

Chaos Engineering with Go

Building resilient systems through controlled chaos

Kapil Kumar Khandelwal

Mohit Garg



www.bpbonline.com

First Edition 2025

Copyright © BPB Publications, India

ISBN: 978-93-55519-641

All Rights Reserved. No part of this publication may be reproduced, distributed or transmitted in any form or by any means or stored in a database or retrieval system, without the prior written permission of the publisher with the exception to the program listings which may be entered, stored and executed in a computer system, but they can not be reproduced by the means of publication, photocopy, recording, or by any electronic and mechanical means.

LIMITS OF LIABILITY AND DISCLAIMER OF WARRANTY

The information contained in this book is true to correct and the best of author's and publisher's knowledge. The author has made every effort to ensure the accuracy of these publications, but publisher cannot be held responsible for any loss or damage arising from any information in this book.

All trademarks referred to in the book are acknowledged as properties of their respective owners but BPB Publications cannot guarantee the accuracy of this information.

To View Complete
BPB Publications Catalogue
Scan the QR Code:



Dedicated to

To my parents

Sh. Mahesh Chand and Smt. Rukmani Devi

whose unwavering support and encouragement have been my foundation

To my wife, **Ankita**

whose love and understanding have been my strength.

And to my son, **Divit**, who inspires me every day

With all my love and gratitude

– Kapil Kr. Khandelwal

To my parents

Sh. Dev Raj Garg and Smt. Late Nisha Garg

for their unwavering belief in me.

To my wonderful wife, **Vatika Garg**

whose love and patience have made this journey possible.

And to my children

Avika Garg and Yash Garg

who are my greatest source of joy and motivation

– Mohit Garg

About the Authors

- **Kapil Kumar Khandelwal** is a passionate tech and product enthusiast with over a decade of experience. He launched his illustrious career after achieving his B.Tech degree and later earned an executive MBA in Leadership and Finance. He also holds multiple certifications in quality, product, and cloud.

Kapil has worked in leading service and product firms, contributing to mission-critical products and unlocking the potential of Chaos Engineering in real-world scenarios. His unquenchable curiosity in Chaos Engineering has propelled him to become a highly sought-after authority in the field.

When he is not immersed in tech challenges, you can find Kapil exploring hiking trails, learning new tools, and building connections within the tech and product community through insightful networking events.

- **Mohit Garg** completed his Bachelor of Technology in Information Technology in 2010, followed by an MBA in Leadership and Finance. With over 13 years of experience in the IT industry, he has contributed to and engaged with leading product and service-based companies. His vast experience has been instrumental in writing this book, providing practical insights and deep understanding of Chaos Engineering.

Throughout his career, Mohit has remained active in learning new tools and technologies, constantly expanding his skill set and building new connections. This dedication to continuous improvement and networking has kept him at the forefront of industry trends and best practices.

Outside of his professional life, Mohit enjoys spending quality time with his family. He loves playing with his children, Avika and Yash, exploring music, watching movies, and taking long drives. These hobbies provide him with a refreshing break from work and keep him motivated.

Mohit's dedication to his profession, combined with his personal interests and family life, has contributed to a well-rounded perspective, making this book a valuable resource for anyone looking to explore Chaos Engineering.

About the Reviewers

- ❖ **Daniel Moreira Cardoso** is a seasoned Senior Software Engineer with several years of experience in software development. He has proficiency in a variety of programming languages and technologies, including Typescript, Kotlin, Golang, PostgreSQL, Apache Kafka, Kubernetes, Google Cloud Platform, AWS, Datadog, Next.js, and React. His expertise lies in developing solutions for financial domains, payments, municipal public sectors, and sales platforms, significantly improving the performance, availability, reliability, and resilience of backend systems. In addition to his technical abilities, Daniel has experience leading technical teams and contributing to major projects that enhanced business efficiency and customer satisfaction. He is known for his proactive problem-solving approach and dedication to automating processes and integrating innovative solutions to meet dynamic business and client needs. Notably, Daniel has made significant contributions to optimizing tooling for load tests and implementing Chaos Engineering practices, which have improved the resilience and reliability of the systems he has worked with.
- ❖ **Syed** is a pioneer and subject matter expert (SME) in application modernization and cloud computing. He is a holder of multiple industry certifications in Azure, AWS, GCP, Oracle, PMP, and CSM. With more than a decade of expertise in product development, architecture, scalability, and modernization, Syed has held positions with Oracle, the Big 4 technology advisory companies like PwC and Ernst & Young, as well as other contemporary application enterprises. Syed has also published articles in the leading Journal of Cloud Computing. Syed is a skilled and devoted technology management executive who employs more than 15 years of expertise in the design and development of high-demand On Premise / Commercial of the Shelf Software / SaaS solutions. He possesses specialized knowledge in designing and constructing highly scalable and highly available systems. In order to assist them in establishing and managing their businesses by taking control of their frequently cumbersome applications, Syed has provided consulting services to some of the world's top corporations. Syed assists businesses in modernizing their applications by emphasizing scale and availability, aiding cloud migration, DevOps transformations, and risk-based problem identification. Syed offers expert guidance, strategy, and thought leadership to assist in solving application modernization issues. Syed promotes change by collaborating with enterprises of all sizes and executive teams.

Acknowledgements

I would like to express my deepest gratitude to all those who have supported and inspired me throughout the journey of writing this book.

To my family, your unwavering love and encouragement have been my foundation. A special thanks to my parents, whose blessings and support have shaped my path, and to my beloved wife, Ankita, and son, Divit, for their endless patience and understanding.

To my extended family and especially my sisters, your constant support and belief in me have been a source of great motivation.

I am immensely grateful to BPB Publications, whose expertise and dedication have been instrumental in bringing this book to life. Your professional guidance and support have been invaluable.

I would also like to acknowledge the valuable contributions of my colleagues and co-workers during my years in the tech industry. Your teachings and feedback have been crucial to my growth.

A special thank you to my close friend and co-author, Mohit Garg. Your collaboration, insights, and hard work have been pivotal in shaping this book. It has been an incredible journey working alongside you.

Finally, to everyone who has contributed in any way to the creation of this book, thank you. Your support, encouragement, and contributions have made this possible.

- Kapil Kr. Khandelwal

Creating this book has been a remarkable journey, and I am incredibly thankful for the support of many individuals. First, I would like to thank my parents, Sh. Dev Raj Garg and Smt. Late Nisha Garg. Your unwavering love and support have been my guiding light.

To my wonderful wife, Vatika Garg, your endless encouragement and understanding have been vital to this project. To my children, Avika Garg and Yash Garg, your energy and enthusiasm have been a constant source of inspiration.

To my extended family, especially my sister, thank you for your support and encouragement. To my friends, your belief in me has been a driving force throughout this journey

I extend my heartfelt thanks to BPB Publications for providing me with this incredible opportunity. Their professional guidance and trust have been paramount in bringing this book to life. Their support has been instrumental in every step of this journey.

I am also deeply grateful to my close friend and co-author, Kapil Kr. Khandelwal. Your collaboration, hard work, and insights have been essential to this project. Working with you has been an enriching experience, and I value your friendship and partnership immensely.

Lastly, I am incredibly thankful to all those who have supported me along the way.

- Mohit Garg

Preface

Chaos Engineering with Go is a comprehensive guide crafted for both beginners and experienced professionals aiming to enhance system reliability and minimize downtime. This book serves a diverse audience, including **Site Reliability Engineers (SREs)**, DevOps engineers, tech professionals, and students, offering valuable insights into building resilient systems through controlled chaos.

Readers should have a basic understanding of programming fundamentals, the Go programming language, the **Software Development Life Cycle (SDLC)**, cloud-native concepts, distributed systems, and DevOps principles. For those with prior experience in Chaos Engineering or Go programming, this book provides advanced topics and practical applications to deepen their knowledge.

Discover innovative techniques for chaos fault injection, explore the latest chaos testing tools, and design effective chaos experiments to strengthen your infrastructure against unexpected failures. Learn about resilience patterns, fault tolerance strategies, and security Chaos Engineering. Real-world case studies demonstrate the practical application of these concepts in various scenarios.

With a focus on best practices and lessons learned, this book equips you with the knowledge and tools needed to embrace chaos, ensuring robust and reliable systems in an ever-evolving digital landscape. Whether you are a seasoned engineer or a curious beginner, Chaos Engineering with Go is your indispensable companion on the journey to engineering resilience.

Chapter 1: Exploring the Essence of Chaos Engineering - This chapter is an essential part of the book Chaos Engineering with Go as it lays the groundwork for those who are unfamiliar with the topic. Using relevant examples, it offers a crucial comprehension of the underlying concepts, evolution, and advantages. It explores the mentality and cultural elements required for firms to successfully adopt Chaos Engineering.

Chapter 2: Chaos Engineering Concepts -This chapter discusses topics such as principles of Chaos Engineering, and how it differs from functional and non-functional testing. Understanding these principles allows readers to execute Chaos Engineering activities in a methodical and structured manner.

Chapter 3: Revision with Go - This chapter introduces the Go programming language, which is essential for implementing Chaos Engineering with Go. It covers Go's syntax, principles, and key ideas to ensure readers have a solid understanding before exploring specific Chaos Engineering topics.

Managing packages and dependencies in Go is crucial for using existing Chaos Engineering tools and creating innovative solutions. Common tasks in Chaos Engineering include simulating concurrent events, injecting errors, and monitoring system reactions. Understanding concurrency and goroutines allows readers to leverage Go's robust concurrency features for effective chaos experiments. Additionally, mastering Go's error management techniques is vital for properly handling errors in Chaos Engineering experiments.

Chapter 4: Fault Tolerance and Resilience Patterns - In this chapter, you will understand fault tolerance in microservices architecture and its crucial role in maintaining system reliability. This chapter outlines the fundamental principles for implementing fault tolerance and laying the groundwork for resilient systems. Explore key resilience patterns such as timeouts, retries, and circuit breakers, examining how they mitigate failures. Learn effective strategies for testing fault tolerance to ensure robustness in distributed systems. Through a compelling case study of a distributed system built in Go, discover third-party libraries for fault injection and gain valuable lessons and best practices for constructing resilient distributed systems in Go.

Chapter 5: Chaos Fault Injection Techniques - This chapter focuses on the practical application of Chaos Engineering in real-world scenarios. It provides detailed guidance on using Go to introduce faults and simulate various stress conditions. Topics include a range of failure scenarios such as process crashes, disk space issues, high CPU usage, out-of-memory situations, component failures, and load balancer changes. Readers will learn how to replicate these failures and monitor system responses effectively.

The chapter equips readers with the knowledge and techniques to create automated chaos experiments using Go, covering topics like error automation and dependency testing. It also serves as a guide to leveraging Go's libraries, tools, and frameworks for successful fault injection and testing system resilience.

Chapter 6: Chaos Testing Tools - This chapter offers an in-depth exploration of essential chaos testing tools, including Chaos Monkey, Toxiproxy, Chaos Toolkit, and Chaos Mesh. Readers gain insights into their functionalities and applications in fault injection testing. Step-by-step instructions for installing Chaos Mesh on Ubuntu facilitate practical implementation. Explore the specific faults that can be injected via Chaos Mesh, such as

Pod and DNS failures, along with stress experiments. Learn about Chaos Mesh's limitations and emerging industry trends in chaos engineering tools like Gremlin and Litmus Chaos, helping readers select the right tools for their chaos engineering endeavors.

Chapter 7: Chaos Experiment Design - This chapter stresses how crucial it is to approach the planning and execution of experiments in chaos engineering in a logical and systematic manner. In order to successfully prepare, carry out, and analyze chaos experiments, it offers readers a structured framework. It aids readers in understanding the general objectives of each chaotic experiment as well as the particular system characteristics that researchers want to test, failure scenarios that they hope to model, and specific failure scenarios that they want to simulate. To evaluate the system's behavior during chaos experiments, this chapter assists readers in defining relevant metrics and observability measurements.

Chapter 8: Chaos with Emerging Tech Stack - Explore cutting-edge chaos engineering techniques in this chapter, featuring the powerful Azure Chaos Studio—an indispensable tool within the Azure ecosystem. Uncover the intricacies of high availability testing through chaos engineering methodologies, ensuring your systems are robust and resilient. Dive into the realm of Artificial Intelligence, discovering how AI elevates chaos engineering to new heights, introducing automated scenario generation and adaptive testing. This chapter provides a succinct guide to leveraging Azure Chaos Studio and AWS fault simulator, mastering high availability testing, and seamlessly integrating AI into chaos engineering for unparalleled system resilience and innovation.

Chapter 9: Essence of Observability in Distributed System - In this chapter, we will explore the essence of observability and its crucial role in understanding distributed systems. It is a foundational guide for engineers, covering key components and implications in modern infrastructures. We will also discuss application monitoring approaches and tools like Dynatrace, emphasizing the importance of logging with the ELK Stack and best practices for Go. Tracing is highlighted, focusing on its relevance and a step-by-step guide to implementing OpenTelemetry. By the end, readers will grasp observability principles and practical tools essential for building resilient, observable distributed systems.

Chapter 10: Observability in Chaos Engineering - This chapter emphasizes the vital role of observability in chaos engineering. It explains how observability provides real-time visibility into system performance during chaos experiments, identifying vulnerabilities before they escalate. Readers will explore various observability tools tailored for chaos engineering and learn to select the most suitable ones for their needs. Key metrics are categorized into resilience, infrastructure, and application metrics, highlighting their

importance in evaluating system resilience. By effectively logging and tracing chaos experiments, engineers can reconstruct system behavior and derive actionable insights, equipping them to build resilient systems through controlled chaos.

Chapter 11: Security Chaos Engineering Overview - Explore Security Chaos Engineering (SCE) essentials, from foundational principles to practical application. Learn how to apply chaos engineering techniques to fortify security resilience, utilizing specialized tools tailored for SCE. Join us as we navigate the intersection of Chaos Engineering and cybersecurity, forging stronger defenses through deliberate disruption.

Chapter 12: Case Studies: Chaos Engineering in Action - This chapter presents real-world case studies demonstrating how Chaos Engineering is applied across various sectors to enhance system resilience and reliability. These examples serve to inspire readers to implement Chaos Engineering techniques within their organizations, highlighting successful implementations in specific domains.

Chapter 13: Best Practices and Lessons Learned - This chapter consolidates insights from real-world chaos engineering experiences, offering valuable information for readers. It analyzes common pitfalls and mistakes to help avoid them on the path to resilient systems. The chapter discusses effectively integrating Chaos Engineering into the Software Development Life Cycle (SDLC) as a preventive strategy.

Code Bundle and Coloured Images

Please follow the link to download the *Code Bundle* and the *Coloured Images* of the book:

<https://rebrand.ly/oxg6aji>

The code bundle for the book is also hosted on GitHub at

<https://github.com/bpbpublications/Chaos-Engineering-with-Go>.

In case there's an update to the code, it will be updated on the existing GitHub repository.

We have code bundles from our rich catalogue of books and videos available at **<https://github.com/bpbpublications>**. Check them out!

Errata

We take immense pride in our work at BPB Publications and follow best practices to ensure the accuracy of our content to provide with an indulging reading experience to our subscribers. Our readers are our mirrors, and we use their inputs to reflect and improve upon human errors, if any, that may have occurred during the publishing processes involved. To let us maintain the quality and help us reach out to any readers who might be having difficulties due to any unforeseen errors, please write to us at :

errata@bpbonline.com

Your support, suggestions and feedbacks are highly appreciated by the BPB Publications' Family.

Did you know that BPB offers eBook versions of every book published, with PDF and ePub files available? You can upgrade to the eBook version at www.bpbonline.com and as a print book customer, you are entitled to a discount on the eBook copy. Get in touch with us at :

business@bpbonline.com for more details.

At **www.bpbonline.com**, you can also read a collection of free technical articles, sign up for a range of free newsletters, and receive exclusive discounts and offers on BPB books and eBooks.

Piracy

If you come across any illegal copies of our works in any form on the internet, we would be grateful if you would provide us with the location address or website name. Please contact us at **business@bpbonline.com** with a link to the material.

If you are interested in becoming an author

If there is a topic that you have expertise in, and you are interested in either writing or contributing to a book, please visit **www.bpbonline.com**. We have worked with thousands of developers and tech professionals, just like you, to help them share their insights with the global tech community. You can make a general application, apply for a specific hot topic that we are recruiting an author for, or submit your own idea.

Reviews

Please leave a review. Once you have read and used this book, why not leave a review on the site that you purchased it from? Potential readers can then see and use your unbiased opinion to make purchase decisions. We at BPB can understand what you think about our products, and our authors can see your feedback on their book. Thank you!

For more information about BPB, please visit **www.bpbonline.com**.

Join our book's Discord space

Join the book's Discord Workspace for Latest updates, Offers, Tech happenings around the world, New Release and Sessions with the Authors:

<https://discord.bpbonline.com>



Table of Contents

1. Exploring the Essence of Chaos Engineering	1
Introduction.....	1
Structure.....	1
Objectives	2
Overview of Chaos Engineering	3
Benefits of Chaos Engineering.....	3
Challenges and considerations.....	4
History of Chaos Engineering	5
Importance of resilience in distributed system.....	6
Elevating system reliability.....	8
Chaos Engineering examples	10
Conclusion.....	12
Points to remember	12
2. Chaos Engineering Concepts.....	13
Introduction.....	13
Structure.....	13
Objectives	13
Chaos Engineering Principles	14
<i>Hypothesis-driven experiments.....</i>	<i>15</i>
<i>Application of hypothesis-driven experiments</i>	<i>16</i>
<i>Minimize Blast Radius.....</i>	<i>16</i>
<i>Application of minimize blast radius</i>	<i>16</i>
<i>Applying real-world events.....</i>	<i>17</i>
<i>Application of real-world events</i>	<i>17</i>
<i>Automate experiments</i>	<i>18</i>
<i>Application of automated experiments.....</i>	<i>18</i>
<i>Continuous iteration</i>	<i>19</i>
<i>Application of continuous iterations.....</i>	<i>19</i>
<i>Capture and learn from data</i>	<i>19</i>
Difference between Chaos Engineering and Testing.....	21
Chaos Engineering vs. Performance Testing.....	23

Conclusion.....	25
Points to remember	25
Multiple choice questions	26
<i>Answer key</i>	28
3. Revision with Go	29
Introduction.....	29
Structure.....	29
Objectives	30
Overview of Chaos Engineering with Go.....	30
<i>Implementing Go in Chaos Engineering: Real-world example</i>	32
Introduction of Go programming language.....	32
Setting up Go development environment	33
<i>Go installation at Windows Environment</i>	33
<i>Go installation at Mac Environment</i>	36
Go fundamentals and syntax.....	37
<i>Hello to Go!!!</i>	38
<i>Variables and constants</i>	39
<i>Shorthand variable</i>	39
Control structure	40
<i>Conditional statements</i>	40
<i>Looping statements</i>	41
<i>Switch statements</i>	41
Data types in Go	42
<i>Primitive data types</i>	42
<i>Composite data types</i>	43
<i>Special data types</i>	44
<i>Zero values</i>	45
<i>Type conversion</i>	46
<i>Type inference</i>	46
<i>Arrays and slices</i>	50
<i>Structs in Go</i>	51
Pointers in Go.....	52
First class functions	53
Packages and dependencies.....	54
Concurrency and Goroutines	56

<i>Synchronization</i>	58
<i>Goroutines vs. Thread</i>	60
Error handling	60
<i>Function returning an error</i>	60
Conclusion.....	62
Points to remember	62
Multiple choice questions	62
<i>Answers</i>	64
4. Fault Tolerance and Resilience Patterns	65
Introduction.....	65
Structure.....	65
Objectives	66
Fault tolerance in distributed system	66
<i>Strategy of fault tolerance</i>	66
<i>Factors to consider in fault tolerance</i>	68
<i>Difference in high availability and fault tolerance</i>	69
Resilience patterns.....	70
<i>Revealing the definition of Resilience</i>	70
Rate limiter.....	71
<i>Bucket algorithms</i>	72
<i>Leaky bucket algorithm</i>	73
<i>Simple rate counters</i>	73
<i>Rate limiting in distributed system</i>	75
Bulkhead	78
<i>Implementation with Bulkhead Pattern</i>	79
Circuit Breaker	80
<i>Circuit Breaker implementation in Go</i>	81
Timeout.....	84
Retry.....	85
Fault tolerant system built in Go.....	87
<i>High level architecture of system</i>	87
<i>Workflow</i>	89
<i>Test strategy</i>	89
Best practices for distributed systems in Go	92
Conclusion.....	94

Multiple choice questions	94
<i>Answers</i>	96
5. Chaos Fault Injection Techniques.....	97
Introduction.....	97
Structure.....	97
Objectives	98
Overview of Fault injection.....	98
Killing processes.....	99
Resource exhaustion technique	101
<i>Setting up resource exhaustion tests</i>	102
<i>Utilizing disk space in resource consumption</i>	104
<i>Mitigation strategies for resource exhaustion faults</i>	106
Component failure: Dependency testing	107
<i>Risk assessment in Dependency Testing</i>	107
<i>Steps of risk assessment</i>	108
<i>Tools and frameworks</i>	109
<i>Fallback mechanisms: Ensuring resilience</i>	110
<i>Steps to perform Dependencies Failure</i>	111
<i>Techniques of dependency testing</i>	113
<i>Simulating database failures</i>	113
<i>Testing 3rd-Party API Reliability</i>	113
<i>Load Testing Message Queues</i>	114
<i>Testing Cloud Service Failures</i>	114
<i>Testing External Storage Failures</i>	117
<i>Benefits of dependency testing</i>	118
Adjusting load balancers.....	119
<i>Techniques of load balancer failures</i>	120
<i>Simulating Increased Traffic Load</i>	120
<i>Testing Failure of Backend Servers</i>	121
<i>Load Balancing Algorithm Changes</i>	122
<i>Benefits of load balancer adjusting</i>	123
<i>Further insights</i>	124
Error injection in databases and storage systems.....	124
<i>Types of Errors to Inject</i>	125
<i>Simulate database failures</i>	126

<i>Introducing data corruption in storage</i>	126
<i>Database node or server failures</i>	127
<i>Further insights</i>	128
Simulating network failures and latency issues	128
<i>Network Failure Scenarios</i>	131
<i>Further insights</i>	132
Conclusion.....	133
Multiple choice questions	133
<i>Answers</i>	134
6. Chaos Testing Tools	135
Introduction.....	135
Structure.....	135
Objectives	136
Overview of Chaos Testing tools	136
<i>Chaos Monkey</i>	136
<i>Limitations of Chaos Monkey</i>	137
<i>ToxiProxy</i>	138
<i>Benefits of using Toxiproxy</i>	139
<i>Limitations of Toxiproxy</i>	140
<i>Chaos Toolkit</i>	140
<i>Experiment design and implementation</i>	141
<i>Chaos Mesh</i>	143
<i>Components of Chaos Mesh</i>	144
<i>Features of Chaos Mesh</i>	144
Chaos Mesh: A deep exploration	146
<i>Installation of Chaos Mesh on Ubuntu</i>	146
<i>Common tips and tricks for troubleshooting</i>	150
Faults injected by Chaos Mesh	151
<i>Pod failure by Chaos Mesh</i>	151
<i>Objective of Pod failure experiment</i>	151
<i>Commencing the Pod failure experiment</i>	151
<i>Result analysis</i>	155
<i>DNS failure experiment by Chaos Mesh</i>	156
<i>Objective of DNS failure experiment</i>	157
<i>Initiating the DNS Failure Experiment</i>	157

<i>Result analysis</i>	159
<i>Stress experiment by Chaos Mesh</i>	160
<i>Objective of stress experiment by Chaos Mesh</i>	161
<i>Embarking on the stress experiment</i>	161
<i>Result analysis</i>	163
<i>Limitation of Chaos Mesh</i>	163
Gremlin: Chaos Engineering as a platform	163
<i>Types of faults</i>	164
<i>Features of Gremlin</i>	164
<i>Getting started with Gremlin</i>	165
<i>Pros and cons of Gremlin</i>	170
<i>Pros</i>	171
<i>Cons</i>	171
<i>Case study: Strengthening e-commerce platform with Gremlin</i>	171
<i>Comparison of Gremlin with other tools</i>	172
Choosing the right Chaos Engineering tools.....	173
Conclusion.....	173
Points to remember	174
Multiple choice questions	175
<i>Answers</i>	176
7. Chaos Experiment Design	177
Introduction.....	177
Structure.....	177
Objectives	178
Defining Chaos Engineering experiments.....	178
Identifying system weaknesses and failure scenarios	179
What do you break first.....	180
<i>Known Knowns</i>	181
<i>Known Unknowns</i>	181
<i>Unknown Knowns</i>	182
<i>Unknown Unknowns</i>	182
Designing hypotheses for chaos experiments.....	183
<i>Key considerations for hypothesis design</i>	183
<i>Example of hypothesis design</i>	184
Define the scope for chaos experiment	185

<i>Legal and compliance consideration</i>	186
<i>Sample experiment documentation</i>	187
<i>Best practices for scope documentation</i>	188
<i>Define scope in Chaos Mesh</i>	188
<i>Target Kubernetes resources</i>	189
<i>Duration and timing</i>	189
Recovery and roll-back procedure of Chaos Experiment	190
<i>Recovery procedures</i>	192
<i>Rollback procedures</i>	192
Identifying the metrics of chaos experiment.....	195
<i>Chaos experiment metrics</i>	195
<i>Tools for clarity</i>	197
<i>Choosing the right metrics: Align with your goals</i>	197
Executing controlled chaos experiments	198
Review the result of chaos experiment	199
Chaos Engineering in continuous integration and deployment	202
<i>Challenges and risks of Chaos Engineering in CI/CD</i>	203
Conclusion.....	204
Points to remember	204
Multiple choice questions	205
<i>Answers</i>	206
8. Chaos with Emerging Tech Stack.....	207
Introduction.....	207
Structure.....	207
Objectives	208
High Availability testing by Chaos Engineering	208
<i>High Availability</i>	208
<i>Benefits of High Availability</i>	209
<i>Techniques to Achieve High Availability</i>	209
<i>Server clustering</i>	209
<i>Load balancing</i>	210
<i>Data replication</i>	211
<i>Failover mechanisms</i>	211
<i>Geographic redundancy</i>	211
<i>High Availability Testing Strategy for Systems</i>	211

Initial verification.....	211
Failure scenarios.....	212
Scalability testing.....	212
Disaster Recovery integration.....	213
Testing of HA use cases in Chaos Engineering	213
Chaos Engineering tools for HA testing	214
Applying AI in Chaos Engineering	215
Machine Learning	216
Labeled or unlabeled dataset	216
Use of Machine Learning in Chaos Engineering	217
Test distributed system using ML and Chaos Mesh	218
Challenges in applying Machine Learning to Chaos Engineering	220
Future trends and opportunities	221
Azure Chaos Studio	222
How Chaos Studio works	223
Faults supported by Chaos Studio	223
Service-direct faults.....	223
Agent-based faults.....	224
AWS Fault Injection Services	224
Role of AWS FIS in Chaos Engineering	224
Supported fault types	225
Benefits of using AWS FIS.....	225
Conclusion.....	226
Points to remember	227
Multiple choice questions	227
Answers.....	228
9. Essence of Observability in Distributed System.....	229
Introduction.....	229
Structure.....	229
Objectives	230
Essence of observability	230
Components of observability.....	231
Metrics	232
Error rate metrics	232
Latency metrics.....	232

<i>Recovery time metrics</i>	232
Logs	233
<i>Limitation of logs</i>	233
Traces.....	233
Observability in distributed systems.....	234
Implementation of monitoring in distributed systems.....	236
<i>Approaches of application monitoring</i>	236
<i>Machine Learning and anomaly detection in distributed system monitoring</i> ..	237
<i>Distributed system monitoring with Dynatrace</i>	238
Logging implementation in a distributed system	241
<i>ELK Stack as Centralized Logging System</i>	241
<i>Best practices of logging implementation</i>	244
<i>Logging in Go Programming Language</i>	245
Tracing in distributed system	248
<i>Distributed tracing tools</i>	249
<i>Implementation of Open Telemetry</i>	249
<i>Instrument the Go application by Open Telemetry</i>	250
<i>Data sampling in Telemetry</i>	251
<i>Importance of data sampling</i>	252
<i>Configuring sampling in OpenTelemetry</i>	252
<i>Outcomes and benefits of distributed tracing</i>	253
Difference between distributed tracing and logging	255
Conclusion.....	255
Points to remember	255
Multiple choice questions	256
<i>Answers</i>	258
10. Observability in Chaos Engineering.....	259
Introduction.....	259
Structure.....	259
Objectives	260
Role of observability in Chaos.....	260
Observability tools for Chaos Engineering	262
<i>Elastic search and Kibana</i>	263
<i>Prometheus and Grafana</i>	263
<i>Dynatrace</i>	264

<i>Splunk</i>	264
<i>New Relic</i>	264
<i>Datadog</i>	265
<i>Challenges in integrating monitoring tools</i>	269
<i>Strategies for cohesive monitoring</i>	270
Metrics used in Chaos Experiment	271
<i>Resilience metrics</i>	271
<i>MTTR</i>	272
<i>MTTD</i>	272
<i>Recovery Time Objective</i>	273
<i>Availability</i>	274
<i>Auto scaling turnaround time</i>	274
<i>Service Level Objectives Compliance</i>	275
<i>Failure Injection Rate</i>	275
<i>Infrastructures metrics</i>	275
<i>Importance of granularity of infrastructure metrics</i>	276
<i>CPU utilization</i>	276
<i>Memory usage</i>	277
<i>Disk Input/Output</i>	278
<i>Network traffic</i>	279
<i>Latency</i>	279
<i>Security monitoring</i>	280
<i>Load balancer metrics</i>	281
<i>Container metrics</i>	281
<i>Application metrics</i>	281
<i>Service health</i>	282
<i>Response time</i>	282
<i>Error rate</i>	282
<i>Throughput</i>	282
<i>Active threads</i>	282
<i>Queue length</i>	282
<i>Best practices for collecting and storing metrics</i>	283
Logging and tracing of Chaos experiment	283
<i>Strategies for trace implementation</i>	284
Conclusion	285

Points to remember	286
Multiple choice questions	286
<i>Answers</i>	288
11. Security Chaos Engineering Overview	289
Introduction.....	289
Structure.....	289
Objectives	290
Fundamentals of Security Chaos Engineering	290
Types of attacks for Security Chaos Engineering	291
<i>Denial-of-Service</i>	291
<i>Injection attacks</i>	292
<i>Privilege escalation</i>	293
<i>Brief overview of Metasploit Framework</i>	294
<i>Data breaches</i>	294
Design experiment of SCE	295
Best practices for Security Chaos Engineering	297
Case study: Security Chaos Engineering	299
Conclusion.....	301
Points to remember	301
Multiple choice questions	301
<i>Answers</i>	302
12. Case Studies: Chaos Engineering in Action	303
Introduction.....	303
Structure.....	303
Objectives	304
LinkedOut: Failure Injection Framework	304
<i>Overview of Case Study</i>	305
<i>Objectives of LinkedOut</i>	305
<i>LinkedOut overview</i>	305
<i>Triggering mechanisms</i>	305
<i>Web application</i>	306
<i>Impact on services and underlying technologies</i>	306
<i>Enhancing service resilience</i>	306
<i>Detecting bottlenecks</i>	307

<i>Validating resilience strategies</i>	307
<i>Technology integration</i>	307
<i>Observability and monitoring</i>	307
<i>Lessons learned</i>	308
<i>Future considerations</i>	310
The Walmart Global Tech case study.....	310
<i>Addressing challenges: Training and tooling</i>	311
<i>DocRX</i>	311
<i>Impact on services and underlying technologies</i>	312
<i>Unified monitoring across hybrid clouds</i>	312
<i>Enhanced resiliency through pre-checks</i>	312
<i>Real-time diagnostics and automated analysis</i>	313
<i>Proactive monitoring with Slack-bot integration</i>	313
<i>Resiliency levels and prerequisites</i>	314
<i>Key takeaways and future perspectives</i>	314
Capital One's journey in Chaos Engineering.....	315
<i>Chaos Engineering at Capital One</i>	315
<i>Testing in production</i>	316
<i>Addressing failure at scale</i>	316
<i>Lower latency and higher capacity</i>	316
<i>Risk mitigation and monitoring</i>	317
<i>Lessons learned</i>	317
<i>Mind-set matters</i>	317
<i>Causes of failures require widespread coverage</i>	317
<i>Observability is a pre-requisite</i>	317
<i>Integrating Chaos Engineering into DevOps practices</i>	318
Conclusion.....	319
Points to remember.....	319
Multiple choice questions.....	320
<i>Answers</i>	322
13. Best Practices and Lessons Learned	323
Introduction.....	323
Structure.....	323
Objectives.....	324
Common mistakes in Chaos Engineering.....	324

<i>Lack of clear objectives</i>	324
<i>Inadequate planning and communication</i>	325
<i>Overlooking safety measures</i>	325
<i>Excessive disruption</i>	325
<i>Ignoring customer impact</i>	326
<i>Insufficient monitoring and observability</i>	326
<i>Inconsistent execution</i>	327
<i>Not learning from failures</i>	327
<i>Isolating Chaos Engineering</i>	328
<i>Not evolving with the system</i>	328
Incorporating Chaos Engineering into the SDLC.....	329
<i>Advantages of Chaos Engineering into the SDLC</i>	330
Challenges and limitations of Chaos Engineering.....	331
Best practices for Chaos Engineering.....	332
Conclusion.....	334
Points to remember.....	335
Multiple choice questions.....	335
<i>Answers</i>	336
Index	337-348

CHAPTER 1

Exploring the Essence of Chaos Engineering

Introduction

The first chapter of this book will cover groundwork for those who are unfamiliar with the topic of Chaos Engineering. Using relevant examples, it offers a crucial comprehension of the underlying concepts, evolution, and advantages.

This chapter outlines the fundamental concepts underlying Chaos Engineering, such as definition, aims, and benefits, to ensure that the readers comprehend its purpose. It explores the mentality and cultural elements required for an organization to successfully adopt Chaos Engineering.

Structure

In this chapter, we will discuss the following topics:

- Overview of Chaos Engineering
- Benefits of Chaos Engineering
- Challenges and considerations
- History of Chaos Engineering
- Importance of resilience in distributed system

- Elevating system reliability
- Chaos Engineering examples

Objectives

The objective of the chapter titled *Exploring the Essence of Chaos Engineering* is to provide readers with an in-depth understanding of the fundamental principles, concepts, and practices that constitute Chaos Engineering. It aims to introduce readers to the core elements involved in Chaos Engineering and shed light on its significance in improving system resilience and reliability. By the end of the chapter, readers should have a clear grasp of what Chaos Engineering entails and how it can be applied to enhance system robustness and user experience.

In the industrial revolution of digitization, the world has undergone a rapid transformation, heavily reliant on technology in almost every aspect of our lives. With the advent of the internet, our global community has become more interconnected than ever before, driving industries towards the online platform. However, this technological shift is still relatively new, and both emerging and developed economies are continuously striving to optimize the infrastructure and ecosystem required to support these online businesses. As a consequence, uncertainties arise, leading to a higher prevalence of failures.

We frequently encounter headlines reporting challenges in accessing mobile and online banking services, instances of bank websites being unavailable or not functioning correctly, and occurrences of *Service Unavailable* messages. Such unpredictability has become a regular occurrence in today's fast-paced digital landscape.

These outages and failures predominantly occur in complex and distributed systems, where multiple components may fail simultaneously, increasing the impact of the problem. Identifying and rectifying these bugs can take anywhere from minutes to hours, depending on the complexities of the system architecture. Consequently, these incidents not only result in revenue losses for the affected companies, but also lose the customer trust as well.

Although these systems are designed to handle individual failures, the complexities of large and chaotic systems pose the risk of cascading failures, leading to severe and extended outages. The challenge lies in enhancing the resilience and reliability of these complex systems to withstand the ever-increasing demands of the digital era.

Standard application testing is not capable of identifying errors in a distributed system. Companies must find more inventive ways to regularly test micro services. Chaos Engineering is a tactic that is gaining notice.

Overview of Chaos Engineering

Chaos Engineering is a specialized field within software engineering that strives to achieve a specific goal to improve system resilience and reliability by intentionally inducing controlled failures in a system. It emerged as a response to the growing complexity of modern distributed systems, where even minor disruptions can have catastrophic consequences.

Understanding the concepts behind Chaos Engineering allows readers to develop a systematic approach to designing and conducting experiments that simulate real-world failures and stress scenarios. This proactive approach is used to identify bottlenecks, performance issues, and scalability limitations; enabling reader to fine-tune systems for optimal performance. It builds confidence in their system's ability to handle unexpected events, and ultimately create more robust and dependable applications.

Chaos Engineering is not only a technical practice; it requires a mind-set shift and a culture of embracing failure as a means to improve system resilience. Chaos Engineering extends its influence across diverse roles and domains, encompassing SRE teams, DevOps engineers, system administrators, and enterprise architects. It involves cross-functional collaboration and effective communication between teams. This promotes collaboration, facilitates knowledge sharing, and aligns the team towards a common goal of building resilient systems.

Chaos Engineering has become an essential practice in modern software engineering to ensure system reliability and resilience. By embracing controlled chaos and learning from failures, organizations can build robust, fault-tolerant systems that provide a more sensible experience for users and withstand the challenges of today's complex computing landscape.

Benefits of Chaos Engineering

Chaos Engineering offers a range of benefits that contribute to the overall reliability, resilience, and customer experience of software systems. Here are some key advantages of adopting Chaos Engineering practices:

- **Early fault detection:** Chaos Engineering enables early detection of potential weaknesses and vulnerabilities within a system, reducing the likelihood of major outages or service disruptions in production.
- **Enhanced resilience:** By subjecting a system to controlled failures, Chaos Engineering helps developers build a more resilient architecture, capable of self-healing from unexpected events.
- **Improved incident response:** Chaos experiments serve as intense training for our incident response team, refining troubleshooting, communication, and recovery skills in a controlled setting. This results in faster, more efficient reactions to

real-world problems, refining important skills like as root cause analysis and collaboration. Measuring response metrics allows you to identify areas for improvement and continuously optimizes your incident response plan to adapt to changing systems.

- **Enhancing trust and better customer service:** Chaos Engineering, by intentionally inducing failures in systems, helps build resilience and trust in services. Chaos experiments expose weaknesses before they become critical issues. This translates to fewer unexpected outages and smoother customer experiences, fostering trust and loyalty. When outages occur, practiced incident response teams can quickly diagnose and resolve issues, minimizing downtime and impact on customers. Proactively communicating potential weaknesses and planned chaos experiments demonstrates transparency and commitment to reliability. This builds trust with customers and fosters a sense of partnership. By minimizing outages, resolving issues quickly, and communicating openly, Chaos Engineering ultimately leads to increased customer satisfaction and loyalty. For instance, Spotify, a music streaming service, uses Chaos Engineering to ensure high availability. This translates to reliable music access for their users, contributing to their high customer satisfaction scores.
- **Cost savings:** Chaos Engineering enables companies to uncover weaknesses in infrastructure, optimize cloud costs, reduce mean time to recovery, prevent data loss, and streamline development processes. For instance, by simulating failures, organizations can fine-tune systems, avoiding costly downtime and ensuring efficient resource allocation, leading to significant financial savings.

By prioritizing the above benefits and leveraging them effectively, companies can build stronger relationships with customers, protect their brand reputation, mitigate risks, drive efficiency and innovation, and attract key talent and partners, all of which contribute to long-term success and profitability.

Challenges and considerations

Implementing Chaos Engineering, while highly beneficial, comes with its own set of challenges and considerations. It is important to be aware of these factors to ensure successful integration and effective practice within an organization. Here are some challenges and considerations to keep in mind:

- **Safety first:** In Chaos Engineering, challenges are intentionally introduced during testing, and occasionally, these challenges might be unnecessary. While Chaos Engineering is beneficial, it should be carried out with caution. Safety measures must be in place to prevent experiments from causing severe disruptions to critical systems.
- **Start small:** Organizations that are new to Chaos Engineering should begin with simple experiments and gradually increase complexity as their understanding and expertise grow.

- **Communication and collaboration:** Chaos Engineering requires close collaboration between developers, operations, and other stakeholders to ensure that all parties are aware of the experiments and potential risks involved.
- **Observability:** One of the most common issues engineers run across while implementing chaotic engineering is their inability to keep track of the observation. When there are gaps in monitoring, the team finds it difficult to understand the entire effects of the issue on the system. The teams are unable to locate the problem's root cause, which does not solve the issue and makes it complicated.
- **Impact analysis:** Experimenters should carefully assess the potential impact of each experiment to avoid unintended consequences and maintain system stability.

History of Chaos Engineering

Beginning in the early 2000s, Chaos Engineering has a long and interconnected history that is closely associated with the emergence of complex distributed systems and the need for improved reliability in technology infrastructure. The concept was pioneered by engineers at Netflix, who faced the challenges of running a vast and rapidly growing online streaming service that required high availability and fault tolerance.

The following key milestones highlight the evolution of Chaos Engineering:

- **Netflix's start:** In the mid-2000s, Netflix started transitioning from a DVD rental service to an online streaming platform. As the scale of their service expanded, they encountered various technical challenges, including outages and system failures. The traditional approaches of testing and monitoring were not enough to ensure the robustness of their distributed systems.

In 2011, Netflix introduced the Simian Army, a suite of tools designed to create controlled disruptions within their production environment. One of the most notable components was Chaos Monkey, which randomly terminated virtual machine instances. This forced engineering teams to build applications and infrastructure that could withstand such failures.

The rationale behind *Chaos Monkey*, according to former VP of Product Engineering at Netflix *John Ciancutti*, is that,

If we aren't constantly testing our ability to succeed despite failure, then it isn't likely to work when it matters most—in the event of an unexpected outage.

- **Wider recognition:** Chaos Engineering gained wider recognition in the software engineering community as Netflix shared its experience and success stories at technology conferences and in blog posts. Other companies became interested in applying similar principles to their own systems.