after the meeting. Assembling this list will require brainstorming and choosing the best potential actions to either prevent the issue from occurring or enable faster detection or recovery. Other ways to improve the systems may also be included.
- Our goal is to identify the smallest number of incremental steps to achieve the desired outcomes, as opposed to “Big Bang” changes, which not only take longer to implement, but delay the improvements we need.
- We will also generate a separate list of lower-priority ideas and assign an owner. If similar problems occur

in the future, these ideas may serve as the foundation for crafting future countermeasures.

- Those at the meeting will reach an agreement on the incident metrics and their organizational impact. For example, we may choose to measure our incidents by the following metrics:
 - Event severity: How severe was this issue? This directly relates to the impact on the service and our customers.
 - Total downtime: How long were customers unable to use the service to any degree?
 - Time to detect: How long did it take for us or our systems to know there was a problem?
 - Time to resolve: How long after we knew there was a problem did it take for us to restore service?

Bethany Macri from Etsy observed, “Blamelessness in a post-mortem does not mean that no one takes responsibility. It means that we want to find out what the circumstances were that allowed the person making the change or who introduced the problem to do this. What was the larger environment? ... The idea is that by removing blame, you remove fear, and by removing fear, you get honesty.”¹⁷

Appendix 9: The Simian Army

After the 2011 AWS US-East outage, Netflix had numerous discussions about engineering their systems to automatically deal with failure. These discussions have evolved into a service called “Chaos Monkey.”¹⁸

Since then, Chaos Monkey has evolved into a whole family of tools, known internally as the “Netflix Simian

Army,” to simulate increasingly catastrophic levels of failures.¹⁹

- **Chaos Gorilla:** simulates the failure of an entire AWS availability zone.
- **Chaos Kong:** simulates failure of entire AWS regions, such as North America or Europe.

Other member of the Simian Army now include:

- **Latency Monkey:** induces artificial delays or downtime in their RESTful client-server communication layer to simulate service degradation and ensure that dependent services respond appropriately.
- **Conformity Monkey:** finds and shuts down AWS instances that don't adhere to best-practices (e.g., when instances don't belong to an auto-scaling group or when there is no escalation engineer email address listed in the service catalog).
- **Doctor Monkey:** taps into health checks that run on each instance and finds unhealthy instances and proactively shuts them down if owners don't fix the root cause in time.
- **Janitor Monkey:** ensures that their cloud environment is running free of clutter and waste; searches for unused resources and disposes of them.
- **Security Monkey:** an extension of Conformity Monkey; finds and terminates instances with security violations or vulnerabilities, such as improperly configured AWS security groups.

Appendix 10: Transparent Uptime

Lenny Rachitsky wrote about the benefits of what he called “transparent uptime:”²⁰

Your support costs go down as your users are able to self-identify system wide problems without calling or emailing your support department. Users will no longer have to guess whether their issues are local or global and can more quickly get to the root of the problem before complaining to you.

You are better able to communicate with your users during downtime events, taking advantage of the broadcast nature of the internet versus the one-to-one nature of email and the phone. You spend less time communicating the same thing over and over and more time resolving the issue.

You create a single and obvious place for your users to come to when they are experiencing downtime. You save your users’ time currently spent searching forums, Twitter, or your blog.

Trust is the cornerstone of any successful SaaS adoption. Your customers are betting their business and their livelihoods on your service or platform. Both current and prospective customers require confidence in your service. Both need to know they won’t be left in the dark, alone and uninformed, when you run into trouble. Real time insight into unexpected events is the best way to build this trust. Keeping them in the dark and alone is no longer an option.

It’s only a matter of time before every serious SaaS provider will be offering a public health dashboard. Your users will demand it.

BIBLIOGRAPHY

- “A Conversation with Werner Vogels.” *ACM Queue* 4, no. 4 (2006): 14–22.
<https://queue.acm.org/detail.cfm?id=1142065>.
- Adler, Paul. “Time-and-Motion Regained.” *Harvard Business Review*, January–February 1993. <https://hbr.org/1993/01/time-and-motion-regained>.
- Agile Alliance. “Information Radiators.” Glossary. Accessed May 31, 2016.
<https://www.agilealliance.org/glossary/information-radiators/>.
- ALICE. “Pair Programming.” Wiki page. Updated April 4, 2014.
http://euler.math.uga.edu/wiki/index.php?title=Pair_programming.
- Allspaw, John. “Blameless PostMortems and a Just Culture.” *Code as Craft* (blog), Etsy, May 22, 2012. <http://codeascraft.com/2012/05/22/blameless-postmortems/>.
- Allspaw, John. “Convincing Management that Cooperation and Collaboration Was Worth It.” *Kitchen Soap* (blog), January 5, 2012.
<http://www.kitchensoap.com/2012/01/05/convincing-management-that-cooperation-and-collaboration-was-worth-it/>.
- Allspaw, John. “Counterfactual Thinking, Rules, and the Knight Capital Accident.” *Kitchen Soap* (blog), October 29, 2013.
<http://www.kitchensoap.com/2013/10/29/counterfactuals-knight-capital/>.
- Allspaw, John interviewed by Jenn Webb. “Post-Mortems, Sans Finger-Pointing.” *The O’Reilly Radar Podcast*. Podcast audio, 30:34. August 21, 2014.
<http://radar.oreilly.com/2014/08/postmortems-sans-finger-pointing-the-oreilly-radar-podcast.html>
- Amazon Web Services, “Summary of the Amazon DynamoDB Service Disruption and Related Impacts in the US-East Region.” Amazon Web Services. Accessed May 28, 2016. <https://aws.amazon.com/message/5467D2/>.
- Anderson, David J. *Kanban: Successful Evolutionary Change for Your Technology Business*. Sequim, WA: Blue Hole Press, 2010.
- Anderson, David J., and Dragos Dumitriu. *From Worst to Best in 9 Months: Implementing a Drum-Buffer-Rope Solution in Microsoft’s IT Department*. Microsoft Corporation, 2005.
- Antani, Snehal. “IBM Innovate DevOps Keynote.” Posted by IBM DevOps, June 12, 2014. YouTube video, 47:57. <https://www.youtube.com/watch?v=s0M1P05-6Io>.
- Arbuckle, Justin. “What Is ArchOps: Chef Executive Roundtable.” 2013.
- Ashman, David. “DOES14—David Ashman—Blackboard Learn—Keep Your Head in the Clouds.” Posted by DevOps Enterprise Summit 2014, October 28, 2014. YouTube video, 30:43. <https://www.youtube.com/watch?v=SSmixnMpsI4>.

- Associated Press. "Number of Active Users at Facebook over the Years," Yahoo! News. May 1, 2013. <https://www.yahoo.com/news/number-active-users-facebook-over-230449748.html?ref=gs>.
- Atwood, Jeff. "Pair Programming vs. Code Reviews." *Coding Horror* (blog), November 18, 2013. <http://blog.codinghorror.com/pair-programming-vs-code-reviews/>.
- Atwood, Jeff. "Software Branching and Parallel Universes." *Coding Horror* (blog), October 2, 2007. <http://blog.codinghorror.com/software-branching-and-parallel-universes/>.
- Axelos. *ITIL Service Transition*. Belfast, Ireland: The Stationary Office, 2011.
- Ayers, Zach, and Joshua Cohen. "Andon Cords in Development Teams—Driving Continuous Learning." Presentation at the DevOps Enterprise Summit, Las Vegas, 2019. <https://videolibrary.doesvirtual.com/?video=504281981>.
- Azzarello, Domenico, Frédéric Debruyne, and Ludovica Mottura. "The Chemistry of Enthusiasm: How Engaged Employees Create Loyal Customers," Bain & Company, May 4, 2012. <https://www.bain.com/insights/the-chemistry-of-enthusiasm>.
- Bahri, Sami. "Few Patients-in-Process and Less Safety Scheduling; Incoming Supplies Are Secondary." *The W. Edwards Deming Institute Blog*, August 22, 2013. <https://blog.deming.org/2013/08/fewer-patients-in-process-and-less-safety-scheduling-incoming-supplies-are-secondary/>.
- Barr, Jeff. "EC2 Maintenance Update." *AWS News Blog*. Amazon Web Services, September 25, 2014. <https://aws.amazon.com/blogs/aws/ec2-maintenance-update/>.
- Basu, Biswanath, Rakesh Goyal, and Jennifer Hansen. "Biz & Tech Partnership Towards 10 'No Fear Releases' Per Day," presentation at DevOps Enterprise Summit, Las Vegas, 2020. <https://videolibrary.doesvirtual.com/?video=468711236>.
- Bazaarvoice, Inc. Announces Its Financial Results for the Fourth Fiscal Quarter and Fiscal Year Ended April 30, 2012." *Bazaar Voice*, June 6, 2012. <http://investors.bazaarvoice.com/releasedetail.cfm?ReleaseID=680964>.
- Beck, Kent. "Slow Deployment Causes Meetings." Facebook, November 19, 2015. https://www.facebook.com/notes/kent-beck/slow-deployment-causes-meetings/1055427371156793?_rdr=p.
- Beck, Kent, Mike Beedle, Arie van Bennekum, Alastair Cockburn, Ward Cunningham, Martin Fowler, James Grenning, et al. "Twelve Principles of Agile Software." *Agile Manifesto*, 2001. <http://agilemanifesto.org/principles.html>.
- Besnard, Denis and Erik Hollnagel. "Some Myths about Industrial Safety." Paris: Centre De Recherche Sur Les Risques Et Les Crises Mines Working Paper Series 2012. ParisTech, Paris, France, December 2012. http://gswong.com/?wpfb_dl=31.

- Betz, Charles. *Architecture and Patterns for IT Service Management, Resource Planning, and Governance: Making Shoes for the Cobbler's Children*. Witham, MA: Morgan Kaufmann, 2011.
- Beyond Lean. "The 7 Wastes (Seven forms of Muda)." The 7 Wastes Explained. Accessed July 28, 2016. <http://www.beyondlean.com/7-wastes.html>.
- Big Fish Games. "Big Fish Celebrates 11th Consecutive Year of Record Growth." Pressroom. January 28, 2014. <http://pressroom.bigfishgames.com/2014-01-28-Big-Fish-Celebrates-11th-Consecutive-Year-of-Record-Growth>.
- Bland, Mike. "DOES15—Mike Bland—Pain Is Over, If You Want It." Posted by Gene Kim to [slideshare.net](http://www.slideshare.net), November 18, 2015. Slideshow. <http://www.slideshare.net/ITRevolution/does15-mike-bland-pain-is-over-if-you-want-it-55236521>.
- Bland, Mike. "Fixits, or I Am the Walrus," *Mike Bland* (blog). Mike Bland, October 4, 2011. <https://mike-bland.com/2011/10/04/fixits.html>.
- Bosworth, Andrew. "Building and Testing at Facebook." Facebook, August 8, 2012. <https://www.facebook.com/notes/facebook-engineering/building-and-testing-at-facebook/10151004157328920>.
- Boubez, Toufic. "Simple Math for Anomaly Detection Toufic Boubez—Metafor Software—Monitorama PDX 2014-05-05," Posted by tboubez to [slideshare.net](http://www.slideshare.net), May 6, 2014. Slideshow. <http://www.slideshare.net/tboubez/simple-math-for-anomaly-detection-toufic-boubez-metafor-software-monitorama-pdx-20140505>.
- Brooks, Jr., Frederick P. *The Mythical Man-Month: Essays on Software Engineering, Anniversary Edition*. Upper Saddle River, NJ: Addison-Wesley, 1995.
- Buchanan, Leigh. "The Wisdom of Peter Drucker from A to Z." *Inc.*, November 19, 2009. <http://www.inc.com/articles/2009/11/drucker.html>.
- Buhr, Sarah. "Etsy Closes Up 86 Percent on First Day of Trading." *Tech Crunch*, April 16, 2015. <http://techcrunch.com/2015/04/16/etsy-stock-surges-86-percent-at-close-of-first-day-of-trading-to-30-per-share/>.
- Burrows, Mike. "The Chubby Lock Service for Loosely-Coupled Distributed Systems." Paper presented at OSDI 2006: Seventh Symposium on Operating System Design and Implementation, November 2006. <http://static.googleusercontent.com/media/research.google.com/en//archive/chubby-osdi06.pdf>.
- Cagan, Marty. *Inspired: How to Create Products Customers Love*. Saratoga, CA: SVPG Press, 2008.
- Campbell-Pretty, Em. "DOES14—Em Campbell-Pretty—How a Business Exec Led Agile, Lead, CI/CD." Posted by DevOps Enterprise Summit, April 20, 2014. YouTube video, 29:47. <https://www.youtube.com/watch?v=-4pIMMTbtwE>.
- Canahuati, Pedro. "Growing from the Few to the Many: Scaling the Operations Organization at Facebook." Filmed December 16, 2013 for QCon. Video,

- 39:39. <http://www.infoq.com/presentations/scaling-operations-facebook>.
- Chacon, Scott. "GitHub Flow." *Scott Chacon* (blog), August 31, 2011. <http://scottchacon.com/2011/08/31/github-flow.html>.
- Chakrabarti, Arup. "Common Ops Mistakes." Filmed presentation at Heavy Bit Industries, June 3, 2014. Video, 36:12. <http://www.heavybit.com/library/video/common-ops-mistakes/>.
- Chan, Jason. "OWASP AppSecUSA 2012: Real World Cloud Application Security." Posted by Christiaan008, December 10, 2012. Youtube video, 37:45. <https://www.youtube.com/watch?v=daNA0jXDvYk>.
- Chandola, Varun , Arindam Banerjee, and Vipin Kumar. "Anomaly Detection: A Survey." *ACM Computing Surveys* 41, no. 3 (July 2009): 15. <http://doi.acm.org/10.1145/1541880.1541882>.
- Chapman, Janet, and Patrick Eltridge. "On A Mission: Nationwide Building Society," presentation at DevOps Enterprise Summit, London, 2020. <https://videolibrary.doesvirtual.com/?video=432109857>.
- Chuvakin, Anton. "LogLogic/Chuvakin Log Checklist," republished with permission, 2008, <http://juliusdavies.ca/logging/llcl.html>.
- Clanton, Ross, and Michael Ducy interviewed by Courtney Kissler and Jason Josephy. "Continuous Improvement at Nordstrom." *The Goat Farm*, season 1, episode 17. Podcast audio, 53:18. June 25, 2015. <http://goatcan.do/2015/06/25/the-goat-farm-episode-7-continuous-improvement-at-nordstrom/>.
- Clanton, Ross and Heather Mickman. "DOES14—Ross Clanton and Heather Mickman—DevOps at Target." Posted by DevOps Enterprise Summit 2014, October 29, 2014. YouTube video, 29:20. <https://www.youtube.com/watch?v=exrjV9V9vhY>.
- Claudius, Jonathan. "Attacking Cloud Services with Source Code." Posted by Jonathan Claudius to speakerdeck.com, April 16, 2013. Slideshow. <https://speakerdeck.com/claudijd/attacking-cloud-services-with-source-code>.
- Clemm, Josh. "LinkedIn Started Back in 2003—Scaling LinkedIn—A Brief History." Posted by Josh Clemm to slideshare.net, November 9, 2015. Slideshow. http://www.slideshare.net/joshclemm/how-linkedin-scaled-a-brief-history/3-LinkedIn_started_back_in_2003.
- Cockcroft, Adrian, Cory Hicks, and Greg Orzell. "Lessons Netflix Learned from the AWS Outage." *The Netflix Tech Blog*, April 29, 2011. <http://techblog.netflix.com/2011/04/lessons-netflix-learned-from-aws-outage.html>.
- Cockcroft, Adrian, interviewed by Michael Ducy and Ross Clanton. "Adrian Cockcroft of Battery Ventures." *The Goat Farm* season 1, episode 8. Podcast audio, July 31, 2015. <http://goatcan.do/2015/07/31/adrian-cockcroft-of-battery-ventures-the-goat-farm-episode-8/>.

- Cockcroft, Adrian. "Monitorama—Please, No More Minutes, Milliseconds, Monoliths or Monitoring Tools." Posted by Adrian Cockcroft to [slideshare.net](https://www.slideshare.net), May 5, 2014. Slideshow.
<http://www.slideshare.net/adriancockcroft/monitorama-please-no-more>.
- Collins, Justin, Alex Smolen, and Neil Matatall. "Putting to your Robots to Work V1.1." Posted by Neil Matatall to [slideshare.net](https://www.slideshare.net), April 24, 2012. Slideshow.
<http://www.slideshare.net/xplodersuv/sf-2013-robots/>.
- Conrad, Ben, and Matt Hyatt. "Saving the Economy from Ruin (with a Hyperscale Paas)," presentation at the 2021 DevOps Enterprise Summit-Europe Virtual.
<https://videolibrary.doesvirtual.com/?video=550704128>.
- Conway, Melvin E. "How Do Committees Invent?" Mel Conway. Originally published in *Datamation* magazine, April 1968.
<http://www.melconway.com/research/committees.html>.
- Cook, Scott. "Leadership in an Agile Age: An Interview with Scott Cook." By Cassie Divine. Intuit, April 20, 2011.
<https://web.archive.org/web/20160205050418/http://network.intuit.com/2011/04/20/leadership-in-the-agile-age/>.
- Corman, Josh and John Willis. "Immutable Awesomeness—Josh Corman and John Willis at DevOps Enterprise Summit 2015." Posted by Sonatype, October 21, 2015. YouTube video, 34:25. <https://www.youtube.com/watch?v=-S8-lrm3iV4>.
- Cornago, Fernando, Vikalp Yadav, and Andreia Otto. "From 6-Eye Principle to Release at Scale - adidas Digital Tech 2021," presentation at DevOps Enterprise Summit-Europe, 2021. <https://videolibrary.doesvirtual.com/?video=524020857>.
- Cox, Jason. "Disney DevOps: To Infinity and Beyond." Presented at DevOps Enterprise Summit, San Francisco, 2014.
- Cundiff, Dan, Levi Geinert, Lucas Rettif. "Crowdsourcing Technology Governance," presentation at DevOps Enterprise Summit, San Francisco, 2018. <https://videolibrary.doesvirtual.com/?video=524020857>.
- Cunningham, Ward. "Ward Explains Debt Metaphor," c2. Last updated January 22, 2011. <http://c2.com/cgi/wiki?WardExplainsDebtMetaphor>.
- Daniels, Katherine. "Devopsdays Minneapolis 2015—Katherine Daniels—DevOps: The Missing Pieces." Posted by DevOps Minneapolis, July 13, 2015. YouTube video, 33:26. <https://www.youtube.com/watch?v=LNJkVw93yTU>.
- Davis, Jennifer and Katherine Daniels. *Effective DevOps: Building a Culture of Collaboration, Affinity, and Tooling at Scale*. Sebastopol, CA: O'Reilly Media, 2016.
- "Decreasing False Positives in Automated Testing." Posted by Sauce Labs to [slideshare.net](https://www.slideshare.net), March 24, 2015. Slideshow.
<http://www.slideshare.net/saucelabs/decreasing-false-positives-in-automated-testing>.

- DeGrandis, Dominica. "DOES15—Dominica DeGrandis—The Shape of Uncertainty." Posted by DevOps Enterprise Summit, November 5, 2015. Youtube video, 22:54. <https://www.youtube.com/watch?v=Gp05i0d34gg>.
- Dekker, Sidney. "DevOpsDays Brisbane 2014—Sidney Decker—System Failure, Human Error: Who's to Blame?" Posted by info@devopsdays.org, 2014. Vimeo video, 1:07:38. <https://vimeo.com/102167635>.
- Dekker, Sidney. *Just Culture: Balancing Safety and Accountability*. Lund University, Sweden: Ashgate Publishing, 2007.
- Dekker, Sidney. *The Field Guide to Understanding Human Error*. Lund University, Sweden: Ashgate Publishing, 2006.
- DeLuccia, James, Jeff Gallimore, Gene Kim, and Byron Miller. *DevOps Audit Defense Toolkit*. Portland, OR: IT Revolution, 2015. <http://itrevolution.com/devops-and-auditors-the-devops-audit-defense-toolkit>.
- "DevOps Culture: How to Transform," [Cloud.Google.com](https://cloud.google.com/architecture/devops/devops-culture-transform), accessed August 26, 2021. <https://cloud.google.com/architecture/devops/devops-culture-transform>.
- Dickerson, Chad. "Optimizing for Developer Happiness." *Code As Craft* (blog), Etsy, June 6, 2011. <https://codeascraft.com/2011/06/06/optimizing-for-developer-happiness/>.
- Dignan, Larry. "Little Things Add Up." *Baseline*, October 19, 2005. <http://www.baselinemag.com/c/a/Projects-Management/Profiles-Lessons-From-the-Leaders-in-the-iBaselinei500/3>.
- Douglas, Jake. "Deploying at Github." *The GitHub Blog*. GitHub, August 29, 2012. <https://github.com/blog/1241-deploying-at-github>.
- Dweck, Carol. "Carol Dweck Revisits the 'Growth Mindset.'" *Education Week*, September 22, 2015. <http://www.edweek.org/ew/articles/2015/09/23-/caroldweck-revisits-the-growth-mindset.html>.
- Edmondson, Amy C. "Strategies for Learning from Failure." *Harvard Business Review*, April 2011. <https://hbr.org/2011/04/strategies-for-learning-from-failure>.
- Edwards, Damon. "DevOps Kaizen: Find and Fix What Is Really Behind Your Problems." Posted by dev2ops to [slideshare.net](https://www.slideshare.net/dev2ops/dev-ops-kaizen-damon-edwards). Slideshow., May 4, 2015. <http://www.slideshare.net/dev2ops/dev-ops-kaizen-damon-edwards>.
- Exner, Ken. "Transforming Software Development." Posted by Amazon Web Services, April 10, 2015. YouTube video, 40:57. <https://www.youtube.com/watch?v=YCrhemssYuI>.
- Figureau, Brice. "The 10 Commandments of Logging." *Masterzen's Blog*, January 13, 2013. <http://www.masterzen.fr/2013/01/13/the-10-commandments-of-logging/>.

- Fitz, Timothy. "Continuous Deployment at IMVU: Doing the Impossible Fifty Times a Day." *Timothy Fitz* (blog), February 10, 2009. <http://timothyfitz.com/2009/02/10/continuous-deployment-at-imvu-doing-the-impossible-fifty-times-a-day/>.
- Forsgren, Nicole, Bas Alberts, Kevin Backhouse, and Grey Baker. *The 2020 State of the Octoverse*. GitHub, 2020. <https://octoverse.github.com/static/github-octoverse-2020-security-report.pdf>.
- Forsgren, Nicole, Jez Humble, and Gene Kim. *Accelerate: State of DevOps 2018*. DORA and Google Cloud, 2018. <https://lp.google-mkto.com/rs/248-TPC-286/images/DORA-State%20of%20DevOps.pdf>.
- Forsgren, Nicole, Jez Humble, Nigel Kersten, and Gene Kim. "2014 State Of DevOps Findings! Velocity Conference." Posted by Gene Kim to [slideshare.net](https://www.slideshare.net), June 30, 2014. Slideshow. <http://www.slideshare.net/realgenekim/2014-state-of-devops-findings-velocity-conference>.
- Forsgren, Nicole, Dustin Smith, Jez Humble, and Jessie Frazelle. *Accelerate State of DevOps 2019*. DORA and Google Cloud, 2019. <https://services.google.com/fh/files/misc/state-of-devops-2019.pdf>.
- Forsgren, Nicole, Margaret-Anne Storey, Chandra Maddila, Thomas Zimmermann, Brain Houck, and Jenna Butler. "The SPACE of Developer Productivity." *ACM Queue* 19, no. 1 (2021): 1–29. <https://queue.acm.org/detail.cfm?id=3454124>.
- Fowler, Chad. "Trash Your Servers and Burn Your Code: Immutable Infrastructure and Disposable Components." *Chad Fowler* (blog), June 23, 2013. <http://chadfowler.com/2013/06/23/immutable-deployments.html>.
- Fowler, Martin. "Continuous Integration." *Martin Fowler* (blog), May 1, 2006. <http://www.martinfowler.com/articles/continuousIntegration.html>.
- Fowler, Martin. "Eradicating Non-Determinism in Tests." *Martin Fowler* (blog), April 14, 2011. <http://martinfowler.com/articles/nonDeterminism.html>.
- Fowler, Martin. "StranglerFigApplication." *Martin Fowler* (blog), June 29, 2004. <http://www.martinfowler.com/bliki/StranglerApplication.html>.
- Fowler, Martin. "TestDrivenDevelopment." *Martin Fowler* (blog), March 5, 2005. <http://martinfowler.com/bliki/TestDrivenDevelopment.html>.
- Fowler, Martin. "TestPyramid." *Martin Fowler* (blog), May 1, 2012. <http://martinfowler.com/bliki/TestPyramid.html>.
- Freakonomics. "Fighting Poverty With Actual Evidence: Full Transcript." *Freakonomics* blog, November 27, 2013. <http://freakonomics.com/2013/11/27/fighting-poverty-with-actual-evidence-full-transcript/>.
- Furtado, Adam, and Lauren Knausenberger. "The Air Force's Digital Journey in 12 Parsecs or Less." Presentation at DevOps Enterprise Summit, London, 2020. <https://videolibrary.doesvirtual.com/?video=467489046>.

- Gaekwad, Karthik. "Agile 2013 Talk: How DevOps Change Everything." Posted by Karthik Gaekwad to [slideshare.net](http://www.slideshare.net/karthequian/howdevopschangeseverythingagile2013karthikgaekwad/), August 7, 2013. Slideshow. <http://www.slideshare.net/karthequian/howdevopschangeseverythingagile2013karthikgaekwad/>.
- Galbreath, Nick. "Continuous Deployment—The New #1 Security Feature, from BSildesLA 2012." Posted by Nick Galbreath to [slideshare.net](http://www.slideshare.net/nickgsuperstar/continuous-deployment-the-new-1-security-feature), August 16, 2012. Slideshow. <http://www.slideshare.net/nickgsuperstar/continuous-deployment-the-new-1-security-feature>.
- Galbreath, Nick. "DevOpsSec: Applying DevOps Principles to Security, DevOpsDays Austin 2012." Posted by Nick Galbreath to [slideshare.net](http://www.slideshare.net/nickgsuperstar/devopssec-apply-devops-principles-to-security), April 12, 2012. Slideshow. <http://www.slideshare.net/nickgsuperstar/devopssec-apply-devops-principles-to-security>.
- Galbreath, Nick. "Fraud Engineering, from Merchant Risk Council Annual Meeting 2012." Posted by Nick Galbreath to [slideshare.net](http://www.slideshare.net/nickgsuperstar/fraud-engineering), May 3, 2012. Slideshow. <http://www.slideshare.net/nickgsuperstar/fraud-engineering>.
- Gallagher, Sean. "When 'Clever' Goes Wrong: How Etsy Overcame Poor Architectural Choices." *Arstechnica*, October 3, 2011. <http://arstechnica.com/business/2011/10/when-clever-goes-wrong-how-etsy-overcame-poor-architectural-choices/>.
- Gardner, Tom. "Barnes & Noble, Blockbuster, Borders: The Killer B's Are Dying." *The Motley Fool*, July 21, 2010. <http://www.fool.com/investing/general/2010/07/21/barnes-noble-blockbuster-borders-the-killer-bs-are.aspx>.
- Geer, Dan and Joshua Corman. "Almost Too Big to Fail." *login: The Usenix Magazine* 39, no. 4 (August 2014): 66–68. https://www.usenix.org/system/files/login/articles/15_geer_0.pdf.
- Gertner, Jon. *The Idea Factory: Bell Labs and the Great Age of American Innovation*. New York: Penguin Books, 2012.
- GitHub. "Etsy's Feature Flagging API Used for Operational Rampups and A/B testing." *Etsy/feature*. Last updated January 4, 2017. <https://github.com/etsy/feature>.
- GitHub. "Library for Configuration Management API." *Netflix/archaius*. Last updated December 4, 2019. <https://github.com/Netflix/archaius>.
- Golden, Bernard. "What Gartner's Bimodal IT Model Means to Enterprise CIOs." *CIO Magazine*, January 27, 2015. <http://www.cio.com/article/2875803/cio-role/what-gartner-s-bimodal-it-model-means-to-enterprise-cios.html>.
- Goldratt, Eliyahu M. *Beyond the Goal: Eliyahu Goldratt Speaks on the Theory of Constraints (Your Coach in a Box)*. Prince Frederick, MD: Gildan Media, 2005.
- Google App Engine Team. "Post-Mortem for February 24, 2010 Outage." Google App Engine website, March 4, 2010. <https://groups.google.com/forum/#!topic/google-appengine/p2QKJ0OSLc8>.

- Govindarajan, Vijay, and Chris Trimble. *The Other Side of Innovation: Solving the Execution Challenge*. Boston, MA: Harvard Business Review, 2010, Kindle.
- Gruver, Gary. "DOES14—Gary Gruver—Macy's—Transforming Traditional Enterprise Software Development Processes." Posted by DevOps Enterprise Summit 2014, October 29, 2014. YouTube video, 27:24. <https://www.youtube.com/watch?v=-HSSGiYXA7U>.
- Gruver, Gary, and Tommy Mouser. *Leading the Transformation: Applying Agile and DevOps Principles at Scale*. Portland, OR: IT Revolution Press, 2015.
- Gupta, Prachi. "Visualizing LinkedIn's Site Performance." LinkedIn Engineering blog, June 13, 2011. <https://engineering.linkedin.com/25/visualizing-linkedins-site-performance>.
- Hammant, Paul. "Introducing Branch by Abstraction." *Paul Hammant's Blog*, April 26, 2007. http://paulhammant.com/blog/branch_by_abstraction.html.
- Hastings, Reed. "Netflix Culture: Freedom and Responsibility." Posted by Reed Hastings to slideshare.net, August 1, 2009. Slideshow. <http://www.slideshare.net/reed2001/culture-1798664>.
- Hendrickson, Elisabeth. "DOES15—Elisabeth Hendrickson—Its All About Feedback." Posted by DevOps Enterprise Summit, November 5, 2015. YouTube video, 34:47. <https://www.youtube.com/watch?v=r2BFTXBundQ>.
- Hendrickson, Elisabeth. "On the Care and Feeding of Feedback Cycles." Posted by Elisabeth Hendrickson to slideshare.net, November 1, 2013. Slideshow. <http://www.slideshare.net/ehendrickson/care-and-feeding-of-feedback-cycles>.
- Hodge, Victoria, and Jim Austin. "A Survey of Outlier Detection Methodologies." *Artificial Intelligence Review* 22, no. 2 (October 2004): 85–126. http://www.geo.upm.es/postgrado/CarlosLopez/papers/Hodge+Austin_OutlierDetection_AIRE381.pdf.
- Holmes, Dwayne. "How A Hotel Company Ran \$30B of Revenue in Containers," presentation at DevOps Enterprise Summit, Las Vegas, 2020. <https://videolibrary.doesvirtual.com/?video=524020857>.
- "How I Structured Engineering Teams at LinkedIn and AdMob for Success." First Round Review, 2015. <http://firstround.com/review/how-i-structured-engineering-teams-at-linkedin-and-admob-for-success/>.
- Hrenko, Michael. "DOES15—Michael Hrenko—DevOps Insured By Blue Shield of California." Posted by DevOps Enterprise Summit, November 5, 2015. YouTube video, 42:24. <https://www.youtube.com/watch?v=NlgrOT24UDw>.
- Huang, Gregory T. "Blackboard CEO Jay Bhatt on the Global Future of Edtech." *Xconomy*, June 2, 2014. <http://www.xconomy.com/boston/2014/06/02/blackboard-ceo-jay-bhatt-on-the-global-future-of-edtech/>.
- Humble, Jez. "What is Continuous Delivery?" Continuous Delivery (website), accessed May 28, 2016, <https://continuousdelivery.com/>.

- Humble, Jez, and David Farley. *Continuous Delivery: Reliable Software Releases through Build, Test, and Deployment Automation*. Upper Saddle River, NJ: Addison-Wesley, 2011.
- Humble, Jez, Joanne Molesky, and Barry O'Reilly. *Lean Enterprise: How High Performance Organizations Innovate at Scale*. Sebastopol, CA: O'Reilly Media, 2015.
- "IDC Forecasts Worldwide IT Spending to Grow 6% in 2012, Despite Economic Uncertainty." *Business Wire*, September 10, 2012.
<http://www.businesswire.com/news/home/20120910005280/en/IDC-Forecasts-Worldwide-Spending-Grow-6-2012>.
- Immelt, Jeff. "GE CEO Jeff Immelt: Let's Finally End the Debate over Whether We Are in a Tech Bubble." *Business Insider*, December 9, 2015.
<http://www.businessinsider.com/ceo-of-ge-lets-finally-end-the-debate-over-whether-we-are-in-a-tech-bubble-2015-12>.
- Intuit, Inc. "2012 Annual Report: Form 10-K." July 31, 2012.
http://s1.q4cdn.com/018592547/files/doc_financials/2012/INTU_2012_7_31_10K_r230_at_09_13_12_FINAL_and_Camera_Ready.pdf.
- Jacobson, Daniel, Danny Yuan, and Neeraj Joshi. "Scrier: Netflix's Predictive Auto Scaling Engine." *The Netflix Tech Blog*, November 5, 2013.
<http://techblog.netflix.com/2013/11/scrier-netflixs-predictive-auto-scaling.html>.
- Jenkins, Jon. "Velocity 2011: Jon Jenkins, 'Velocity Culture.'" Posted by O'Reilly, June 20, 2011. YouTube video, 15:13. <https://www.youtube.com/watch?v=dxk8b9rSKOo>.
- "Jeremy Long: The (Application) Patching Manifesto," YouTube video, 41:17, posted by LocoMocoSec: Hawaii Product Security Conference, May 17, 2018.
<https://www.youtube.com/watch?v=qVVZrTRJ290>.
- JGFLL. Review of *The Phoenix Project: A Novel About IT, DevOps, and Helping Your Business Win*, by Gene Kim, Kevin Behr, and George Spafford. Amazon review, March 4, 2013. <http://www.amazon.com/review/R1KSSPTEGLWJ23>.
- Johnson, Kimberly H., Tim Judge, Christopher Porter, and Ramon Richards. "How Fannie Mae Uses Agility to Support Homeowners and Renters," presentation at DevOps Enterprise Summit, Las Vegas, 2020.
<https://videolibrary.doesvirtual.com/?video=467488997>.
- Jones, Angie. "3 Ways to Get Test Automation Done Within Your Sprints." *TechBeacon*. Accessed February 15, 2021. <https://techbeacon.com/app-dev-testing/3-ways-get-test-automation-done-within-your-sprints>.
- Kash, Wyatt. "New Details Released on Proposed 2016 IT Spending." *FedScoop*, February 4, 2015. <http://fedscoop.com/what-top-agencies-would-spend-on-it-projects-in-2016>.
- Kastner, Erik. "Quantum of Deployment." *Code as Craft* (blog). Etsy, May 20, 2010.
<https://codeascraft.com/2010/05/20/quantum-of-deployment/>.

- Kersten, Mik. "Project to Product: From Stories to Scenius," Tasktop blog, November 21, 2018, <https://www.tasktop.com/blog/project-product-stories-scenius/>.
- Kersten, Mik. *Project to Product: How to Survive and Thrive in the Age of Digital Disruption with the Flow Framework*. Portland, Oregon: IT Revolution Press, 2018.
- Kersten, Nigel, IT Revolution, and PwC. *2015 State of DevOps Report*. Portland, OR: Puppet Labs, 2015. https://puppet.com/resources/white-paper/2015-state-of-devops-report?_ga=1.6612658.168869.1464412647&link=blog.
- Kim, Gene. "The Amazing DevOps Transformation of the HP LaserJet Firmware Team (Gary Gruver)." IT Revolution blog, February 13, 2014. <http://itrevolution.com/the-amazing-devops-transformation-of-the-hp-laserjet-firmware-team-gary-gruver/>.
- Kim, Gene. "Organizational Learning and Competitiveness: Revisiting the 'Allspaw/Hammond 10 Deploys Per Day at Flickr' Story." IT Revolution blog, December 13, 2014. <http://itrevolution.com/organizational-learning-and-competitiveness-a-different-view-of-the-allspawhammond-10-deploys-per-day-at-flickr-story/>.
- Kim, Gene. "State of DevOps: 2020 and Beyond," IT Revolution blog, March 1, 2021. <https://itrevolution.com/state-of-devops-2020-and-beyond/>.
- Kim, Gene. "The Three Ways: The Principles Underpinning DevOps." IT Revolution blog, August 22, 2012. <http://itrevolution.com/the-three-ways-principles-underpinning-devops/>.
- Kim, Gene, Kevin Behr, and George Spafford. *The Visible Ops Handbook: Implementing ITIL in 4 Practical and Auditable Steps*. Eugene, OR: IT Process Institute, 2004.
- Kim, Gene, Gary Gruver, Randy Shoup, and Andrew Phillips. "Exploring the Uncharted Territory of Microservices." Xebia Labs. Webinar, February 20, 2015. <https://xebialabs.com/community/webinars/exploring-the-uncharted-territory-of-microservices/>.
- Kissler, Courtney. "DOES14—Courtney Kissler—Nordstrom—Transforming to a Culture of Continuous Improvement." Posted by DevOps Enterprise Summit 2014, October 29, 2014. YouTube video, 29:59. <https://www.youtube.com/watch?v=0ZAcsrZBSlo>.
- Kohavi, Ron, Thomas Crook, and Roger Longbotham. "Online Experimentation at Microsoft." Paper presented at the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Paris, France, 2009. http://www.exp-platform.com/documents/exp_dmcastudies.pdf.
- Krishnan, Kripa. "Kripa Krishnan: 'Learning Continuously From Failures' at Google." Posted by Flowcon, November 11, 2014. YouTube video, 21:35. <https://www.youtube.com/watch?v=KqqS3wgQum0>.

- Krishnan, Kripa. "Weathering the Unexpected." *Communications of the ACM* 55, no. 11 (November 2012): 48–52.
<http://cacm.acm.org/magazines/2012/11/156583-weathering-the-unexpected/abstract>.
- Kumar, Ashish. "Development at the Speed and Scale of Google." PowerPoint presented at QCon, San Francisco, CA, 2010.
https://qconsf.com/sf2010/dl/qcon-sanfran-2010/slides/AshishKumar_DevelopingProductsattheSpeedandScaleofGoogle.pdf.
- Leibman, Maya, and Ross Clanton. "DevOps: Approaching Cruising Altitude." Presentation at DevOps Enterprise Summit, Las Vegas, 2020,
<https://videolibrary.doesvirtual.com/?video=550704282>.
- Letuchy, Eugene. "Facebook Chat." Facebook, May 3, 2008.
http://www.facebook.com/note.php?note_id=14218138919&id=944554719.
- Lightbody, Patrick. "Velocity 2011: Patrick Lightbody, 'From Inception to Acquisition.'" Posted by O'Reilly, June 17, 2011. YouTube video, 15:28.
<https://www.youtube.com/watch?v=ShmPod8JecQ>.
- Limoncelli, Tom. "Python Is Better than Perl6." Everything SysAdmin blog, January 10, 2011. <http://everythingsysadmin.com/2011/01/python-is-better-than-perl6.html>.
- Limoncelli, Tom. "SRE@Google: Thousands Of DevOps Since 2004." USENIX Association Talk, NYC. Posted by USENIX, January 12, 2012. Youtube video, 45:57. <http://www.youtube.com/watch?v=iluTnhdTzK0>.
- Limoncelli, Tom. "Stop Monitoring Whether or Not Your Service Is Up!" Everything SysAdmin blog, November 27, 2013.
<http://everythingsysadmin.com/2013/11/stop-monitoring-if-service-is-up.html>.
- Limoncelli, Tom. "Yes, You Can Really Work from HEAD." Everything SysAdmin blog, March 15, 2014. <http://everythingsysadmin.com/2014/03/yes-you-really-can-work-from-head.html>.
- Lindsay, Jeff. "Consul Service Discovery with Docker." Progridium blog, August 20, 2014. <http://progridium.com/blog/2014/08/20/consul-service-discovery-with-docker>.
- Loura, Ralph, Olivier Jacques, and Rafael Garcia. "DOES15—Ralph Loura, Olivier Jacques, & Rafael Garcia—Breaking Traditional IT Paradigms to ..." Posted by DevOps Enterprise Summit, November 16, 2015. YouTube video, 31:07.
https://www.youtube.com/watch?v=q9nNqqie_sM.
- Lublinsky, Boris. "Versioning in SOA." *The Architecture Journal*, April 2007.
<https://msdn.microsoft.com/en-us/library/bb491124.aspx>.
- Lund University. "Just Culture: Balancing Safety and Accountability." Human Factors & System Safety website, November 6, 2015.
<http://www.humanfactors.lth.se/sidney-dekker/books/just-culture/>.

- Luyten, Stefan. "Single Piece Flow: Why Mass Production Isn't the Most Efficient Way of Doing 'Stuff.'" *Medium* (blog), August 8, 2014. <https://medium.com/@stefanluyten/single-piece-flow-5d2c2bec845b#.9o7sn74ns>.
- Macri, Bethany. "Morgue: Helping Better Understand Events by Building a Post Mortem Tool—Bethany Macri." Posted by info@devopsdays.org, October 18, 2013. Vimeo video, 33:34. <http://vimeo.com/77206751>.
- Malpass, Ian. "DevOpsDays Minneapolis 2014—Ian Malpass, Fallible Humans." Posted by DevOps Minneapolis, July 20, 2014. YouTube video, 35:48. <https://www.youtube.com/watch?v=5NY-SrQFrBU>.
- Malpass, Ian. "Measure Anything, Measure Everything." *Code as Craft* (blog). Etsy, February 15, 2011. <http://codeascraft.com/2011/02/15/measure-anything-measure-everything/>.
- Mangot, Dave, and Karthik Rajan. "Agile.2013.effecting.a.devops.transformation.at.salesforce." Posted by Dave Mangot to slideshare.net, August 12, 2013. Slideshow. <http://www.slideshare.net/dmangot/agile2013effectingadevopstransformationatsalesforce>.
- Marsh, Dianne. "Dianne Marsh: 'Introducing Change while Preserving Engineering Velocity.'" Posted by Flowcon, November 11, 2014. YouTube video, 17:37. <https://www.youtube.com/watch?v=eW3ZxY67fnc>.
- Martin, Karen, and Mike Osterling. *Value Stream Mapping: How to Visualize Work and Align Leadership for Organizational Transformation*. New York: McGraw Hill, 2013.
- Maskell, Brian. "What Does This Guy Do? Role of Value Stream Manager." *Maskell* (blog), July 3, 2015. <http://blog.maskell.com/?p=2106http://www.lean.org/common/display/?o=221>.
- Masli, Adi., Vernon J. Richardson, Marcia Widenmier Watson, and Robert W. Zmud. "Senior Executives' IT Management Responsibilities: Serious IT-Related Deficiencies and CEO/CFO Turnover." *MIS Quarterly* 40, no. 3 (2016): 687–708. <https://doi.org/10.25300/misq/2016/40.3.08>.
- Massachusetts Institute of Technology. "Creating High Velocity Organizations." Course Descriptions. Accessed May 30, 2016. <http://executive.mit.edu/openenrollment/program/organizational-development-high-velocity-organizations>.
- Mathew, Reena, and Dave Mangot. "DOES14—Reena Mathew and Dave Mangot—Salesforce." Posted by ITRevolution to slideshare.net, October 29, 2014. Slideshow. <http://www.slideshare.net/ITRevolution/does14-reena-matthew-and-dave-mangot-salesforce>.
- Mauro, Tony. "Adopting Microservices at Netflix: Lessons for Architectural Design." *NGINX* (blog), February 19, 2015.

<https://www.nginx.com/blog/microservices-at-netflix-architectural-best-practices/>.

McDonnell, Patrick. “Continuously Deploying Culture: Scaling Culture at Etsy—Velocity Europe 2012.” Posted by Patrick McDonnell to [slideshare.net](https://www.slideshare.net), October 4, 2012. Slideshow. <http://www.slideshare.net/mcdonnps/continuously-deploying-culture-scaling-culture-at-etsy-14588485>.

McKinley, Dan. “Why MongoDB Never Worked Out at Etsy.” *Dan McKinley* (blog), December 26, 2012. <http://mcfunley.com/why-mongodb-never-worked-out-at-etsy>.

Mell, Peter, and Timothy Grance. *The NIST Definition of Cloud Computing: Recommendations of the National Institute of Standards and Technology*. Washington, DC: National Institute of Standards and Technology, 2011.

Messeri, Eran. “What Goes Wrong When Thousands of Engineers Share the Same Continuous Build?” Presented at the GOTO Conference, Aarhus, Denmark, October 2, 2013.

Metz, Cade. “Google Is 2 Billion Lines of Code—and It’s All in One Place.” *Wired*, September 16, 2015. <http://www.wired.com/2015/09/google-2-billion-lines-codeand-one-place/>.

Metz, Cade. “How Three Guys Rebuilt the Foundation of Facebook.” *Wired*, June 10, 2013. <http://www.wired.com/wiredenterprise/2013/06/facebook-hhvm-saga/all/>.

Mickman, Heather, and Ross Clanton. “DOES15—Heather Mickman & Ross Clanton—(Re)building an Engineering Culture: DevOps at Target.” Posted by DevOps Enterprise Summit, November 5, 2015. YouTube video, 33:39. <https://www.youtube.com/watch?v=7s-VbB1fG5o>.

Milstein, Dan. “Post-Mortems at HubSpot: What I Learned from 250 Whys.” *HubSpot* (blog), June 1, 2011. <http://product.hubspot.com/blog/bid/64771/Post-Mortems-at-HubSpot-What-I-Learned-From-250-Whys>.

Morgan, Timothy Prickett. “A Rare Peek Into The Massive Scale of AWS.” *Enterprise AI*, November 14, 2014. <http://www.enterprisetech.com/2014/11/14/rare-peek-massive-scale>.

Morrison, Erica. “How We Turned Our Company’s Worst Outage into a Powerful Learning Opportunity.” Presentation at DevOps Enterprise Summit, London, 2020. <https://videolibrary.doesvirtual.com/?video=431872263>.

Moore, Geoffrey A., and Regis McKenna. *Crossing the Chasm: Marketing and Selling High-Tech Products to Mainstream Customers*. New York: HarperCollins, 2009.

Mueller, Ernest. “2012—A Release Odyssey.” Posted by Ernest Mueller to [slideshare.net](https://www.slideshare.net), March 12, 2014. Slideshow. <http://www.slideshare.net/mxyzplk/2012-a-release-odyssey>.

- Mueller, Ernest. "Business Model Driven Cloud Adoption: What NI Is Doing in the Cloud." posted by Ernest Mueller to [slideshare.net](http://www.slideshare.net/mxyzplk/business-model-driven-cloud-adoption-what-ni-is-doing-in-the-cloud), June 28, 2011. Slideshow. <http://www.slideshare.net/mxyzplk/business-model-driven-cloud-adoption-what-ni-is-doing-in-the-cloud>.
- Mueller, Ernest. "DOES15—Ernest Mueller—DevOps Transformations at National Instruments and ..." Posted by DevOps Enterprise Summit, November 5, 2015. YouTube video, 34:14. <https://www.youtube.com/watch?v=6Ry40h1UAYE>.
- Mulkey, Jody. "DOES15—Jody Mulkey—DevOps in the Enterprise: A Transformation Journey." Posted by DevOps Enterprise Summit, November 5, 2015. YouTube video, 28:22. <https://www.youtube.com/watch?v=USYrDaPEFtM>.
- Nagappan, Nachiappan, E. Michael Maximilien, Thirumalesh Bhat, and Laurie Williams. "Realizing Quality Improvement through Test Driven Development: Results and Experiences of Four Industrial Teams." *Empire Software Engineering* 13 (2008): 289–302. http://research.microsoft.com/en-us/groups/ese/nagappan_tdd.pdf.
- Naraine, Ryan. "Twilio, HashiCorp Among Codecov Supply Chain Hack Victims," SecurityWeek, May 10, 2021. <https://www.securityweek.com/twilio-hashicorp-among-codecov-supply-chain-hack-victims>.
- Nationwide. *2014 Annual Report*. 2014 <https://www.nationwide.com/about-us/nationwide-annual-report-2014.jsp>.
- Nielsen, Jonas Klit. "8 Years with LinkedIn—Looking at the Growth." Mind Jumpers blog, May 10, 2011. <http://www.mindjumpers.com/blog/2011/05/linkedin-growth-infographic/>.
- Netflix. Letter to Shareholders, January 19, 2016. http://files.shareholder.com/downloads/NFLX/2432188684x0x870685/C6213FF9-5498-4084-A0FF-74363CEE35A1/Q4_15_Letter_to_Shareholders_-_COMBINED.pdf.
- Newland, Jesse. "ChatOps at GitHub." Posted on speakerdeck.com, February 7, 2013. Slideshow. <https://speakerdeck.com/jnewland/chatops-at-github>.
- North, Dan. "Ops and Operability." Posted to speakerdeck.com, February 25, 2016. Slideshow. <https://speakerdeck.com/tastapod/ops-and-operability>.
- "NUMMI." *This American Life* episode 403, March 26, 2010. Radio. <http://www.thisamericanlife.org/radio-archives/episode/403/transcript>.
- Nygard, Michael T. *Release It!: Design and Deploy Production-Ready Software*. Raleigh, NC: Pragmatic Bookshelf, 2007, Kindle.
- O'Donnell, Glenn. "DOES14—Glenn O'Donnell—Forrester—Modern Services Demand a DevOps Culture Beyond Apps." Posted by DevOps Enterprise Summit 2014, November 5, 2014. YouTube video, 12:20. https://www.youtube.com/watch?v=pvPWKuO4_48.

- O'Reilly, Barry. "How to Implement Hypothesis-Driven Development." Barry O'Reilly blog, October 21, 2013. <http://barryoreilly.com/explore/blog/how-to-implement-hypothesis-driven-development/>.
- Osterweil, Leon. "Software Processes Are Software Too." Paper presented at International Conference on Software Engineering, Monterey, CA, 1987. <http://www.cs.unibo.it/cianca/wwwpages/ids/lettura/Osterweil.pdf>.
- OWASP. "OWASP Cheat Sheet Series." Updated March 2, 2016. https://www.owasp.org/index.php/OWASP_Cheat_Sheet_Series.
- Özil, Giray. "Ask a programmer to review 10 lines of code." Twitter, February 27, 2013. <https://twitter.com/girayozil/status/306836785739210752>.
- Pal, Tapabrata. "DOES15—Tapabrata Pal—Banking on Innovation & DevOps." Posted by DevOps Enterprise Summit, January 4, 2016. YouTube video, 32:57. <https://www.youtube.com/watch?v=bbWFCKGhxOs>.
- Parikh, Karan. "From a Monolith to Microservices + REST: The Evolution of LinkedIn's Architecture." Posted by Karan Parikh to slideshare.net, November 6, 2014. Slideshow. <http://www.slideshare.net/parikhk/restli-and-deco>.
- "Paul O'Neill." *Forbes*, October 11, 2001. <http://www.forbes.com/2001/10/16/poneill.html>.
- Paul, Ryan. "Exclusive: A Behind-the-Scenes Look at Facebook Release Engineering." *Ars Technica*, April 5, 2012. <http://arstechnica.com/business/2012/04/exclusive-a-behind-the-scenes-look-at-facebook-release-engineering/1/>.
- PCI Security Standards Council. "Glossary." Glossary of terms (website). Accessed May 30, 2016. https://www.pcisecuritystandards.org/pci_security/glossary.
- PCI Security Standards Council. *Payment Card Industry (PCI) Data Security Standards: Requirements and Security Assessment Procedures*, Version 3.1. PCI Security Standards Council, 2015, Section 6.3.2. https://webcache.googleusercontent.com/search?q=cache:hpRe2COzdzAJ:https://www.pcisecuritystandards.org/documents/PCI_DSS_v3-1_SAQ_D_Merchant_rev1-1.docx+&cd=2&hl=en&ct=clnk&gl=us.
- Pepitone, Julianne. "Amazon EC2 Outage Downs Reddit, Quora." *CNN Money*, April 22, 2011. http://money.cnn.com/2011/04/21/technology/amazon_server_outage/index.htm.
- Perrow, Charles. *Normal Accidents: Living with High-Risk Technologies*. Princeton, NJ: Princeton University Press, 1999.
- Plastic SCM. "Version Control History." History of version control. Accessed May 31, 2016. <https://www.plasticscm.com/version-control-history>.
- Pomeranz, Hal. "Queue Inversion Week." *Righteous IT*, February 12, 2009. <https://righteousit.wordpress.com/2009/02/12/queue-inversion-week/>.

- Poppendieck, Mary, and Tom Poppendieck. *Implementing Lean Software: From Concept to Cash*. Upper Saddle River, NJ: Addison-Wesley, 2007.
- Potvin, Rachel, and Josh Levenber. “Why Google Stores Billions of Lines of Code in a Single Repository.” *Communications of the ACM* 59, no.7 (July 2016): 78–87. <https://cacm.acm.org/magazines/2016/7/204032-why-google-stores-billions-of-lines-of-code-in-a-single-repository/fulltext>.
- “Post Event Retrospective—Part 1.” *Rally Blogs*, accessed May 31, 2016. <https://www.rallydev.com/blog/engineering/post-event-retrospective-part-i>.
- Protalinski, Emil. “Facebook Passes 1.55B Monthly Active Users and 1.01B Daily Active Users.” *Venture Beat*, November 4, 2015. <http://venturebeat.com/2015/11/04/facebook-passes-1-55b-monthly-active-users-and-1-01-billion-daily-active-users/>.
- Prugh, Scott. “Continuous Delivery.” Scaled Agile Framework. Updated February 14, 2013, <http://www.scaledagileframework.com/continuous-delivery/>.
- Prugh, Scott. “DOES14: Scott Prugh, CSG—DevOps and Lean in Legacy Environments.” Posted by DevOps Enterprise Summit to [slideshare.net](https://www.slideshare.net), November 14, 2014. Slideshow. <http://www.slideshare.net/DevOpsEnterpriseSummit/scott-prugh>.
- Prugh, Scott, and Erica Morrison. “DOES15—Scott Prugh & Erica Morrison—Conway & Taylor Meet the Strangler (v2.0).” Posted by DevOps Enterprise Summit, November 5, 2015. YouTube video, 29:39. <https://www.youtube.com/watch?v=tKdIHCL0DUg>.
- Prugh, Scott, and Erica Morrison. “When Ops Swallows Dev,” presentation at DevOps Enterprise Summit 2016. <https://videolibrary.doesvirtual.com/?video=524430639>.
- Puppet Labs and IT Revolution Press. *2013 State of DevOps Report*. Portland, OR: Puppet Labs, 2013. <http://www.exin-library.com/Player/eKnowledge/2013-state-of-devops-report.pdf>.
- Ratchitsky, Lenny. “7 Keys to a Successful Public Health Dashboard.” *Transparent Uptime*, December 1, 2008. <http://www.transparentuptime.com/2008/11/rules-for-successful-public-health.html>.
- Raymond, Eric S. “Conway’s Law.” Eric Raymond. Accessed May 31, 2016. <http://catb.org/~esr/jargon/>.
- Rembetsy, Michael, and Patrick McDonnell. “Continuously Deploying Culture: Scaling Culture at Etsy.” Posted by Patrick McDonnell.bl to [slideshare.net](https://www.slideshare.net), October 4, 2012. Slideshow. <http://www.slideshare.net/mcdonnps/continuously-deploying-culture-scaling-culture-at-etsy-14588485>.
- Ries, Eric. *The Lean Startup: How Today’s Entrepreneurs Use Continuous Innovation to Create Radically Successful Businesses*. New York: Random House, 2011. Audiobook.

- Ries, Eric. "Work in Small Batches." *Startup Lessons Learned* (blog), February 20, 2009. <http://www.startuplessonslearned.com/2009/02/work-in-small-batches.html>.
- Robbins, Jesse. "GameDay: Creating Resiliency Through Destruction—LISA11." Posted by Jesse Robbins to [slideshare.net](http://www.slideshare.net), December 7, 2011. Slideshow. <http://www.slideshare.net/jesserobbins/ameday-creating-resiliency-through-destruction>.
- Robbins, Jesse. "Hacking Culture at VelocityConf." posted by Jesse Robbins to [slideshare.net](http://www.slideshare.net), June 28, 2012. Slideshow. <http://www.slideshare.net/jesserobbins/hacking-culture-at-velocityconf>.
- Robbins, Jesse, Kripa Krishnan, John Allspaw, and Tom Limoncelli. "Resilience Engineering: Learning to Embrace Failure." *ACM Queue* 10, no. 9 (September 13, 2012). <https://queue.acm.org/detail.cfm?id=2371297>.
- Roberto, Michael, Richard M. J. Bohmer, and Amy C. Edmondson. "Facing Ambiguous Threats." *Harvard Business Review*, November 2006. <https://hbr.org/2006/11/facing-ambiguous-threats>.
- Rossi, Chuck. "Release Engineering and Push Karma: Chuck Rossi." Facebook, April 5, 2012. <https://www.facebook.com/notes/facebook-engineering/release-engineering-and-push-karma-chuck-rossi/10150660826788920>.
- Rossi, Chuck. "Ship early and ship twice as often." Facebook, August 3, 2012. <https://www.facebook.com/notes/facebook-engineering/ship-early-and-ship-twice-as-often/10150985860363920>.
- Rother, Mike. *Toyota Kata: Managing People for Improvement, Adaptiveness and Superior Results*. New York: McGraw Hill, 2010. Kindle.
- Rubinstein, Joshua S., David E. Meyer, and Jeffrey E. Evans. "Executive Control of Cognitive Processes in Task Switching." *Journal of Experimental Psychology: Human Perception and Performance* 27, no. 4 (2001): 763–797. <http://www.umich.edu/~bcalab/documents/RubinsteinMeyerEvans2001.pdf>.
- Senge, Peter M. *The Fifth Discipline: The Art & Practice of the Learning Organization*. New York: Doubleday, 2006.
- Sharwood, Simon. "Are Your Servers PETS or CATTLE?" *The Register*, March 18 2013. http://www.theregister.com/2013/03/18/servers_pets_or_cattle_cern/.
- Shingo, Shigeo. *A Study of the Toyota Production System: From an Industrial Engineering Viewpoint*. London: Productivity Press, 1989.
- Shinn, Bill. "DOES15—Bill Shinn—Prove it! The Last Mile for DevOps in Regulated Organizations." Posted by Gene Kim to [slideshare.net](http://www.slideshare.net), November 20, 2015. Slideshow. <http://www.slideshare.net/ITRevolution/does15-bill-shinn-prove-it-the-last-mile-for-devops-in-regulated-organizations>.

- Shook, John. "Five Missing Pieces in Your Standardized Work (Part 3 of 3)." Lean Enterprise Institute, October 27, 2009. <http://www.lean.org/shook/DisplayObject.cfm?o=1321>.
- Shoup, Randy. "Exploring the Uncharted Territory of Microservices." Posted by XebiaLabs, Inc., February 20, 2015. YouTube video, 56:50. <https://www.youtube.com/watch?v=MRa21icSIQk>.
- Shoup, Randy. "The Virtuous Cycle of Velocity: What I Learned About Going Fast at eBay and Google by Randy Shoup." Posted by Flowcon, December 26, 2013. YouTube video, 30:05. <https://www.youtube.com/watch?v=EwLBoRyXTOI>.
- Skinner, Chris. "Banks Have Bigger Development Shops than Microsoft." *Chris Skinner's Blog*, September 9, 2011. <http://thefinanser.com/2011/09/banks-have-bigger-development-shops-than-microsoft.html/>.
- Smart, Jonathan, Zsolt Berend, Myles Ogilvie, and Simon Rohrer. *Sooner Safer Happier: Antipatterns and Patterns for Business Agility*. Portland, OR: IT Revolution, 2020.
- Snyder, Ross. "Scaling Etsy: What Went Wrong, What Went Right." Posted by Ross Snyder to slideshare.net, October 5, 2011. Slideshow. <http://www.slideshare.net/beamrider9/scaling-etsy-what-went-wrong-what-went-right>.
- Snyder, Ross. "Surge 2011—Scaling Etsy: What Went Wrong, What Went Right." Posted by OmniTiSurge Conference, December 23, 2011. YouTube video, 37:17. <https://www.youtube.com/watch?v=eenrfm50mXw>.
- Sonatype. *2015 State of the Software Supply Chain Report: Hidden Speed Bumps on the Way to "Continuous."* Fulton, MD: Sonatype, Inc., 2015. http://cdn2.hubspot.net/hubfs/1958393/White_Papers/2015_State_of_the_Software_Supply_Chain_Report-.pdf?t=1466775053631.
- Sonatype. *2019 State of the Software Supply Chain Report*. 2019. <https://www.sonatype.com/resources/white-paper-state-of-software-supply-chain-report-2019>.
- Sonatype. *2020 State of the Software Supply Chain Report*. 2020. <https://www.sonatype.com/resources/white-paper-state-of-the-software-supply-chain-2020>.
- Sowell, Thomas. *Basic Economics*, Fifth Edition. New York: Basic Books, 2014.
- Sowell, Thomas. *Knowledge and Decisions*. New York: Basic Books, 1980.
- Spear, Steven J. *The High-Velocity Edge: How Market Leaders Leverage Operational Excellence to Beat the Competition*. New York: McGraw Hill Education, 2009.
- Srivastava, Shivam, Kartik Trehan, Dilip Wagle, and Jane Wang. "Developer Velocity: How Software Excellence Fuels Business Performance." *McKinsey*, April 20, 2020. <https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/developer-velocity-how-software-excellence-fuels-business-performance>.

- Staats, Bradley, and David M. Upton. "Lean Knowledge Work." *Harvard Business Review*, October 2011. <https://hbr.org/2011/10/lean-knowledge-work>.
- Strear, Chris. "Leadership Lessons Learned from Improving Flow in Hospital Settings using Theory of Constraints." Presentation at DevOps Enterprise Summit, Europe, 2021. <https://videolibrary.doesvirtual.com/?video=550704199>.
- Stehr, Nico, and Reiner Grundmann. *Knowledge: Critical Concepts*, vol. 3. London: Routledge, 2005.
- Sterling, Bruce. "Scenius, or Communal Genius." *Wired*, June 16, 2008. <https://www.wired.com/2008/06/scenius-or-comm/#:~:text=His%20actual%20definition%20is%3A%20%22Scenius,scenius%2C%20you%20act%20like%20genius>.
- Stillman, Jessica. "Hack Days: Not Just for Facebookers." *Inc.*, February 3, 2012. <http://www.inc.com/jessica-stillman/hack-days-not-just-for-facebookers.html>.
- Sussman, Noah and Laura Beth Denker. "Divide and Conquer." *Code as Craft* (blog). Etsy, April 20, 2011. <https://codeascraft.com/2011/04/20/divide-and-concur/>.
- Sussna, Jeff. "From Design Thinking to DevOps and Back Again: Unifying Design & Operations." Posted by William Evans, June 5, 2015. Vimeo video, 21:19. <https://vimeo.com/129939230>.
- Takeuchi, Hirotaka, and Ikujiro Nonaka. "New Product Development Game." *Harvard Business Review* (January 1986): 137-146.
- Taleb, Nicholas. *Antifragile: Things That Gain from Disorder* (Incerto). New York: Random House, 2012.
- Target. "All About Target." A Bullseye View. Accessed June 9, 2016. <https://corporate.target.com>.
- Temple-Raston, Dina. "A 'Worst Nightmare' Cyberattack: The Untold Story of the SolarWinds Hack," NPR, April 16, 2021. <https://www.npr.org/2021/04/16/985439655/a-worst-nightmare-cyberattack-the-untold-story-of-the-solarwinds-hack>.
- Thomas, John and Ashish Kumar. "Welcome to the Google Engineering Tools Blog." Google Engineering Tools blog, posted May 3, 2011. <http://google-engtools.blogspot.com/2011/05/welcome-to-google-engineering-tools.html>.
- Townsend, Mark L. Review of *The Phoenix Project: A Novel About IT, DevOps, and Helping Your Business Win*, by Gene Kim, Kevin Behr, and George Spafford. Amazon review, March 2, 2013. http://uedata.amazon.com/gp/customer-reviews/R1097DFODM12VD/ref=cm_cr_getr_d_rvw_ttl?ie=UTF8&ASIN=B00VATFAMI.
- Treynor, Ben. "Keys to SRE." Presented at Usenix SREcon14, Santa Clara, CA, May 30, 2014. <https://www.usenix.org/conference/srecon14/technical-sessions/presentation/keys-sre>.

- Tucci, Linda. "Four Pillars of PayPal's 'Big Bang' Agile Transformation." *TechTarget*, August 2014. <http://searchcio.techtarget.com/feature/Four-pillars-of-PayPals-big-bang-Agile-transformation>.
- Turnbull, James. *The Art of Monitoring*. Seattle, WA: Amazon Digital Services, 2016. Kindle.
- Twitter Engineering. "Hack Week @ Twitter." Twitter blog, January 25, 2012. <https://blog.twitter.com/2012/hack-week-twitter>.
- Van Den Elzen, Scott. Review of *The Phoenix Project: A Novel About IT, DevOps, and Helping Your Business Win*, by Gene Kim, Kevin Behr, and George Spafford. Amazon review, March 13, 2013. http://uedata.amazon.com/gp/customer-reviews/R2K95XEH5OL3Q5/ref=cm_cr_getr_d_rvw_ttl?ie=UTF8&ASIN=B00VATFAMI.
- Van Leeuwen, Evelijn and Kris Buytaert. "DOES15—Evelijn Van Leeuwen and Kris Buytaert—Turning Around the Containership." Posted by DevOps Enterprise Summit, December 21, 2015. YouTube video, 30:28. <https://www.youtube.com/watch?v=0GI4AMKvPc>.
- Van Kemande, Ron. "Nothing Beats Engineering Talent: The Agile Transformation at ING." Presented at the DevOps Enterprise Summit, London, UK, June 30–July 1, 2016.
- Vance, Ashlee. "Inside Operation InVersion, the Code Freeze that Saved LinkedIn." *Bloomberg*, April 11, 2013. <http://www.bloomberg.com/news/articles/2013-04-10/inside-operation-inversion-the-code-freeze-that-saved-linkedin>.
- Vance, Ashlee. "LinkedIn: A Story About Silicon Valley's Possibly Unhealthy Need for Speed." *Bloomberg*, April 30, 2013. <http://www.bloomberg.com/articles/2013-04-29/linkedin-a-story-about-silicon-valleys-possibly-unhealthy-need-for-speed>.
- Vault. "Nordstrom, Inc." Company Profile. Accessed March 30, 2021. <http://www.vault.com/company-profiles/retail/nordstrom,-inc/company-overview.aspx>.
- Velasquez, Nicole Forsgren, Gene Kim, Nigel Kersten, and Jez Humble. *2014 State of DevOps Report*. Portland, OR: Puppet Labs, IT Revolution Press, and ThoughtWorks, 2014. <https://services.google.com/fh/files/misc/state-of-devops-2014.pdf>.
- Verizon Wireless. *2014 Data Breach Investigations Report*. Verizon Enterprise Solutions, 2014. https://dti.delaware.gov/pdfs/rp_Verizon-DBIR-2014_en_xg.pdf.
- Verizon Wireless. *2021 Data Breach Investigations Report*. Verizon, 2021. <https://enterprise.verizon.com/resources/reports/2021-data-breach-investigations-report.pdf>.
- "VPC Best Configuration Practices." Flux7 blog, January 23, 2014. <http://blog.flux7.com/blogs/aws/vpc-best-configuration-practices>.

- Walsh, Mark. "Ad Firms Right Media, AdInterax Sell to Yahoo." *MediaPost*, October 18, 2006. <http://www.mediapost.com/publications/article/49779/ad-firms-right-media-adinterax-sell-to-yahoo.html>.
- Wang, Kendrick. "Etsy's Culture Of Continuous Experimentation and A/B Testing Spurs Mobile Innovation." Apptimize blog, January 30, 2014. <http://apptimize.com/blog/2014/01/etsy-continuous-innovation-ab-testing/>.
- "Weekly Top 10: Your DevOps Flavor." Electric Cloud, April 1, 2016. <http://electric-cloud.com/blog/2016/04/weekly-top-10-devops-flavor/>.
- West, David. "Water scrum-fall is-reality_of_agile_for_most." Posted by harsoft to [slideshare.net](http://www.slideshare.net), April 22, 2013. Slideshow. <http://www.slideshare.net/harsoft/water-scrumfall-isrealityofagileformost>.
- Westrum, Ron. "A Typology of Organisation Culture." *BMJ Quality & Safety* 13, no. 2 (2004): ii22–ii27. doi:10.1136/qshc.2003.009522.
- Westrum, Ron. "The Study of Information Flow: A Personal Journey." *Proceedings of Safety Science* 67 (August 2014): 58–63. https://www.researchgate.net/publication/261186680_The_study_of_information_flow_A_personal_journey.
- "What Happens to Companies That Get Hacked? FTC Cases." Posted by SuicidalSnowman to Giant Bomb forum, July 2012. <http://www.giantbomb.com/forums/off-topic-31/what-happens-to-companies-that-get-hacked-ftc-case-540466/>.
- "When will Google permit languages other than Python, C++, Java and Go to be used for internal projects?" Quora forum. Accessed May 29, 2016. <https://www.quora.com/When-will-Google-permit-languages-other-than-Python-C-Java-and-Go-to-be-used-for-internal-projects/answer/Neil-Kandalgaonkar>.
- "Which programming languages does Google use internally?" Quora forum. Accessed May 29, 2016. <https://www.quora.com/Which-programming-languages-does-Google-use-internally>.
- Wickett, James. "Attacking Pipelines—Security Meets Continuous Delivery." Posted by James Wickett to [slideshare.net](http://www.slideshare.net), June 11, 2014. Slideshow. <http://www.slideshare.net/wickett/attacking-pipelinessecurity-meets-continuous-delivery>.
- Wiggins, Adams. "The Twelve-Factor App." 12Factor, January 30, 2012. <http://12factor.net/>.
- Wikipedia. "Direct Marketing." Wikipedia. Updated May 28, 2016. https://en.wikipedia.org/wiki/Direct_marketing.
- Wikipedia. "Imposter Syndrome." *Wikipedia*. Updated November 17, 2020. [https://en.wikipedia.org/wiki/Impostor_syndrome#:~:text=Impostor%20syndrome%20\(also%20known%20as,exposed%20as%20a%20%22fraud%22](https://en.wikipedia.org/wiki/Impostor_syndrome#:~:text=Impostor%20syndrome%20(also%20known%20as,exposed%20as%20a%20%22fraud%22).

- Wikipedia. "Kaizen." Wikipedia. Updated May 12, 2016.
<https://en.wikipedia.org/wiki/Kaizen>.
- Wikipedia. "Kolmogorov–Smirnov Test." Wikipedia. Updated May 19, 2016.
http://en.wikipedia.org/wiki/Kolmogorov–Smirnov_test.
- Wikipedia. "Telemetry." Wikipedia. Updated May 5, 2016.
<https://en.wikipedia.org/wiki/Telemetry>.
- Willis, John. "Docker and the Three Ways of DevOps Part 1: The First Way—Systems Thinking." Docker blog, May 26, 2015.
<https://blog.docker.com/2015/05/docker-three-ways-devops/>.
- Winslow, Michael, Tamara Ledbetter, Adam Zimman, John Esser, Tim Judge, Carmen DeArdo. *Change in a Successful Organization: Avoid Complacency by Making a Case for Continuous Improvement*. Portland, OR: IT Revolution, 2020.
<https://myresources.itrevolution.com/id006657108/Change-in-a-Successful-Organization>.
- Wiseman, Ben. *2021 Work Trend Index: Annual Report: The Next Great Disruption Is Hybrid Work—Are We Ready?* Microsoft, March 22, 2021. https://ms-worklab.azureedge.net/files/reports/hybridWork/pdf/2021_Microsoft_WTI_Report_March.pdf
- Womack, Jim. *Gemba Walks*. Cambridge, MA: Lean Enterprise Institute, 2011). Kindle.
- Wong, Bruce, and Christos Kalantzis. "A State of Xen—Chaos Monkey & Cassandra." *The Netflix Tech Blog*, October 2, 2014.
<http://techblog.netflix.com/2014/10/a-state-of-xen-chaos-monkey-cassandra.html>.
- Wong, Eric. "Eric the Intern: The Origin of InGraphs." LinkedIn Engineering blog, June 30, 2011. <http://engineering.linkedin.com/32/eric-intern-origin-ingraphs>.
- Womack, James P., and Daniel T. Jones. *Lean Thinking: Banish Waste and Create Wealth in Your Corporation*. New York: Free Press, 2010.
- Zhao, Haiping. "HipHop for PHP: Move Fast." Facebook, February 2, 2010.
<https://www.facebook.com/notes/facebook-engineering/hiphop-for-php-move-fast/280583813919>.
- Zia, Mossadeq, Gabriel Ramírez, and Noah Kunin. "Compliance Masonry: Building a Risk Management Platform, Brick by Brick." 18F blog, April 15, 2016.
<https://18f.gsa.gov/2016/04/15/compliance-masonry-building-a-risk-management-platform/>.

NOTES

Note from the Publisher

1. Kim, “State of DevOps: 2020 and Beyond.”

Preface

1. Branden Williams, personal correspondence with the authors, 2015.
2. Christopher Little, personal correspondence with Gene Kim, 2010.

Introduction

1. Goldratt, *Beyond the Goal*.
2. Immelt, “Let’s Finally End the Debate.”
3. “Weekly Top 10: Your DevOps Flavor,” *Electric Cloud*.
4. Goldratt, *Beyond the Goal*.
5. Spear, *The High-Velocity Edge*, [Chapter 3](#).
6. Christopher Little, personal correspondence with Gene Kim, 2010.
7. Skinner, “Banks Have Bigger Development Shops than Microsoft.”
8. Stehr and Grundmann, *Knowledge*, 139.
9. Masli et al., “Senior Executive’s IT Management Responsibilities.”
10. “IDC Forecasts Worldwide IT Spending to Grow 6%,” Business Wire.
11. Kersten, IT Revolution, and PwC, *2015 State of DevOps Report*.
12. Azzarello, Debruyne, and Mottura, “The Chemistry of Enthusiasm.”
13. Brooks, *The Mythical Man-Month*.
14. Kim et al., “Exploring the Uncharted Territory of Microservices.”
15. Kersten, IT Revolution, and PwC, *2015 State of DevOps Report*.
16. Jenkins, “Velocity Culture”; Exner, “Transforming Software Development.”
17. Goldratt, *Beyond the Goal*.
18. JGFLL, review of *The Phoenix Project*; Townsend, review of *The Phoenix Project*; Van Den Elzen, review of *The Phoenix Project*.

Part I Introduction

1. Beck, et al., “Twelve Principles of Agile Software.”
2. Rother, *Toyota Kata*, [Part III](#).

Chapter 1

1. Martin and Osterling, *Value Stream Mapping*, [Chapter 1](#).
2. Martin and Osterling, *Value Stream Mapping*, [Chapter 3](#).
3. Martin and Osterling, *Value Stream Mapping*, [Chapter 3](#).
4. Kersten, *Project to Product*.
5. Forsgren, Humble, and Kim, *Accelerate 2018*.
6. Leibman and Clanton, “DevOps: Approaching Cruising Altitude.”
7. Leibman and Clanton, “DevOps: Approaching Cruising Altitude.”
8. Leibman and Clanton, “DevOps: Approaching Cruising Altitude.”
9. Leibman and Clanton, “DevOps: Approaching Cruising Altitude.”
10. Leibman and Clanton, “DevOps: Approaching Cruising Altitude.”
11. Leibman and Clanton, “DevOps: Approaching Cruising Altitude.”
12. Leibman and Clanton, “DevOps: Approaching Cruising Altitude.”
13. Leibman and Clanton, “DevOps: Approaching Cruising Altitude.”
14. Leibman and Clanton, “DevOps: Approaching Cruising Altitude.”

Chapter 2

1. Rubinstein, Meyer, and Evans, “Executive Control of Cognitive Processes in Task Switching.”
2. DeGrandis, “DOES15—Dominica DeGrandis—The Shape of Uncertainty.”
3. Bahri, “Few Patients-In-Process and Less Safety Scheduling.”
4. Meeting between David J. Andersen and team at Motorola with Daniel S. Vacanti, February 24, 2004; story retold at USC CSSE Research Review with Barry Boehm in March 2004.
5. Womack and Jones, *Lean Thinking*. [Chapter 1](#).
6. Ries, “Work in Small Batches.”
7. Goldratt, *Beyond the Goal*.
8. Goldratt, *The Goal*, “Five Focusing Steps.”
9. Shingo, *A Study of the Toyota Production System*.
10. Poppendieck and Poppendieck, *Implementing Lean Software*, 74.
11. Poppendieck and Poppendieck, *Implementing Lean Software*, [Chapter 4](#).
12. Edwards, “DevOps Kaizen.”
13. Strear, “Leadership Lessons Learned From Improving Flow.”
14. Strear, “Leadership Lessons Learned From Improving Flow.”
15. Strear, “Leadership Lessons Learned From Improving Flow.”

Chapter 3

1. Perrow, *Normal Accidents*.
2. Dekker, *The Field Guide to Understanding Human Error*.
3. Spear, *The High-Velocity Edge*, [Chapter 8](#).
4. Spear, *The High-Velocity Edge*, [Chapter 8](#).
5. Senge, *The Fifth Discipline*, [Chapter 5](#).
6. “NUMMI,” *This American Life*.
7. Hendrickson, “DOES15—Elisabeth Hendrickson—Its All About Feedback.”
8. Hendrickson, “DOES15—Elisabeth Hendrickson—Its All About Feedback.”
9. Spear, *The High-Velocity Edge*, [Chapter 1](#).
10. Spear, *The High-Velocity Edge*, [Chapter 4](#).
11. Ayers and Cohen, “Andon Cords in Development Teams.”
12. Ayers and Cohen, “Andon Cords in Development Teams.”
13. Ayers and Cohen, “Andon Cords in Development Teams.”
14. Jeff Gallimore, personal correspondence with the authors, 2021.
15. Sowell, *Knowledge and Decisions*, 222.
16. Sowell, *Basic Economics*.
17. Gary Gruver, personal correspondence with Gene Kim, 2014.

Chapter 4

1. Adler, “Time-and-Motion Regained.”
2. Dekker, *The Field Guide to Understanding Human Error*, [Chapter 1](#).
3. Dekker, “Just Culture: Balancing Safety and Accountability.”
4. Westrum, “The Study of Information Flow.”
5. Westrum, “A Typology of Organisation Culture.”
6. Velasquez et al., *2014 State of DevOps Report*.
7. Macri, “Morgue.”
8. Spear, *The High-Velocity Edge*, [Chapter 1](#).
9. Senge, *The Fifth Discipline*, [Chapter 1](#).
10. Rother, *Toyota Kata*, 12.
11. Mike Orzen, personal correspondence with Gene Kim, 2012.
12. “Paul O’Neill,” *Forbes*.
13. Spear, *The High-Velocity Edge*, [Chapter 4](#).

14. Spear, *The High-Velocity Edge*, [Chapter 4](#).
15. Spear, *The High-Velocity Edge*, [Chapter 4](#).
16. Spear, *The High-Velocity Edge*, [Chapter 4](#).
17. Taleb, *Antifragile*.
18. Womack, *Gemba Walks*, Kindle location 4113.
19. Rother, *Toyota Kata*, [Part IV](#).
20. Rother, *Toyota Kata*, Conclusion.
21. Winslow et al., Change in a Successful Organization.
22. Gertner, *The Idea Factory*.
23. Kersten, *Project to Product*; Kersten, “Project to Product: From Stories to Scenius.”
24. Brian Eno, as quoted in Sterling, “Scenius, or Communal Genius.”
25. Gertner, *The Idea Factory*.
26. Gertner, *The Idea Factory*.

Chapter 5

1. Rembetsy and McDonnell, “Continuously Deploying Culture.”
2. “Nordstrom, Inc.,” Vault (website).
3. Kissler, “DOES14—Courtney Kissler—Nordstrom.”
4. Gardner, “Barnes & Noble, Blockbuster, Borders.”
5. Kissler, “DOES14—Courtney Kissler—Nordstrom.”
6. Kissler, “DOES14—Courtney Kissler—Nordstrom” [Alterations to quote made by Courtney Kissler via personal correspondence with Gene Kim, 2016.]
7. Kissler, “DOES14—Courtney Kissler—Nordstrom” [Alterations to quote made by Courtney Kissler via personal correspondence with Gene Kim, 2016.]
8. Kissler, “DOES14—Courtney Kissler—Nordstrom” [Alterations to quote made by Courtney Kissler via personal correspondence with Gene Kim, 2016.]
9. Kissler, “DOES14—Courtney Kissler—Nordstrom” [Alterations to quote made by Courtney Kissler via personal correspondence with Gene Kim, 2016.]
10. Kissler, “DOES14—Courtney Kissler—Nordstrom” [Alterations to quote made by Courtney Kissler via personal correspondence with Gene Kim, 2016.]
11. Mueller, “Business Model Driven Cloud Adoption.”
12. Unpublished calculation by Gene Kim after the 2014 DevOps Enterprise Summit.
13. Kersten, IT Revolution, and PwC, *2015 State of DevOps Report*.
14. Prugh, “DOES14: Scott Prugh, CSG.”

15. Rembetsy and McDonnell, "Continuously Deploying Culture."
16. Golden, "What Gartner's Bimodal IT Model Means to Enterprise CIOs."
17. Furtado and Knausenberger, "The Air Force's Digital Journey in 12 Parsecs or Less."
18. Furtado and Knausenberger, "The Air Force's Digital Journey in 12 Parsecs or Less."
19. Furtado and Knausenberger, "The Air Force's Digital Journey in 12 Parsecs or Less."
20. Furtado and Knausenberger, "The Air Force's Digital Journey in 12 Parsecs or Less."
21. Furtado and Knausenberger, "The Air Force's Digital Journey in 12 Parsecs or Less."
22. Furtado and Knausenberger, "The Air Force's Digital Journey in 12 Parsecs or Less."
23. Furtado and Knausenberger, "The Air Force's Digital Journey in 12 Parsecs or Less."
24. Golden, "What Gartner's Bimodal IT Model Means to Enterprise CIOs."
25. Golden, "What Gartner's Bimodal IT Model Means to Enterprise CIOs."
26. Kersten, IT Revolution, and PwC, *2015 State of DevOps Report*.
27. Scott Prugh, personal correspondence with Gene Kim, 2014.
28. Moore and McKenna, *Crossing the Chasm*, 11.
29. Tucci, "Four Pillars of PayPal's 'Big Bang' Agile Transformation."
30. Fernandez and Spear, "Creating High Velocity Organizations."
31. Van Kemande, "Nothing Beats Engineering Talent."
32. Leibman and Clanton, "DevOps: Approaching Cruising Altitude."
33. Leibman and Clanton, "DevOps: Approaching Cruising Altitude."
34. Leibman and Clanton, "DevOps: Approaching Cruising Altitude."
35. Leibman and Clanton, "DevOps: Approaching Cruising Altitude."
36. Leibman and Clanton, "DevOps: Approaching Cruising Altitude."
37. Leibman and Clanton, "DevOps: Approaching Cruising Altitude."
38. Leibman and Clanton, "DevOps: Approaching Cruising Altitude."
39. Conrad and Hyatt, "Saving the Economy from Ruin (with a Hyperscale PaaS)."
40. Conrad and Hyatt, "Saving the Economy from Ruin (with a Hyperscale PaaS)."
41. Conrad and Hyatt, "Saving the Economy from Ruin (with a Hyperscale PaaS)."
42. Conrad and Hyatt, "Saving the Economy from Ruin (with a Hyperscale PaaS)."
43. Conrad and Hyatt, "Saving the Economy from Ruin (with a Hyperscale PaaS)."

44. Conrad and Hyatt, "Saving the Economy from Ruin (with a Hyperscale PaaS)."
45. Conrad and Hyatt, "Saving the Economy from Ruin (with a Hyperscale PaaS)."
46. Conrad and Hyatt, "Saving the Economy from Ruin (with a Hyperscale PaaS)."
47. Conrad and Hyatt, "Saving the Economy from Ruin (with a Hyperscale PaaS)."
48. Conrad and Hyatt, "Saving the Economy from Ruin (with a Hyperscale PaaS)."
49. Buchanan, "The Wisdom of Peter Drucker from A to Z."

Chapter 6

1. Kissler, "DOES14—Courtney Kissler—Nordstrom."
2. Clanton and Ducy, interview of Courtney Kissler and Jason Josephy, "Continuous Improvement at Nordstrom."
3. Clanton and Ducy, interview of Courtney Kissler and Jason Josephy, "Continuous Improvement at Nordstrom."
4. Clanton and Ducy, interview of Courtney Kissler and Jason Josephy, "Continuous Improvement at Nordstrom."
5. Maskell, "What Does This Guy Do? Role of Value Stream Manager."
6. Edwards, "DevOps Kaizen."
7. Govindarajan and Trimble, *The Other Side of Innovation*.
8. Govindarajan and Trimble, *The Other Side of Innovation*, [Part I](#).
9. Cagan, *Inspired*, 12
10. Cagan, *Inspired*, 12.
11. Vance, "LinkedIn."
12. Clemm, "LinkedIn Started Back in 2003."
13. Nielsen, "8 Years with LinkedIn."
14. Clemm, "LinkedIn Started Back in 2003."
15. Parikh, "From a Monolith to Microservices + REST."
16. Clemm, "LinkedIn Started back in 2003."
17. Vance, "LinkedIn."
18. "How I Structured Engineering Teams at LinkedIn and AdMob for Success," *First Round Review*.
19. Vance, "Inside Operation InVersion."
20. Vance, "LinkedIn."
21. Clemm, "LinkedIn Started Back in 2003."
22. "How I Structured Engineering Teams," *First Round Review*.
23. Christopher Little, personal correspondence with Gene Kim, 2011.

24. Ryan Martens, personal correspondence with Gene Kim, 2013.

Chapter 7

1. Conway, “How Do Committees Invent?”
2. Conway, “How Do Committees Invent?”
3. Raymond, “Conway’s Law.”
4. Buhr, “Etsy Closes Up 86 Percent on First Day of Trading.”
5. Snyder, “Scaling Etsy.”
6. Snyder, “Scaling Etsy.”
7. Gallagher, “When ‘Clever’ Goes Wrong.”
8. Snyder, “Scaling Etsy.”
9. Snyder, “Scaling Etsy.”
10. Forsgren et al., *Accelerate State of DevOps 2019*.
11. Snyder, “Scaling Etsy.”
12. Snyder, “Surge 2011.”
13. Snyder, “Surge 2011.”
14. Snyder, “Surge 2011.”
15. McDonnell, “Continuously Deploying Culture.”
16. Fernandez and Spear, “Creating High Velocity Organizations.”
17. Adrian Cockcroft, personal correspondence with Gene Kim, 2014.
18. “A Conversation with Werner Vogels.”
19. Forsgren, Humble, and Kim, *Accelerate State of DevOps 2018*; Forsgren et al., *Accelerate State of DevOps 2019*..
20. Spear, *The High-Velocity Edge*, [Chapter 8](#).
21. Rother, *Toyota Kata*, 250.
22. Mulkey, “DOES15—Jody Mulkey.”
23. Mulkey, “DOES15—Jody Mulkey.”
24. Canahuati, “Growing from the Few to the Many.”
25. Spear, *The High-Velocity Edge*, [Chapter 1](#).
26. Prugh, “*Continuous Delivery*.”
27. Prugh, “*Continuous Delivery*.”
28. Prugh, “*Continuous Delivery*.”
29. Dweck, “Carol Dweck Revisits the ‘Growth Mindset.’”
30. Cox, “Disney DevOps.”

31. John Lauderbach, personal conversation with Gene Kim, 2001.
32. Mauro, “Adopting Microservices at Netflix”; Wiggins, “The Twelve-Factor App.”
33. Shoup, “Exploring the Uncharted Territory of Microservices.”
34. Humble, O’Reilly, and Molesky, *Lean Enterprise*, [Part III](#).
35. Hastings, “Netflix Culture.”
36. Dignan, “Little Things Add Up.”
37. Mickman and Clanton, “DOES15—Heather Mickman & Ross Clanton.”
38. Mickman and Clanton, “DOES15—Heather Mickman & Ross Clanton.”
39. Mickman and Clanton, “DOES15—Heather Mickman & Ross Clanton.”
40. Mickman and Clanton, “DOES15—Heather Mickman & Ross Clanton.”
41. Mickman and Clanton, “DOES15—Heather Mickman & Ross Clanton.”
42. Mickman and Clanton, “DOES15—Heather Mickman & Ross Clanton.”
43. Mickman and Clanton, “DOES15—Heather Mickman & Ross Clanton.”
44. Mickman and Clanton, “DOES15—Heather Mickman & Ross Clanton.”

Chapter 8

1. “Big Fish Celebrates 11th Consecutive Year of Record Growth,” Big Fish Games (website).
2. Paul Farrall, personal correspondence with Gene Kim, January 2015.
3. Paul Farrall, personal correspondence with Gene Kim, 2014.
4. Paul Farrall, personal correspondence with Gene Kim, 2014.
5. Ernest Mueller, personal correspondence with Gene Kim, 2014.
6. Edwards, “DevOps Kaizen.”
7. Marsh, “Dianne Marsh ‘Introducing Change while Preserving Engineering Velocity.’”
8. Cox, “Disney DevOps.”
9. Daniels, “Devopsdays Minneapolis 2015—Katherine Daniels—DevOps: The Missing Pieces.”
10. Ernest Mueller, personal correspondence with Gene Kim, 2015.
11. Takeuchi and Nonaka, “New Product Development Game.”
12. Chapman and Eltridge. “On A Mission: Nationwide Building Society.”
13. Chapman and Eltridge. “On A Mission: Nationwide Building Society.”
14. Chapman and Eltridge. “On A Mission: Nationwide Building Society.”
15. Chapman and Eltridge. “On A Mission: Nationwide Building Society.”

16. Chapman and Eltridge. “On A Mission: Nationwide Building Society.”
17. Chapman and Eltridge. “On A Mission: Nationwide Building Society.”
18. Chapman and Eltridge. “On A Mission: Nationwide Building Society.”

Chapter 9

1. Campbell-Pretty, “DOES14—Em Campbell-Pretty—How a Business Exec Led Agile, Lead, CI/CD.”
2. Campbell-Pretty, “DOES14—Em Campbell-Pretty—How a Business Exec Led Agile, Lead, CI/CD.”
3. Campbell-Pretty, “DOES14—Em Campbell-Pretty—How a Business Exec Led Agile, Lead, CI/CD.”
4. Campbell-Pretty, “DOES14—Em Campbell-Pretty—How a Business Exec Led Agile, Lead, CI/CD.”
5. Campbell-Pretty, “DOES14—Em Campbell-Pretty—How a Business Exec Led Agile, Lead, CI/CD.”
6. Campbell-Pretty, “DOES14—Em Campbell-Pretty—How a Business Exec Led Agile, Lead, CI/CD.”
7. Campbell-Pretty, “DOES14—Em Campbell-Pretty—How a Business Exec Led Agile, Lead, CI/CD.”
8. “Version Control History,” Plastic SCM (website).
9. Davis and Daniels, *Effective DevOps*, 37.
10. Velasquez et al., *2014 State of DevOps Report*.
11. Sharwood, “Are Your Servers PETS or CATTLE?”
12. Chan, “OWASP AppSecUSA 2012.”
13. Fowler, “Trash Your Servers and Burn Your Code.”
14. Willis, “Docker and the Three Ways of DevOps [Part 1](#).”
15. Forsgren et al., *2020 State of the Octoverse*.
16. Holmes, “How A Hotel Company Ran \$30B of Revenue in Containers.”
17. Holmes, “How A Hotel Company Ran \$30B of Revenue in Containers.”
18. Holmes, “How A Hotel Company Ran \$30B of Revenue in Containers.”
19. Holmes, “How A Hotel Company Ran \$30B of Revenue in Containers.”
20. Holmes, “How A Hotel Company Ran \$30B of Revenue in Containers.”
21. Holmes, “How A Hotel Company Ran \$30B of Revenue in Containers.”
22. Holmes, “How A Hotel Company Ran \$30B of Revenue in Containers.”

Chapter 10

1. Gary Gruver, personal correspondence with Gene Kim, 2014.
2. Bland, "DOES15—Mike Bland—Pain Is Over, If You Want It."
3. Bland, "DOES15—Mike Bland—Pain Is Over, If You Want It."
4. "Imposter Syndrome," *Wikipedia*."
5. Bland, "DOES15—Mike Bland—Pain Is Over, If You Want It."
6. Bland, "DOES15—Mike Bland—Pain Is Over, If You Want It."
7. Bland, "DOES15—Mike Bland—Pain Is Over, If You Want It."
8. Potvin and Levenber, "Why Google Stores Billions of Lines of Codes in a Single Repository."
9. Messeri, "What Goes Wrong When Thousands of Engineers Share the Same Continuous Build?"
10. Messeri, "What Goes Wrong When Thousands of Engineers Share the Same Continuous Build?"
11. Potvin and Levenber, "Why Google Stores Billions of Lines of Codes in a Single Repository."
12. Potvin and Levenber, "Why Google Stores Billions of Lines of Codes in a Single Repository"; Messeri, "What Goes Wrong When Thousands of Engineers Share the Same Continuous Build?"
13. Jez Humble and David Farley, personal correspondence with Gene Kim, 2012.
14. Humble and Farley, *Continuous Delivery*, 3.
15. Humble and Farley, *Continuous Delivery*, 188.
16. Humble and Farley, *Continuous Delivery*, 258.
17. Fowler, "Continuous Integration."
18. Fowler, "Test Pyramid."
19. Fowler, "Test Driven Development."
20. Nagappan et al., "Realizing Quality Improvement through Test Driven Development."
21. Hendrickson, "On the Care and Feeding of Feedback Cycles."
22. "Decreasing False Positives in Automated Testing."; Fowler, "Eradicating Non-determinism in Tests."
23. Gruver, "DOES14—Gary Gruver—Macy's—Transforming Traditional Enterprise Software Development Processes."
24. Jones, "3 Ways to Get Test Automation Done Within Your Sprints."
25. Shoup, "The Virtuous Cycle of Velocity."
26. West, "Water scrum-fall is-reality_of_agile_for_most."
27. Forsgren et al., *Accelerate: State of DevOps 2019*.

28. Forsgren, Humble, and Kim, *Accelerate: State of DevOps 2018*.

Chapter 11

1. Kim, “The Amazing DevOps Transformation of the HP LaserJet Firmware Team.”
2. Kim, “The Amazing DevOps Transformation of the HP LaserJet Firmware Team.”
3. Kim, “The Amazing DevOps Transformation of the HP LaserJet Firmware Team.”
4. Kim, “The Amazing DevOps Transformation of the HP LaserJet Firmware Team.”
5. Kim, “The Amazing DevOps Transformation of the HP LaserJet Firmware Team.”
6. Gruver and Mouser, *Leading the Transformation*, 60.
7. Gary Gruver, personal communication with the authors, 2016.
8. Kim, “The Amazing DevOps Transformation of the HP LaserJet Firmware Team.”
9. Kim, “The Amazing DevOps Transformation of the HP LaserJet Firmware Team.”
10. Kim, “The Amazing DevOps Transformation of the HP LaserJet Firmware Team.”
11. Kim, “The Amazing DevOps Transformation of the HP LaserJet Firmware Team.”
12. Atwood, “Software Branching and Parallel Universes.”
13. Cunningham, “Ward Explains Debt Metaphor.”
14. Mueller, “2012: A Release Odyssey.”
15. “Bazaarvoice, Inc. Announces Its Financial Results,” Bazaar Voice (website).
16. Mueller, “DOES15—Ernest Mueller—DevOps Transformations At National Instruments.”
17. Mueller, “DOES15—Ernest Mueller—DevOps Transformations At National Instruments.”
18. Mueller, “DOES15—Ernest Mueller—DevOps Transformations At National Instruments.”
19. Mueller, “DOES15—Ernest Mueller—DevOps Transformations At National Instruments.”
20. Mueller, “DOES15—Ernest Mueller—DevOps Transformations At National Instruments”
21. Kersten, IT Revolution, and PwC, *2015 State of DevOps Report*.

22. Brown, et al., *State of DevOps Report*; Forsgren et al., *State of DevOps Report 2017*.

Chapter 12

1. Rossi, “Release Engineering and Push Karma.”
2. Paul, “Exclusive: A Behind-the-Scenes Look at Facebook Release Engineering.”
3. Rossi, “Release Engineering and Push Karma.”
4. Paul, “Exclusive: a Behind-the-Scenes Look at Facebook Release Engineering.”
5. Rossi, “Ship early and ship twice as often.”
6. Beck, “Slow Deployment Causes Meetings.”
7. Prugh, “DOES14: Scott Prugh, CSG—DevOps and Lean in Legacy Environments.”
8. Prugh, “DOES14: Scott Prugh, CSG—DevOps and Lean in Legacy Environments.”
9. Prugh, “DOES14: Scott Prugh, CSG—DevOps and Lean in Legacy Environments.”
10. Prugh, “DOES14: Scott Prugh, CSG—DevOps and Lean in Legacy Environments.”
11. Prugh, “DOES14: Scott Prugh, CSG—DevOps and Lean in Legacy Environments.”
12. Prugh, “DOES14: Scott Prugh, CSG—DevOps and Lean in Legacy Environments.”
13. Puppet Labs and IT Revolution Press, *2013 State of DevOps Report*.
14. Prugh and Morrison, “DOES15—Scott Prugh & Erica Morrison—Conway & Taylor Meet the Strangler (v2.0).”
15. Prugh and Morrison, “DOES15—Scott Prugh & Erica Morrison—Conway & Taylor Meet the Strangler (v2.0).”
16. Prugh and Morrison, “DOES15—Scott Prugh & Erica Morrison—Conway & Taylor Meet the Strangler (v2.0).”
17. Tim Tischler, personal conversation with Gene Kim, FlowCon 2013.
18. Puppet Labs and IT Revolution Press, *2013 State of DevOps Report*.
19. Forsgren et al. *Accelerate: State of DevOps 2019*.
20. Dickerson, “Optimizing for Developer Happiness.”
21. Sussman and Denker, “Divide and Conquer.”
22. Sussman and Denker, “Divide and Conquer.”
23. Sussman and Denker, “Divide and Conquer.”
24. Sussman and Denker, “Divide and Conquer.”

25. Sussman and Denker, “Divide and Conquer.”
26. Kastner, “Quantum of Deployment.”
27. Fitz, “Continuous Deployment at IMVU.”
28. Fitz, “Continuous Deployment at IMVU”; Hrenko, “DOES15—Michael Hrenko—DevOps Insured By Blue Shield of California.”
29. Humble and Farley, *Continuous Delivery*, 265.
30. Ries, *The Lean Startup*.
31. Bosworth, “Building and testing at Facebook”; “Etsy’s Feature Flagging,” GitHub (website).
32. Allspaw, “Convincing Management.”
33. Rossi, “Release Engineering and Push Karma.”
34. Protalinski, “Facebook Passes 1.55B Monthly Active Asers.”
35. Protalinski, “Facebook Passes 1.55B Monthly Active Users.”
36. Letuchy, “Facebook Chat.”
37. Letuchy, “Facebook Chat.”
38. Letuchy, “Facebook Chat.”
39. Jez Humble, personal correspondence with Gene Kim, 2014.
40. Jez Humble, personal correspondence with Gene Kim, 2014.
41. Jez Humble, personal correspondence with Gene Kim, 2014.
42. Forsgren, Humble, and Kim, *Accelerate State of DevOps 2018*; Forsgren et al., *Accelerate State of DevOps 2019*.
43. Prugh and Morrison, “When Ops Swallows Dev.”
44. Prugh and Morrison, “When Ops Swallows Dev.”
45. Prugh and Morrison, “When Ops Swallows Dev.”
46. Prugh and Morrison, “When Ops Swallows Dev.”
47. Prugh and Morrison, “When Ops Swallows Dev.”
48. Prugh and Morrison, “When Ops Swallows Dev.”

Chapter 13

1. Humble, “What is *Continuous Delivery*?”
2. Kim et al., “Exploring the Uncharted Territory of Microservices.”
3. Kim et al., “Exploring the Uncharted Territory of Microservices.”
4. Kim et al., “Exploring the Uncharted Territory of Microservices.”
5. Shoup, “From Monolith to Microservices.”
6. Betz, *Architecture and Patterns for IT Service Management*, 300.

7. Shoup, “From Monolith to Micro-services.”
8. Shoup, “From Monolith to Micro-services.”
9. Shoup, “From Monolith to Micro-services.”
10. Vogels, “A Conversation with Werner Vogels.”
11. Vogels, “A Conversation with Werner Vogels.”
12. Vogels, “A Conversation with Werner Vogels.”
13. Vogels, “A Conversation with Werner Vogels.”
14. Jenkins, “Velocity Culture.”
15. Exner, “Transforming Software Development.”
16. Fowler, “Strangler Fig Application.”
17. Lublinsky, “Versioning in SOA.”
18. Hammant, “Introducing Branch by Abstraction.”
19. Fowler, “Strangler Fig Application.”
20. Huang, “Blackboard CEO Jay Bhatt on the Global Future of Edtech.”
21. Ashman, “DOES14—David Ashman—Blackboard Learn—Keep Your Head in the Clouds.”
22. Ashman, “DOES14—David Ashman—Blackboard Learn—Keep Your Head in the Clouds.”
23. Ashman, “DOES14—David Ashman—Blackboard Learn—Keep Your Head in the Clouds.”
24. Ashman, “DOES14—David Ashman—Blackboard Learn—Keep Your Head in the Clouds.”
25. Ashman, “DOES14—David Ashman—Blackboard Learn—Keep Your Head in the Clouds.”
26. Forsgren et al., *State of DevOps Report 2017*.
27. Forsgren, Humble, and Kim, *Accelerate: State of DevOps 2018*; Forsgren et al., *Accelerate State of DevOps 2019*.

Chapter 14

1. Kim, Behr, and Spafford, *The Visible Ops Handbook*, Introduction.
2. Kim, Behr, and Spafford, *The Visible Ops Handbook*, Introduction.
3. Kim, Behr, and Spafford, *The Visible Ops Handbook*, Introduction.
4. “Telemetry,” *Wikipedia*.
5. Rembetsy and McDonnell, “Continuously Deploying Culture.”
6. Rembetsy and McDonnell, “Continuously Deploying Culture.”
7. John Allspaw, personal conversation with Gene Kim, 2014.

8. Malpass, “Measure Anything, Measure Everything.”
9. Kersten, IT Revolution, and PwC, *2015 State of DevOps Report*.
10. Forsgren et al., *Accelerate: State of DevOps 2019*.
11. Turnbull, *The Art of Monitoring*, Introduction.
12. Cockcroft, “Monitorama.”
13. Prugh, “DOES14: Scott Prugh, CSG—DevOps and Lean in Legacy Environments.”
14. Figureau, “The 10 Commandments of Logging.”
15. Dan North, personal correspondence with Gene Kim, 2016.
16. Chuvakin, “LogLogic/Chuvakin Log Checklist.”
17. Kim, Behr, and Spafford, *The Visible Ops Handbook*, Introduction.
18. North, “Ops and Operability.”
19. John Allspaw, personal correspondence with Gene Kim, 2011.
20. Agile Alliance, “Information Radiators.”
21. Ernest Mueller, personal correspondence with Gene Kim, 2014.
22. Gupta, “Visualizing LinkedIn’s Site Performance.”
23. Wong, “Eric the Intern.”
24. Wong, “Eric the Intern.”
25. Wong, “Eric the Intern.”
26. Ed Blankenship, personal correspondence with Gene Kim, 2016.
27. Burrows, “The Chubby Lock Service for Loosely-Coupled Distributed Systems.”
28. Lindsay, “Consul Service Discovery with Docker.”
29. Mulkey, “DOES15—Jody Mulkey—DevOps in the Enterprise: A Transformation Journey.”
30. Forsgren et al., *Accelerate: State of DevOps 2019*.

Chapter 15

1. Netflix Letter to Shareholders.
2. Roy Rapoport, personal correspondence with Gene Kim, 2014.
3. Hodge and Austin, “A Survey of Outlier Detection Methodologies.”
4. Roy Rapoport, personal correspondence with Gene Kim, 2014.
5. Roy Rapoport, personal correspondence with Gene Kim, 2014.
6. Roy Rapoport, personal correspondence with Gene Kim, 2014.
7. Boubez, “Simple Math for Anomaly Detection.”

8. Limoncelli, “Stop Monitoring Whether or Not Your Service Is Up!”
9. Boubez, “Simple Math for Anomaly Detection.”
10. Dr. Nicole Forsgren, personal correspondence with Gene Kim, 2015.
11. Jacobson, Yuan, and Joshi, “Scryer: Netflix’s Predictive Auto Scaling Engine.”
12. Jacobson, Yuan, and Joshi, “Scryer: Netflix’s Predictive Auto Scaling Engine.”
13. Jacobson, Yuan, and Joshi, “Scryer: Netflix’s Predictive Auto Scaling Engine.”
14. Chandola, Banerjee, and Kumar, “Anomaly Detection: A Survey.”
15. Tarun Reddy, personal interview with Gene Kim, Rally headquarters, Boulder, CO, 2014.
16. “Kolmogorov-Smirnov Test,” Wikipedia.
17. Boubez, “Simple Math for Anomaly Detection.”
18. Boubez, “Simple Math for Anomaly Detection.”

Chapter 16

1. Walsh, “Ad Firms Right Media.”
2. Nick Galbreath, personal conversation with Gene, 2013.
3. Galbreath, “Continuous Deployment.”
4. Galbreath, “Continuous Deployment.”
5. Galbreath, “Continuous Deployment.”
6. Canahuati, “Growing from the Few to the Many.”
7. Lightbody, “From Inception to Acquisition.”
8. Chakrabarti, “Common Ops Mistakes.”
9. Sussna, “From Design Thinking to DevOps and Back Again.”
10. Anonymous, personal conversation with Gene Kim, 2005.
11. Limoncelli, “SRE@Google.”
12. Treynor, “Keys to SRE.”
13. Limoncelli, “SRE@Google.”
14. Limoncelli, “SRE@Google.”
15. Limoncelli, “SRE@Google.”
16. Tom Limoncelli, personal correspondence with Gene Kim, 2016.
17. Tom Limoncelli, personal correspondence with Gene Kim, 2016.

Chapter 17

1. Humble, O’Reilly and Molesky, *Lean Enterprise*, [Part II](#).
2. Intuit, Inc., “2012 Annual Report.”

3. Cook, "Leadership in an Agile Age."
4. Cook, "Leadership in an Agile Age."
5. Cook, "Leadership in an Agile Age."
6. "Direct Marketing," Wikipedia.
7. "Fighting Poverty With Actual Evidence: Full Transcript," Freakonomics (blog).
8. Kohavi, Crook, and Longbotham, "Online Experimentation at Microsoft."
9. Kohavi, Crook, and Longbotham, "Online Experimentation at Microsoft."
10. Jez Humble, personal correspondence with Gene Kim, 2015.
11. Wang, "Etsy's Culture Of Continuous Experimentation."
12. O'Reilly, "How to Implement Hypothesis-Driven Development."
13. Kim, "Organizational Learning and Competitiveness."
14. Kim, "Organizational Learning and Competitiveness."
15. Kim, "Organizational Learning and Competitiveness."
16. Kim, "Organizational Learning and Competitiveness."
17. Kim, "Organizational Learning and Competitiveness."
18. Kim, "Organizational Learning and Competitiveness."
19. Kim, "Organizational Learning and Competitiveness."

Chapter 18

1. Chacon, "GitHub Flow."
2. Douglas, "Deploying at GitHub."
3. Allspaw, "Counterfactual Thinking, Rules, and the Knight Capital Accident."
4. Allspaw, "Counterfactual Thinking, Rules, and the Knight Capital Accident."
5. Staats and Upton, "Lean Knowledge Work."
6. Forsgren et al., *Accelerate State of DevOps 2019*.
7. Forsgren et al., *Accelerate State of DevOps 2019*.
8. Velasquez et al., *2014 State of DevOps Report*.
9. Randy Shoup, personal interview with Gene Kim, 2015.
10. Özil, "Ask a programmer."
11. Cornago, Yadav, and Otto, "From 6-Eye Principle to Release at Scale - adidas Digital Tech 2021."
12. Cornago, Yadav, and Otto, "From 6-Eye Principle to Release at Scale - adidas Digital Tech 2021."

13. Cornago, Yadav, and Otto, "From 6-Eye Principle to Release at Scale - adidas Digital Tech 2021."
14. Cornago, Yadav, and Otto, "From 6-Eye Principle to Release at Scale - adidas Digital Tech 2021."
15. Cornago, Yadav, and Otto, "From 6-Eye Principle to Release at Scale - adidas Digital Tech 2021."
16. Cornago, Yadav, and Otto, "From 6-Eye Principle to Release at Scale - adidas Digital Tech 2021."
17. Cornago, Yadav, and Otto, "From 6-Eye Principle to Release at Scale - adidas Digital Tech 2021."
18. Cornago, Yadav, and Otto, "From 6-Eye Principle to Release at Scale - adidas Digital Tech 2021."
19. Cornago, Yadav, and Otto, "From 6-Eye Principle to Release at Scale - adidas Digital Tech 2021."
20. Özil, "Ask a programmer to review 10 lines of code."
21. Messeri, "What Goes Wrong When Thousands of Engineers Share the Same Continuous Build?"
22. Thomas and Kumar, "Welcome to the Google Engineering Tools Blog."
23. Kumar, "Development at the Speed and Scale of Google."
24. Randy Shoup, personal correspondence with Gene Kim, 2014.
25. Atwood, "Pair Programming vs. Code Reviews."
26. Atwood, "Pair Programming vs. Code Reviews."
27. "Pair Programming," ALICE Wiki page.
28. "Pair Programming," ALICE Wiki page.
29. Hendrickson, "DOES15—Elisabeth Hendrickson—Its All About Feedback."
30. Hendrickson, "DOES15—Elisabeth Hendrickson—Its All About Feedback."
31. Hendrickson, "DOES15—Elisabeth Hendrickson—Its All About Feedback."
32. Hendrickson, "DOES15—Elisabeth Hendrickson—Its All About Feedback."
33. Hendrickson, "DOES15—Elisabeth Hendrickson—Its All About Feedback."
34. Ryan Tomayko, personal correspondence with Gene Kim, 2014.
35. Ryan Tomayko, personal correspondence with Gene Kim, 2014.
36. Ryan Tomayko, personal correspondence with Gene Kim, 2014.
37. Ryan Tomayko, personal correspondence with Gene Kim, 2014.
38. Ryan Tomayko, personal correspondence with Gene Kim, 2014.
39. Cockcroft, Ducey, and Clanton, "Adrian Cockcroft of Battery Ventures."
40. Pal, "DOES15—Tapabrata Pal—Banking on Innovation & DevOps."

41. Cox, “Disney DevOps.”
42. Clanton and Mickman, “DOES14—Ross Clanton and Heather Mickman—DevOps at Target.”
43. Clanton and Mickman, “DOES14—Ross Clanton and Heather Mickman—DevOps at Target.”
44. Clanton and Mickman, “DOES14—Ross Clanton and Heather Mickman—DevOps at Target.”
45. Clanton and Mickman, “DOES14—Ross Clanton and Heather Mickman—DevOps at Target.”
46. Clanton and Mickman, “DOES14—Ross Clanton and Heather Mickman—DevOps at Target.”
47. John Allspaw and Jez Humble, personal correspondence with Gene Kim, 2014.

Chapter 19

1. Spear, *The High-Velocity Edge*, [Chapter 1](#).
2. Spear, *The High-Velocity Edge*, [Chapter 1](#).
3. Pepitone, “Amazon EC2 Outage Downs Reddit, Quora.”
4. Morgan, “A Rare Peek into the Massive Scale of AWS.”
5. Cockcroft, Hicks, and Orzell, “Lessons Netflix Learned from the AWS Outage.”
6. Cockcroft, Hicks, and Orzell, “Lessons Netflix Learned from the AWS Outage.”
7. Cockcroft, Hicks, and Orzell, “Lessons Netflix Learned from the AWS Outage.”
8. Dekker, “Just Culture,” 152.
9. Dekker, “DevOpsDays Brisbane 2014—Sidney Decker—System Failure, Human Error: Who’s to Blame?”
10. Allspaw, “Post-Mortems, Sans Finger-Pointing.”
11. Allspaw, “Blameless PostMortems and a Just Culture.”
12. Malpass, “DevOpsDays Minneapolis 2014—Ian Malpass, Fallible Humans.”
13. Milstein, “Post-Mortems at HubSpot: What I Learned from 250 Whys.”
14. Randy Shoup, personal correspondence with Gene Kim, 2014.
15. Google, “Post-Mortem for February 24, 2010 Outage”; Amazon Web Services, “Summary of the Amazon DynamoDB Service Disruption and Related Impacts in the US-East Region.”
16. Macri, “Morgue.”
17. Macri, “Morgue.”
18. Forsgren, Humble, and Kim, *Accelerate: State of DevOps 2018*.

19. Edmondson, “Strategies for Learning from Failure.”
20. Spear, *The High-Velocity Edge*, [Chapter 4](#).
21. Spear, *The High-Velocity Edge*, [Chapter 4](#).
22. Spear, *The High-Velocity Edge*, [Chapter 3](#).
23. Roberto, Bohmer, and Edmondson, “Facing Ambiguous Threats.”
24. Roberto, Bohmer, and Edmondson, “Facing Ambiguous Threats.”
25. Roberto, Bohmer, and Edmondson, “Facing Ambiguous Threats.”
26. Roy Rapoport, personal correspondence with Gene Kim, 2012.
27. Roy Rapoport, personal correspondence with Gene Kim, 2012.
28. Roy Rapoport, personal correspondence with Gene Kim, 2012.
29. Nygard, *Release It!*, [Part I](#).
30. Barr, “EC2 Maintenance Update.”
31. Wong and Kalantzis, “A State of Xen—Chaos Monkey & Cassandra.”
32. Wong and Kalantzis, “A State of Xen—Chaos Monkey & Cassandra.”
33. Wong and Kalantzis, “A State of Xen—Chaos Monkey & Cassandra.”
34. Roy Rapoport, personal correspondence with Gene Kim, 2015.
35. Adrian Cockcroft, personal correspondence with Gene Kim, 2012.
36. Robbins, “GameDay.”
37. Robbins et al., “Resilience Engineering.”
38. Robbins et al., “Resilience Engineering.”
39. Robbins et al., “Resilience Engineering.”
40. Robbins et al., “Resilience Engineering.”
41. Robbins et al., “Resilience Engineering.”
42. Krishman, “‘Learning Continuously From Failures’ at Google.”
43. Krishnan, “Weathering the Unexpected.”
44. Krishnan, “Weathering the Unexpected.”
45. Morrison, “How We Turned Our Company’s Worst Outage into a Powerful Learning Opportunity.”
46. Widely attributed to Peter Senge.

Chapter 20

1. Newland, “ChatOps at GitHub.”
2. Newland, “ChatOps at GitHub.”
3. Mark Imbriaco, personal correspondence with Gene Kim, 2015.

4. Newland, “ChatOps at GitHub.”
5. Newland, “ChatOps at GitHub.”
6. Newland, “ChatOps at GitHub.”
7. Newland, “ChatOps at GitHub.”
8. Osterweil, “Software Processes are Software Too.”
9. Arbuckle, “What Is ArchOps: Chef Executive Roundtable.”
10. Arbuckle, “What Is ArchOps: Chef Executive Roundtable.”
11. Arbuckle, “What Is ArchOps: Chef Executive Roundtable.”
12. Metz, “Google Is 2 Billion Lines of Code—and It’s All in One Place.”
13. Metz, “Google Is 2 Billion Lines of Code—and It’s All in One Place.”
14. Metz, “Google Is 2 Billion Lines of Code—and It’s All in One Place.”
15. Messeri, “What Goes Wrong When Thousands of Engineers Share the Same Continuous Build?”
16. Randy Shoup, personal correspondence with Gene Kim, 2014.
17. Limoncelli, “Yes, You Can Really Work from HEAD.”
18. Forsgren et al., *Accelerate: State of DevOps 2019*.
19. Forsgren et al., *Accelerate: State of DevOps 2019*.
20. Mell and Grance, *The NIST Definition of Cloud Computing*, 6.
21. Forsgren et al., *Accelerate: State of DevOps 2019*.
22. Loura, Jacques, and Garcia, “DOES15—Ralph Loura, Olivier Jacques, & Rafael Garcia—Breaking Traditional IT Paradigms.”
23. Rembetsy and McDonnell, “Continuously Deploying Culture.”
24. Rembetsy and McDonnell, “Continuously Deploying Culture.”
25. McKinley, “Why MongoDB Never Worked Out at Etsy.”
26. Cundiff, Geinert, and Rettig, “Crowdsourcing Technology Governance.”
27. Cundiff, Geinert, and Rettig, “Crowdsourcing Technology Governance.”
28. Cundiff, Geinert, and Rettig, “Crowdsourcing Technology Governance.”
29. Cundiff, Geinert, and Rettig, “Crowdsourcing Technology Governance.”
30. Cundiff, Geinert, and Rettig, “Crowdsourcing Technology Governance.”
31. Cundiff, Geinert, and Rettig, “Crowdsourcing Technology Governance.”

Chapter 21

1. “Kaizen,” Wikipedia.
2. Spear, *The High-Velocity Edge*, [Chapter 8](#).

3. Spear, *The High-Velocity Edge*, [Chapter 8](#).
4. Mickman and Clanton, “(Re)building an Engineering Culture.”
5. Ravi Pandey, personal correspondence with Gene Kim, 2015.
6. Mickman and Clanton, “(Re)building an Engineering Culture.”
7. Pomeranz, “Queue Inversion Week.”
8. Spear, *The High-Velocity Edge*, [Chapter 3](#).
9. Stillman, “Hack Days.”
10. Associated Press, “Number of Active Users at Facebook over the Years.”
11. Zhao, “HipHop for PHP.”
12. Metz, “How Three Guys Rebuilt the Foundation of Facebook.”
13. Metz, “How Three Guys Rebuilt the Foundation of Facebook.”
14. Steve Farley, personal correspondence with Gene Kim, January 5, 2016.
15. Gaekwad, “Agile 2013 Talk.”
16. O’Donnell, “DOES14—Glenn O’Donnell—Forrester—Modern Services Demand a DevOps Culture Beyond Apps.”
17. Smart et al., *Sooner Safer Happier*, 314.
18. Nationwide, *2014 Annual Report*.
19. Steve Farley, personal correspondence with Gene Kim, 2016.
20. Pal, “DOES15—Tapabrata Pal—Banking on Innovation & DevOps.”
21. Pal, “DOES15—Tapabrata Pal—Banking on Innovation & DevOps.”
22. Tapabrata Pal, personal correspondence with Gene Kim, 2015.
23. “All About Target,” Target (website).
24. Mickman and Clanton, “(Re)building an Engineering Culture.”
25. Van Leeuwen and Buytaert, “DOES15—Evelijn Van Leeuwen and Kris Buytaert—Turning Around the Containership.”
26. Mickman and Clanton, “(Re)building an Engineering Culture.”
27. “DevOps Culture: How to Transform.”
28. Bland, “DOES15—Mike Bland—Pain Is Over, If You Want It.”
29. Bland, “DOES15—Mike Bland—Pain Is Over, If You Want It.”
30. Bland, “DOES15—Mike Bland—Pain Is Over, If You Want It.”
31. Bland, “DOES15—Mike Bland—Pain Is Over, If You Want It.”
32. Bland, “DOES15—Mike Bland—Pain Is Over, If You Want It.”
33. Bland, “DOES15—Mike Bland—Pain Is Over, If You Want It.”
34. Bland, “DOES15—Mike Bland—Pain Is Over, If You Want It.”

35. Bland, “DOES15—Mike Bland—Pain Is Over, If You Want It.”
36. Bland, “Fixits, or I Am the Walrus.”
37. Bland, “Fixits, or I Am the Walrus.”
38. Bland, “Fixits, or I Am the Walrus.”

Chapter 22

1. Wickett, “Attacking Pipelines—Security Meets *Continuous Delivery*.”
2. Wickett, “Attacking Pipelines—Security Meets *Continuous Delivery*.”
3. Pal, “DOES15—Tapabrata Pal—Banking on Innovation & DevOps.”
4. Justin Arbuckle, personal interview with Gene Kim, 2015.
5. Justin Arbuckle, personal interview with Gene Kim, 2015.
6. Antani, “IBM Innovate DevOps Keynote.”
7. Galbreath, “DevOpsSec: Applying DevOps Principles to Security, DevOpsDays Austin 2012.”
8. Galbreath, “DevOpsSec: Applying DevOps Principles to Security, DevOpsDays Austin 2012.”
9. “OWASP Cheat Sheet Series,” OWASP (website).
10. Collins, Smolen, and Matatall, “Putting to your Robots to Work V1.1.”
11. “What Happens to Companies That Get Hacked? FTC Cases,” Giant Bomb forum.
12. “What Happens to Companies That Get Hacked? FTC Cases,” Giant Bomb forum.
13. Collins, Smolen, and Matatall, “Putting to your Robots to Work V1.1.”
14. Twitter Engineering, “Hack Week @ Twitter.”
15. Twitter Engineering, “Hack Week @ Twitter.”
16. Corman and Willis, “Immutable Awesomeness—Josh Corman and John Willis at DevOps Enterprise Summit 2015.”
17. Forsgren et al., *2020 State of the Octoverse*.
18. Forsgren et al., *2020 State of the Octoverse*.
19. Verison, *2014 Data Breach Investigations Report*.
20. Verizon, *2021 Data Breach Investigations Report*, 20.
21. Sonatype, *2019 State of the Software Supply Chain Report*.
22. Sonatype, *2019 State of the Software Supply Chain Report*.
23. Sonatype, *2019 State of the Software Supply Chain Report*.
24. “Jeremy Long: The (Application)Patching Manifesto.”

25. “Jeremy Long: The (Application)Patching Manifesto.”
26. Sonatype, *2019 State of the Software Supply Chain Report*.
27. Sonatype, *2019 State of the Software Supply Chain Report*.
28. Sonatype, *2019 State of the Software Supply Chain Report*.
29. Sonatype, *2019 State of the Software Supply Chain Report*.
30. Sonatype, *2020 State of the Software Supply Chain Report*.
31. Sonatype, *2020 State of the Software Supply Chain Report*.
32. Geer and Corman, “Almost Too Big to Fail.”
33. Forsgren et al., *2020 State of the Octoverse*.
34. Temple-Raston, “A ‘Worst Nightmare’ Cyberattack.”
35. Naraine, “Twilio, HashiCorp Among Codecov Supply Chain Hack Victims.”
36. Kash, “New Details Released on Proposed 2016 IT Spending.”
37. Bland, “DOES15—Mike Bland—Pain Is Over, If You Want It.”
38. Bland, “DOES15—Mike Bland—Pain Is Over, If You Want It.”
39. Zia, Ramírez, and Kunin, “Compliance Masonry.”
40. Marcus Sachs, personal correspondence with Gene Kim, 2010.
41. “VPC Best Configuration Practices,” Flux7 blog.
42. Galbreath, “Fraud Engineering, from Merchant Risk Council Annual Meeting 2012.”
43. Galbreath, “DevOpsSec.”
44. Galbreath, “DevOpsSec.”
45. Galbreath, “DevOpsSec.”
46. Galbreath, “DevOpsSec.”
47. Galbreath, “DevOpsSec.”
48. Galbreath, “DevOpsSec.”
49. Claudius, “Attacking Cloud Services with Source Code.”
50. Johnson, et. al., “How Fannie Mae Uses Agility to Support Homeowners and Renters.”
51. Johnson, et. al., “How Fannie Mae Uses Agility to Support Homeowners and Renters.”
52. Johnson, et. al., “How Fannie Mae Uses Agility to Support Homeowners and Renters.”
53. Johnson, et. al., “How Fannie Mae Uses Agility to Support Homeowners and Renters.”

54. Johnson, et. al., “How Fannie Mae Uses Agility to Support Homeowners and Renters.”
55. Johnson, et. al., “How Fannie Mae Uses Agility to Support Homeowners and Renters.”
56. Kimberly Johnson, personal correspondence with the authors, 2021.

Chapter 23

1. Axelos, *ITIL Service Transition*, 48.
2. Axelos, *ITIL Service Transition*, 48 and 68.
3. Matthew and Mangot, “DOES14—Reena Mathew and Dave Mangot—Salesforce.”
4. Mangot and Rajan, “Agile.2013.effecting.a.dev ops.transformation.at.salesforce.”
5. Mangot and Rajan, “Agile.2013.effecting.a.dev ops.transformation.at.salesforce.”
6. Mangot and Rajan, “Agile.2013.effecting.a.dev ops.transformation.at.salesforce.”
7. Matthew and Mangot, “DOES14—Reena Mathew and Dave Mangot—Salesforce.”
8. Matthew and Mangot, “DOES14—Reena Mathew and Dave Mangot—Salesforce.”
9. Matthew and Mangot, “DOES14—Reena Mathew and Dave Mangot—Salesforce.”
10. Matthew and Mangot, “DOES14—Reena Mathew and Dave Mangot—Salesforce.”
11. Matthew and Mangot, “DOES14—Reena Mathew and Dave Mangot—Salesforce.”
12. Bill Massie, personal correspondence with Gene Kim, 2014.
13. “Glossary,” PCI Security Standards Council website.
14. PCI Security Standards Council, *Payment Card Industry (PCI) Data Security Stands*.
15. Bill Massie, personal correspondence with Gene Kim, 2014.
16. Bill Massie, personal correspondence with Gene Kim, 2014.
17. Bill Massie, personal correspondence with Gene Kim, 2014.
18. Basu, Goyal, and Hansen, “Biz & Tech Partnership Towards 10 'No Fear Releases' Per Day.”
19. Basu, Goyal, and Hansen, “Biz & Tech Partnership Towards 10 'No Fear Releases' Per Day.”

20. Basu, Goyal, and Hansen, “Biz & Tech Partnership Towards 10 'No Fear Releases' Per Day.”
21. Basu, Goyal, and Hansen, “Biz & Tech Partnership Towards 10 'No Fear Releases' Per Day.”
22. Basu, Goyal, and Hansen, “Biz & Tech Partnership Towards 10 'No Fear Releases' Per Day.”
23. Basu, Goyal, and Hansen, “Biz & Tech Partnership Towards 10 'No Fear Releases' Per Day.”
24. Shinn, “DOES15—Bill Shinn—Prove it! The Last Mile for DevOps in Regulated Organizations.”
25. Shinn, “DOES15—Bill Shinn—Prove it! The Last Mile for DevOps in Regulated Organizations.”
26. Shinn, “DOES15—Bill Shinn—Prove it! The Last Mile for DevOps in Regulated Organizations.”
27. Shinn, “DOES15—Bill Shinn—Prove it! The Last Mile for DevOps in Regulated Organizations.”
28. Shinn, “DOES15—Bill Shinn—Prove it! The Last Mile for DevOps in Regulated Organizations.”
29. Shinn, “DOES15—Bill Shinn—Prove it! The Last Mile for DevOps in Regulated Organizations.”
30. Shinn, “DOES15—Bill Shinn—Prove it! The Last Mile for DevOps in Regulated Organizations.”
31. Shinn, “DOES15—Bill Shinn—Prove it! The Last Mile for DevOps in Regulated Organizations.”
32. Shinn, “DOES15—Bill Shinn—Prove it! The Last Mile for DevOps in Regulated Organizations.”
33. DeLuccia, Gallimore, Kim, and Miller, DevOps Audit Defense Toolkit.
34. Mary Smith (a pseudonym), personal correspondence with Gene Kim, 2013.
35. Mary Smith (a pseudonym), personal correspondence with Gene Kim, 2013

Conclusion

1. Robbins, “Hacking Culture at VelocityConf.”

Afterword

1. Forsgren et al., *2020 State of the Octoverse*.
2. Wiseman, *2021 Work Trend Index: Annual Report*.
3. Forsgren et al., “The SPACE of Developer Productivity.”

Appendix

1. Ries, *The Lean Startup*.
2. Beck et al., “Twelve Principles of Agile Software.”
3. Humble and Farley, *Continuous Delivery*.
4. Fitz, “Continuous Deployment at IMVU.”
5. Rother, *Toyota Kata*, Introduction.
6. Rother, *Toyota Kata*, Introduction.
7. Ries, *The Lean Startup*.
8. Kim, Behr, and Spafford, *The Phoenix Project*, 365.
9. Besnard and Hollnagel, *Some Myths about Industrial Safety*, 3.
10. Besnard and Hollnagel, *Some Myths about Industrial Safety*, 4.
11. Besnard and Hollnagel, *Some Myths about Industrial Safety*, 6.
12. Besnard and Hollnagel, *Some Myths about Industrial Safety*, 8.
13. Besnard and Hollnagel, *Some Myths about Industrial Safety*, 9.
14. Besnard and Hollnagel, *Some Myths about Industrial Safety*, 11.
15. Shook, “Five Missing Pieces in Your Standardized Work (Part 3 of 3).”
16. “Post Event Retrospective—Part 1,” Rally Blogs.
17. Macri, “Morgue.”
18. Cockcroft, Hicks, and Orzell, “Lessons Netflix Learned.”
19. Cockcroft, Hicks, and Orzell, “Lessons Netflix Learned.”
20. Rachitsky, “7 Keys to a Successful Public Health Dashboard.”

Index

Note: Figures are indicated with *f*; footnotes are indicated with *n*; tables are indicated with *t*.

SYMBOLS

%C/A percent complete and accurate), 11, 85*f*, 95

NUMBERS

2/4 outage, 318–320

18F team, 369–371

A

A/B testing

conclusions on, 280

feature planning and, 278

feature testing and, 276–277

feature toggles and, 196–197

history of, 275

need for, 273

our release and, 277–278

practical guide to, 276

Accelerate: The Science of Lean and DevOps, xln, 14, 219

acceptance stage, 152*f*, 153

acceptance test-driven development (ATDD), 159

acceptance tests, 36, 155–156

accountability, as cultural pillar, 76

Adams, Keith, 338

Adidas, 286–287

Agile

DevOps as continuation of, xxv

Infrastructure Movement, 5–6, 410

Manifesto, 4–5, 410

Agile Software Development with Scrum, 122*n*

Aisen Seiki Global, 52

Alcoa, 50, 313

alert fatigue, 247

Algra, Ingrid, 343*n*

Allspaw, John, xxiii, 5, 6, 197, 234, 283, 284, 298, 308, 319, 360, 385*n*, 410

Amazon

architecture transformation at, 210, 212–213

continuous delivery at, 200–201

DevOps myths and, xxiv–xxv

market-oriented teams at, 102

“master of disaster” at, 316, 317, 399

post-mortem, 311*n*

service-oriented architecture (SOA) at, 109–110

two-pizza teams at, 110, 111

Amazon Auto Scaling (AAS), 251–252

Amazon Reboot of 2014, Great, 315–316

Amazon Web Services (AWS)

CloudWatch, 391

outage, 305–306, 315

security solutions architect at, 389

ambiguous threats, 313–314

American Airlines

brownfield transformation, 68

case study, 15–18, 74–77

new vocabulary at, 74–77*t*

anchoring bias, 388

Anderson, David J., 22

Andon button, 38*n*

Andon cord

description of, 37–39

Excella, 39–41*f*

for low-risk releases, 181

work stoppage and, 416–417*f*

anomaly detection

advanced, 255–257

anomaly detection (*continued*)

- defined, 253
- Kolmogorov-Smirnov test, 254, 255, 256*f*, 257
- smoothing for, 253–254

Antani, Snehal, 354

antifragility, 52

APIs

- cleanly defined, 214
- enablement, 112–114
- loosely coupled architecture and, 209
- microservice-based architecture driven by, 388
- self-service, 230
- versioned, 214

application logging telemetry, 231–233

application performance monitors, 235*n*

application-based release patterns

- dark launching, 190, 197–199
- defined, 190
- feature toggles, 190, 195–197

Arbuckle, Justin, 324, 354

Architectural Review Board (ARB), 333

architecture

- Amazon, 212–213
- Blackboard Learn, 215–217*f*
- conclusions on, 218
- downward spiral in, 208–209
- eBay, 207–208
- loosely coupled, 26–27, 209, 217
- monoliths versus microservices, 210–212
- overly tight, 26–27
- service-oriented, 109, 210
- strangler fig application pattern and, 70, 208, 213–217, 218

Architecture and Patterns for IT, 208

Art of Monitoring, The, 228, 230*f*

Ashman, David, 215, 216, 217*f*

ATM cash machines, 392

Atwood, Jeff, 170, 171, 293

Audit Defense Toolkit, DevOps, 391

auditors and compliance officers

- ATM cash machines and, 392
- PCI compliance at Etsy, 385–387
- proving compliance, 389–391
- tension between IT and, 389

Austin, Jim, 245

automated testing

- Andon cord and, 163–165
- categories of tests, 155–156
- conclusions on, 166
- constraints and, 26
- deployment pipeline infrastructure, 151–154
- essential components of, 165–166
- fast and reproducible tests, 166
- at Google, 148–151
- green build and, 154, 163, 166
- ideal testing pyramid, 157f–158
- need for, 147–148
- non-functional requirements and, 162–163, 328
- observability and, 147
- performance testing environment, 161–162
- reducing reliance on manual tests, 160–161
- research supporting, 165–166
- running tests in parallel, 158f–159
- test-driven development (TDD), 159, 161, 327

automation, DevOps and, xxvi

Ayers, Zack, 39, 41f

B

bad apple theory, 307

Baker, Bill, 141

banks, as IT companies, xxxv

batch sizes, reducing, 9, 22–24*f*, 409

Bazaarvoice, 173–175, 176

Beck, Kent, 159, 178

Beedle, Mike, 122*n*

Behr, Kevin, *xlii*, 225, 233*n*, 415*f*

Bell Labs, 54–55

Besnard, Denis, 416

Betz, Charles, 208

Beyond The Phoenix Project, 57

Big Fish Games, 115–117

big-bang approach, 73

bimodal IT, 71, 72

Blackboard Learn, 215–217*f*

blameless post-mortems

- CSG case study, 318–320
- defined, 308
- inviting Ops engineers to, 117, 123–124
- for organizational learning, *xxxix*, 48–49
- publishing reports of, 311–312
- sample agenda, 418–419
- scheduling of, 308–310

Bland, Mike, 148, 149, 344, 345, 370

Blank, Steve, 411

Blankenship, Ed, 240

blitz, improvement, 335, 336

blue-green deployment pattern

- description of, 190, 191*f*–192
- for point-of-sale system, 193

Bohmer, Richard M. J., 313

Booch, Grady, 6*n*

Bosu, Biswanath, 388

bottlenecks

- DevOps and, 405–406
- generalists and, 106*t*

- handoffs, queues, and, 414
- Boubez, Toufic, 249, 250f, 251f, 255, 256
- bounded contexts, 109
- Bouwman, Jan-Joost, 343n
- Brakeman, 359, 362f–363
- branches, feature, 170
- branching by abstraction, 214n
- branching strategies, 167, 170–171
- Brittain, Mike, 241f, 262f
- Brook, Frederick, xli
- Brooke’s Law, 286
- brownfield projects
 - case study, 69–71
 - defined, 67
 - greenfield projects versus, 66–69
 - technical debt and, 67
- Building Blocks, Blackboard Learn, 215–217f
- Building the Future: Big Teaming for Audacious Innovation*, 312
- bureaucratic organizations, 42, 47, 48t
- bureaucratic processes, cutting, 42–43, 296–297, 299
- Burgess, Mark, 6n
- burnout, decreased, xl, 15, 175
- business relationship manager, 116
- Buytaert, Kris, 343n
- C**
- Cagan, Marty, 90
- Campbell-Pretty, Em, 135–136
- Canahuati, Pedro, 105, 263
- canary release pattern, 190, 194f–195
- canary release test, 177n
- Capital One
 - biz and tech partnership at, 387–389
 - case study, 342–343, 387–389
 - Got Goo? program, 296

- internal conferences, 342–343
- cardholder data breaches, 364, 371
- cardholder data environment (CDE), 385–386
- case studies
 - Adidas, 286–287
 - Amazon, 212–213
 - American Airlines, 15–18, 74–77
 - ATM cash machines, 392
 - Bazaarvoice, 173–175, 176
 - Bell Labs, 54–55
 - Blackboard Learn, 215–217*f*
 - Capital One, 342–343, 387–389
 - CSG International, 181–183, 201–206, 318–320
 - Dixons Retail, 193
 - Etsy, 186–188, 332, 373–375, 385–387
 - Excella, 39–41*f*
 - Facebook, 198–199
 - Fannie Mae, 376–378
 - Federal Government agencies, 369–371
 - Google, 269–271, 290–291
 - hospital system, 29–32
 - hotel company, 143–144
 - Kessel Run refueling system, 66–69
 - LinkedIn, 91–93, 237–238
 - Nationwide Building Society, 124–127
 - Nationwide Insurance, 342
 - Netflix, 251–253
 - Pivotal Labs, 293–294
 - [Salesforce.com](https://www.salesforce.com), 383–384
 - Target, 112–114, 333–334, 342, 343
 - tax collection agency for UK, 77–80
 - Twitter, 360–363
 - Yahoo! Answers, 278–280
- Chacon, Scott, 281, 282*f*

Chakrabarti, Arup, 263

Change, John Shook's Model of, 205

change advisory board (CAB), 380, 382

change approval processes

- case studies, 383–384, 385–387, 387–389
- dangers of, 283–284
- at Etsy, 385–387
- normal changes, 380, 381–382
- at [Salesforce.com](https://www.salesforce.com), 383–384
- security and compliance in, 379–380
- standard changes, 379–381, 383–384
- three categories of changes, 379–380
- urgent changes, 380

change control failure, 283

change freezes, 292

Chaos Gorilla, 420

Chaos Kong, 420

Chaos Monkey, 52, 55, 306–307, 315

Chapman, Brent, 347

Chapman, Janet, 124, 125, 126*f*

chat rooms

- announcing changes with, 288
- drawbacks of, 94–95
- Hubot at GitHub, 321–323
- organizational knowledge and, 321–322
- shared goals and, 94
- as watercooler, 322

Chuvakin, Anton A., 232

Clanton, Ross, 74, 75*f*, 76, 77*t*, 297, 335, 336, 343

Claudius, Jonathan, 375

Clemm, Josh, 91–93

cloud computing, five characteristics of, 330–331

cloud native, 306

Cloud System Administration, The Practice of, 325

[Cloud.gov](#), 370, 371

cluster immune system, 190, 195*n*

coaching kata, 53

Cockcroft, Adrian, xxxit, 102*n*, 231, 296

code

- infrastructure as, 6*n*
- libraries, 356
- maintainability, 326
- repositories, 355–357
- signing, 359–360
- Code Climate, 359
- code commits
 - automated tests on, 160, 166
 - daily, 172
 - gated commits, 172
 - Google, 150, 290
 - Pivotal Labs, 294
 - security and, 357
 - strangler fig application pattern, 215–217*f*

code reviews. *See also* change approval processes

- ATM systems and, 392
- change reviews versus, 289*n*
- defined, 288
- e-mail pass around, 290
- forms of, 289–290
- Google, 290–291
- guidelines for, 288–289
- learning-based culture and, 339
- “over the shoulder,” 290
- pair programming, 289, 292–294
- Pivotal Labs, 293–294
- requesting, 282
- separation of duty and, 184, 384–385, 386
- size of change and, 289

- tool-assisted, 290
- unauthorized access and, 375, 376
- Codecov security breach, 368–369
- Cohen, Joshua, 39, 41*f*
- Collins, Justin, 360, 361
- Columbia* space shuttle, 313
- commit stage, 152*f*, 153
- Common Vulnerabilities and Exposures (CVE), 365
- communities of practice, 343–345
- compliance officers and auditors
 - ATM cash machines and, 392
 - Payment Card Industry Data Security Standards (PCI DSS), 385
 - PCI compliance at Etsy, 385–387
 - proving compliance, 389–391
 - separation of duty and, 379, 384–387
 - tension between IT and, 389
- complicated-subsystem teams, 112
- conferences
 - internal, 342–343
 - sharing experiences from DevOps, 341–342
- Conformity Monkey, 420
- Conrad, Ben, 78
- constraint identification, 25–27
- Constraints, Theory of, 4, 412–413
- Consul, 242*n*
- containers, 143–144, 152, 153*n*
- contextual inquiry, 264
- continual learning and experimentation, 3, 45–56
- continuous delivery. *See also* deployment process; low-risk releases, enabling
 - continuous deployment versus, 199–201
 - defined, 133, 200
 - elite performance and, 201
 - infrastructure monitoring and, 243
 - low-risk releases and, 199–201

Continuous Delivery: Reliable Software Releases Through Build, Test, Deployment Automation, xxiv, 151, 193, 199, 219

Continuous Delivery Movement, 6, 410–411

continuous integration (CI) and testing

- defined, 36, 151*n*
- deployment pipeline infrastructure and, 151–154

continuous integration practices

- Andon cord and, 163–165
- at Bazaarvoice, 173–175, 176
- catching errors early, 157–163
- reliable automated validation test suite, 154–156
- three capabilities required for, 154
- trunk-based development, 167–176

convergence of DevOps, 409–412

Conway, Melvin, 97

Conway's Law

conclusions about, 114

- defined, 61, 97–98
- at Etsy, 98–100, 108
- organizational archetypes and, 100–101
- Target case study and, 112–114
- team boundaries in accordance with, 108
- two-pizza team and, 110–111

Cook, Richard, 58, 319

Cook, Scott, 274

core conflict in IT operations, xxxii–xxxiii, 412*f*–413

Corman, Josh, 353, 363, 368, 412

Cornago, Fernando, 286

costs of IT-related deficiencies, xxxvi–xxxvii

COTS software, 417–418

counterfactual thinking, 283*n*, 310

COVID-19 pandemic

- call centers during, 126–127
- remote work and, 108*n*

- UK's financial support package, 77–80
- Cox, Jason, 107, 119–120, 296
- crowdsourcing technology governance, 333–334
- CSG International
 - blameless post-mortem, 318–320
 - brownfield transformation, 68
 - case study, 181–183, 201–206, 318–320
 - daily deployments at, 181–183
 - generalists, 106–107
 - organizational transformation, 201–206
- culture, organizational
 - high-trust, xxxix, 45, 48, 150
 - importance of, 46
 - just culture, 47, 307–308
 - learning-based, 48–49, 339
 - safety culture, 46–49
 - three pillars of, 75–76
 - three types of, 47, 48*t*
- culture of causality, 225, 233*n*
- Cundiff, Dan, 333
- Cunningham, Ward, xxxii, 171
- customer acquisition funnels, 240, 273, 275, 278
- customers, two types of, 43

D

- daily work, improvement of, 49–50
- daily work of development
 - at Big Fish Games, 115–117
 - conclusions on, 127–128
 - embedded Ops engineers in, 116, 119–120
 - Ops liaisons in, 116, 117, 120–121
 - shared services in, 117–118
 - team rituals in, 117, 121–124
- dark launch
 - defined, 190, 197–198

- Facebook Chat, 198–199
- dashboard
 - Adidas, 287
 - creating a simple, 237n
 - daily work and, 234
 - Etsy, 227
 - public health, 421
- Data Breach Investigation Report (DBIR), 364
- database changes, dealing with, 192
- database syntax error, 374
- Davenport, Shawn, 295n
- Debois, Patrick, xxiii, 5, 6, 405–406, 410
- DEBUG level, 232
- dedicated release engineer, 116
- defects, as waste, 28
- DeGrandis, Dominica, 22, 58
- Dekker, Sidney, 34, 47, 58, 307, 347, 395
- (Delicate) Art of Bureaucracy, The*, 299
- Deming, W. Edwards, 38, 54
- demonstrations, compliance by, 354
- dependency scanning, 359
- deploy code button, 186
- Deployinator console, Etsy, 187–188f
- deployment lead time, 8–11f, 409
- deployment pipeline foundations
 - Andon cord and, 163–165
 - containers, 143–144, 152, 153n
 - defined, 6, 151, 152f
 - definition of “done,” 144–145
 - Enterprise Data Warehouse, 135–137
 - goal of deployment pipeline, 153
 - infrastructure, 151–154
 - on-demand creation of test environments, 137–138
 - rebuilding instead of repairing infrastructure, 141–143

- single repository of truth, 139–141, 150, 324–327
- deployment pipeline protection
 - ATMs and production telemetry, 392
 - auditors and, 389–391
 - Capital One, 387–389
 - change advisory board (CAB) and, 380, 381, 382
 - change approval processes and, 379–380
 - Etsy, 385–387
 - for normal changes, 380, 381–382
 - separation of duty, 379, 384–387
 - for standard changes, 379–381, 383–384
- deployment process
 - Andon cord and, 181
 - automating, 179–181
 - continuous deployment, 199–201
 - CSG International, 181–183, 201–206
 - decouple deployments from releases, 189–199
 - Dixons Retail, 193
 - Etsy, 186–188
 - Facebook, 177–179
 - release versus, 189
- deployment process (*continued*)
 - self-service deployments, 184–188
 - smoke testing deployments, 180, 187
- deploys per day per developer, xxxviii, xxxixf–xl
- destructive testing, 383–384
- Dev team rituals, 121–124
- Dev tests, 36
- developer productivity
 - DevOps practices and, xlii
 - measuring, 401–403
 - shared services for, 117–119
 - development, daily work of
 - at Big Fish Games, 115–117

- conclusions on, 127–128
- embedded Ops engineers in, 116, 119–120
- Ops liaisons in, 116, 117, 120–121
- shared services in, 117–118
- team rituals in, 117, 121–124

development, hypothesis-driven

- at Intuit, 273–275
- need for, 273, 280

development, test-driven (TDD). *See also* automated testing

- building incrementally with, 161
- defined, 159
- shared libraries and, 327
- study on, 159*n*

development, trunk-based

- adopting practices of, 172
- at Bazaarvoice, 173–175, 176
- benefits of, 175
- conclusions on, 175–176
- gated commits and, 172
- for HP's LaserJet Firmware division, 168–170
- need for, 167, 171

DevOps

- breaking downward spiral with, xxxv–xxxvii
- business value of, xxxvii–xxxix
- convergence of, 409–412
- core conflict in IT operations and, xxxii–xxxiii, 412*f*–413
- developer productivity and, xli*f*–xlii
- downward spiral in IT and, xxxiii–xxxvii, 413*t*–414*t*
- history of, 3–6
- myths about, xxiv–xxv
- outcomes created by, xxviii, xxix
- principles underpinning, 12–14

DevOps Days. *See* DevOpsDays

DevOps Enterprise Summit, xxxix, 341–342, 343

DevOps journeys. *See* case studies

DevOpsDays, xxiii, xxiv, 6, 341, 343*n*, 353, 355, 410

Dickerson, Chad, 99

Dignan, Larry, 110, 111

direct response marketing, 275

Disaster Recovery Program (DiRT), Google's, 317

Disney, 107, 119–120, 296

Dixons Retail, 193

Doctor Monkey, 420

Dojo, DevOps, 335–336

Dojos, *Getting Started With*, 58

“done,” modifying definition of, 144–145, 172

DORA, *State of DevOps Reports* from, 5, 14, 57, 103, 140, 165, 166, 175, 185, 201, 217, 227, 243, 285, 312, 326, 331, 343

downward spiral in IT

description of, xxxiii–xxxvii, 413*t*–414*t*

DevOps for breaking, xxxvii–xxxix

Drucker, Peter, 75, 80

Dunbar, Robin, 111*n*

Dunbar's number, 111*n*

Dweck, Carol, 107

dynamic analysis, 354

E

early adopters

- defined, 72*f*, 73
- finding, 73–74

eBay, 90, 207–208

economic costs of IT-related deficiencies, xxxvi–xxxviii

Edmondson, Amy C., 312, 313, 347

Edwards, Damon, 10*f*, 28, 118

Eli Lilly, 312

Eltridge, Patrick, 124, 125, 126*f*, 127

email pass around code review, 290

emergency change advisory board (ECAB), 380

- employee burnout, lower rates of, xl, 175
- enabling teams, 112
- Encasement Strategy, 70
- Eno, Brian, 54, 55
- Enterprise Data Warehouse, 135–137
- envelope game, simulation of, 23, 24*f*
- environment-based release patterns
 - blue-green deployments, 190, 191*f*–192, 193
 - canary releases, 190, 194*f*–195
 - cluster immune systems, 190, 195
 - defined, 190
- ERROR level, 232
- ethics of DevOps, xxxvii–xxxix
- Etsy
 - blameless post-mortems, 310, 311, 419
 - brownfield transformation, 68–69
 - case studies, 186–188, 332, 373–375, 385–387
 - cluster immune system, 195*n*
 - continuous delivery at, 201
 - Conway’s Law and, 98–100, 108
 - designated Ops liaison at, 120–121
 - DevOps myths and, xxiv–xxv
 - DevOps transformation at, 226–227
 - experimentation, 277–278
 - functional orientation, 101*n*, 104
 - LAMP stack, 226
 - learning-based culture, 308
 - Morgue tool, 311–312
 - organizational learning, 49
 - PCI compliance at, 385–387
 - PHP run-time warning, 262*f*
 - retrospective meetings, 310, 311, 419
 - self-service deployment, 186–188
 - separation of duty, 385–387

- standardizing technology stack at, 332
- StatsD metric library, 234, 235*f*
- transformation projects, 63

Evans, Eric J., 109

Evans, Jason, 338

event router, 229, 230*f*

Excella, Andon cord at, 39–41*f*

Expanding Pockets of Greatness: Spreading DevOps Horizontally in Your Organization, 129

experimentation, rapid

- A/B testing, 273, 275–278
- customer acquisition funnel and, 240, 273, 275, 278
- at Etsy, 277–278
- at Intuit, 273–275
- need for, 273
- at TurboTax, 274
- at Yahoo! Answers, 278–280

exploratory testing, 36

Explore It!: Reduce Risk and Increase Confidence with Exploratory Testing, 36, 219

extra features, 28

extra processes, 28

F

Facebook

- canary release pattern, 194*f*–195
- case study, 198–199
- code deployment, 177–179
- continuous delivery, 201
- experimentation, 279
- feedback, 105
- Gatekeeper, 196*n*, 199
- hackathon, 338

Facebook Chat, dark launch, 198–199

fail fasts, 316*n*

failure parties, 312

failures

blameless post-mortem at CSG, 318–320

calculated risk-taking and, 314–315

game days to rehearse, 316–318

no fear of, 55

publishing post-mortem reports, 311–312

redefining, 314–315

rehearsing and injecting, 315–316

retrospective meetings after, 308–310

weak failure signals, 313–314

fallbacks, 316*n*

Fannie Mae, 376–378

Farley, David, 6, 151, 152*f*, 156, 158*f*, 193, 199, 214, 410

Farley, Steve, 339, 342

Farr, Will, 295*n*

Farrall, Paul, 115–116

fast and reproducible tests, 166

FATAL level, 232

Fearless Organization, The, 347

feature branches, 170

feature flags. *See* feature toggles

feature freezes, 175

feature toggles, 195–197, 277

features

extra, 28

planning, 278

testing, 276–277, 279

user excitement and, 241*f*

Federal Government agencies, 369–371

feedback. *See also* The Second Way (principles of feedback)

Andon cord and, 37–41

customer observation, 264–265

cycle times, 37*f*

fast and constant, 10, 13–14, 33

- optimizing for downstream work centers, 43–44
- quality control closer to the source, 42–43
- pager rotation duties, 263–264
- principles of, 3, 13–14, 33
- production telemetry and, 261–262
- safe deployment of code and, 259–261
- safety within complex systems, 33–34
- seeing problems as they occur, 35–36, 244
- self-management by developers, 265–271
- stakeholder, 36
- swarming, 37–41
- types and cycle times, 36–37*f*
- user, 36, 37*f*

Fernandez, Roberto, 73, 100

Fifth Discipline, The, 35, 49

First Way, The. *See* The First Way (principles of flow)

Fitz, Tim, 6, 192*n*, 199, 411

Five Dysfunctions of a Team: A Leadership Fable, 347

fix forward, 262

fixed mindset, 107

fixits, 345

Flickr, 197–198, 278, 360

flow, principles of, 3. *See also* The First Way (principles of flow)

flow metrics, 11–12

focusing steps, five, 26, 32

following work downstream, 264–265

Forsgren, Nicole, xxxix–xl, 5, 14, 140, 228*f*, 363, 401–403, 404

Fowler, Martin, 156, 157, 213, 214, 220, 306*n*

fraud, defined, 373

Fryman, James, 295*n*

full-stack engineer, 106, 386

functional-oriented organizations

- defined, 100
- DevOps outcomes in, 103–104

- market orientation versus, 100–101, 103*f*
- problems of, 101–102
- funding services and products, 107–108

Furtado, Adam, 69, 70

G

Gaekwad, Karthik, 339

Galbreath, Nick, 259–261, 355, 373–375

Gall’s Law, 70

Gallimore, Jeff, 41

game days, 52, 316–318

Ganglia, 226, 229

gated commits, 172

Gatekeeper, Facebook’s, 196*n*, 199

Gauntlet security tool, 353, 357

Gaussian distribution, 247*f*, 249, 253

GE Capital, 324, 354

Geer, Dan, 368

Geinert, Levi, 333

General Electric, CEO of, xxxi–xxxii

General Motors manufacturing plant, 35, 38, 45

generalists, 105–107

generative organizations, 47–48*t*, 57

Gertner, Jon, 54, 55

GitHub

- functional orientation, 101*n*, 104
- Hubot, 321–323
- Octoverse Report*, 57, 368, 401
- peer review, 281–283
- pull request processes, 295–296
- vulnerability timeline, 368

GitHub Flow, 282

The Goal: A Process of Ongoing Improvement, xlii, xliii, 30, 406

goals

- global, 19, 21

improvement, 88

Goldratt, Eliyahu M., xxxiii, xlii, 25, 26, 32, 406

Goldratt's Theory of Constraint, 25–26, 32

Google

- architecture, 209*f*–210
- automated testing, 148–151
- case study, 269–271, 290–291
- code reviews, 290–291
- continuous delivery, 200–201
- DevOps myths and, xxiv–xxv
- disaster recovery program, 317–318
- grouplets, 344, 345
- launch and hand-off readiness review, 269–271
- retrospective documents, 311
- service-oriented architectures (SOA), 109–110
- shared source code repository, 300, 324–326
- Testing on the Toilet newsletter, 149*n*, 344
- Web Server team, 148–151

Google Cloud Datastore, 209*f*, 210

Got Goo? program, 296

Gothelf, Jeff, 411

Govindarajan, Vijay, 86, 87

Goyal, Rakesh, 387, 388

Grafana, 79, 234, 254, 255

Gramm-Leach-Bailey Act, 389

Graphite, 226, 234, 235*f*, 236, 254, 255, 322, 374*f*

Gray, Jim, 212

green build, 154, 163, 166

greenfield vs. brownfield services, 66–69

grouplets, 344, 345

growth mindset, 107

Gruver, Gary, 43, 148, 160, 168, 169, 170

guardrails, 79

Gupta, Prachi, 237, 238

H

- Haber, Eben, 5
- hackathon
 - defined, 337*n*
 - Facebook, 338
- Hamilton, James, 306*n*
- Hammant, Paul, 214*n*
- Hammond, Paul, xxiii, 5, 6, 360, 410
- hand back mechanism, 268–269
- hand-off readiness review (HRR), 269, 270, 271*f*
- handoffs, 20, 24–25, 414–415
- hardships and waste, 27–29
- healthcare organizations
 - generative cultures and, 47–48*t*
 - HIPAA requirements, 390–391
 - hospital case study, 29–32
- helplessness, learned, xxxvi
- Hendrickson, Elizabeth, 36, 160, 219, 293–294, 300
- heroics, 10, 28–29, 170
- High-Velocity Edge, The*, xxxv
- HIPAA, 390–391
- HipHop virtual machine project (HHVM), 338
- history
 - A/B testing, 275
 - DevOps, 3–6
 - software delivery, xxx–xxx*it*
- HMRC tax collection agency, 77–80
- Hodge, Victoria J., 245
- holdouts, identifying, 74
- Hollnagel, Erik, 416
- Holmes, Dwayne, 143, 144
- hospital system case study, 29–32
- HP LaserJet, 69
- HP's LaserJet Firmware division, 168–170, 176

HSBC bank, xxxv*n*

Hubot, at GitHub, 321–323

Humble, Jez, xxii–xxiii, xl, 6, 151, 152*f*, 156, 158*f*, 191*f*, 199, 200, 207, 219, 273, 277, 284, 404–405, 406, 410

Hyatt, Matt, 78, 79, 80

hybrid schedules, 108*n*

hypothesis-driven development

- at Intuit, 273–275
- need for, 273, 280

I

Idea Factory: Bell Labs and the Great Age of American Innovation, The, 54

ideal testing pyramid, 157*f*–158

Imbriaco, Mark, 322

Immelt, Jeffrey, xxxi–xxxii

immersive learning opportunities, 16, 18

immutable infrastructure, 142

immutable services, 214

imposter syndrome, 148*n*, 310

improvement blitz, 335, 336

improvement goals, 88

improvement kata, 6

improvement of daily work, 49–50

INFO level, 232

information radiators, 236, 237, 241

information security. *See also* deployment pipeline protection

- 18F team, 369–371
- bad paths, 358
- Brakeman, 359, 362*f*–363
- change approval processes, 379–380
- code signing, 359–360
- data breaches, 364, 368–369, 371
- defect tracking, 355
- dependency scanning, 359
- deployment pipeline and, 357, 375–376

- DevOps and, 353
- dynamic analysis, 354
- early-stage product demonstrations, 354
- Etsy, 373–375, 385–387
- Fannie Mae, 376–378
- Gauntlt security tool, 353, 357
- happy path, 358
- Open Web Application Security Project (OWASP), 359*n*, 360
- OWASP Dependency Check, 368*n*
- OWASP ZAP, 358*f*, 359
- Payment Card Industry Data Security Standards (PCI DSS), 385
- post-mortems and, 355
- preventive security controls, 355–357
- production telemetry and, 371–373
- Rugged DevOps, 353
- sad and bad paths, 358
- separation of duty, 379, 384–387
- shared source code repositories and, 355–357
- shifting security left, 376–378
- silo, 353
- software supply chain and, 363–369
- source code integrity and code signing, 359–360
- SQL injection attacks, 373, 374*f*
- static analysis, 358*f*–359
- Twitter case study, 360–363

Infosec. *See* information security

Infosec team, 83

infrastructure

- centralized telemetry, 227–231
- changes, 383–384
- as code, 6*n*
- deployment pipeline, 151–154
- immutable, 142
- metrics, 242–243

- rebuilding instead of repairing, 141–143
- ING technology organization, 74
- innovators and early adopters
 - defined, 72*f*–73
 - finding, 73–74
- integration tests, 156
- Intuit, 273–275
- IT operations
 - core conflict in, xxxii–xxxiii, 412*f*–413
 - developer productivity, xli*f*–xlii
 - DevOps and, xxvi
 - downward spiral in, xxxv–xxxix, 413*t*–414*t*
 - impact of DevOps on, xxxvii–xxxix
- ITIL, xxvii, 139, 263*n*, 285, 286, 379, 380*n*
- ITIL CMBD, 242*n*

J

- Jacob, Adam, 6*n*
- Jacobson, Daniel, 252*f*, 254*f*
- Janitor Monkey, 420
- Java
 - automation, 161
 - Bazaarvoice, 173
 - dependency scanning, 359
 - EAR and WAR files, 152
 - LinkedIn, 91
 - Maven Central, 364
 - ORM, 99*n*
- JavaScript
 - application logging, 231*n*
 - client software level, 239
 - CSG, 318
 - eBay, 207*n*
 - Facebook, 199
 - libraries, 326

NPM, 364

- open-source dependencies, 363
- StatsD, 234

Jenkins, 152, 153, 180, 187, 205, 322, 358*f*, 377

JIRA, 355, 377, 381, 382, 386

Johnson, Kimberly, 376, 377–378

Jones, Angie, 161

Jones, Daniel T., 23

Joshi, Neeraj, 252*f*, 254*f*

just culture, 47, 307–308

Just Culture, 347

K

kaizen blitz, 49, 335, 336

Kalantzis, Christos, 315–316

Kanban: Successful Evolutionary Change for Your Technology Business, 22

kanban boards

- example, 20*f*–21
- Ops work on, 124
- shared goals and, 94
- Toyota Production System, 4

Kandogan, Eser, 5

Kanies, Luke, xxiii, 6*n*

Kastner, Erik, 187, 188*f*

Kelly, Mervin, 54, 55

Kersten, Mik, 12, 54

Kessel Run mid-air refueling system, 69–71

Kim, Gene, xxi–xxii, xl, xlii, 13*f*, 47*n*, 54, 57, 58, 225, 233*n*, 265, 284, 295*n*, 300, 353*n*, 364, 385*n*, 403–404, 406, 415*f*

Kissler, Courtney, 63, 64, 65, 66, 81–82, 83

Knight Capital failure, 283

Kohavi, Ron, 276, 277

Kolmogorov-Smirnov test, 254, 255, 256*f*, 257

Krishnan, Kripa, 317–318

Kumar, Ashish, 291*f*

L

laggards (skeptics), 72*f*, 73

LAMP stack

 DevOps myth and, xxvi

 Etsy, 226

large batch sizes

 merges of, 170–171

 small versus, 9, 22–24*f*, 409

Latency Monkey, 420

latent defects, 317

Lauderbach, John, 108*n*

launch guidance, 266–267

launch readiness review (LRR), 269, 270, 271*f*

Lead Architecture Review Board (LARB), 296, 297

lead time

 defined, 409

 focus on, 8, 18

 Lean Movement and, 409

 of minutes, 10–11*f*

 process time versus, 9*f*

 queue size and, 22, 415*f*

 of three months, 10*f*

leaders

 role of, 52–54

 vocabulary for, 76–77*t*

Lean Enterprise: How High Performance Organizations Innovate at Scale, 278

Lean Manufacturing, xxx, 7, 8, 9

Lean Movement

 description of, 3, 409

 missing element in, 6

Lean Startup, The, 411

Lean UX movement, 411

LeanKit, 381

learned helplessness, xxxvi

learning-based culture

- ASREDS learning loop, 340, 341*f*
- communities of practice, 343–345
- conclusions on, 346
- DevOps conferences, 341–342
- Etsy, 49
- grouplets, 344, 345
- importance of, 46
- improvement blitz, 335, 336
- internal conferences, 342–343
- just culture, 47, 307–308
- rituals to pay down technical debt, 336–339
- Teaching Thursday, 339
- thirty-day challenge at Target, 335–336
- safety culture, 46–49
- trust and, xli, 45, 48, 150

Leibman, Maya, 15, 16, 17, 75*f*, 77*t*

Lencioni, Patrick, 347

Lesiecki, Nick, 344

Letuchy, Eugene, 198, 199

Levenberg, Josh, 149, 300

liaisons, Ops

- assigning, 117, 120–121
- purpose of, 115–117
- team rituals and, 121–124
- two types of, 116

libraries, shared, 326, 327

Lightbody, Patrick, 263

limiting work in progress (WIP), 7, 21–22

Limoncelli, Tom, 248, 270, 271, 325–326

LinkedIn

- case study, 91–93, 237–238
- former monolithic architecture, 210
- Operation Inversion, 91–93

- self-service metrics, 237–238
- Little, Christopher, xxxv, xxvi, 94
- logging infrastructure, 231*n*
- logging telemetry, 231–233
- logs
 - centralized, 229
 - in monitoring framework, 230*f*–231
- Lonestar Application Security Conference, 353
- Long, Jeremy, 365
- long-lived private branches, 170
- loosely coupled architecture, 26–27, 209, 217
- Love, Paul, 353*n*
- low-risk releases, architecture for
 - at Amazon, 212–213
 - at Blackboard Learn, 215–217*f*
 - conclusion on, 218
 - downward spiral and, 208–209
 - at eBay, 207–208
 - loosely coupled architecture, 26–27, 209, 217
 - monoliths versus microservices, 210–212
 - service-oriented architecture, 109, 210
 - strangler fig application pattern and, 70, 208, 213–217, 218
- low-risk releases, enabling
 - Andon cord and, 181
 - application-based release patterns, 190, 195–199
 - automating deployment process, 179–181
 - conclusion on, 206
 - continuous delivery and, 199–201
 - at CSG International, 181–183, 201–206
 - dark launching, 197–199
 - environment-based release patterns, 190, 191–195
 - at Etsy, 186–188, 201
 - at Facebook, 177–179, 201
 - feature toggles, 195–197

self-service deployments, 184–188

Luyten, Stefan, 24*f*

M

Macri, Bethany, 49, 419

Magill, Stephen, 364

Maglio, Paul, 5

Making Work Visible: Exposing Time Theft to Optimize Work & Flow, 22, 58

Malpass, Ian, 227, 235*f*, 310

Mangot, Dave, 383, 384

manual tests

 automating, 160–161

 high reliance on, 10

manufacturing value stream, 7

market-oriented organizations, 100–101

market-oriented teams, 102–103

Marsh, Dianne, 118

Martens, Ryan, 94

Martin, Karen, 7, 9*n*, 11

mass marketing, 275

Massie, Bill, 385–387

Matatall, Neil, 360, 361

Mathew, Reena, 383, 384

matrix-oriented organizations, 100

Maurer, Dan, 274

Maven Central repository, 364, 365

Maximilien, E. Michael, 159*n*

McChrystal, Stanley, 347

McDonnell, Patrick, 226

McKinley, Dan, 332

mean time to recover. *See* MTTR

means and standard deviations, 246–248

Measuring Software Quality whitepaper, 58

Mediratta, Bharat, 149, 344

mentors, test certified, 345

Messeri, Eran, 150, 290, 325

Metasploit, 359, 369

metrics

- actionable business, 240–241*f*
- application and business, 240–242
- flow, 11–12
- infrastructure, 242–243
- library (StatsD), 234
- MTTR and, 227
- production, 234–235
- self-service, 236–238

Metz, Cade, 338

Mickman, Heather, 112–114, 297, 343

microservices

- monoliths versus, 210–212
- pros and cons of, 211*t*
- shifting to, 212–213

Microsoft Excel, 255

Microsoft Office Framework (MOF) Study, 225

Milstein, Dan, 310

Minimum viable product (MVP), 388

mistakes. *See also* failures

- examining, 308–310
- in learning-based culture, 307–308
- redefining failure, 314–315

Model of Change, John Shook's, 205

monitoring frameworks, 227–231

monolithic architectures

- defined, 210, 211*t*
- shifting from, 212–213

Moore, Geoffrey A., 72

Morgue tool, Etsy's, 311–312

Morrison, Erica, 183*f*, 204, 206, 318

motion waste, 28

MTTR

- batch size and, 292
- culture of causality and, 233*n*
- daily deployments and, 183*f*
- defined, 16
- fact-based problem-solving and, 234
- high performers and, 186*f*, 227, 228*f*
- low-risk changes and, 381
- Morgue tool for recording, 311
- optimizing for, 261*n*
- tracking metrics and, 227

Mueller, Ernest, 117*n*, 122, 173–175, 237

Mulkey, Jody, 104–105

Multichannel Digital Tax Platform (MDTP), 78–80

multitasking, 21–22

multivariate testing, defined, 276

MySQL, xxvi, xxxit, 226, 296, 332

Mythical Man-Month, The, xli

myths

- DevOps, xxiv–xxvi
- industrial safety, 416*t*

N

Nagappan, Nachi, 159*n*

NASA, 231, 313, 314

National Institute of Standards and Technology (NIST), 330, 331, 389

National Vulnerability Database, 368

Nationwide Building Society, 124–127

Nationwide Insurance, 339, 342

Naval Reactors (NR), 51

.NET, xxviii, 363

Netflix

- Archaius library, 196*n*
- auto-scaling capacity, 251–253
- case study, 251–253

Chaos Monkey, 52, 55, 306–307, 315, 420

- cluster immune system, 195*n*
- DevOps myths and, xxiv–xxv
- as digital-native company, 15
- market-oriented teams, 102
- shared services, 118
- Simian Army, 369*n*, 420
- telemetry, 245–246

Newland, Jesse, 321, 323

Nmap, 359, 369

non-functional requirements (NFRs), 89, 90*f*, 162–163, 328

Nordstrom, 15, 63–66

normal changes, 380, 381–382

North, Dan, 191*f*, 193, 232

Nygaard, Michael, 315, 347

O

object relational mapping (ORM), 99*n*

observability, and testing, 147

Octoverse Report, 57, 368, 401

O’Donnell, Glenn, 340

O’Neill, Paul, 50, 313

Open Web Application Security Project (OWASP), 359*n*, 360

Operation Desert Shield, 189*n*

opinion of the team, 40

opinionated platform, 79

Ops engineers on service teams, embedding, 119–120

Ops liaisons

- assigning, 117, 120–121
- purpose of, 115–117
- team rituals and, 121–124
- two types of, 116

O’Reilly, Barry, 278

organizational archetypes, 100–101

organizational culture

- importance of, 46
- just culture, 47, 307–308
- safety culture, 46–49
- three pillars of, 75–76
- three types of, 47, 48*t*
- trust and, xli, 45, 48, 150

organizational goals, and technology choices, 329–330

organizational knowledge

- chat rooms for capturing, 321–323
- easy-to-reuse, 323–324
- Ops user stories and, 328–329
- recommend_tech program at Target, 333–334
- source code repository for, 139–141, 150, 324–327

organizational learning

- ASREDS learning loop, 340, 341*f*
- communities of practice, 343–345
- conclusions on, 346
- DevOps conferences, 341–342
- Etsy, 49
- grouplets, 344, 345
- improvement blitz, 335, 336
- internal conferences, 342–343
- rituals to pay down technical debt, 336–339
- Teaching Thursday, 339
- thirty-day challenge at Target, 335–336

ORM (object relational mapping), 99*n*

Orzen, Mike, 49

Osterling, Mike, 7, 9*n*, 11

Otto, Andreia, 287

outages

- 2/4 outage, 318–320
- Adidas, 286
- Amazon Web Services (AWS), 305–306, 315
- culture of blame around, 233

- culture of causality and, 225
- mastering, 347
- Netflix, 314–315
- outlier detection, 245–246
- outsourcing contracts, IT, 102*n*
- “over the shoulder” code review, 290
- OWASP (Open Web Application Security Project), 359*n*, 360
- OWASP Dependency Check, 359, 365, 368*n*
- OWASP ZAP, 358*f*, 359
- Özil, Giray, 289

P

- pager rotation duties, 263–264
- pair programmed rollback, 164*n*
- pair programming, 289, 292–294
- pairing hours, 293*n*
- Pais, Manuel, 111, 117, 129
- Pal, Tapabrata, 296, 342–343, 353*n*
- Pandey, Ravi, 336
- Parker, Doug, 76–77
- Paroski, Drew, 338
- partially done work, 28
- passion, as cultural piller, 75
- passwords, OWASP guidance on, 360
- pathological organizations, 47, 48*t*
- Payment Card Industry Data Security Standards (PCI DSS), 385, 386
- PayPal, 73*n*
- PDCA (Plan, Do, Check, Act) cycle, 38, 54
- peer review
 - of changes, 288–290
 - code review basics, 288–291
 - email pass around, 290
 - GitHub, 281–283
 - “over the shoulder” code review, 290
 - pair programming, 289, 292–294

- pull requests, 281–283, 295–296
 - tool-assisted code review, 290
- performance testing environment, 161–162
- Perrow, Charles, 34
- Phoenix Project, The*, xxii, xlii, xxxix, 12, 404, 406, 413, 414
- PHP
 - Conway’s Law, 98, 99
 - Etsy, 226, 262*f*, 332
 - Facebook, 177*n*, 178, 199, 338
 - LAMP stack and DevOps, xxvi
 - as representative technology, xxxit
- Pinterest, 114
- Pivotal Labs, 293–294
- planning horizons, short, 88–89
- platform teams, 112
- Poppendieck, Mary and Tom, 27
- Porter, Chris, 376, 377
- post-mortems, blameless
 - CSG case study, 318–320
 - defined, 308
 - inviting Ops engineers to, 117, 123–124
 - for organizational learning, xxxix, 48–49
 - publishing reports of, 311–312
 - sample agenda, 418–419
 - scheduling of, 308–310
- Potvin, Rachel, 149, 300, 324, 326
- powerlessness, employees and, xxxvi
- pretotyping, 275*n*
- preventive security controls, 355–357
- problem-solving
 - guided by telemetry, 225, 233–234
 - seeing problems as they occur, 35–36, 244
 - swarming, 37–41
- process time, lead time versus, 9*f*

- product owner, 83
- production telemetry. *See also* telemetry; telemetry analysis
 - ATM systems and, 392
 - contextual inquiry, 264
 - in daily work, 234–235
 - feature toggles, 195–197, 262, 277
 - fix forward, 262
 - hand-off readiness review (HRR), 269, 270, 271*f*
 - information security and, 371–373
 - launch guidance, 266–267
 - launch readiness review (LLR), 270, 271*f*
 - LinkedIn, 237–238
 - pager rotation duties, 263–264
 - roll back, 262
 - service hand-back mechanism, 268*f*–269
 - site reliability engineers (SREs), 269–271
 - UX observation, 264–265
 - productivity, developer
 - DevOps practices for, *xlif*
 - measuring, 401–403
 - shared services for, 117–120
- Project to Product*, 12
- prototypes, creating, 275*n*
- Prugh, Scott, 72, 106–107, 181, 182, 183, 201–204, 231
- psychological safety, 347
- pull requests
 - defined, 281
 - evaluating, 295–296
 - GitHub, 281–283
- Puppet Labs, xxiii, xxxix, *xlif*, 5, 14, 140, 185, 217
- Python
 - Etsy, 332
 - ORM, 99*n*
 - PyPi for, 364

Zenoss and, 238

Q

quality controls, ineffective, 42–43

queue size, controlling, 22, 414–415*f*

Quicken, 274*n*

Quora, 278, 306

R

Rachitsky, Lenny, 420

Rajan, Karthik, 383

Rally, 94, 253, 381

Rapoport, Roy, 245, 246, 314

Rational Unified Process, 4

Raymond, Eric S., 97

rebooting servers, 225

Red Hat, xxxit, 138

Reddit, 306

Reddy, Tarum, 253

release, deployment versus, 189

Release It! Design and Deploy Production-Ready Software, 315, 347

release managers, 83

release patterns

application-based, 190, 195–199

environment-based, 190, 191–195

releases, architecture for low-risk. *See* architecture

releases, enabling low-risk

Andon cord and, 181

application-based release patterns, 190, 195–199

automating deployment process, 179–181

conclusion on, 206

continuous delivery and, 199–201

at CSG International, 181–183, 201–206

dark launching, 197–199

environment-based release patterns, 190, 191–195

at Etsy, 186–188, 201

- at Facebook, 177–179, 201
- feature toggles, 195–197
- self-service deployments, 184–188
- Rembetsy, Michael, 63, 226, 332
- remote work, Covid-19 pandemic and, 108*n*
- request for change (RFC) form, 380, 382
- resilience engineering, defined, 316
- resilience patterns, 51–52
- resilient organizations
 - calculated risk-taking in, 314–315
 - CSG case study, 318–320
 - description of, 305
 - experimental model for, 313
 - game days in, 316–318
 - just, learning culture in, 307–308, 320
 - Netflix example, 305–307
 - post-mortem reports in, 311–312
 - retrospective meetings at, 308–310
- retrospectives
 - CSG case study, 318–320
 - defined, 308
 - Ops engineers at, 117, 123–124
 - for organizational learning, xxxix, 48–49
 - publishing reports of, 311–312
 - sample agenda, 418–419
 - scheduling of, 308–310
- Rettif, Lucas, 333
- review and coordination processes
 - Adidas, 286–287
 - bureaucracy and, 42–43, 296–297
 - code review basics, 288–291
 - conclusions on, 297–298
 - coordination and scheduling, 288
 - dangers of change approval processes, 283–284

- dangers of manual testing and change freezes, 292
- dangers of overly controlling changes, 284–285
- GitHub, 281–283
- Google, 290–291
- pair programming, 289, 292–294
- Pivotal Labs, 293–294
- pull requests, 281–283, 295–296
- Rhoades, Lacy, 277–278
- Rice, David, 412
- Richardson, Vernon, xxxviii
- Ries, Eric, 6*n*, 24, 195*n*, 411
- Right Media, 259–261
- risk-taking, calculated, 314–315
- rituals
 - Dev team, 121–124
 - technical debt and, 336–339
- Robbins, Jesse, 316, 317, 399, 410
- Roberto, Michael, 313
- roll back, 262
- Rossi, Chuck, 177, 178, 198
- Rother, Mike, 6, 49, 53, 54, 104, 411
- Rouster, 383
- Ruby, 152, 231*n*, 234, 359, 363, 364
- Ruby on Rails, xxxix, 99*n*, 143, 191*n*, 362
- Rugged Computing movement, 412
- Rugged DevOps, 353
- Rushgrove, Gareth, 358*f*

S

- Sachs, Marcus, 371
- safety
 - in complex systems, 33–34
 - myths about, 416*t*
- safety culture, 46–49
- Safety Differently: Human Factors for a New Era*, 395

[Salesforce.com](https://www.salesforce.com), 383–384

Sarbanes-Oxley Act, 267

Savoia, Alberto, 275*n*

scenius, 54–55

Schmidt, Eric, 69

Schwaber, Ken, 122*n*

Schwartz, Mark, 299

Scott, Kevin, 92, 93

Scrum, 122*n*, 410

Scryer, 251–253

Second Way, The. *See* The Second Way (principles of feedback)

security. *See* information security

security and compliance, DevOps compatibility with, xxv

Security Chaos Engineering, 142*n*

Security Monkey, 420

selflessness, as cultural pillar, 75

Senge, Peter, 35, 49, 320

separation of duty, 379, 384–387

service-oriented architectures (SOAs), 109, 210

settling period, 243

Shafer, Andrew Clay, xxiv, 5, 410

Shared Operations Team (SOT), 181–182*n*

shared services, 117–119

Shewhart, Walter, 54

Shewhart cycle, 38, 54

Shingo, Shigeo, 27

Shinn, Bill, 389–391

Shook, John, 205

Shortridge, Kelly, 142*n*

Shoup, Randy, xli, 109–110, 164, 207, 208, 209*f*, 210, 211*f*, 289, 291, 311, 326

silent majority, 74

siloization, 105

silos

- breaking down, 124–127

- functional teams in, 125, 126*f*
- learning, 340
- Simian Army, 369*n*, 420
- single repository of truth
 - as deployment pipeline foundation, 139–141, 150
 - for global improvement, 324–327, 334
 - at Google, 324–326
- site reliability engineers (SREs), 269–271
- Skelton, Matthew, 111, 117, 129
- small batch sizes, 7, 9, 22–24*f*, 409
- small teams, 110–111
- Smart, Jonathan, 125, 341*f*, 403
- smoke testing deployments, 180, 187
- Smollen, Alex, 360, 361
- smoothing, 253
- Snover, Jeffrey, xxxii
- Snyder, Ross, 98, 99, 100
- software, COTS, 417–418
- software delivery. *See also* deployment pipeline foundations; low-risk release, enabling
 - faster, xxxiiit
 - history of, xxx–xxxit
- software supply chain, security of, 363–369
- software type, myth about DevOps and, xxvi
- SolarWinds security breach, 368
- Sonatype Nexus Lifecycle, 368*n*
- Sonatype State of the Software Supply Chain Report, 364, 365, 366
- Sooner Safer Happier*, 125, 340, 395, 403
- source code repository, shared
 - as deployment pipeline foundation, 139–141, 150
 - for global improvement, 324–327, 334
 - at Google, 324–326
 - preventive security controls and, 355–357
- SPACE framework, 403

Spafford, George, xliv, 225, 233*n*, 353*n*, 415*f*

Spear, Steven J., xxxvii, 34, 38, 49, 50, 58, 105, 305, 313, 335, 338

specialists, 105–107

Spotify, 15

spring cleanings, 337

sprint planning boards, 20-21

Sprouter, 98–100*n*, 108

SQL injection attacks, 373, 374*f*

Stack Exchange, 278, 293

stakeholder feedback, 36

standard changes

- description of, 379–380
- at [Salesforce.com](https://salesforce.com), 383–384

standups, daily, 117, 120, 122, 410

Starbucks, 15

startups, DevOps for, xxiv

State of DevOps Reports, xln, xlii, 5, 14, 15, 57, 67, 71, 103, 140, 165, 166, 175, 185, 201, 217, 227, 243, 284, 285, 312, 326, 331, 343

State of the Octoverse report, GitHub's, 57, 368, 401

static analysis, 358*f*–359

statistical analysis software, Tableau, 253

StatsD metric library, 234, 235*f*

Stillman, Jessica, 338

Stoneham, Jim, 278–279, 280

strangler fig application pattern

- at Blackboard Learn, 215–217*f*
- blog explaining, 220
- defined, 70, 208, 213–215, 218

stream-aligned teams, 112

Strear, Chris, 29–32

Sussman, Noah, 187

Sussna, Jeff, 265*n*

swarming

- case study, 39–41

purpose of, 37–39
systems of engagement, 71, 72
systems of record, 71, 72

T

Tableau, 253

Taleb, Nassim Nicholas, 52

Tang, Diane, 276

Target

- API enablement, 112–114
- approval processes, 296–297
- case study, 112–114, 333–334, 342, 343
- internal conferences, 342, 343
- recommend_tech program, 333–334
- thirty-day challenge, 335–336
- as traditional enterprise, 15

task switching, 28

tax collection agency, UK's, 77–80

Team of Teams: New Rules of Engagement for a Complex World, 347

team opinion, requesting, 40

team rituals, 121–124

Team Topologies: Organizing Business and Technology Teams for Fast Flow, 111, 117, 129

teams

- 18F team, 369–371
- Conway's Law and, 108
- dedicated team, 86–88
- four types of, 111–112
- functional teams in silos, 125, 126*f*
- generalists on, 105–107
- Google Web Server (GWS) team, 148–151
- long-lived, 126*f*
- platform team, 112
- Shared Operations Team (SOT), 181–182*n*
- stream-aligned, 112

- two-pizza team, 110–111
 - in value stream, 83
- technical debt
 - brownfield transformations and, 67
 - defined, xxxiv–xxxv, 171, 205
 - reducing, 89–93, 165
 - rituals to pay down, 336–339
 - test automation and, 161, 165
- technology choices
 - Etsy, 332
 - organizational goals and, 329–330
 - Target, 333–334
- Technology Enterprise Adoption Process (TEAP), 296, 297
- technology value stream. *See also* value stream; value stream selection
 - description of, 8–12
 - invisible work in, 19–20
- telemetry. *See also* production telemetry; telemetry analysis
 - application and business metrics, 240–242
 - application logging, 231–233
 - centralized telemetry infrastructure, 227–231
 - conclusions on, 244
 - culture of causality, 225, 233*n*
 - daily work and, 234–235
 - defined, 225–226
 - detection of problems with, 36
 - Etsy, 226–227
 - gaps, 239–240
 - information radiators, 236, 237, 241
 - infrastructure metrics, 242–243
 - monitoring frameworks, 227–231
 - Netflix, 245–246
 - problem-solving guided by, 233–234
 - security-related, 371–375
 - self-service access to, 236–237

- self-service metrics at LinkedIn, 237–238
- settling period, 243–244
- StatsD metric library, 234, 235*f*
- telemetry analysis
 - alerting systems, 248–249, 250*f*, 255*f*
 - anomaly detection techniques, 253–257
 - Gaussian distribution, 247*f*, 249
 - Kolmogorov-Smirnov test, 254, 255, 256*f*, 257
 - means and standard deviations, 246–248
 - non-Gaussian distribution, 249–251*f*
 - outlier detection, 245–246
 - Scryer tool at Netflix, 251–253
 - smoothing, 253–254
 - Tableau software, 253
- Terhorst-North, Dan, 191*f*, 193, 232
- test setup and run, 26
- test-driven development (TDD), 159, 161, 327
- testing, A/B
 - conclusions on, 280
 - feature toggles and, 196–197
 - history of, 275
 - integrating into feature planning, 278
 - integrating into feature testing, 276–277
 - integrating into our release, 277–278
 - need for, 273
 - practical guide to, 276
- testing, automated
 - Andon cord and, 163–165
 - categories of tests, 155–156
 - conclusions on, 166
 - deployment pipeline infrastructure, 151–154
 - essential components of, 165–166
 - fast and reproducible tests, 166
 - at Google, 148–151

- green build and, 154, 163, 166
- ideal testing pyramid, 157f–158
- need for, 147–148
- non-functional requirements and, 162–163, 328
- observability and, 147
- performance testing environment, 161–162
- reducing reliance on manual tests, 160–161
- research supporting, 165–166
- running tests in parallel, 158f–159
- test-driven development (TDD), 159, 161, 327

testing, operations, and security, as everyone’s job, 104–105

Testing on the Toilet newsletter, 149n, 344

The First Way (principles of flow)

- case study, 29–32
- constraint identification, 25–27
- general description of, 13f, 19
- handoff reduction, 24–25
- kanban boards, 4, 20f–21, 94, 124
- limiting work in progress (WIP), 7, 10, 21–22
- making work visible, 19–21, 93, 94, 124
- reducing batch sizes, 22–24f
- small batch sizes, 7, 9, 22–24f
- waste and hardship elimination, 27–29

The Second Way (principles of feedback)

- case study, 39–41
- description of, 13f–14
- feedback types and cycle times, 36–37f
- optimizing for downstream work centers, 43–44
- quality control closer to the source, 42–43
- safety within complex systems, 33–34
- seeing problems as they occur, 35–36
- swarming, 37–39, 41

The Third Way (continual learning and experimentation)

- case study, 54–55

- conclusions on, 55–56
- description of, 13*f*, 14
- global knowledge, 51
- improvement of daily work, 49–50
- just culture, 47, 307–308
- leader’s role, 52–54
- resilience patterns, 51–52
- safety culture for, 46–49

The Three Ways

- case study, 15–18
- defined, 12–14
- high-level principles of, xlv
- research supporting, 14–15

Theory of Constraints, 4, 412–413

Third Way, The. *See* The Third Way (continual learning and experimentation)

Thoughtworks’ *Tech Radar*, 57

Thrasher, Paula, 129

Three Mile Island, 34

Three Ways, The. *See* The Three Ways

ticket, defined, 381*n*

Ticketmaster/LiveNation, 104, 243

Tischler, Tim, 184

Tomayko, Ryan, 295, 296

tool-assisted code review, 290

Toyota Kata: Managing People for Improvement, Adaptiveness and Superior Results, 6, 49, 104, 411

Toyota Kata Movement, description of, 4, 6, 49, 54, 411

Toyota Production System

- Andon cord, 37–39, 163, 416–417*f*
- core belief in, 285
- improvement blitz, 335
- improvement kata, 6
- information radiators, 236, 237, 241
- Lean Movement and, 409

- techniques, 4
- transparent uptime, 311*n*, 420–421
- Treynor Sloss, Ben, 269
- Trimble, Chris, 86, 87
- trunk-based development
 - adopting practices of, 172
 - at Bazaarvoice, 173–175, 176
 - benefits of, 175
 - conclusions on, 175–176
 - gated commits and, 172
 - for HP's LaserJet Firmware division, 168–170
 - need for, 167, 171
- Trustworthy Online Controlled Experiments: A Practical Guide to A/B Testing, 276
- TurboTax, 274, 275
- Turnbull, James, 228, 230*f*
- Twitter
 - architecture, 210
 - experimentation, 279
 - Fail Whale error page, 360
 - static security testing, 360–363
- two-pizza team, 110–111

U

- Unicorn Project, The*, 219, 333*n*
- unit tests, 155
- universality of the solution, xlii–xliii
- urgent changes, 380
- US Air Force, 69–71
- US Federal Government agencies, 369–371
- US Navy, 34*n*, 51
- user acceptance testing (UAT), 152
- user feedback, 36, 37*f*
- user research, 273, 275*n*, 277. *See also* A/B testing
- user stories, 328–329

UX movement, Lean, 411

UX observation, 264–265

V

value stream. *See also* value stream selection

defined, 7, 18

manufacturing, 7

selection of, 61, 63

technology, 8–12

value stream manager, 83

value stream mapping

%C/A metrics in, 11, 85, 95

case studies, 15–18, 91–93

conclusions on, 85

creating a map, 84–86

defined, 81, 84–86

dedicated transformation team and, 86–88, 95

example of a value stream map, 85*f*

goals and, 88

Lean movement and, 4

members of value stream, 83

at Nordstrom, 81–83

technical debt and, 89–93

Value Stream Mapping: How to Visualize Work and Align Leadership for Organizational Transformation, 7

value stream selection

American Airlines' new vocabulary, 74–77*t*

brownfield transformation of refueling system, 69–71

case studies, 69–71, 74–77, 77–80

expanding scope of initiative, 73–77

greenfield vs. brownfield services, 66–69

innovators and early adopters, 72–74

Nordstrom's DevOps journey, 63–66

systems of engagement, 71, 72

systems of record, 71, 72

tax collection agency's journey, 77–80

van Kemenade, Ron, 74

Van Leeuwen, Evelijn, 343*n*

Vance, Ashlee, 92

Velocity Conference, 410

Verizon data breach, 364, 371

version control system, 139–141

versioned APIs, 214

Vincent, John, 247

visibility of automated test failures, 164

visibility of work, 19–21, 93, 94

Visible Ops Handbook, The, 225

Visible Ops Security, 353*n*

visual work boards, 20*f*–21. *See also* kanban boards

vocabulary, using a new, 76–77*t*

Vogels, Werner, 110, 212

Vulnerabilities and Exposures, Common (CVEs), 364, 365

Vulnerability Database, National, 368

W

Walker, Jason, 333

Wang, Kendrick, 277

WARN level, 232

waste and hardship elimination, 27–29

water-Scrum-fall anti-pattern, 165*n*

weak failure signals, 313–314

Westrum, Ron, 47, 48*t*, 57

Westrum Organizational Typology Model, 48*t*

Wickett, James, 353

Williams, Branden, xxv

Williams, Jeff, 412

Williams, Laurie, 159*n*, 293

Willis, John, xxiii–xxiv, 3, 5, 57, 406–407, 409, 410

Wolberg, Kirsten, 73*n*

Womack, James P., 23, 53

Wong, Bruce, 316

Wong, Eric, 238

work, visibility of, 19–21, 93, 94, 124

work in process (WIP), 7, 10, 21–22

workplace safety, at Alcoa, 50

X

Xu, Ya, 276

Y

Yadav, Yikalp, 286

Yahoo! Answers, 278–280

Yuan, Danny, 252*f*, 254*f*

Z

Zenoss, 229, 238

Zhao, Haiping, 338

ZooKeeper, 242

Zuckerberg, Mark, 338

ACKNOWLEDGMENTS

Gene Kim

Thank you to my wife, Margueritte, and my sons, Reid, Parker, and Grant, for putting up with me being in deadline mode for over five years. And to my parents for helping me become a nerd early in life. Thank you to my amazing co-authors, as well as the IT Revolution team who made this book possible: Anna Noak and especially to Leah Brown for all her heroic work pulling this second edition together!

I am so grateful to everyone who made the first edition possible: John Allspaw (Etsy), Alanna Brown (Puppet), Adrian Cockcroft (Battery Ventures), Justin Collins (Brakeman Pro), Josh Corman (Atlantic Council), Jason Cox (The Walt Disney Company), Dominica DeGrandis (LeanKit), Damon Edwards (DTO Solutions), Dr. Nicole Forsgren (Chef), Gary Gruver, Sam Guckenheimer (Microsoft), Elisabeth Hendrickson (Pivotal Software), Nick Galbreath (Signal Sciences), Tom Limoncelli (Stack Exchange), Chris Little, Ryan Martens, Ernest Mueller (AlienVault), Mike Orzen, Christopher Porter (CISO, Fannie Mae), Scott Prugh (CSG International), Roy Rapoport (Netflix), Tarun Reddy (CA/Rally), Jesse Robbins (Orion Labs), Ben Rockwood (Chef), Andrew Shafer (Pivotal), Randy Shoup (Stitch Fix), James Turnbull (Kickstarter), and James Wickett (Signal Sciences).

I also want to thank the many people whose incredible DevOps journeys we studied, including Justin Arbuckle, David Ashman, Charlie Betz, Mike Bland, Dr. Toufic Boubez, Em Campbell-Pretty, Jason Chan, Pete Cheslock, Ross Clanton, Jonathan Claudius, Shawn Davenport, James DeLuccia, Rob England, John Esser, James Fryman, Paul Farrall, Nathen Harvey, Mirco Hering, Adam Jacob, Luke

Kanies, Kaimar Karu, Nigel Kersten, Courtney Kissler, Bethany Macri, Simon Morris, Ian Malpass, Dianne Marsh, Norman Marks, Bill Massie, Neil Matatall, Michael Nygard, Patrick McDonnell, Eran Messeri, Heather Mickman, Jody Mulkey, Paul Muller, Jesse Newland, Dan North, Dr. Tapabrata Pal, Michael Rembetsy, Mike Rother, Paul Stack, Gareth Rushgrove, Mark Schwartz, Nathan Shimek, Bill Shinn, JP Schneider, Dr. Steven Spear, Laurence Sweeney, Jim Stoneham, and Ryan Tomayko.

And I am so profoundly grateful for the many reviewers who gave us fantastic feedback that shaped this book: Will Albenzi, JT Armstrong, Paul Auclair, Ed Bellis, Daniel Blander, Matt Brender, Alanna Brown, Branden Burton, Ross Clanton, Adrian Cockcroft, Jennifer Davis, Jessica DeVita, Stephen Feldman, Martin Fisher, Stephen Fishman, Jeff Gallimore, Becky Hartman, Matt Hatch, William Hertling, Rob Hirschfeld, Tim Hunter, Stein Inge Morisbak, Mark Klein, Alan Kraft, Bridget Kromhaut, Chris Leavory, Chris Leavoy, Jenny Madorsky, Dave Mangot, Chris McDevitt, Chris McEniry, Mike McGarr, Thomas McGonagle, Sam McLeod, Byron Miller, David Mortman, Chivas Nambiar, Charles Nelles, John Osborne, Matt O’Keefe, Manuel Pais, Gary Pedretti, Dan Piessens, Brian Prince, Dennis Ravenelle, Pete Reid, Markos Rendell, Trevor Roberts, Jr., Frederick Scholl, Matthew Selheimer, David Severski, Samir Shah, Paul Stack, Scott Stockton, Dave Tempero, Todd Varland, Jeremy Voorhis, and Branden Williams.

And several people gave me an amazing glimpse of what the future of authoring with modern toolchains looks like, including Andrew Odewahn (O’Reilly Media) who let us use the fantastic Chimera reviewing platform, James Turnbull (Kickstarter) for his help creating my first publishing

rendering toolchain, and Scott Chacon (GitHub) for his work on GitHub Flow for authors.

Jez Humble

Creating this book has been a labor of love, for Gene in particular. It's an immense privilege and pleasure to have worked with Gene and my other co-authors, John and Pat, along with Todd, Anna, Robyn and the editorial and production team at IT Revolution preparing this work—thank you. I also want to thank Nicole Forsgren whose work with Gene, Alanna Brown, Nigel Kersten and I on the PuppetLabs/DORA *State of DevOps Report* over the last three years has been instrumental in developing, testing and refining many of the ideas in this book. My wife, Rani, and my two daughters, Amrita and Reshmi, have given me boundless love and support during my work on this book, as in every part of my life. Thank you. I love you. Finally, I feel incredibly lucky to be part of the DevOps community, which almost without exception walks the talk of practicing empathy and growing a culture of respect and learning. Thanks to each and every one of you.

Patrick Debois

I would like to thank those who were on this ride, much gratitude to you all.

John Willis

First and foremost, I need to acknowledge my saint of a wife for putting up with my crazy career. It would take another book to express how much I learned from my co-authors Patrick, Gene and Jez. Other very important influencers and advisers in my journey are Mark Hinkle, Mark Burgess, Andrew Clay Shafer, and Michael Cote. I also want to give a shout out to Adam Jacob for hiring me at Chef and giving me the freedom to explore, in the early days, this thing we call

Devops. Last but definitely not least is my partner in crime, my *Devops Cafe* cohost, Damon Edwards.

Nicole Forsgren

I'm grateful to Jez and Gene, whose collaboration on the *State of DevOps Reports* (and later DORA) provided a wonderful research lab and foundation for work. The partnership wouldn't have been possible without the insights of Alanna Brown, who spearheaded the initiative and who we were able to work with on those first few reports. I continue to be grateful for those who believe in women—and believe women—with ambition, ideas, and opinions. Some notable people in my life are my parents (sorry for the stress!), my dissertation advisors Suzie Weisband and Alexandra Durcikova (much love and respect to you for taking a chance on me and my pivot), Xavier Velasquez (you've always believed my plans will work), and of course my lady backchannel (thank you for love, support, and for space to rage so it doesn't end up on my Twitter timeline). And as always, Diet Coke.

ABOUT THE AUTHORS

Gene Kim is a bestselling author, researcher, multiple award-winning CTO, and founder of IT Revolution. His books include *The Phoenix Project*, *The Unicorn Project*, and *Accelerate*. Since 2014, he has been the founder and organizer of DevOps Enterprise Summit, studying the technology transformations of large, complex organizations.

Jez Humble has co-authored several bestselling books on software, including Jolt Award-winner *Continuous Delivery* and Shingo Publication Award winner *Accelerate*. He works at Google and teaches at the University of California, Berkeley.

Patrick Debois is Director of DevOps Relations and Advisor at Snyk. He works to bridge the gap between projects and operations by using Agile techniques in development, project management, and system administration.

John Willis is Senior Director of Global Transformation Office at RedHat. He has worked in the IT management industry for more than thirty-five years. He is coauthor of *Beyond The Phoenix Project* and host the Profound podcast.

Nicole Forsgren, PhD, is Partner at Microsoft Research leading the Developer Velocity Lab. She is author of the Shingo Publication Award-winning book *Accelerate: The Science of Lean Software and DevOps* and is best known as lead investigator on the largest DevOps studies to date. She has been a successful entrepreneur (with an exit to Google), professor, performance engineer, and sysadmin. Her work has been published in several peer-reviewed journals.