

AWS Cloud Engineer Guide

Building scalable cloud solutions with AWS

Sizwe Molefe



www.bpbonline.com

First Edition 2025

Copyright © BPB Publications, India

ISBN: 978-93-65899-757

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:



www.bpbonline.com

Kup ksi k

Dedicated to

*The woman who raised me
and continues to inspire, my mother:*

Joyce Fikile Vezi

About the Author

Sizwe Molefe has been working in Information Technology for more than 14 years, having been in many roles in numerous projects as a technical point of contact, technical leader, project co-ordinator, project manager, and delivering projects. The projects delivered include but are not limited to SAP ERP Migrations, Microsoft Exchange, Active Directory, Database, and Cloud Migrations, and quite a few more others. Sizwe has worked for multi-international companies, the likes of MRI Systems, Zando a Jumia Group subsidiary, and Amazon. Currently, he is a Solutions Architect at Amazon Web Services. He is an Information Technology graduate and holds multiple AWS certifications. Furthermore, he is also an accredited AWS Redshift Subject Matter Expert. He also runs a research and consulting startup called Zuluchain, which has a primary focus on Blockchain and Web3.0 technologies. He is a Web3.0 enthusiast and actively participates in the Web3.0 community.

About the Reviewer

Dhirendra Kumar is a Kubernetes cloud Architect, presently working in the USA with 20 years of experience in Kubernetes (certified) as a platform engineer. His interests include multi-cloud, hybrid-cloud, distributed systems, and security-oriented cloud engineering. He is currently focused on and working with cloud-native solutions and tools, including AWS, Kubernetes, and Terraform. He is also focused on Infrastructure as Code and monitoring the correct thing with tools such as Dynatrace, Prometheus, and ELK Stack, with a passion for automating everything and avoiding doing anything that resembles manual work.

I would like to thank my wife, *Geethashri Ananda*, and my parents for always supporting me and motivating me to contribute to society.

Acknowledgement

I would like to express my gratitude to my family and friends for the support and cheering throughout the writing of this book.

I am grateful for my wife's (Luyanda) support as well as my two kids, Refiloe and Dithapelo.

I am highly grateful to my mother, who raised and put me through high school and tertiary school single-handedly after my father passed on in my younger days.

I give thanks to BPB for the guidance and expertise in helping me make this book a finished product. The journey has been a long one, filled with learnings and collaboration with editors and technical editors.

Last but not least, I would like to thank the readers who have shown interest in the contents of this book by purchasing it. Your support goes a long way.

Preface

The journey to cloud computing can be an overwhelming and complex experience. This is especially true for companies that have had their IT infrastructure hosted on-premise for years. The idea of data migration and having data hosted in the cloud can sound daunting. In that sense, the cloud journey requires that one has a comprehensive understanding of the many cloud services that cloud service providers offer.

This book is designed as a guide on how to make use of the different AWS Cloud Services. Here, we focus on the core services that you will likely cross paths with in your AWS cloud journey. These services include EC2, S3, VPC, RDS, CloudFormation, and a lot more. These are some of the core services likely to be used in migration cases.

In the course of this book, you will gain knowledge on AWS Services and their key features. You will learn how to use each service, different use cases, and gain in-depth knowledge by performing knowledge checks on each chapter. The book will give you a theoretical idea on how to design a cloud infrastructure. You will also learn about the best practices of each service described and examples.

This book is intended for students pursuing IT/computer science, solutions architects, developers, and IT professionals who are new to the AWS cloud. This book will help them expand their horizon when it comes to AWS and improve their skills in building scalable and resilient environments on AWS.

With this book, you will gain the knowledge and skills to become a proficient AWS Cloud Engineer using critical AWS Services. I hope you will find it informative and helpful.

Chapter 1: Creating an AWS Environment – This chapter provides a comprehensive introduction to creating and managing an AWS environment. Here, we will cover key concepts that include setting up an AWS Account, using the Free Trial, Account Security best practices, installing AWS CLI and Boto3. This chapter also introduces key concepts and tools that will be expanded upon later, such as **Elastic Compute Cloud (EC2)**.

Chapter 2: Amazon Elastic Compute Cloud – This chapter covers the introduction to EC2. This includes the definition and history of EC2, step-by-step instructions on creating EC2 instances through the AWS Console, SDK, and CloudFormation. We will take the user through practical, hands-on guidance for setting up and managing EC2 instances, emphasizing best practices for security and data management. We will also introduce

the key concepts and related AWS services that integrate with EC2, providing a solid foundation for readers to build upon in their cloud computing journey.

Chapter 3: Amazon Virtual Private Cloud - This chapter explores **Amazon Virtual Private Cloud (Amazon VPC)**, a fundamental AWS service that allows users to create and manage their own virtual networks within the cloud. It emphasizes the core aspects of VPC, such as security, scalability, and customization, enabling users to build isolated, controlled network environments. The key topics covered include the configuration of IP address ranges, the creation of subnets, routing tables, and network gateways, alongside an exploration of the tools provided by AWS to secure and optimize cloud deployments. This chapter is essential for understanding how to maintain robust and flexible network architectures while leveraging the cloud's vast resources.

Chapter 4: Amazon S3: Simple Storage Service – In this chapter, we will explore Amazon S3, a fundamental AWS service that offers users object storage, a type of data storage. Here, the reader will understand the key concepts of S3, which include storage classes, S3 versioning, S3's high level of durability, scalability, and security, alongside the integration of S3 with other AWS services. Learning about these key topics will help the reader gain the necessary skills to fully utilize S3 as a storage service.

Chapter 5: Amazon API Gateway – This chapter provides the reader with a comprehensive overview of Amazon API Gateway. This also includes and covers API Gateway integration with Lambda functions. The key concepts covered include passing API keys to API Gateway, integrating Lambda functions with API Gateway, and deploying and invoking API endpoints. The readers will get to learn and understand with real-life scenarios the deploying and invoking of REST API endpoints and building serverless applications using API Gateway.

Chapter 6: AWS Database Services – This chapter provides a high-level understanding of the various AWS Database Services and explores the key concepts of database administration and monitoring. The readers will learn about the numerous advantages of AWS RDS, a fully managed relational database service that simplifies the process of setting up, operating, and scaling relational databases in the AWS cloud. This chapter also allows the user to learn about DynamoDB, a fully managed, highly scalable, and high-performance NoSQL database service that provides consistent, single-digit millisecond latency at any scale. The readers will also learn about Amazon Aurora, a fully managed, high-performance, and cost-effective relational database service compatible with MySQL and PostgreSQL.

Chapter 7: Elastic Load Balancing and Auto Scaling – In this chapter, the readers will explore two powerful services provided by AWS that are essential for building scalable and highly available applications: **Elastic Load Balancing (ELB)** and Auto Scaling. We will cover key concepts for each service, which for ELB include features and benefits (scaling, load distribution, health checks, sticky sessions, and more) and configuring ELB (defining listener rules, creating target groups, and configuring routing). For Auto Scaling, the readers will learn about automatically adjusting the capacity of infrastructure based on application demand and how this helps with optimal performance.

Chapter 8: Amazon Route 53 - The readers will learn about the fundamental AWS service known as Route53. This is a highly scalable and reliable **Domain Name System (DNS)** service offered by AWS. The readers will learn about the Route53 services, which is more than just translating domain names to IP addresses. They will also learn and understand the advanced routing policies, traffic management options, implementing health checks and failover configurations, managing domain registration and settings, integrating Route 53 with other AWS services, and many more.

Chapter 9: Decouple Applications – This chapter will provide a fundamental approach to designing and building modern, scalable, and flexible systems. It involves separating components and minimizing dependencies to achieve a loosely coupled architecture. The readers will explore the AWS services, which include AWS Lambda, Amazon SQS, Amazon SNS, Amazon Kinesis, and others. They will learn how data can be managed independently from application logic, providing the freedom to scale and evolve data-intensive applications.

Chapter 10: CloudFormation – In this chapter, we will explore the capabilities and the benefits of AWS CloudFormation, a service that allows users to automate and manage their AWS infrastructure in a scalable and efficient manner. The readers will learn key concepts that take them through fundamental concepts such as templates, stacks, resources, and parameters. Templates are the JSON or YAML files that define the desired state of the users' AWS resources. We will also take the readers through the various methods of deploying CloudFormation templates. Additionally, this chapter will provide hands-on examples and walkthroughs to demonstrate the practical application of CloudFormation.

Chapter 11: AWS Monitoring – This chapter will explore the critical AWS services used for monitoring. We will focus primarily on Amazon CloudWatch, which is the central monitoring. Additionally, we will explore other essential monitoring services like AWS CloudTrail and AWS Config, which offer visibility into API activity and resource

configuration changes. We will also look into the key features of these products, practical examples, and step-by-step instructions to help readers effectively leverage the AWS monitoring services

Chapter 12: AWS Security and Encryption – Security is paramount in the cloud, and at AWS, security is one of the highest priority topics when it comes to services offered. In this chapter, the readers will cover AWS Security and Encryption, focusing on various objectives that are pivotal in securing AWS environments. The key security concepts covered in this chapter include IAM, KMS, SSM, Secrets Manager, CloudHSM, Shield, WAF, GuardDuty, Inspector, and Macie.

Chapter 13: AWS Containers – In this chapter, we will explore the concept of AWS Container Service. The readers will learn about Amazon ECS and Amazon EKS. We will uncover how these services empower organizations to orchestrate and deploy containers at scale. The readers will also discover AppRunner, a service catered to accelerate code to scalable web applications. By the end of this chapter, they will gain the knowledge and insights needed to embrace the container revolution and harness the full potential of this transformative technology.

Chapter 14: Automating Deployments with CI/CD in AWS – This chapter looks into the world of DevOps and is designed to help the reader gain insights into the rise of agile methodologies. We will start by exploring core CI/CD concepts more deeply, including the differences between continuous integration vs. delivery vs. deployment. The readers will also learn about CI/CD in the context of Kubernetes and containerized applications, followed by demonstration patterns leveraging Docker, Kubernetes concepts, and AWS container tools. We will provide guidelines, best practices, and troubleshooting advice gleaned from real-world experience with CI/CD in AWS across various applications and teams.

Chapter 15: AWS Cloud Migrations – In this chapter, we will discover the various fundamental AWS Services that play a major role in migration to the cloud. These services include but not limited to AWS **Database Migration Service (DMS)**, a specialized tool for replicating on-premises databases to AWS with minimal downtime. We will explore AWS DataSync, a software-based transfer acceleration tool optimized for Nas, filesystems, and S3. The readers will learn about the methodical cloud migration components, which include components, dependencies, waves, and landing zones. They will also gain comprehensive technical knowledge and strategies required to migrate organizations to the AWS cloud successfully.

Code Bundle and Coloured Images

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

<https://rebrand.ly/fytde1y>

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

<https://github.com/bpbpublications/AWS-Cloud-Engineer-Guide>.

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. Creating an AWS Environment	1
Introduction.....	1
Structure.....	1
Objectives	2
Creating an AWS account and environment	2
<i>Creating an IAM user</i>	<i>3</i>
<i>Walkthrough IAM roles and policies</i>	<i>7</i>
IAM roles.....	7
IAM policies	8
<i>Account creation best practices</i>	<i>8</i>
Installing AWS CLI.....	9
Installing Boto3.....	11
Leveraging SDK tools to create resources in your new AWS environment.....	12
Automating AWS resource creation with CloudFormation.....	13
Creating an EC2 instance using CloudFormation	14
Knowledge check	15
<i>Answers.....</i>	<i>16</i>
Conclusion.....	16
 2. Amazon Elastic Compute Cloud	 17
Introduction.....	17
Structure.....	18
Objectives	18
Creating an EC2 instance through the console	18
<i>SDK tools</i>	<i>20</i>
<i>CloudFormation templates</i>	<i>22</i>
<i>Amazon Machine Images</i>	<i>25</i>
Controlling traffic to virtual machines using security groups.....	25
Sharing data volumes between machines using EFS.....	26
Elastic Block Storage and backing up data with snapshots	28

Knowledge check	31
<i>Answers</i>	31
Conclusion.....	32
3. Amazon Virtual Private Cloud	33
Introduction.....	33
Structure.....	33
Objectives	34
Key features of Amazon VPC	34
Monitoring VPC with AWS CloudWatch	35
<i>Example use case: Monitoring NAT gateway traffic</i>	36
Implementing virtual private cloud	37
Defining a virtual private network.....	39
NAT gateway	41
Network access control lists.....	43
<i>Knowledge check</i>	45
<i>Answers</i>	46
Security groups.....	46
<i>Knowledge check</i>	49
<i>Answers</i>	50
Ethernet network interface.....	50
Transitive Cross-VPC Connections	53
Knowledge check	54
<i>Answers</i>	54
Conclusion.....	55
4. Amazon S3: Simple Storage Service.....	57
Introduction.....	57
Structure.....	58
Objectives	58
<i>Amazon S3 features</i>	58
<i>S3 uploads and downloads</i>	59
S3 lifecycle policies.....	60

<i>Knowledge check</i>	62
<i>Answers</i>	62
Archive S3 objects.....	62
<i>Knowledge check</i>	64
<i>Answers</i>	65
Observing S3 storage	65
<i>Knowledge check</i>	66
<i>Answers</i>	67
Application-specific access to S3.....	67
<i>Knowledge check</i>	68
<i>Answers</i>	69
Creating and restoring EC2 backups.....	69
<i>Knowledge check</i>	70
<i>Answers</i>	71
Replication data and replication rules on S3 bucket	71
<i>Knowledge check</i>	73
<i>Answers</i>	74
Conclusion.....	74
5. Amazon API Gateway	75
Introduction.....	75
Structure.....	76
Objectives	76
Amazon API Gateway and integration with Lambda	76
Passing API keys to API Gateway in detail	84
Deploying and invoking API endpoint.....	86
Knowledge check	90
<i>Answers</i>	91
Conclusion.....	91
6. AWS Database Services	93
Introduction.....	93
Structure.....	94

Objectives	94
Importance of database services in the cloud	94
<i>Growth of digital data</i>	94
<i>Data-driven decision making</i>	95
<i>Agility and speed</i>	95
<i>Cost efficiency</i>	95
<i>Innovation and collaboration</i>	95
AWS database service portfolio	96
<i>Amazon Relational Database Service</i>	96
<i>Supported database engines</i>	96
<i>Example use cases</i>	97
<i>Amazon DynamoDB</i>	100
<i>Key features</i>	101
<i>Use cases</i>	101
<i>Amazon Aurora</i>	104
<i>Key features</i>	104
<i>Use cases</i>	105
<i>Scenario</i>	107
<i>Problem statement</i>	107
<i>Solution</i>	108
<i>Amazon Aurora Serverless</i>	113
<i>How Aurora Serverless works</i>	113
<i>Enabling REST access to Aurora Serverless</i>	113
<i>Benefits of Aurora Serverless</i>	114
<i>Using Aurora Serverless</i>	114
<i>Amazon DocumentDB</i>	115
<i>Use cases of Amazon DocumentDB</i>	115
<i>Architecture of Amazon DocumentDB</i>	116
<i>Real life scenario for Amazon DocumentDB</i>	117
<i>Amazon ElastiCache</i>	119
<i>How ElastiCache works</i>	120
<i>Use cases of ElastiCache</i>	120
<i>Real life scenario for Amazon ElastiCache</i>	121

Customer success stories with Amazon ElastiCache	122
Amazon Redshift	122
How Redshift works	122
Use cases of Redshift.....	124
Real life scenario using Redshift	125
Redshift Spectrum	128
Additional features and capabilities of AWS Database Services.....	131
Simplified database administration	131
IAM authentication with RDS database.....	132
How IAM database authentication works.....	132
Benefits of IAM database authentication	133
Setting up IAM database authentication	133
Leveraging RDS Proxy	133
How RDS Proxy works	133
Benefits of using RDS Proxy.....	134
Using RDS Proxy	135
RDS password rotation	135
Why rotate database passwords.....	135
Rotating RDS Database passwords	136
Best practices for RDS password rotation.....	136
Knowledge check	137
Answers.....	137
Conclusion.....	138
7. Elastic Load Balancing and Auto Scaling.....	139
Introduction.....	139
Structure.....	140
Objectives	140
Understanding load balancing	140
Types of Elastic Load Balancers.....	141
Classic Load Balancer	141
Application Load Balancer	142
Network Load Balancer	142

Key features and benefits of ELB	142
<i>Setting up an Elastic Load Balancer</i>	143
Cross-zone load balancing	145
Elastic Load Balancer Connection Draining.....	146
<i>Real life use case scenario using Elastic Load Balancing</i>	148
<i>Understanding the concept of Auto Scaling</i>	150
<i>Benefits of Auto Scaling</i>	151
Auto Scaling groups.....	151
<i>Creating and managing Auto Scaling groups</i>	152
<i>Scaling policies on metrics and conditions</i>	152
<i>Integration with Elastic Load Balancing for dynamic scaling</i>	153
Scaling strategies	153
Advanced Auto Scaling features.....	155
<i>Real life use case scenario using Auto Scaling</i>	156
<i>Real-time scenario</i>	158
<i>Creating ELB using the AWS Console</i>	159
<i>Creating an Elastic Load Balancer using the AWS CLI</i>	159
<i>Creating Auto Scaling using the AWS Console</i>	160
<i>Creating Auto Scaling using the AWS CLI</i>	161
Knowledge check	161
<i>Answers</i>	163
Conclusion.....	164
8. Amazon Route 53	165
Introduction.....	165
Structure.....	166
Objectives	166
Understanding DNS fundamentals.....	166
<i>Common misunderstood topics on DNS</i>	168
Getting started with Route 53.....	169
<i>Getting started with Route 53 using the AWS CLI</i>	172
Routing policies and traffic management.....	175
<i>Real-life example</i>	176

Health checks and failover.....	177
<i>Real-life example</i>	178
Domain registration and management	179
<i>Real-life example</i>	180
Route 53 integration with other AWS services.....	181
<i>Real-life example</i>	182
Best practices and troubleshooting techniques.....	183
<i>Real-life example</i>	184
Knowledge check	185
<i>Answers</i>	185
Conclusion.....	186
9. Decouple Applications.....	187
Introduction.....	187
Structure.....	188
Objectives	188
Exploring AWS decoupling services	188
<i>Configuring AWS decoupling services</i>	193
<i>AWS Lambda</i>	193
Amazon Simple Queue Service.....	197
Amazon Simple Notification Service.....	202
<i>Delay Queues in SNS and SQS</i>	205
Amazon EventBridge.....	205
<i>Use cases</i>	206
Knowledge check	208
<i>Answers</i>	208
Conclusion.....	209
10. CloudFormation	211
Introduction.....	211
Structure.....	211
Objectives	212
Fundamental concepts of CloudFormation.....	212

<i>Writing CloudFormation templates</i>	213
Intrinsic functions.....	213
AWS CloudFormation Designer	217
CloudFormation artifacts	219
<i>Stack Outputs</i>	219
<i>Stack Events</i>	220
<i>Resource metadata</i>	221
<i>CloudFormation registry</i>	222
<i>CloudFormation drift detection</i>	223
Deployment process using CloudFormation	224
<i>Stack creation</i>	224
<i>Stack updates</i>	225
<i>Stack failures</i>	227
<i>Stack deletion</i>	228
Best practices for CloudFormation templates	229
Challenges working with CloudFormation.....	231
Knowledge check	232
<i>Answers</i>	232
Conclusion.....	233
11. AWS Monitoring	235
Introduction.....	235
Structure.....	235
Objectives	236
Fundamental concepts of Amazon CloudWatch	236
Basic monitoring.....	237
<i>Metrics collected by basic monitoring</i>	237
<i>Monitoring intervals and data retention</i>	238
Detailed monitoring.....	239
<i>Increased monitoring granularity</i>	239
<i>Additional metrics</i>	239
<i>Enabling detailed monitoring</i>	240
<i>Use cases for detailed monitoring</i>	240

AWS CloudWatch alarms	241
<i>Setting up CloudWatch alarms</i>	241
<i>Alarm use cases</i>	242
<i>Alarm notifications and integration</i>	243
<i>Practical scenario using AWS CloudWatch</i>	243
CloudWatch Metrics Insights	248
Grafana integration	249
Setting up CloudWatch real user monitoring	249
Fundamental of AWS CloudTrail	251
<i>AWS CloudTrail event history</i>	252
<i>Practical scenario using AWS CloudTrail</i>	252
Fundamental concepts of AWS Config	255
<i>Practical scenario using AWS Config</i>	256
Key considerations and challenges with monitoring	259
<i>CloudWatch</i>	259
<i>CloudTrail</i>	260
<i>AWS Config</i>	261
AWS monitoring best practices	263
<i>AWS CloudWatch</i>	263
<i>AWS CloudTrail</i>	264
<i>AWS Config</i>	264
Knowledge check	266
<i>Answers</i>	267
Conclusion	267
12. AWS Security and Encryption	269
Introduction	269
Structure	269
Objectives	270
AWS share responsibility	270
<i>AWS responsibilities</i>	271
<i>Customer responsibilities</i>	271
<i>Use case example: E-commerce company</i>	272