

High Performance SRE

*Automation, error budgeting, RPAs, SLOs, and
SLAs with site reliability engineering*

Anchal Arora Mishra



www.bpbonline.com

First Edition 2024

Copyright © BPB Publications, India


ISBN: 978-93-55516-718

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.



Dedicated to

My husband

Ashutosh Mishra

whose belief in my potential has always made me aim higher

About the Author



Anchal Arora Mishra brings an extensive amount of experience as a Site Reliability Engineer from Walmart Global Tech to the creative sphere of technology. Anchal is not only adept in maintaining system reliability but also possesses a robust background in both the theoretical and practical aspects of Cloud Computing, DevOps, and SRE.

Anchal holds a post-graduate degree in Computer Applications from the prestigious VIT University in Vellore, ensuring a solid academic foundation behind the hands-on expertise. A lifelong learner and advocate for continuous improvement, Anchal has acquired key industry certifications, including CCNA and AWS, and is an active member of Toastmasters International, reflecting a commitment to technical and communication excellence. An active contributor to the SRE community, Anchal has also taken the stage at VIT University, sharing insights on DevOps in the Mobile World, demonstrating her dedication to nurturing the next generation of tech professionals.

About the Reviewer



Naresh Kumar Miryala, a distinguished engineering leader at Meta, possesses an extensive background in cloud and platform engineering honed over nearly two decades in the field. His deep understanding of both technical and business intricacies empowers him to provide innovative solutions spanning diverse domains such as database systems, large-scale backend infrastructure, multi-cloud environments (AWS/GCP/OCI/Azure), automation, cloud infrastructure, DevOps, Kubernetes, and Elasticsearch.

Having previously contributed to esteemed organizations like Oracle Corp and Computer Sciences Corporations, Naresh played a pivotal role in migrating or implementing Oracle technologies for over 50 organizations globally, many of which are Fortune 500 entities. His impact spans across various industries, including pharmaceuticals, retail, banking, and gold mining companies worldwide.

Naresh is highly experienced in cloud migrations, particularly involving databases (Oracle/Exadata/MySQL/Postgres) and applications (EBS/Fusion/EPM/GTM). He played a pivotal role in ensuring their seamless execution for large and complex deployments globally.

Naresh's affiliations include membership in IEEE, AIM leadership council, and fellowship at RSA. He holds certifications as a professional in Multi Cloud and Data platforms, and actively engages as a blogger, tech reviewer, and frequent speaker in international conferences.

Acknowledgement

In creating this book on **site reliability engineering (SRE)**, I have been fortunate enough to receive the support, guidance, and inspiration from various individuals and organizations whose contributions have been invaluable.

Firstly, I extend my deepest gratitude to the pioneers of the SRE field, whose innovative work and foundational principles have shaped this domain. Their insights and experiences have been a guiding light in the exploration of SRE concepts. Special mention goes to Google SRE work here, as Google was the first mover in this space.

I owe my gratitude to the seniors at my workspace whose mentorship and expertise have been instrumental in shaping my understanding and approach. My appreciation also goes to the SRE teams at Walmart, Amazon, Google, and Netflix, whose real-world experiences and challenges have provided critical case studies and practical examples that enrich this book. The willingness to share their journeys has been invaluable in drafting the work of this book.

To the editorial and production teams at BPB, thank you for your unwavering support and guidance throughout the publishing process. Your professionalism and expertise have been pivotal in bringing this book to fruition.

My heartfelt thanks to my husband and family for their endless encouragement and understanding, especially when writing consumed much of my attention. Your support has been my pillar of strength.

Thank you all for making this journey enriching and possible.

Preface

Welcome to this exploration of **site reliability engineering (SRE)**, a discipline that resides at the crossroads of software engineering and systems operations. This book is designed to guide you through the principles, practices, and philosophies that define SRE, a field pioneered by Google and now embraced by organizations around the world.

The genesis of this book lies in my journey as a software engineer and systems administrator. Like many in this field, I was intrigued by the challenges of maintaining the reliability of increasingly complex systems. The advent of SRE provided a framework that transformed how we approach these challenges, blending traditional IT operations with the innovation and speed of software engineering.

In writing this book, my goal is to demystify SRE for a broad audience. Whether you are an experienced practitioner, an aspiring SRE, or simply curious about the field, this book aims to provide valuable insights and practical guidance. It covers a range of topics from the foundational principles of SRE, such as **service level objectives (SLOs)** and error budgets, to advanced practices like chaos engineering and incident management.

This book also delves into the cultural and organizational aspects of SRE. Implementing SRE is not just about adopting new tools and practices; it is about fostering a culture that values reliability, accountability, and continuous improvement. To this end, I have included case studies and real-world examples illustrating how various organizations have successfully integrated SRE principles into their operations.

I am grateful to the many SRE professionals who have shared their knowledge and experiences, contributing to the rich tapestry of insights presented in this book. Their practical advice and real-world examples have been invaluable in illustrating the application of SRE in diverse environments.

As the field of SRE continues to evolve, this book aims to be a comprehensive guide and starting point for further exploration. The world of technology is ever-changing, and the practices of SRE will continue to adapt to new challenges and opportunities.

I invite you to join me on this journey through the dynamic and exciting world of SRE. Whether you are looking to implement SRE practices in your organization or simply seeking to understand this rapidly growing field, I hope this book provides you with the knowledge and inspiration you need.

Chapter 1: Introduction to Site Reliability Engineer - This chapter introduces SRE, explaining its origins at Google, its importance in modern software engineering, and how it bridges the gap between software development and operations.

Chapter 2: DevOps to Site Reliability Engineering - This chapter delves into the principles and practices of designing inherently reliable systems, including discussions on redundancy, fault tolerance, and the balance between reliability and cost.

Chapter 3: Monitoring - This chapter covers the strategies for monitoring complex distributed systems, emphasizing the selection of key performance indicators and tools to maintain visibility over system health.

Chapter 4: Incident Management and Risk Mitigation - This chapter discusses the frameworks and protocols for responding to system incidents, managing risks, and minimizing user impact, highlighting the importance of preparation and training.

Chapter 5: Error Budgets - This chapter introduces the concept of error budgets, a quantitative measure that balances the pace of innovation with the need for system stability and reliability.

Chapter 6: SLI/SLO/SLA - Here, the book breaks down **service level indicators (SLIs)**, **service level objectives (SLOs)**, and **service level agreements (SLAs)**, explaining how they guide reliability work and business decisions.

Chapter 7: Capacity Planning - The focus of this chapter is on predicting future system load and ensuring that infrastructure can handle growth in demand without sacrificing performance or reliability.

Chapter 8: On-call and First-response - This chapter discusses the human elements of SRE, including the organization of on-call rotations, the responsibilities of first responders, and best practices for incident response.

Chapter 9: RCA and Post-mortem - Root cause analysis (RCA) and post-mortem culture are explored in this chapter, detailing how teams can learn from failures to prevent future incidents and improve system reliability.

Chapter 10: Chaos Engineering - The chapter describes chaos engineering practices, where systems are deliberately subjected to stress in controlled environments to uncover weaknesses and improve resilience.

Chapter 11: Artificial Intelligence for Site Reliability Engineering - This chapter examines the emerging role of **artificial intelligence (AI)** in **site reliability engineering (SRE)**, including how AI can automate incident response, anomaly detection, and predictive analysis.

Chapter 12: Case Studies - Real-world case studies provide insights into how various organizations have implemented SRE principles, the challenges they faced, and the outcomes of their efforts.

Coloured Images

Please follow the link to download the
Coloured Images of the book:

<https://rebrand.ly/3b107k9>

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. Introduction to Site Reliability Engineer	1
Introduction.....	1
Structure.....	2
Objectives	2
Historical context and origin of the SRE role	3
Type of DevOps teams in different companies	3
Roles and responsibilities of SRE	4
Bridging the gap between development and operations	5
<i>Maintaining system and service reliability</i>	5
Importance of SRE in the modern tech ecosystem	5
Skills and knowledge for SRE	6
<i>Necessary technical skills</i>	6
<i>Soft skill requirements</i>	6
Culture of SREs and DevOps.....	7
<i>Understanding DevOps</i>	7
<i>SRE's role in promoting the DevOps culture</i>	7
<i>Effect on the process of making and delivering software</i>	8
Importance of SRE in the digital age	8
<i>Effect of service downtime on businesses</i>	8
<i>SRE's role in reducing and preventing downtime</i>	8
<i>Prospects and developments for SREs in the future</i>	9
Career path and professional development	9
<i>Starting point and prerequisites for becoming an SRE</i>	10
<i>Continuous learning and upskilling</i>	10
<i>Career progression for SREs</i>	10
<i>Evolving SRE role</i>	11
Conclusion.....	11
Multiple choice questions	12
<i>Answers</i>	12

2. DevOps to Site Reliability Engineering.....	13
Introduction.....	13
Structure.....	14
Objectives	14
DevOps to site reliability engineering.....	14
Need for site reliability engineering.....	15
<i>Site reliability engineering team structure</i>	<i>15</i>
<i>Site reliability engineering discipline.....</i>	<i>18</i>
<i>Unspoken commitments.....</i>	<i>18</i>
Site reliability engineering engagement model	20
<i>Site reliability engineering implements DevOps</i>	<i>20</i>
Site reliability engineering strategy adoption.....	23
Site reliability engineering challenges.....	25
Site reliability engineering best practices.....	27
Site reliability engineering best practices tools.....	27
Conclusion.....	30
Multiple choice questions	30
<i>Answers.....</i>	<i>31</i>
3. Monitoring.....	33
Introduction.....	33
Structure.....	33
Objectives	34
Need for monitoring	34
Pillars of monitoring	38
<i>Latency</i>	<i>39</i>
<i>Errors.....</i>	<i>39</i>
<i>Saturation.....</i>	<i>40</i>
Threshold monitoring.....	40
Monitoring and observability.....	41
Application monitoring.....	46
Monitoring best practices.....	47
<i>Examples of monitoring and observability tools</i>	<i>47</i>

Conclusion.....	50
Multiple choice questions	50
<i>Answers</i>	51
4. Incident Management and Risk Mitigation	53
Introduction.....	53
Structure.....	54
Objectives	54
Purpose of incident management	54
<i>More about software risks</i>	55
Incident prioritization.....	55
Incident severity level.....	56
<i>Use of severity level</i>	56
<i>Difference between severity and priority</i>	56
<i>Defining incident severity levels</i>	57
Incident response planning.....	58
Risks to consider.....	58
Analyzing the risks	59
<i>Production incident lifecycle</i>	60
<i>Cost of reliability</i>	60
<i>Response plan</i>	61
Best practices to reduce production incidents.....	62
Risk and mitigation.....	63
Best practices for risk mitigation.....	64
Conclusion.....	65
Multiple choice questions	66
<i>Answers</i>	67
5. Error Budgets	69
Introduction.....	69
Structure.....	70
Objectives	70
Purpose of error budgets	71

Defining error budgets	71
<i>Error budget equation</i>	73
<i>Prioritizing development over end-user experience</i>	74
Relation of error budgets with SLI and SLO	75
Benefits to setting the proper error budgets	76
Outage policies	77
Action items if the error budget is exceeded	77
Best practices to get the correct error budgets	78
Conclusion	79
Multiple choice questions	79
<i>Answers</i>	80
6. SLI/SLO/SLA	81
Introduction	81
Structure	81
Objectives	82
Introduction to service level management	83
<i>Overview of service level management</i>	83
<i>Key components of SLM: SLI, SLO, and SLA</i>	83
<i>Benefits of implementing an SLM program</i>	84
Understanding service level indicators	84
<i>Purpose of SLIs</i>	85
<i>Types of SLIs and their use cases</i>	85
<i>Key features of selecting appropriate SLI</i>	86
<i>Importance of SLIs</i>	87
Setting service level objectives	87
<i>Purpose of SLOs</i>	87
<i>Setting up appropriate SLOs</i>	87
Creating service level agreements	89
<i>Purpose of SLAs</i>	89
<i>Components of SLA</i>	90
<i>Negotiations of SLA</i>	91
Implementing and managing the SLM program	91

<i>Steps for implementing the SLM program</i>	92
<i>Best practices for managing SLIs, SLOs, and SLAs</i>	92
<i>Common challenges in setting up correct SLA</i>	93
<i>Role of technology in automating SLM</i>	93
Case studies and real-world examples.....	95
<i>Netflix</i>	95
<i>Adobe</i>	95
<i>LinkedIn</i>	96
Conclusion.....	96
Multiple choice questions	97
<i>Answers</i>	98
7. Capacity Planning	99
Introduction.....	99
Structure.....	100
Objectives	101
Importance of capacity planning	101
<i>Principles of capacity management</i>	101
Understanding resource requirements.....	102
<i>Identifying key resources</i>	102
<i>Analyzing historical usage data</i>	103
<i>Forecasting future usage patterns</i>	103
Capacity analysis.....	104
<i>Capacity analysis to determine workload resources</i>	105
<i>Trade-offs between performance, availability, and cost</i>	105
Scaling strategies	106
<i>Choosing the right scaling strategy</i>	107
<i>Considerations for auto-scaling and load balancing</i>	108
Monitoring and alerting	109
<i>Setting up monitoring tools</i>	109
<i>Defining alerting thresholds for key metrics</i>	110
<i>Strategies for proactive capacity planning</i>	111
Capacity planning in the cloud	113

<i>Understanding cloud resource allocation</i>	113
<i>Leveraging cloud provider tools</i>	114
Capacity planning for disaster recovery	115
<i>Disaster recovery capacity needs</i>	115
<i>Developing disaster recovery capacity plans</i>	116
<i>Disaster recovery plans and capacity</i>	117
Conclusion.....	118
Multiple choice questions	118
<i>Answers</i>	119
8. On-call and First-response	121
Introduction.....	121
Structure.....	121
Objectives	122
Understanding on-call	123
<i>Types of on-call rotations</i>	123
<i>Key responsibilities of on-call engineers</i>	123
First response processes	124
<i>Common steps in first response processes</i>	124
<i>Best practices for first response</i>	125
<i>Preparing for on-call and first-response</i>	125
<i>Importance of proactive preparation</i>	125
<i>Key tools and resources for on-call engineers</i>	126
<i>Strategies for reducing stress and avoiding burnout</i>	126
Communicating during incidents.....	127
<i>Importance of effective communication</i>	127
<i>Best practices for communicating with stakeholders</i>	127
<i>Tools for effective incident communication</i>	128
Incident review and post-mortems.....	128
<i>Incidents and post-mortems</i>	128
<i>Common post-mortem processes and best practices</i>	129
<i>Preventing incidents with post-mortems</i>	129
Case studies.....	130

Google.....	130
Amazon	131
Atlassian.....	132
Netflix.....	133
Conclusion.....	134
Multiple choice questions	135
Answers.....	136
9. RCA and Post-mortem.....	137
Introduction.....	137
Structure.....	138
Objectives	139
Root cause analysis	139
<i>Understanding the RCA process.....</i>	<i>139</i>
<i>Problem identification</i>	<i>140</i>
<i>Data collection</i>	<i>140</i>
<i>Root cause identification</i>	<i>140</i>
<i>Implementing solutions</i>	<i>140</i>
<i>Reviewing the efficiency of the solutions</i>	<i>140</i>
<i>Various methods of RCA.....</i>	<i>141</i>
<i>The five whys.....</i>	<i>141</i>
<i>Fishbone/Ishikawa diagrams</i>	<i>141</i>
<i>Fault tree analysis.....</i>	<i>141</i>
<i>Role of RCA in problem-solving and actions</i>	<i>142</i>
Post-mortem.....	142
<i>How to conduct a post-mortem</i>	<i>142</i>
<i>Gathering data and information</i>	<i>142</i>
<i>Analyzing the incident.....</i>	<i>142</i>
<i>Identifying actions for improvement.....</i>	<i>143</i>
<i>Implementing changes</i>	<i>143</i>
<i>Role of a blameless post-mortem.....</i>	<i>143</i>
<i>Role of post-mortem in learning and improvement</i>	<i>143</i>
<i>Real-world examples of effective post-mortems.....</i>	<i>143</i>

<i>Challenges and pitfalls in conducting post-mortems</i>	143
Relationship between RCA and post-mortem	144
<i>RCA feeds into the post-mortem process</i>	144
<i>RCA and post-mortem: Synergies and differences</i>	144
<i>Optimizing incident management</i>	144
Future trends.....	145
<i>Applying AI and ML to RCA and post mortem</i>	145
<i>Post-mortem best practices</i>	146
Conclusion.....	147
Multiple choice questions	147
<i>Answers</i>	148
10. Chaos Engineering	149
Introduction.....	149
Structure.....	150
Objectives	151
Principles of chaos engineering.....	151
<i>Building a hypothesis</i>	152
<i>Introducing real-world events</i>	152
<i>Observing the system</i>	152
<i>Verifying the hypothesis</i>	152
<i>Incremental complexity</i>	152
Role of chaos engineering in SRE.....	153
Key concepts in chaos engineering.....	154
<i>Blast radius</i>	154
<i>Failure injection</i>	154
<i>Steady-state</i>	154
<i>Observability and monitoring</i>	154
<i>Chaos experiments</i>	155
<i>Game days</i>	155
Preparing for chaos engineering	155
<i>Setting objectives and metrics</i>	155
<i>Building an observability infrastructure</i>	156

<i>Establishing a strong incident response strategy</i>	157
Implementing chaos testing.....	157
Tools and technologies for chaos engineering	158
<i>Chaos toolkit</i>	159
<i>Gremlin</i>	159
<i>Chaos Monkey</i>	159
Case studies on chaos engineering	160
<i>Netflix</i>	160
<i>Amazon</i>	161
<i>Google</i>	162
Future of chaos engineering	164
Conclusion.....	164
Multiple choice questions	165
<i>Answers</i>	166
11. Artificial Intelligence for Site Reliability Engineering	167
Introduction.....	167
Structure.....	168
Objectives	169
Role of AI in transforming SRE processes	169
Automated testing and quality assurance.....	170
<i>Role of AI in test case generation and automation</i>	170
<i>Role of AI in testing</i>	171
Intelligent debugging.....	171
<i>AI techniques for code analysis and issue identification</i>	172
<i>Real-time insights and suggestions for issue resolution</i>	172
<i>Impact of intelligent debugging on system stability</i>	172
Predictive maintenance	172
<i>AI for maintenance and upgrades</i>	173
<i>Predicting potential failures and resource depletion</i>	173
<i>Predictive maintenance and resource optimization</i>	173
Code generation and augmentation.....	174
<i>Code snippets and faster development</i>	174

<i>AI-assisted code review for improved code quality</i>	174
<i>Enhanced development and coding practices</i>	174
Performance optimization.....	174
<i>Monitoring and analysis</i>	175
<i>Bottleneck detection and root cause analysis</i>	175
<i>Automated performance tuning</i>	175
<i>Predictive and adaptive scaling</i>	175
<i>User experience optimization</i>	176
Anomaly detection and security.....	176
<i>AI for anomaly detection</i>	176
<i>Leveraging AI to prevent security threats</i>	177
<i>Enhancing system security and maintaining data</i>	177
Continuous integration and deployment.....	178
<i>Automation of CI/CD processes using AI</i>	178
<i>AI-driven code analysis and release management</i>	178
<i>Software delivery and development</i>	179
Natural language processing for SRE.....	179
<i>Role of NLP in processing requirements</i>	179
<i>Tools for requirement analysis</i>	180
<i>Sentiment analysis and user feedback</i>	180
Future trends and challenges.....	181
<i>Potential challenges and ethical considerations</i>	181
<i>Future of AI in SRE</i>	181
Conclusion.....	182
<i>Multiple choice questions</i>	182
<i>Answers</i>	183
12. Case Studies	185
Introduction.....	185
Structure.....	185
Objectives.....	186
Google.....	186
<i>Background and difficulties</i>	186

Google's software reliability engineering model	186
Netflix.....	188
Background and difficulties	188
Netflix's software reliability engineering methodology	189
Core ideas of Netflix's SRE strategy	189
Spotify	191
Background and difficulties	191
Spotify's software reliability engineering approach	191
LinkedIn.....	193
Background and difficulties	193
Journey of LinkedIn's software reliability engineering	193
Amazon.....	195
Background and challenges	196
SRE at Amazon.....	196
Conclusion.....	198
Index	199-207

CHAPTER 1

Introduction to Site Reliability Engineer

Introduction

As the digital world evolves and the demand for uninterrupted and seamless service delivery increases, the need for specialized roles to ensure this consistency becomes paramount. One such role is that of the **site reliability engineer (SRE)**. It was first introduced by *Google* to bridge the gap between the development and operations teams; the role of SRE has revolutionized how organizations handle their digital infrastructure and service reliability.

The SRE role embodies the principles of **Infrastructure as Code (IaC)**, seeking to apply software engineering methods to operations problems. The role is pivotal, focusing on developing highly scalable and robust software systems and ensuring their resilience under varying levels of user demand.

As we delve into this chapter, we aim to elucidate the nuances of the SRE role, which involves understanding their objectives, responsibilities, and day-to-day activities. This role is essential to facilitate proactive and reactive problem-solving to optimize system performance and uptime. SRE does not only manage incident response and system troubleshooting but also works toward the proactive prevention of such incidents. This is achieved by constantly analyzing system trends and identifying areas of potential failure. Therefore, an SRE serves as the key guardian of system stability, focusing on reducing organizational chaos and promoting a culture of stability, reliability, and efficiency.

Moreover, SREs foster collaboration between different teams in an organization, eliminating the conventional silos between development and operations teams. By encouraging shared responsibility for the reliability and quality of the services, they help to drive the DevOps culture in organizations, ensuring faster and more stable delivery of features to users.

In addition, we will touch upon the necessary skills and knowledge for a successful career in SRE. They need a strong foundation in computer science principles and programming, coupled with a deep understanding of system design and architecture. Furthermore, they must possess soft skills, including problem-solving, communication, and collaboration, to liaise effectively with various teams and drive improvements.

Understanding the SRE's role is pivotal in today's technological landscape. In an era where companies run on digital platforms, and downtime can result in significant financial and reputational losses, the role of an SRE is crucial. So, sit back and dive into this chapter to unravel the complexities and understand the impact of the SRE's role in creating reliable and resilient digital infrastructures.

Structure

In this chapter, we will cover the following topics:

- Historical context and origin of the SRE role
- Type of DevOps teams in different companies
- Roles and responsibilities of an SRE
- Importance of SRE in the modern tech ecosystem
- Skills and knowledge for SRE
- Culture of SRE and DevOps
- Importance of SRE in the digital age
- Career path and professional development

Objectives

A major goal of this chapter is to familiarize readers with the SRE function, an integral part of the current technological ecology. We aim to know everything about the SRE position, from how it was created to its effect on an organization's bottom line, and to make SRE less mysterious by supplying straightforward terminology and practical examples. Their primary responsibilities include system architecture, implementation, problem-solving, incident management, and team coordination, which will be covered in detail for the reader. By breaking down these responsibilities, we hope to give the reader a comprehensive picture of an SRE's typical workday. In addition, this section of the book

explains why SREs are significant for the success of the DevOps culture by bridging the gap between the development and operations teams. We will do our best to illustrate how this synergy shortens the time to create stable and reliable software.

This chapter also hopes to clarify the credentials and experience an aspiring SRE should have. We intend to serve as a guide for everybody interested in learning more about it, whether they are aspiring SREs, professionals who deal with SREs, or business leaders who want to learn more about the value SREs provide to their companies. We hope that by the end of this chapter, readers will agree that the function of the SRE in ensuring service dependability and system resilience is more important than ever before in today's era of digital changes and online services.

Historical context and origin of the SRE role

The role of the SRE was first conceived at Google in the early 2000s when the company faced challenges with maintaining its large-scale sites. To manage these complexities, they created a new kind of role that merged the skills of a software engineer and a systems engineer. The main task of this new role was to ensure that Google's services were highly available, efficient, and scalable. Over the years, many other companies have adopted principles and practices developed by Google's SREs, establishing it as a vital discipline in the tech industry.

Type of DevOps teams in different companies

When it comes to DevOps, most organizations operate differently. Many refer to those using **continuous integration** and **continuous delivery (CI/CD)** tools as DevOps professionals. However, developers are adept at using the tools and deploying apps. It is crucial to remember that DevOps is about a blend of best practices, mentality, and technology rather than just CI/CD tools. The duties carried out by DevOps engineers in different organizations will differ. DevOps and SRE engineers are in high demand right now, and job portals are contacting or poaching candidates from various firms to fill the positions. However, delve a little deeper and look at job duties. You will discover that some firms refer to the work of migration engineers, network engineers, support engineers, CI/CD engineers, and system engineers under the title of DevOps/SRE engineers. It is evident that DevOps is not your profession if it requires 24/7 assistance. In later chapters, we will discuss the support provided by SRE engineers.

Application DevOps engineers and platform DevOps engineers are the two different categories of DevOps engineers. Because fewer employees are working in small businesses, it will be challenging to determine the type of position the engineer has. Various situations can be found in large organizations, such as one DevOps team supporting numerous teams concurrently and possessing knowledge of various technologies. All team members use a

single pipeline they design for deployments. The DevOps engineer's responsibility is to safeguard the pipeline and prevent any vulnerabilities from making it into production.

The second kind of team would consist of one DevOps person who would accompany one application team; for example, let us consider a database team, and then one database SRE would be a team component responsible for all SRE-related tasks. These SRE are heavily involved at every step of the process, from the product's conception through its manufacture. They are incredibly knowledgeable about the technology, the item, or the application. SRE is the first person to be called, even before developers, if anything goes wrong in production. They investigate the problem using all available monitoring tools before contacting the necessary team members. By identifying the core cause and implementing best practices, they also ensure this situation does not occur again.

Roles and responsibilities of SRE

The essential tasks and duties of an SRE are as follows:

- **System design and implementation:** An SRE is responsible for designing and implementing highly scalable and robust systems. They work closely with software development teams, providing insights into designing systems that are easy to manage and scale. Their software engineering knowledge and understanding of systems allow them to create designs that can handle high traffic levels and recover quickly from any failure.
- **Incident management:** Another critical role of an SRE is incident management. When a service or system fails or behaves unexpectedly, the SREs are the first person to respond. They are responsible for quickly identifying the issue, mitigating the impact, and resolving the incident. After the incident, they conduct a thorough post-mortem analysis to understand the root cause and prevent similar incidents.
- **Problem-solving and system optimization:** SREs are problem solvers. They constantly monitor the system's performance and identify potential bottlenecks or areas of improvement. By analyzing system trends and using various tools, they proactively solve problems before they become system-wide issues. They also seek to optimize system performance and efficiency, ensuring that resources are utilized best.
- **Change management:** SREs serve a critical role in ensuring that the proposed system and application modifications are properly examined and deployed. They conduct comprehensive risk assessments to foresee and prevent potential negative implications on system reliability while also ensuring that changes follow established norms. They confirm the necessity and safety of each modification by executing pre-release tests and considering security concerns. Following the implementation of these improvements, SREs will continuously monitor the system to detect and address any unforeseen issues, ensuring that system performance