

# EXAM ✓ CRAM

# AWS Certified SysOps Administrator – Associate

(SOA-C02)



CramSheet



Flash  
Cards



Practice  
Tests



MARKO SLUGA  
RICK CRISCI  
WILLIAM “BO” ROTHWELL

FREE SAMPLE CHAPTER |



**EXAM ✓ CRAM**

**AWS Certified  
SysOps  
Administrator –  
Associate  
(SOA-C02)  
Exam Cram**

**Marko Sluga  
Rick Crisci  
William “Bo” Rothwell**



Pearson

## **AWS Certified SysOps Administrator – Associate (SOA-C02) Exam Cram**

Copyright © 2022 by Pearson Education, Inc.

All rights reserved. This publication is protected by copyright, and permission must be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise. For information regarding permissions, request forms, and the appropriate contacts within the Pearson Education Global Rights & Permissions Department, please visit [www.pearson.com/permissions](http://www.pearson.com/permissions).

No patent liability is assumed with respect to the use of the information contained herein. Although every precaution has been taken in the preparation of this book, the publisher and author assume no responsibility for errors or omissions. Nor is any liability assumed for damages resulting from the use of the information contained herein.

ISBN-13: 978-0-13-750958-4

ISBN-10: 0-13-750958-8

Library of Congress Cataloging-in-Publication Data: 2022902124

ScoutAutomatedPrintCode

### **Trademarks**

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Pearson IT Certification cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

### **Warning and Disclaimer**

Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied. The information provided is on an “as is” basis. The authors and the publisher shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book.

### **Special Sales**

For information about buying this title in bulk quantities, or for special sales opportunities (which may include electronic versions; custom cover designs; and content particular to your business, training goals, marketing focus, or branding interests), please contact our corporate sales department at [corpsales@pearsoned.com](mailto:corpsales@pearsoned.com) or (800) 382-3419.

For government sales inquiries, please contact [governmentsales@pearsoned.com](mailto:governmentsales@pearsoned.com).

For questions about sales outside the U.S., please contact [intlcs@pearson.com](mailto:intlcs@pearson.com).

### **Editor-in-Chief**

Mark Taub

### **Director, ITP Product Management**

Brett Bartow

### **Executive Editor**

Nancy Davis

### **Managing Editor**

Sandra Schroeder

### **Development Editor**

Christopher A.  
Cleveland

### **Project Editor**

Mandie Frank

### **Copy Editor**

Chuck Hutchinson

### **Indexer**

Ken Johnson

### **Proofreader**

Donna E. Mulder

### **Technical Editors**

Mark Wilkins  
Ryan Dymek

### **Publishing Coordinator**

Cindy Teeters

### **Designer**

Chuti Prasertsith

### **Compositor**

codeMantra

# Pearson's Commitment to Diversity, Equity, and Inclusion

Pearson is dedicated to creating bias-free content that reflects the diversity of all learners. We embrace the many dimensions of diversity, including but not limited to race, ethnicity, gender, socioeconomic status, ability, age, sexual orientation, and religious or political beliefs.

Education is a powerful force for equity and change in our world. It has the potential to deliver opportunities that improve lives and enable economic mobility. As we work with authors to create content for every product and service, we acknowledge our responsibility to demonstrate inclusivity and incorporate diverse scholarship so that everyone can achieve their potential through learning. As the world's leading learning company, we have a duty to help drive change and live up to our purpose to help more people create a better life for themselves and to create a better world.

Our ambition is to purposefully contribute to a world where

- Everyone has an equitable and lifelong opportunity to succeed through learning
- Our educational products and services are inclusive and represent the rich diversity of learners
- Our educational content accurately reflects the histories and experiences of the learners we serve
- Our educational content prompts deeper discussions with learners and motivates them to expand their own learning (and worldview)

While we work hard to present unbiased content, we want to hear from you about any concerns or needs with this Pearson product so that we can investigate and address them.

- Please contact us with concerns about any potential bias at <https://www.pearson.com/report-bias.html>.

# Credits

## Figures

Figures 2.1–2.3, 4.2–4.7, 4.10–4.13,  
5.1, 5.3, 5.4, 6.1–6.5, 7.1–7.3, 9.1–9.9,  
10.1, 11.1–11.9, 11.11, 12.2, 12.3,  
12.8–12.17, 13.1–13.7, 14.1–14.8,  
15.3–15.5

Figure 4.1

Figure 5.2

Figures 10.2, 10.3

## Credit

Amazon Web Services, Inc

Medium

DeveloperCK

Google LLC

# Contents at a Glance

|                   |   |       |
|-------------------|---|-------|
|                   | Introduction                            | xviii |
| <b>CHAPTER 1</b>  | Introduction to AWS                     | 1     |
| <b>CHAPTER 2</b>  | Monitoring Services in AWS              | 25    |
| <b>CHAPTER 3</b>  | Troubleshooting and Remediation         | 41    |
| <b>CHAPTER 4</b>  | Implementing Scalability and Elasticity | 55    |
| <b>CHAPTER 5</b>  | High Availability and Resilience        | 85    |
| <b>CHAPTER 6</b>  | Backup and Restore Strategies           | 109   |
| <b>CHAPTER 7</b>  | Provisioning Resources                  | 125   |
| <b>CHAPTER 8</b>  | Application Management                  | 137   |
| <b>CHAPTER 9</b>  | Security and Compliance                 | 155   |
| <b>CHAPTER 10</b> | Data Protection at Rest and in Transit  | 171   |
| <b>CHAPTER 11</b> | Networking and Connectivity             | 183   |
| <b>CHAPTER 12</b> | Domains, DNS, and Content Delivery      | 217   |
| <b>CHAPTER 13</b> | Troubleshoot Network Connectivity       | 239   |
| <b>CHAPTER 14</b> | Cost Optimization Strategies            | 259   |
| <b>CHAPTER 15</b> | Performance Optimization                | 271   |
|                   | Glossary                                | 283   |
|                   | Index                                   | 291   |

# Table of Contents

|  |              |
|--|--------------|
| <b>Introduction</b> . . . . .                | <b>xviii</b> |
| <b>CHAPTER 1</b>                             |              |
| <b>Introduction to AWS</b> . . . . .         | <b>1</b>     |
| What Is Cloud Computing? . . . . .           | 2            |
| The AWS Cloud. . . . .                       | 8            |
| AWS Global Architecture . . . . .            | 9            |
| Datacenters . . . . .                        | 9            |
| Availability Zones. . . . .                  | 9            |
| Regions . . . . .                            | 10           |
| Edge Locations . . . . .                     | 10           |
| Accessing AWS . . . . .                      | 11           |
| Six Benefits of the AWS Cloud. . . . .       | 13           |
| AWS Services Overview . . . . .              | 16           |
| Foundation Services . . . . .                | 17           |
| Network Services . . . . .                   | 17           |
| Compute Services . . . . .                   | 18           |
| Storage Services. . . . .                    | 18           |
| Security and Identity Services . . . . .     | 19           |
| End-User Applications Services . . . . .     | 19           |
| Platform Services . . . . .                  | 20           |
| Database Services. . . . .                   | 20           |
| Analytics Services. . . . .                  | 21           |
| Application Services . . . . .               | 21           |
| Other Examples. . . . .                      | 21           |
| Management and Deployment Services . . . . . | 22           |
| What Next? . . . . .                         | 23           |
| <b>CHAPTER 2</b>                             |              |
| <b>Monitoring Services in AWS</b> . . . . .  | <b>25</b>    |
| Metering, Monitoring, and Alerting . . . . . | 26           |
| Metering . . . . .                           | 27           |
| Monitoring . . . . .                         | 27           |
| Alerting . . . . .                           | 27           |
| Security and Compliance. . . . .             | 27           |
| CloudWatch . . . . .                         | 30           |
| Namespaces and Dimensions . . . . .          | 31           |

|   |           |
|---|-----------|
| Metrics . . . . .   | 32        |
| Statistics . . . . .                                      | 33        |
| Percentiles . . . . .                                     | 33        |
| CloudWatch Alarms . . . . .                               | 33        |
| CloudWatch Logs . . . . .                                 | 34        |
| CloudWatch Logs Insights . . . . .                        | 34        |
| CloudTrail . . . . .                                      | 36        |
| What Next? . . . . .                                      | 39        |
| <br>  |           |
| <b>CHAPTER 3</b>  |           |
| <b>Troubleshooting and Remediation . . . . .</b>          | <b>41</b> |
| Responding to Alarms . . . . .                            | 42        |
| Infrastructure Issues . . . . .                           | 43        |
| Application Issues . . . . .                              | 44        |
| Security Issues . . . . .                                 | 44        |
| Error Handling in AWS . . . . .                           | 45        |
| Amazon EventBridge . . . . .                              | 48        |
| Integration with AWS Systems Manager Automation . . . . . | 49        |
| AWS Config . . . . .                                      | 51        |
| What Next? . . . . .                                      | 53        |
| <br>  |           |
| <b>CHAPTER 4</b>  |           |
| <b>Implementing Scalability and Elasticity . . . . .</b>  | <b>55</b> |
| Scaling in the Cloud . . . . .                            | 56        |
| Horizontal vs. Vertical Scaling . . . . .                 | 59        |
| AWS Autoscaling . . . . .                                 | 62        |
| Dynamic Scaling . . . . .                                 | 63        |
| Manual and Scheduled Scaling . . . . .                    | 63        |
| Predictive Scaling . . . . .                              | 64        |
| Caching . . . . .   | 67        |
| Types of Caching . . . . .                                | 68        |
| Client-Side Caching . . . . .                             | 68        |
| Edge Caching . . . . .                                    | 68        |
| Server-Side Caching . . . . .                             | 68        |
| Database Caching . . . . .                                | 69        |
| ElastiCache . . . . .                                     | 71        |
| Memcached . . . . .                                       | 72        |
| Redis . . . . .   | 73        |
| Amazon CloudFront . . . . .                               | 74        |
| CloudFront Security . . . . .                             | 75        |



|                         |    |
|-------------------------|----|
| Read Replicas . . . . . | 78 |
| What Next? . . . . .    | 83 |

## CHAPTER 5

|  |           |
|--|-----------|
| <b>High Availability and Resilience . . . . .</b>                    | <b>85</b> |
| Availability Zones in AWS . . . . .                                  | 88        |
| VPC . . . . .  | 89        |
| Private Subnets . . . . .  | 91        |
| Public Subnets . . . . .   | 91        |
| High Availability with Elastic Load Balancers and Route 53 . . . . . | 94        |
| Highly Available Datastores . . . . .                                | 98        |
| Amazon S3 . . . . .  | 98        |
| S3 Storage Tiers . . . . .   | 99        |
| S3 Security . . . . .  | 100       |
| Highly Available Databases . . . . .                                 | 103       |
| Amazon RDS . . . . .   | 104       |
| Amazon Aurora . . . . .  | 105       |
| Amazon DynamoDB . . . . .  | 106       |
| DynamoDB Global Tables . . . . .                                     | 107       |
| What Next? . . . . .   | 108       |

## CHAPTER 6

|  |            |
|--|------------|
| <b>Backup and Restore Strategies . . . . .</b> | <b>109</b> |
| Backup in the Cloud . . . . .                  | 110        |
| Backups and Snapshots . . . . .                | 112        |
| RPO and RTO . . . . .                          | 113        |
| Disaster Recovery . . . . .                    | 114        |
| Overview of Backup Strategies . . . . .        | 115        |
| Backup and Restore . . . . .                   | 115        |
| Pilot Light . . . . .                          | 115        |
| Warm Standby . . . . .                         | 116        |
| Multisite Active-Active . . . . .              | 116        |
| S3 as a Backup Service . . . . .               | 119        |
| S3 Sync . . . . .                              | 120        |
| AWS DataSync . . . . .                         | 121        |
| Versioning and Life-Cycling Rules . . . . .    | 121        |
| Glacier . . . . .                              | 122        |
| S3 Cross-Region Replication . . . . .          | 123        |
| What Next? . . . . .                           | 124        |

**CHAPTER 7**

**Provisioning Resources . . . . . 125**

- Deployment Tools in AWS . . . . . 126
  - AWS CLI . . . . . 126
- Managing EC2 AMIs . . . . . 129
- AWS CloudFormation . . . . . 129
  - AWS Elastic Beanstalk . . . . . 130
  - CloudFormation StackSets . . . . . 132
- AWS Systems Manager . . . . . 133
- AWS Resource Access Manager . . . . . 134
- What Next? . . . . . 136

**CHAPTER 8**

**Application Management . . . . . 137**

- Lifecycle Management . . . . . 138
  - Deployment Environments . . . . . 138
    - Development . . . . . 139
    - Quality Assurance (QA) . . . . . 139
    - Production . . . . . 139
    - Staging . . . . . 139
    - Blue/Green . . . . . 139
    - Disaster Recovery (DR) . . . . . 140
  - Application Builds . . . . . 140
    - Stable . . . . . 140
    - Long-Term Support (LTS) . . . . . 141
    - Canary . . . . . 141
  - Upgrade Methods . . . . . 142
    - Rolling Upgrades . . . . . 142
    - Active-Passive . . . . . 142
  - Lifecycle Management . . . . . 142
    - Roadmaps . . . . . 142
    - Versioning . . . . . 143
    - Upgrading and Migrating Systems . . . . . 143
    - Deprecations or End of Life . . . . . 144
  - AWS OpsWorks . . . . . 145
- Patching . . . . . 147
  - Security Patches . . . . . 147
    - Hot Fixes . . . . . 148
    - Scheduled Updates . . . . . 148
    - Virtual Patches . . . . . 148

|                                     |     |
|-------------------------------------|-----|
| Signature Updates . . . . .         | 148 |
| Rollups . . . . .                   | 148 |
| Patching Cloud Components . . . . . | 149 |
| VMs . . . . .                       | 149 |
| Virtual Appliances . . . . .        | 149 |
| Networking Components . . . . .     | 149 |
| Applications . . . . .              | 150 |
| Storage Components . . . . .        | 150 |
| Firmware . . . . .                  | 150 |
| Software . . . . .                  | 151 |
| Policies . . . . .                  | 151 |
| N-1 . . . . .                       | 151 |
| Rollbacks . . . . .                 | 152 |
| What Next? . . . . .                | 153 |

## CHAPTER 9

### **Security and Compliance . . . . . 155**

|  |     |
|--|-----|
| Account Management . . . . .               | 156 |
| IAM Essentials . . . . .                   | 156 |
| IAM Users . . . . .                        | 157 |
| IAM User Groups . . . . .                  | 158 |
| IAM Roles . . . . .                        | 158 |
| IAM Policies . . . . .                     | 159 |
| Identity Providers . . . . .               | 161 |
| Password Policy . . . . .                  | 162 |
| Access Keys . . . . .                      | 163 |
| Multifactor Authentication (MFA) . . . . . | 164 |
| Best Practices in IAM . . . . .            | 165 |
| Trusted Advisor . . . . .                  | 168 |
| What Next? . . . . .                       | 169 |

## CHAPTER 10

### **Data Protection at Rest and in Transit . . . . . 171**

|                                      |     |
|--------------------------------------|-----|
| Protecting Data . . . . .            | 172 |
| Encryption . . . . .                 | 172 |
| AWS Key Management Service . . . . . | 173 |
| Data Classification . . . . .        | 174 |
| Certificate Management . . . . .     | 175 |
| AWS Certificate Manager . . . . .    | 177 |
| AWS Secrets Manager . . . . .        | 178 |

|                            |     |
|----------------------------|-----|
| Amazon GuardDuty . . . . . | 179 |
| Amazon Inspector . . . . . | 180 |
| AWS Security Hub . . . . . | 181 |
| What Next? . . . . .       | 182 |

**CHAPTER 11**

**Networking and Connectivity . . . . . 183**

|  |     |
|--|-----|
| The VPC . . . . .                              | 184 |
| Virtual Private Cloud . . . . .                | 184 |
| Elastic Network Interfaces . . . . .           | 185 |
| Elastic IP Address . . . . .                   | 186 |
| Internet Gateway (IGW) . . . . .               | 187 |
| VPC Connectivity . . . . .                     | 189 |
| Subnets . . . . .                              | 189 |
| Route Tables . . . . .                         | 190 |
| NAT Gateway . . . . .                          | 191 |
| VPC Security . . . . .                         | 193 |
| VPC Security by Default . . . . .              | 193 |
| Security Groups . . . . .                      | 194 |
| Network Access Control Lists (NACLs) . . . . . | 194 |
| Public and Private Subnets . . . . .           | 195 |
| VPC Flow Logs . . . . .                        | 195 |
| AWS Network Firewall . . . . .                 | 197 |
| About the AWS Network Firewall . . . . .       | 197 |
| AWS Marketplace . . . . .                      | 198 |
| VPC Endpoints . . . . .                        | 199 |
| Interface Endpoints . . . . .                  | 199 |
| Gateway Endpoints . . . . .                    | 200 |
| VPC Peering . . . . .                          | 202 |
| About VPC Peering . . . . .                    | 202 |
| Transit Gateway . . . . .                      | 203 |
| VPN . . . . .                                  | 206 |
| AWS-Managed VPN . . . . .                      | 206 |
| Software Site-to-Site VPN . . . . .            | 207 |
| Direct Connect . . . . .                       | 209 |
| AWS Direct Connect . . . . .                   | 209 |
| AWS Direct Connect + VPN . . . . .             | 210 |
| Hosted and Dedicated Connections . . . . .     | 210 |

|                               |     |
|-------------------------------|-----|
| AWS WAF . . . . .             | 212 |
| AWS WAF . . . . .             | 212 |
| AWS Shield . . . . .          | 214 |
| AWS Shield Standard . . . . . | 214 |
| AWS Shield Advanced. . . . .  | 215 |
| What Next? . . . . .          | 216 |

## CHAPTER 12

### **Domains, DNS, and Content Delivery . . . . . 217**

|   |     |
|---|-----|
| Route 53 . . . . .                              | 218 |
| Public and Private Zones. . . . .               | 218 |
| Registering a Domain with Route 53 . . . . .    | 219 |
| Route 53 Routing Policies . . . . .             | 222 |
| Simple Routing Policy . . . . .                 | 222 |
| Weighted Routing Policy . . . . .               | 223 |
| Latency-Based Routing Policy . . . . .          | 223 |
| Failover Routing Policy. . . . .                | 224 |
| Geolocation Routing Policy. . . . .             | 225 |
| S3 Static Website Hosting . . . . .             | 227 |
| Configuring S3 Static Website Hosting . . . . . | 227 |
| Amazon CloudFront . . . . .                     | 231 |
| Introduction to CloudFront. . . . .             | 231 |
| Creating a CloudFront Distribution . . . . .    | 231 |
| Time-to-Live (TTL) . . . . .                    | 232 |
| Price Classes . . . . .                         | 232 |
| CloudFront and HTTPs. . . . .                   | 233 |
| Final CloudFront Configuration Tasks . . . . .  | 233 |
| S3 Origin Access Identity. . . . .              | 235 |
| Configuring an Origin Access Identity . . . . . | 235 |
| What Next? . . . . .                            | 237 |

## CHAPTER 13

### **Troubleshoot Network Connectivity . . . . . 239**

|                                    |     |
|------------------------------------|-----|
| VPC Flow Logs . . . . .            | 240 |
| Configuring VPC Flow Logs. . . . . | 240 |
| ELB Access Logs . . . . .          | 245 |
| AWS WAF ACL Logs. . . . .          | 248 |
| AWS WAF, Kinesis, and S3 . . . . . | 248 |
| CloudFront Logs . . . . .          | 251 |

|  |            |
|--|------------|
| Using CloudFront Logs . . . . .                    | 251        |
| Analyzing Standard Log Files . . . . .             | 252        |
| CloudFront Caching Issues . . . . .                | 253        |
| Improving Cache Hit Rate . . . . .                 | 253        |
| HTTP Status Codes from an Origin . . . . .         | 254        |
| Troubleshooting Hybrid and Private Links . . . . . | 256        |
| Troubleshooting Direct Connect . . . . .           | 256        |
| Troubleshooting AWS Managed VPNs . . . . .         | 257        |
| What Next? . . . . .                               | 258        |
| <b>CHAPTER 14</b>                                  |            |
| <b>Cost Optimization Strategies . . . . .</b>      | <b>259</b> |
| Operational Optimization . . . . .                 | 260        |
| Cost Allocation Tags . . . . .                     | 260        |
| Trusted Advisor . . . . .                          | 262        |
| AWS Compute Optimizer . . . . .                    | 263        |
| Cost Explorer . . . . .                            | 265        |
| AWS Budgets and Billing Alarms . . . . .           | 266        |
| Managed Services . . . . .                         | 267        |
| What Next? . . . . .                               | 269        |
| <b>CHAPTER 15</b>                                  |            |
| <b>Performance Optimization . . . . .</b>          | <b>271</b> |
| Optimizing for Performance . . . . .               | 272        |
| Monitoring for Performance Efficiency . . . . .    | 272        |
| Generation . . . . .                               | 274        |
| Aggregation . . . . .                              | 275        |
| Real-Time Processing and Alarming . . . . .        | 275        |
| Storage . . . . .                                  | 275        |
| Analytics . . . . .                                | 275        |
| Optimizing Compute Performance . . . . .           | 275        |
| Enhanced Networking . . . . .                      | 276        |
| Instance Store . . . . .                           | 276        |
| Placement Groups . . . . .                         | 277        |
| Optimizing Datastore Performance . . . . .         | 278        |
| Optimizing Database Performance . . . . .          | 280        |
| What Next? . . . . .                               | 281        |
| <b>Glossary . . . . .</b>                          | <b>283</b> |
| <b>Index . . . . .</b>                             | <b>291</b> |

# About the Authors

**Marko Sluga** has more than 20 years of experience in IT and has had the benefit of witnessing the rise of cloud computing. Marko has worked on a variety of cloud-related projects, from the early stages of SOC, corporate virtualization, and open-source API offerings to modern, fully automated, intelligent, serverless, and cloud-native solutions. Marko has been working with Amazon Web Services (AWS) since the start of the 2010s and holds three associate, two professional, and three specialty AWS certifications. Marko performs training and advising on cloud technologies and strategies, DevOps, and IT system and process optimization to clients from a wide range of companies, including startups, SMBs, enterprise businesses, and Fortune 500 companies. Marko runs his own cloud training, coaching, and consulting practice under the markocloud.com brand. He is the author of the *AWS Certified Developer - Associate (DVA-C01) Cert Guide*.

**Rick Crisci** is the founder and principal trainer for Trainertests.com and is an experienced AWS and VMware instructor. Original content by Rick can be found on a variety of platforms including Pearson, LinkedIn Learning, and Udemy, with more than 30 courses, over 250,000 students, and exceptionally high course review scores. Rick also teaches live hands-on AWS courses regularly for Pearson on O'Reilly.

Prior to becoming an instructor, Rick had over 15 years of real-world experience. Some career highlights include designing high-speed Internet networks in the early 2000s and managing virtualization and networking teams for a financial institution.

In 2017, VMware recognized Rick as the first runner-up for Instructor of the Year for the Americas. This honor was largely due to the glowing feedback from his students. Rick seeks to help students learn by taking complex concepts and providing clear and simple explanations, diagrams, and analogies.

At the impressionable age of 14, **William “Bo” Rothwell** crossed paths with a TRS-80 Micro Computer System (affectionately known as a “Trash 80”). Soon after, the adults responsible for Bo made the mistake of leaving him alone with the TRS-80. He immediately dismantled it and held his first computer class, showing his friends what made this “computer thing” work. Since that experience, Bo’s passion for understanding how computers work and sharing this knowledge with others has resulted in a rewarding career in IT training. His experience includes Cloud, Linux, UNIX, IT security, DevOps, and programming languages such as Perl, Python, Tcl, and BASH. He is the founder and lead instructor of One Course Source, an IT training organization.

# Dedication

*I would like to dedicate this book to my wife and children for their continued support of projects like this book.*

—Marko Sluga

*I would like to dedicate this book to my wife, Jessica, whose constant dedication to helping her students is always an inspiration, and to my children, who give me a purpose to work for. Also, to my mother for giving me confidence to pursue big goals, and to my father who showed me the value of hard work, dedication, and always putting his family before himself.*

—Rick Crisci

*Normally, I use this space to thank my wife, daughter, and parents. While they continue to be my biggest supporters, I feel compelled to dedicate this book to those who embrace lifelong learning. Thank you all for making this book a small part of your learning goals.*

*“There is no end to education. It is not that you read a book, pass an examination, and finish with education. The whole of life, from the moment you are born to the moment you die, is a process of learning.”*

—Jiddu Krishnamurti

—William “Bo” Rothwell

# Acknowledgments

I need to thank Rick and Bo for helping make this book a reality. I would also like to thank Nancy Davis and Chris Cleveland for their generous support during the creation of this title.

—Marko Sluga

Thanks to Marko for giving me the opportunity to help with this book, and to Nancy Davis for keeping this thing going. I’m appreciative of the opportunity to work with such wonderful people.

—Rick Crisci

Thanks to the entire Pearson team for all of the gentle pushes, patience, and dedication to making this book a success.

—William “Bo” Rothwell



# About the Technical Reviewers

**Mark Wilkins** is an electronic engineering technologist with a wealth of experience in designing, deploying, and supporting software and hardware technology in the corporate and small business world. Since 2013, Mark has focused on supporting and designing cloud service solutions with Amazon Web Services, Microsoft Azure, and the IBM Cloud. He is certified in Amazon Web Services (Architecture and SysOps). Mark is also a Microsoft Certified Trainer (MCT) and holds certifications in MCTS, MCSA, Server Virtualization with Windows Server Hyper-V, and Azure Cloud Services. Mark currently develops AWS curriculum on technical aspects of AWS architecture for O'Reilly Media, Pluralsight, and LinkedIn Learning. His published books include *Learning Amazon Web Services*, *AWS Certified Solutions Architect - Associate (SAA-C02) Cert Guide*, *Windows 2003 Registry for Dummies*, and *Administering SMS 3.0, Administering Active Directory*.

As a consultant, **Ryan Dymek** has been building, designing, and improving cloud solutions for some of the largest companies in the world for more than 10 years. Ryan began working with AWS cloud solutions in 2009 and in 2016 added GCP to his portfolio. Ryan transforms teams and organizations using DevOps principles and solid architectural approaches through a model of culture and “people first” philosophies. Ryan has trained more than 10,000 people over the past 6 years on topics such as DevOps engineering, security, cost optimization, performance tuning, and operational excellence.

# We Want to Hear from You!

As the reader of this book, *you* are our most important critic and commentator. We value your opinion and want to know what we're doing right, what we could do better, what areas you'd like to see us publish in, and any other words of wisdom you're willing to pass our way.

We welcome your comments. You can email or write to let us know what you did or didn't like about this book—as well as what we can do to make our books better.

*Please note that we cannot help you with technical problems related to the topic of this book.*

When you write, please be sure to include this book's title and author as well as your name and email address. We will carefully review your comments and share them with the author and editors who worked on the book.

Email: [community@informit.com](mailto:community@informit.com)

# Introduction

Welcome to *AWS Certified SysOps Administrator – Associate (SOA-C02) Exam Cram*. This book will help you get ready to take and pass the AWS Certified SysOps Administrator – Associate (SOA-C02) exam.

This book is designed to remind you of everything you need to know to pass the SOA-C02 certification exam. Each chapter includes a number of practice questions that should give you a reasonably accurate assessment of your knowledge, and, yes, we've provided the answers and their explanations for these questions. Read this book, understand the material, and you'll stand a very good chance of passing the real test.

*Exam Cram* books help you understand and appreciate the subjects and materials you need to know to pass AWS certification exams. *Exam Cram* books are aimed strictly at test preparation and review. They do not teach you everything you need to know about a subject. Instead, the authors streamline and highlight the pertinent information by presenting and dissecting the questions and problems they've discovered that you're likely to encounter on an AWS certification exam.

Let's begin by looking at preparation for the exam.

## How to Prepare for the Exam

This text follows the official exam objectives closely to help ensure your success. The AWS Certified SysOps Administrator – Associate exam covers 6 domains and 14 objectives, and this book is aligned with those domains and objectives. These official objectives from AWS can be found here:

[https://d1.awsstatic.com/training-and-certification/docs-sysops-associate/AWS-Certified-SysOps-Administrator-Associate\\_Exam-Guide.pdf](https://d1.awsstatic.com/training-and-certification/docs-sysops-associate/AWS-Certified-SysOps-Administrator-Associate_Exam-Guide.pdf)

As you examine the numerous exam topics now covered in the exam, resist the urge to panic! This book will provide you with the knowledge (and confidence) that you need to succeed. You just need to make sure you read it and follow the guidance it provides throughout your SysOps Administrator – Associate journey.

## Practice Tests

This book is filled with practice exam questions to get you ready!

- ▶ **CramSaver questions at the beginning of each major topic in each chapter:** These difficult, open-ended questions ensure you really know the material. Some readers use these questions to “test out” of a particular topic.
- ▶ **CramQuizzes at the end of each chapter:** These quizzes provide another chance to demonstrate your knowledge after completing a chapter.

In addition, the book comes with two full practice tests in the Pearson Test Prep software available to you either online or as an offline Windows application. To access the practice exams, please see the instructions in the card inserted in the sleeve in the back of the book. This card includes a unique access code that enables you to activate your exams in the Pearson Test Prep software.

If you are interested in more practice exams than are provided with this book, Pearson IT Certification publishes a Premium Edition eBook and Practice Test product. In addition to providing you with three eBook files (EPUB, PDF, and Kindle), this product provides you with two additional exams’ worth of questions. The Premium Edition version also offers you a link to the specific section in the book that presents an overview of the topic covered in the question, allowing you to easily refresh your knowledge. The insert card in the back of the book includes a special offer for an 80 percent discount off this Premium Edition eBook and Practice Test product, which is an incredible deal.

## Taking a Certification Exam

After you prepare for your exam, you need to register with a testing center. At the time of this writing, the cost to take the AWS Certified SysOps Administrator – Associate exam is \$150 USD for individuals. The test is administered by Pearson VUE testing centers with locations globally or as an online proctored exam.

You will have 180 minutes to complete the exam. The exam consists of a maximum of 65 questions. If you have prepared, you should find that this is plenty of time to properly pace yourself and review the exam before submission.

## Arriving at the Exam Location

As with any examination, arrive at the testing center early (at least 15 minutes). Be prepared! You need to bring two forms of identification (one with a picture). The testing center staff requires proof that you are who you say you are and that someone else is not taking the test for you. Arrive early, because if you are late, you will be barred from entry and will not receive a refund for the cost of the exam.

### ExamAlert

You'll be spending a lot of time in the exam room. Plan on using the full 180 minutes allotted for your exam and surveys. Policies differ from location to location regarding bathroom breaks. Check with the testing center before beginning the exam.

## In the Testing Center

You will not be allowed to take into the examination room study materials or anything else that could raise suspicion that you're cheating. This includes practice test material, books, exam prep guides, or other test aids. The Testing Center will provide you with scratch paper and a pen or pencil. These days, this often comes in the form of an erasable whiteboard.

Examination results are available after the exam. After submitting the exam, you will be notified whether you have passed or failed. The test administrator will also provide you with a printout of your results.

## About This Book

The ideal reader for an *Exam Cram* book is someone seeking certification. However, it should be noted that an *Exam Cram* book is a very easily readable, rapid presentation of facts. Therefore, an *Exam Cram* book is also extremely useful as a quick reference manual.

The book can be read cover to cover, or you may jump across chapters as needed. Because the book chapters align with the exam objectives, some chapters may have slight overlap on topics. Where required, references to the other chapters are provided for you. If you need to brush up on a topic, you can use the index, table of contents, or Table I.1 to find the topics and go to the questions that you need to study. Beyond helping you prepare for the test, we think you'll find this book useful as a tightly focused reference on some of the most important aspects of the AWS Certified SysOps Administrator – Associate (SOA-C02) certification.

This book includes other helpful elements in addition to the actual logical, step-by-step learning progression of the chapters themselves. *Exam Cram* books use elements such as ExamAlerts, tips, notes, and practice questions to make information easier to read and absorb. This text also includes a very helpful glossary to assist you.

**Note**

Reading this book from start to finish is not necessary; this book is set up so that you can quickly jump back and forth to find sections you need to study.

Use the *CramSheet* to remember last-minute facts immediately before the exam. Use the practice questions to test your knowledge. You can always brush up on specific topics in detail by referring to the table of contents and the index. Even after you achieve certification, you can use this book as a rapid-access reference manual.

## Exam Objectives

Table I.1 lists the skills the AWS Certified SysOps Administrator – Associate (SOA-C02) exam measures and the chapter in which the objective is discussed.

TABLE I.1

| Exam Domain                                  | Objective   | Chapter in Book That Covers It |
|--|---|--------------------------------|
| 1.0 Monitoring, Logging, and Remediation     | 1.1 Implement metrics, alarms, and filters by using AWS monitoring and logging services | Chapters 1, 2, 3               |
| 1.0 Monitoring, Logging, and Remediation     | 1.2 Remediate issues based on monitoring and availability metrics                       | Chapters 1, 2, 3               |
| 2.0 Reliability and Business Continuity      | 2.1 Implement scalability and elasticity  | Chapters 1, 4, 5, 6            |
| 2.0 Reliability and Business Continuity      | 2.2 Implement high availability and resilient environments                              | Chapters 1, 4, 5, 6            |
| 2.0 Reliability and Business Continuity      | 2.3 Implement backup and restore strategies   | Chapters 1, 4, 5, 6            |
| 3.0 Deployment, Provisioning, and Automation | 3.1 Provision and maintain cloud resources  | Chapters 1, 4, 7, 8            |
| 3.0 Deployment, Provisioning, and Automation | 3.2 Automate manual or repeatable processes   | Chapters 1, 4                  |

| Exam Domain                           | Objective   | Chapter in Book That Covers It |
|---------------------------------------|---|--------------------------------|
| 4.0 Security and Compliance           | 4.1 Implement and manage security and compliance policies   | Chapters 1, 2, 3, 6, 9         |
| 4.0 Security and Compliance           | 4.2 Implement data and infrastructure protection strategies | Chapters 1, 2, 3, 6, 10        |
| 5.0 Networking and Content Delivery   | 5.1 Implement networking features and connectivity          | Chapters 1, 4, 5, 11           |
| 5.0 Networking and Content Delivery   | 5.2 Configure domains, DNS services, and content delivery   | Chapters 1, 4, 5, 12           |
| 5.0 Networking and Content Delivery   | 5.3 Troubleshoot network connectivity issues                | Chapters 1, 4, 5, 13           |
| 6.0 Cost and Performance Optimization | 6.1 Implement cost optimization strategies                  | Chapters 1, 14                 |
| 6.0 Cost and Performance Optimization | 6.2 Implement performance optimization strategies           | Chapter 15                     |

## The Chapter Elements

Each *Exam Cram* book has chapters that follow a predefined structure. This structure makes *Exam Cram* books easy to read and provides a familiar format for all *Exam Cram* books. The following elements typically are used:

- ▶ Chapter topics
- ▶ CramSavers
- ▶ CramQuizzes
- ▶ ExamAlerts
- ▶ Notes
- ▶ Available exam preparation software practice questions and answers

### Note

Bulleted lists, numbered lists, tables, and graphics are also used where appropriate. A picture can paint a thousand words sometimes, and tables can help to associate different elements with each other visually.

Now let's look at each of the elements in detail.

- ▶ **Chapter topics**—Each chapter contains details of all subject matter listed in the table of contents for that particular chapter. The objective of an

*Exam Cram* book is to cover all the important facts without giving too much detail; it is an exam cram.

- ▶ **CramSavers**—Each chapter kicks off with a short-answer quiz to help you assess your knowledge of the chapter topic. This chapter element is designed to help you determine whether you need to read the whole chapter in detail or merely skim the material and skip ahead to the CramQuiz at the end of the chapter.
- ▶ **CramQuizzes**—Each chapter concludes with a multiple-choice quiz to help ensure that you have gained familiarity with the chapter content.
- ▶ **ExamAlerts**—ExamAlerts address exam-specific, exam-related information. An ExamAlert addresses content that is particularly important, tricky, or likely to appear on the exam. An ExamAlert looks like this:

#### ExamAlert

Make sure you remember the different ways in which you can access a router remotely. Know which methods are secure and which are not.

- ▶ **Notes**—Notes typically contain useful information that is not directly related to the current topic under consideration. To avoid breaking up the flow of the text, they are set off from the regular text.

#### Note

This is a note. You have already seen several notes.

## Other Book Elements

Most of this *Exam Cram* book on SysOps Administrator - Associate follows the consistent chapter structure already described. However, various important elements are not part of the standard chapter format. These elements apply to the book as a whole.

- ▶ **Glossary**—The glossary contains a listing of important terms used in this book with explanations.
- ▶ **CramSheet**—The CramSheet is a quick-reference, tear-out cardboard sheet of important facts useful for last-minute preparation. CramSheets often include a simple summary of the facts that are most difficult to remember.



- ▶ **Companion website**—The companion website for your book allows you to access several digital assets that come with your book, including
  - ▶ Pearson Test Prep software (both online and Windows desktop versions)
  - ▶ Key Terms Flash Cards application
  - ▶ A PDF version of the CramSheet

To access the book's companion website, simply follow these steps:

1. Register your book by going to: [PearsonITCertification.com/register](http://PearsonITCertification.com/register) and entering the ISBN: **9780137509584**.
2. Respond to the challenge questions.
3. Go to your account page and select the **Registered Products** tab.
4. Click the **Access Bonus Content** link under the product listing.

## Pearson Test Prep Practice Test Software

As noted previously, this book comes complete with the Pearson Test Prep practice test software containing two full exams. These practice tests are available to you either online or as an offline Windows application. To access the practice exams that were developed with this book, please see the instructions in the card inserted in the sleeve in the back of the book. This card includes a unique access code that enables you to activate your exams in the Pearson Test Prep software.

## Accessing the Pearson Test Prep Software Online

The online version of this software can be used on any device with a browser and connectivity to the Internet, including desktop machines, tablets, and smartphones. To start using your practice exams online, simply follow these steps:

1. Go to <http://www.PearsonTestPrep.com>.
2. Select **Pearson IT Certification** as your product group.
3. Enter your email/password for your account. If you don't have an account on [PearsonITCertification.com](http://PearsonITCertification.com) or [CiscoPress.com](http://CiscoPress.com), you will need to establish one by going to [PearsonITCertification.com/join](http://PearsonITCertification.com/join).

4. In the **My Products** tab, click the **Activate New Product** button.
5. Enter the access code printed on the insert card in the back of your book to activate your product.
6. The product will now be listed in your My Products page. Click the **Exams** button to launch the exam settings screen and start your exam.

## Accessing the Pearson Test Prep Software Offline

If you wish to study offline, you can download and install the Windows version of the Pearson Test Prep software. There is a download link for this software on the book's companion website, or you can just enter this link in your browser:

<http://www.pearsonitcertification.com/content/downloads/pcpt/engine.zip>

To access the book's companion website and the software, simply follow these steps:

1. Register your book by going to [PearsonITCertification.com/register](http://PearsonITCertification.com/register) and entering the ISBN: **9780137509584**.
2. Respond to the challenge questions.
3. Go to your account page and select the **Registered Products** tab.
4. Click the **Access Bonus Content** link under the product listing.
5. Click the **Install Pearson Test Prep Desktop Version** link under the Practice Exams section of the page to download the software.
6. After the software finishes downloading, unzip all the files on your computer.
7. Double-click the application file to start the installation, and follow the on-screen instructions to complete the registration.
8. When the installation is complete, launch the application and select the **Activate Exam** button on the My Products tab.
9. Click the **Activate a Product** button in the Activate Product Wizard.
10. Enter the unique access code found on the card in the sleeve in the back of your book and click the **Activate** button.
11. Click **Next** and then the **Finish** button to download the exam data to your application.

12. You can now start using the practice exams by selecting the product and clicking the **Open Exam** button to open the exam settings screen.

Note that the offline and online versions will synch together, so saved exams and grade results recorded on one version will be available to you on the other as well.

## Customizing Your Exams

Once you are in the exam settings screen, you can choose to take exams in one of three modes:

- ▶ Study Mode
- ▶ Practice Exam Mode
- ▶ Flash Card Mode

Study Mode allows you to fully customize your exams and review answers as you are taking the exam. This is typically the mode you would use first to assess your knowledge and identify information gaps. Practice Exam Mode locks certain customization options because it is presenting a realistic exam experience. Use this mode when you are preparing to test your exam readiness. Flash Card Mode strips out the answers and presents you with only the question stem. This mode is great for late stage preparation when you really want to challenge yourself to provide answers without the benefit of seeing multiple-choice options. This mode will not provide the detailed score reports that the other two modes will, so it should not be used if you are trying to identify knowledge gaps.

In addition to these three modes, you will be able to select the source of your questions. You can choose to take exams that cover all of the chapters, or you can narrow your selection to just a single chapter or the chapters that make up specific parts in the book. All chapters are selected by default. If you want to narrow your focus to individual chapters, simply deselect all the chapters and then select only those on which you wish to focus in the Objectives area.

You can also select the exam banks on which to focus.

There are several other customizations you can make to your exam from the exam settings screen, such as the time of the exam, the number of questions served up, whether to randomize questions and answers, whether to show the number of correct answers for multiple-choice questions, or whether to serve up only specific types of questions. You can also create custom test banks by selecting only questions that you have marked or questions on which you have added notes.

## Updating Your Exams

If you are using the online version of the Pearson Test Prep software, you should always have access to the latest version of the software as well as the exam data. If you are using the Windows desktop version, every time you launch the software, it will check to see if there are any updates to your exam data and automatically download any changes that were made since the last time you used the software. This requires that you are connected to the Internet at the time you launch the software.

Sometimes, due to many factors, the exam data may not fully download when you activate your exam. If you find that figures or exhibits are missing, you may need to manually update your exams.

To update a particular exam you have already activated and downloaded, simply select the **Tools** tab and select the **Update Products** button. Again, this is an issue only with the desktop Windows application.

If you wish to check for updates to the Pearson Test Prep exam engine software, Windows desktop version, simply select the **Tools** tab and select the **Update Application** button. This will ensure you are running the latest version of the software engine.

## Contacting the Authors

Hopefully, this book provides you with the tools you need to pass the AWS SysOps Administrator - Associate exam. Feedback is appreciated. You can follow and contact the authors on LinkedIn:

**Marko Sluga:**

<https://www.linkedin.com/in/markosluga/>

**Rick Crisci:**

<https://www.linkedin.com/in/rickcrisci/>

**William “Bo” Rothwell:**

<https://www.linkedin.com/in/bo-rothwell/>

*This page intentionally left blank*

## CHAPTER 4

# Implementing Scalability and Elasticity

**This chapter covers the following official AWS Certified SysOps Administrator - Associate (SOA-C02) exam domains:**

- ▶ Domain 2: Reliability and Business Continuity
- ▶ Domain 3: Deployment, Provisioning, and Automation
- ▶ Domain 5: Networking and Content Delivery

(For more information on the official AWS Certified SysOps Administrator - Associate [SOA-C02] exam topics, see the Introduction.)

Ensuring your application's infrastructure is scalable and elastic delivers a double benefit for the application. First, by adding more resources dynamically when required, you can adapt to any amount of traffic to ensure you do not leave any request unanswered. Second, by removing resources, you ensure the application is cost-effective when little or no requests are being received. However, designing a scalable/elastic application is not always an easy task.

In this chapter, we examine scaling, request offloading, and loose coupling as strategies that can enable an application to meet demand while maintaining cost-effectiveness. Ensuring your application is scalable and elastic also builds a good underlying foundation to achieve high availability and resilience, which we discuss in Chapter 5, "High Availability and Resilience."

# Scaling in the Cloud

This section covers the following official AWS Certified SysOps Administrator - Associate (SOA-C02) exam domains:

- ▶ Domain 2: Reliability and Business Continuity
- ▶ Domain 3: Deployment, Provisioning, and Automation

## CramSaver

If you can correctly answer these questions before going through this section, save time by skimming the Exam Alerts in this section and then completing the Cram Quiz at the end of the section.

1. You are operating a forum application that consists of three layers: a web front end on EC2, an application on EC2, and a database layer RDS. The web front end delivers the static forum content and formatting, the application layer stores the session information for each user, and the database layer stores all the user preferences. Assess the scalability of this deployment, identify any issues, and propose a solution.
2. You have been tasked with troubleshooting an image-processing platform. The application resides on a single-layer ECS container deployment that accepts requests for image processing from an incoming S3 bucket and deposits the processed image in an output S3 bucket. Lately, a spike in usage has caused the ECS application to reach its maximum scale. Users using the paid platform have reported that bulk image uploads complete successfully to S3, but some images are never processed. As a result, users are left searching for unprocessed images and need to resubmit them for processing. How could you ensure that the application works as intended?
3. Your developers have updated the forum application as per your previous comments. Your application is now growing, and you have been tasked with ensuring the application maintains scalability to millions of users. To assess the scalability, you have been given more information on the deployment. The web front end uses Apache2 HTTPS servers on Ubuntu Linux on EC2. The application layer runs custom Python code on EC2 that connects to a DynamoDB table to store session data. The database layer is deployed on a Multi-AZ RDS MySQL cluster with two instances (primary and secondary). Assess the scalability of this deployment and identify any potential bottlenecks.
4. After you optimize the application, your clients report highly improved performance. After receiving the latest AWS bill, the CFO has questions about additional cost of the application. While examining the cost report, you find that the application layer seems to still be deployed with a static number of servers like it was before the session state was offloaded. How can you further optimize the application to reduce the cost?

## Answers

1. Answer: Both the web and database layers are scalable. The application layer is limited in elasticity due to the persistence of the session data on the EC2 instances. Session data should be moved off the EC2 instances.
2. Answer: The bulk image uploads seem to exceed the capacity of the ECS cluster. The application needs to be redesigned with a buffer for the image-processing requests. Implementing a message queue service could offload the requests so that the back end can process them in a more predictable manner.
3. Answer: The only issue to identify is with the database layer. A Multi-AZ RDS deployment is only vertically scalable with an upper limit of the maximum size of the RDS instance. The maximum size of the instance could potentially bottleneck the whole forum application.
4. Answer: Find a good metric on which to scale the application layer and implement AWS Autoscaling. The application should shut down some of the instances when usage is low and power them on when traffic increases.

For any application to be made elastic and scalable, you need to consider the sum of the configurations of its components. Any weakness in any layer or service that the application depends on can cause the application scalability and elasticity to be reduced. Any reduction in scalability and elasticity can potentially introduce weaknesses in the high availability and resilience of your application. At the end of the day, any poorly scalable, rigid application with no guarantee of availability can have a tangible impact on the bottom line of any business.

The following factors need to be taken into account when designing a scalable and elastic application:

- ▶ **Compute layer:** The compute layer receives a request and responds to it. To make an application scalable and elastic, you need to consider how to scale the compute layer. Can you scale on a metric-defined amount of CPU, memory, and number of connections, or do you need to consider scaling at the smallest scale—per request? Also consider the best practice for keeping the compute layer disposable. There should never be any persistent data within the compute layer.
- ▶ **Persistent layer:** Where is the data being generated by the application stored? Is the storage layer decoupled from the instances? Is the same (synchronous) data available to all instances, or do you need to account for asynchronous platforms and eventual consistency? You should always ensure that the persistent layer is designed to be scalable and elastic and take into account any issues potentially caused by the replication configuration.



- ▶ **Decoupled components:** Are you scaling the whole application as one, or is each component or layer of the application able to scale independently? You need to always ensure each layer or section of the application can scale separately to achieve maximum operational excellence and lowest cost.
- ▶ **Asynchronous requests:** Does the compute platform need to process every request as soon as possible within the same session, or can you schedule the request to process it at a later time? When requests are allowed to process for a longer amount of time (many seconds, perhaps even minutes or hours), you should always decouple the application with a queue service to handle any requests asynchronously—meaning at a later time. Using a queue can enable you to buffer the requests, ensuring you receive all the requests on the incoming portion of the application and handle the processing with predictable performance on the back end. A well-designed, asynchronously decoupled application should almost never respond with a 500-type HTTP (service issue) error.

Assessing your application from these points of view should give you a rough idea of the scalability and elasticity of the platform. When you have a good idea of the scalability/elasticity, also consider any specific metrics within the defined service-layer agreement (SLA) of the application. After both are defined, assess whether the application will meet the SLA in its current configuration. Make a note if you need to take action to improve the scalability/elasticity and continuously reassess because both the application requirements and defined SLA of the application are likely to change over time.

After the application has been designed to meet the defined SLA of the application, you can make use of the cloud metrics provided in the platform at no additional cost to implement automated scaling and meet the demand in several different ways. We discuss how to implement automation in the AWS Autoscaling later in this section.

### ExamAlert

Remember, one of the crucial factors that enables scalability and elasticity is ensuring your resources are disposable. That means any data is always written outside the processing layer. All databases, files, logs, and any kind of output the application generates should always be decoupled from the processing layer. In the exam, different services might be part of the proposed solution. When analyzing the proposed answer, always ensure the data being output by the application doesn't stay on the EC2 instance for a longer time and that the solution has the lowest cost. For example, one answer could propose that logs be sent to CloudWatch logs via the CloudWatch agent, whereas another answer might propose logs be written to S3 every hour. Although the second solution will probably be more cost-effective, you should consider the first solution if the defined SLA of the application requires that all logs need to be captured.

## Horizontal vs. Vertical Scaling

The general consensus is that there are only two ways (with minor variance, depending on the service or platform) to scale a service or an application:

- ▶ Vertically, by adding more power (more CPU, memory, disk space, network bandwidth) to an application instance
- ▶ Horizontally, by adding more instances to an application layer, thus increasing the power by a factor of the size of the instance added

A great benefit to vertical scaling is that it can be deployed in any circumstance, even without any application support. Because you are maintaining one unit and increasing its size, you can vertically scale to the maximum size of the unit in question. The maximum scaling size of an instance is defined by the maximum size that the service supports. For example, at the time of writing, the maximum instance size supported on EC2 (u-24tb1.metal) offers 448 CPU cores and 24 TB (that's 24,576 GB!) of memory, while the smallest still-supported size (t2.nano) has only 1 CPU core and 0.5 GB of memory. Additionally, there are plenty of instance types and sizes to choose from, which means that you can horizontally scale an application the exact size you need at a certain moment in time. This same fact applies to other instance-based services such as EMR, RDS, and DocumentDB. Figure 4.1 illustrates vertical scaling of instances.

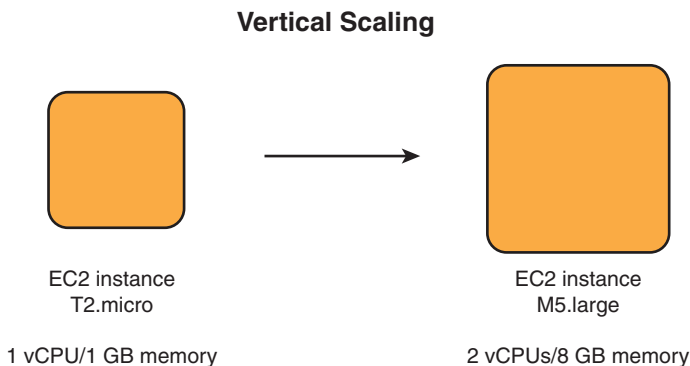


FIGURE 4.1 Vertical scaling

However, when thinking about scalability, you have to consider the drawbacks of vertical scaling. We mentioned maximum size, but one other major drawback is what makes horizontal scaling impractical—a single instance. Because all of AWS essentially operates on EC2 as the underlying design, you can take EC2 as a great example of why a single instance is not the way to go. Each

instance you deploy is deployed on a hypervisor in a rack in a datacenter, and this datacenter can only ever be part of one availability zone. In Chapter 1, “Introduction to AWS,” we defined an availability zone as a fault isolation environment—meaning any failure, whether it is due to a power or network outage or even an earthquake or flood, is always isolated to one availability zone. Although a single instance is vertically scalable, it is by no means highly available, nor does the vertical scaling make the application very elastic, because every time you scale to a different instance type, you need to reboot the instance.

This is where horizontal scaling steps in. With horizontal scaling, you add more instances (scale-out) when traffic to your application increases and remove instances (scale-in) when traffic to your application is reduced. You still need to select the appropriate scaling step, which is, of course, defined by the instance size. When selecting the size of the instances in a scaling environment, always ensure that they can do the job and don’t waste resources. Figure 4.2 illustrates horizontal scaling.




FIGURE 4.2 Horizontal scaling

In the ideal case, the application is stateless, meaning it does not store any data in the compute layer. In this case, you can easily scale the number of instances instead of scaling one instance up or down. The major benefit is that you can scale across multiple availability zones, thus inherently achieving high availability as well as elasticity. This is why the best practice on AWS is to create stateless, disposable instances and decouple the processing from the data and the layers of the application from each other. However, a potential drawback of horizontal scaling is when the application does not support it. The reality is that you will sooner or later come upon a case where you need to support an “enterprise” application, being migrated from some virtualized infrastructure,

that simply does not support adding multiple instances to the mix. In this case you can still use the services within AWS to make the application highly available and recover it automatically in case of a failure. For such instances you can now utilize the AWS EC2 Auto Recovery service, for the instance types that support it, which automatically re-creates the instance in case of an underlying system impairment or failure.

Another potential issue that can prevent horizontal scalability is the requirement to store some data in a stateful manner. Thus far, we have said that you need to decouple the state of the application from the compute layer and store it in a back-end service—for example, a database or in-memory service. However, the scalability is ultimately limited by your ability to scale those back-end services that store the data. In the case of the Relational Database Service (RDS), you are always limited to one primary instance that handles all writes because both the data and metadata within a traditional database need to be consistent at all times. You can scale the primary instance vertically; however, there is a maximum limit the database service will support, as Figure 4.3 illustrates.

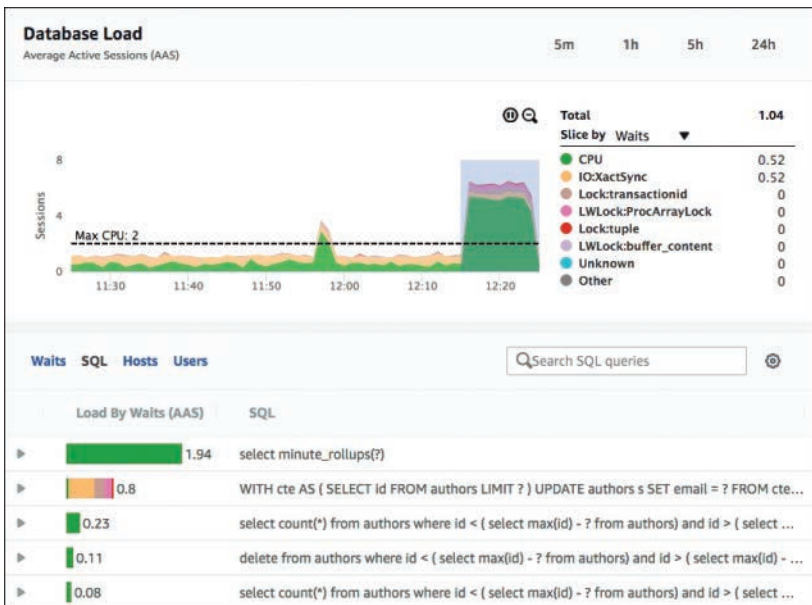


FIGURE 4.3 Database bottlenecking the application

You can also create a Multi-AZ deployment, which creates a secondary, synchronous replica of the primary database in another availability zone; however, Multi-AZ does not make the application more scalable because the replica is inaccessible to any SQL operations and is provided for the sole purpose of high

availability. Another option is adding read replicas to the primary instance to offload read requests. We delve into more details on read replicas later in this chapter and discuss database high availability in Chapter 5.

### ExamAlert

If the exam question is ambiguous about scalability or elasticity, you should still consider these as an important requirement of the application. Unless specific wording indicates cost as the primary or only driver, always choose the answer with the solution that will scale and is designed with elastic best practices in mind.

## AWS Autoscaling

After you design all the instance layers to be scalable, you should take advantage of the AWS Autoscaling service to automate the scale-in and scale-out operations for your application layers based on performance metrics—for example, EC2 CPU usage, network capacities, and other metrics captured in the CloudWatch service.

The AutoScaling service can scale the following AWS services:

- ▶ **EC2:** Add or remove instances from an EC2 AutoScaling group.
- ▶ **EC2 Spot Fleets:** Add or remove instances from a Spot Fleet request.
- ▶ **ECS:** Increase or decrease the number of containers in an ECS service.
- ▶ **DynamoDB:** Increase or decrease the provisioned read and write capacity.
- ▶ **RDS Aurora:** Add or remove Aurora read replicas from an Aurora DB cluster.

To create an autoscaling configuration on EC2, you need the following:

- ▶ **EC2 Launch template:** Specifies the instance type, AMI, key pair, block device mapping, and other features the instance should be created with.
- ▶ **Scaling policy:** Defines a trigger that specifies a metric ceiling (for scaling out) and floor (for scaling in). Any breach of the floor or ceiling for a certain period of time triggers autoscaling.
- ▶ **EC2 AutoScaling group:** Defines scaling limits and the minimum, maximum, and desired numbers of instances. You need to provide a launch configuration and a scaling policy to apply during a scaling event.

## Dynamic Scaling

Traditionally, scaling policies have been designed with dynamic scaling in mind. For example, a common setup would include

- ▶ An AutoScaling group with a minimum of 1 and a maximum of 10 instances
- ▶ A CPU % ceiling of 70 percent for scale-out
- ▶ A CPU % floor of 30 percent for scale-in
- ▶ A breach duration of 10 minutes
- ▶ A scaling definition of +/- 33 percent capacity on each scaling event

The application is now designed to operate at a particular scale between 30 and 70 percent aggregate CPU usage of the AutoScaling group. After the ceiling is breached for 10 minutes, the Autoscaling service adds a third more instances to the AutoScaling group. If you are running one instance, it adds another because it needs to meet 33 percent or more of the capacity. If you are running two instances, it also adds one more; however, at three instances, it needs to add two more instances to meet the rules set out in the scaling policy. When the application aggregate CPU usage falls below 30 percent for 10 minutes, the AutoScaling group is reduced by 33 percent, and the appropriate number of instances is removed each time the floor threshold is breached. Figure 4.4 illustrates dynamic scaling.




FIGURE 4.4 Dynamic scaling

## Manual and Scheduled Scaling

The AutoScaling configuration also has a desired instance count. This feature enables you to scale manually and override the configuration as per the scaling policy. You can set the desired count to any size at any time and resize the

AutoScaling group accordingly. This capability is useful if you have knowledge of an upcoming event that will result in an increase of traffic to your site. You can prepare your environment to meet the demand in a much better way because you can increase the AutoScaling group preemptively in anticipation of the traffic.

You can also set up a schedule to scale if you have a very predictable application. Perhaps it is a service being used only from 9 a.m. to 5 p.m. each day. You simply set the scale-out to happen at 8 a.m. in anticipation of the application being used and then set a scale-in scheduled action at 6 p.m. after the application is not being used anymore. This way you can easily reduce the cost of operating intermittently used applications by over 50 percent. Figure 4.5 illustrates scheduled scaling.

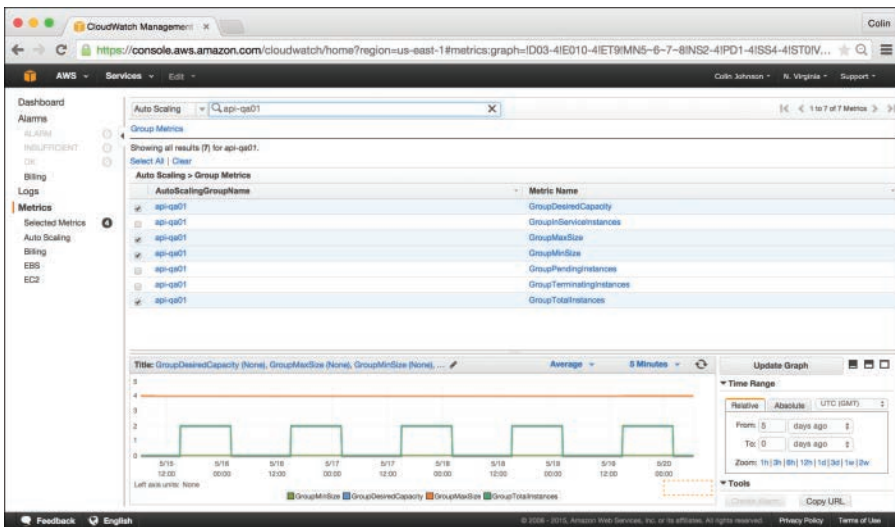


FIGURE 4.5 Scheduled scaling

## Predictive Scaling

Another AutoScaling feature is predictive scaling, which uses machine learning to learn the scaling pattern of your application based on the minimum amount of historical data. The machine learning component then predicts the scaling after reviewing CW data from the previous 14 days to account for daily and weekly spikes as it learns the patterns on a longer time scale. Figure 4.6 illustrates predictive scaling.




FIGURE 4.6 Predictive scaling

## Cram Quiz

Answer these questions. The answers follow the last question. If you cannot answer these questions correctly, consider reading this section again until you can.

- Which of the following are not characteristics of a scalable/elastic application?
  - A. Synchronous request handling in the compute layer
  - B. Session persistence in an external database
  - C. Session persistence in the compute layer
  - D. Asynchronous request offloading to a message queue
- Which of the following are required to enable the application to scale automatically with AWS AutoScaling? (Choose three.)
  - A. EC2 Launch Configuration
  - B. Scaling Policy
  - C. EC2 User Data
  - D. DynamoDB
  - E. CloudWatch Alarm
  - F. AutoScaling Group
- True or False: After you assess that your application is fully scalable and elastic, you only need to maintain the application as is in the cloud.
- True or False: AutoScaling supports only dynamic, scheduled, and predictive scaling.



## Cram Quiz Answers

1. Answer: C is correct. The compute layer should be made stateless. Any persistence in the compute layer hinders scalability and elasticity and potentially causes disruption in the application operation. If an instance in a cluster is lost, all the sessions on the instances are lost with it, meaning all the users connected to that particular instance have to log in and start working with the application from scratch.
  2. Answer: A, B, and F are correct. To create an autoscaling configuration on EC2, you need an EC2 Launch Configuration that defines how to configure the EC2 instances that are launched; a scaling policy that determines the scaling thresholds; and an autoscaling group that determines the minimum, maximum, and desired numbers of instances.
  3. Answer: False. The application should periodically be reassessed for scalability and elasticity because both the application requirements and the SLA might have changed.
  4. Answer: False. Autoscaling also supports manual scaling by setting the desired number of instances in the autoscaling group.
-

# Caching

This section covers the following official AWS Certified SysOps Administrator - Associate (SOA-C02) exam domains:

- ▶ Domain 2: Reliability and Business Continuity
- ▶ Domain 5: Networking and Content Delivery

## CramSaver

If you can correctly answer these questions before going through this section, save time by skimming the Exam Alerts in this section and then completing the Cram Quiz at the end of the section.

1. You have implemented autoscaling on both the web and app tier of your three-tier application, but in times of high read requests, the application seems to be performing slowly or even times out. What could you do to make the application more responsive?
2. You have been tasked with deploying a reliable caching solution that can handle multiple different data types and deliver microsecond to millisecond response performance. Which AWS service would you recommend?
3. True or False: To deliver static content to the user in the fastest possible manner, use a web server with lots of memory and utilize server-side caching.

## Answers

1. Answer: Implement the read cache to offload the database that seems to be bottlenecking the read requests.
2. Answer: ElastiCache Redis would support all the required features.
3. Answer: False. Static content should be delivered via a content delivery network (CDN). In AWS, you can use CloudFront to deliver static content through more than 200 geographically distributed locations across the globe.

Now that you have seen how to create an application that is highly scalable and elastic, you need to also consider the scaling impact on the persistent layer. As mentioned in the previous section, you need to consider the scalability of the persistent layer and include it in the overall assessment of the elasticity of the application. It serves no purpose to make an application highly elastic when the database back end is rigid and introduces a bottleneck for the whole application.

This is where caching comes in as a good way to ensure the data in the application is accessible in the fastest manner possible. When delivering an application from the cloud, you can use several different caching strategies to deliver and reuse frequently used data and offload the need to request the same data over and over again.

A simple analogy to caching is your refrigerator. It takes a minute for you to walk to the fridge (the cache) and grab some milk (cache hit). However, when there is no milk (cache miss) in the fridge, you need to go to the store (the origin), grab a few cartons, and take them home to put them in the fridge. The journey to the store and back is many times longer, so it makes sense to buy items that you frequently use and put them in the fridge. The fridge does have a limited capacity, just like cache usually does, and you will typically find a much greater variety of items at the store.

## Types of Caching

There are several different types of caching that you can use in your application.

### Client-Side Caching

When a client requests the contents of the application from a server, you should ensure that components that are static or change infrequently are reused with client-side caching. Modern browsers have this capability built in, and you can use it by specifying cache control headers within your web server or the service that delivers the content, such as S3.

### Edge Caching

When content is delivered frequently to multiple users, you can employ edge caching or what is more commonly referred to as a content delivery network. In AWS, you can use the Amazon CloudFront service to deliver frequently used content in a highly efficient manner to millions of users around the globe while at the same time offloading multiple same requests off the application or back end.

### Server-Side Caching

When a feature, a module, or certain content stored within the web service is requested frequently, you typically use server-side caching to reduce the need for the server to look for the feature on disk. The first time the feature is requested and the response assembled, the server caches the response in memory so it can be delivered with much lower latency than if it were read from

disk and reassembled each time. There is a limitation to the amount of memory the server has, and of course, server-side caching is traditionally limited to each instance. However, in AWS, you can use the ElastiCache service to provide a shared, network-attached, in-memory datastore that can fulfill the needs of caching any kind of content you would usually cache in memory.

## Database Caching

The last layer of caching is database caching. This approach lets you cache database contents or database responses into a caching service. There are two approaches to database caching:

- ▶ **In-line caching:** This approach utilizes a service that manages the reads and writes to and from the database.
- ▶ **Side-loaded caching:** This approach is performed by an application that is aware of the cache and database as two distinct entities. All reads and writes to and from the cache and the database are managed within the application because both the cache and database are two distinct entities.

An example of an in-line caching solution is the DynamoDB Accelerator (DAX) service. With DAX, you can simply address all reads and writes to the DAX cluster, which is connected to the DynamoDB table in the back end. DAX automatically forwards any writes to DynamoDB, and all reads deliver the data straight from the cache in case of a cache hit or forward the read request to the DynamoDB back end transparently. Any responses and items received from DynamoDB are thus cached in the response or item cache. In this case the application is not required to be aware of the cache because all in-line cache operations are identical to the operations performed against the table itself. Figure 4.7 illustrates DAX in-line caching.




FIGURE 4.7 DAX in-line caching

An example of a sideloaded caching solution is ElastiCache. First, you set up the caching cluster with ElastiCache and a database. The database can be DynamoDB, RDS, or any other database because ElastiCache is not a purpose-built solution like DAX. Second, you have to configure the application to look for any content in the cache first. If the cache contains the content, you get a cache hit, and the content is returned to the application with very low latency. If you get a cache miss because the content is not present in the cache, the application needs to look for the content in the database, read it, and also perform a write to the cache so that any subsequent reads are all cache hits.

There are two traditional approaches to implement sideloaded caching:

- ▶ **Lazy loading:** Data is always written only to the database. When data is requested, it is read from the database and cached. Every subsequent read of the data is read from the cache until the item expires. This is a lean approach, ensuring only frequently read data is in the cache. However, every first read inherently suffers from a cache miss. You can also warm up the cache by issuing the reads you expect to be frequent. Figure 4.8 illustrates lazy loading.

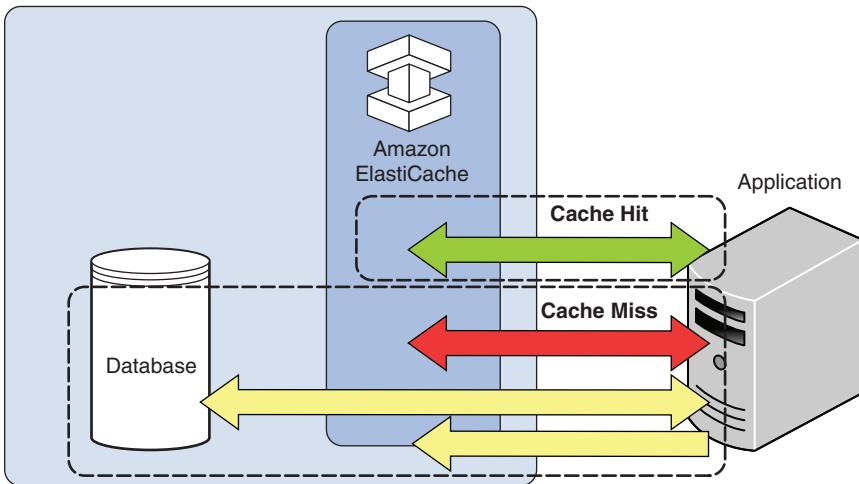


FIGURE 4.8 Lazy loading

- ▶ **Write through:** Data is always written to both the database and the cache. This avoids cache misses; however, this approach is highly intensive on the cache. Figure 4.9 illustrates write-through caching.

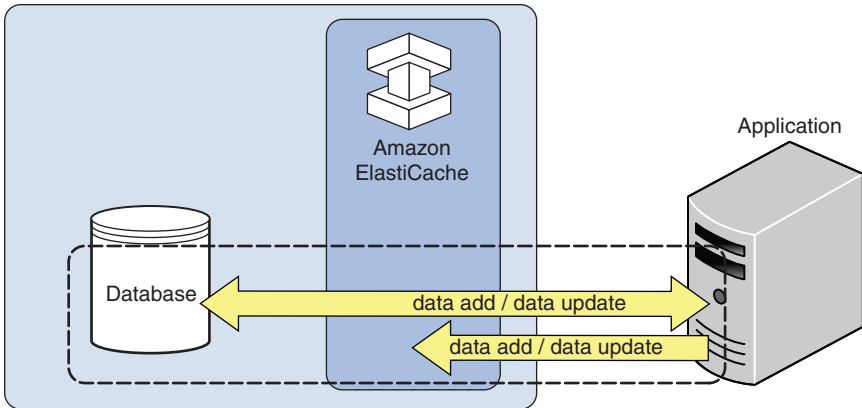


FIGURE 4.9 Write-through caching

Because the application controls the caching, you can implement some items to be lazy loaded but others to be written through. The approaches are in no way mutually exclusive, and you can write the application to perform caching based on the types of data being stored in the database and how frequently those items are expected to be requested.

### ExamAlert

When choosing a caching strategy, always consider the rate of data change and choose the correct time-to-live (TTL) of the data in the cache, to match the rate of change of the data. Data on an e-commerce site such as item descriptions, reviews, and images are unlikely to change frequently, but data such as item stock and price might not be suitable for caching at all. Always keep this in mind when selecting an answer on the exam.

## ElastiCache

ElastiCache is a managed service that can deploy clusters of in-memory data stores. They can be used to perform server-side and database caching.

A typical use case is database offloading with an application-managed side-loaded cache, as described in the previous section. Most applications that work with a database have a high read-to-write ratio—somewhere in the range of 80–90 percent reads to 10–20 percent writes. This means that offloading the reads from the database could save as much as 60–80 percent of the load on the database. For example, if the read-to-write ratio is 90–10 percent, each write (10 percent resources) requires a subsequent read (at least 10 percent

resources), and if the rest of the reads are cached, up to 80 percent of the read load can be offloaded to the cache.

An additional benefit of caching with ElastiCache is that the two engines used, Memcached and Redis, are both in-memory datastores. In comparison with databases where response latencies are measured in milliseconds to seconds, the response times from in-memory databases are measured in microseconds to milliseconds. That can mean that any cached data can be delivered 10, 100, or potentially even 1000 times faster than it would be if the request were read from the database. Figure 4.10 contrasts latency of Redis versus S3.

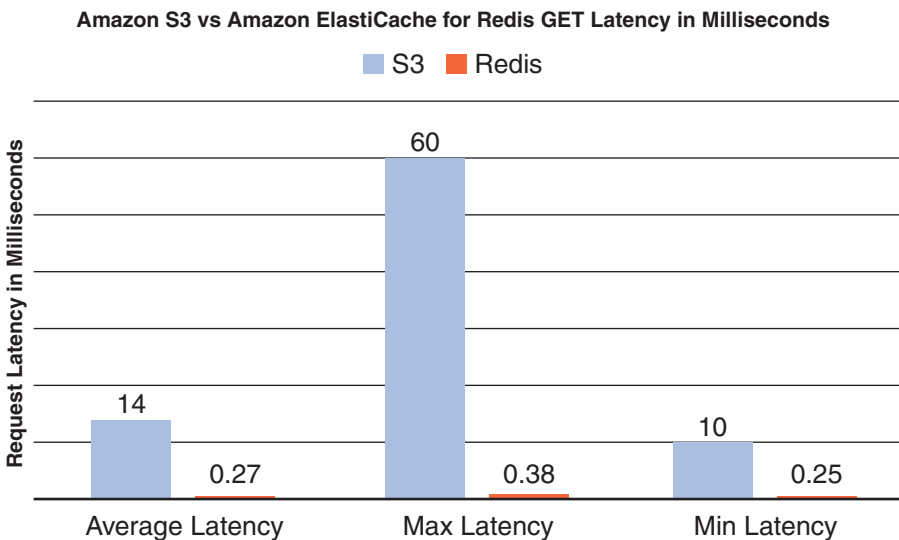


FIGURE 4.10 Redis vs. S3 GET latencies

## Memcached

One of the engines supported by ElastiCache is Memcached, a high-performance, distributed, in-memory key-value store. The service has a simple operational principle in which one key can have one or many nested values. All the data is served out of memory; thus, there is no indexing or data organization in the platform. When using Memcached, you can design either a single instance cache or a multi-instance cache where data is distributed into partitions. These partitions can be discovered by addressing the service and requesting a partition map. There is no resilience or high availability within the cluster because there is no replication of partitions.

Memcached is a great solution for offloading frequent identical database responses. Another use for Memcached is as a shared session store for multiple web instances. However, the lack of resilience within the Memcached design means that any failure of a node requires the application to rebuild the node from persistent database data or for the sessions with the users to be re-established. The benefit of Memcached is that it is linearly read and write scalable because all you need to do is add nodes to the cluster and remap the partitions.

## Redis

The other engine supported by ElastiCache is Redis, a fully-fledged in-memory database. Redis supports much more complex datasets such as tables, lists, hashes, and geospatial data. Redis also has a built-in push messaging feature that can be used for high-performance messaging between services and chat. Redis also has three operational modes that give you advanced resilience, scalability, and elasticity:

- ▶ **Single node:** A single Redis node, not replicated, non-scalable, and non-resilient. It can be deployed only in a single availability zone; however, it does support backup and restore.
- ▶ **Multinode, cluster mode disabled:** A primary read-write instance with up to five read replicas. The read replicas are near synchronous and offer resilience and read offloading for scalability. Multinode clusters with cluster mode disabled can also be deployed in one or more availability zones.
- ▶ **Multinode, cluster mode enabled:** One or more shards of multinode deployments, each with one primary and up to five read replicas. By enabling cluster mode, you retain the resilience and scalability but also add elasticity because you can always reshard the cluster and horizontally add more write capacity. Cluster mode always requires the database nodes to be deployed across multiple availability zones. However, sharding means that multiple primary nodes are responsible for multiple sets of data. Just like with database sharding, this increases write performance in ideal circumstances, but it is not always guaranteed and can add additional complexity to an already-complex solution. Figure 4.11 illustrates Redis multinode cluster mode, both disabled and enabled.



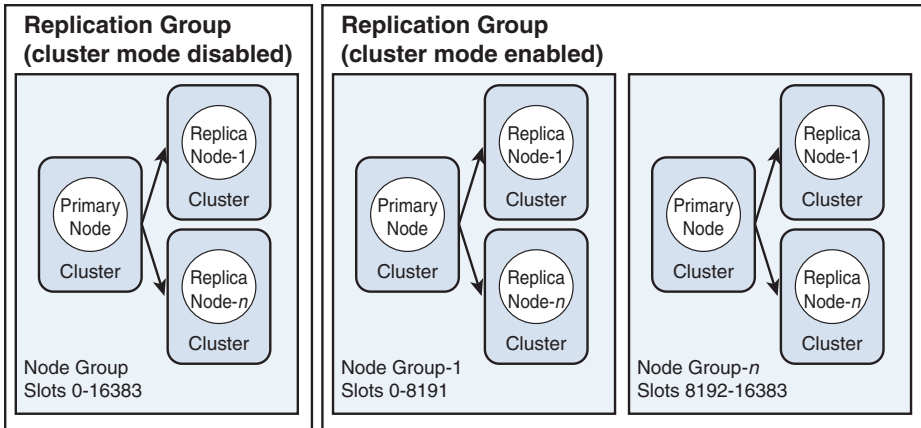


FIGURE 4.11 Redis multinode, cluster mode disabled vs. enabled

The features of Redis give you much more than just a caching service. Many businesses out there use ElastiCache Redis as a fully functional in-memory database because the solution is highly resilient, scalable, elastic, and can even be backed up and restored, just like a traditional database.

### ExamAlert

If an exam question indicates the database is overloaded with reads, caching should be the main strategy to enable scalability of reads. The benefit of caching is that the data in the cache, when configured correctly, is highly likely to be current and synchronous with the database. However, always be careful when scaling the caching cluster because too many caching instances can add additional unnecessary cost to the application.

## Amazon CloudFront

Now that the database and server-side caching are sorted, you need to focus on edge caching. CloudFront is a global content delivery network that can offload static content from the data origin and deliver any cached content from a location that is geographically much closer to the user. This reduces the response latency and makes the application feel as if it were deployed locally, no matter which region or which physical location on the globe the origin resides in.

CloudFront also can terminate all types of HTTP and HTTPS connections at the edge, thus offloading the load to servers or services. CloudFront also works in combination with the Amazon Certificate Manager (ACM) service, which

can issue free HTTPS certificates for your domain that are also automatically renewed and replaced each year.

CloudFront enables you to configure it to accept and terminate the following groups of HTTP methods:

- ▶ **GET and HEAD:** Enables standard caching for documents and headers. Useful for static websites.
- ▶ **GET, HEAD, and OPTIONS:** Enables you to cache OPTIONS responses from an origin server.
- ▶ **GET, HEAD, OPTIONS, PUT, PATCH, POST, and DELETE:** Terminates all HTTP(S) sessions at the CloudFront Edge Location and can increase the performance of both the read and write requests.

One of the better features is the ability to rewrite the caching settings for both the cache held within the CloudFront CDN as well as the client-side caching headers defined within servers. This means that you can customize the time-to-live of both the edge and client-side cache. CloudFront distributions can be configured with the following options for setting TTL:

- ▶ **Min TTL:** Required setting when all HTTP headers are forwarded from the origin server. It defines the minimum cache TTL and also determines the shortest interval for CloudFront to refresh the data from the origin.
- ▶ **Max TTL:** Optional setting. It defines the longest possible cache TTL. It is used to override any cache-control headers defined at the origin server.
- ▶ **Default TTL:** Optional setting. It allows you to define cache behavior for any content where no TTL is defined at the origin server.

## CloudFront Security

CloudFront is also inherently secure against distributed denial-of-service (DDoS) attacks because the content is distributed to more than 200 locations around the globe. An attacker would need to have a massive, globally distributed botnet to be able to attack your application. On top of the benefit of the distributed architecture, CloudFront is also resilient to L3 and L4 DDoS attacks with the use of AWS Shield Standard service. Any CloudFront distribution can also be upgraded to Shield Advanced, a subscription-based service that provides a detailed overview of the state of any DDoS attacks as well as a dedicated 24/7 response team, which can help with custom DDoS mitigation at any OSI layer.

CloudFront also supports integration with the AWS Web Application Firewall (WAF) service that can filter any attempts at web address manipulations, SQL injections, CSS attacks, and common web server vulnerabilities as well as filter traffic based on IP, geography patterns, regular expressions, and methods.

CloudFront also enables you to limit access to content by

- ▶ Restricting access to your application content with signed URLs or cookies
- ▶ Restricting access to content based on geolocation
- ▶ Restricting access to S3 buckets using Origin Access Identity (OAI)

---

## Cram Quiz

Answer these questions. The answers follow the last question. If you cannot answer these questions correctly, consider reading this section again until you can.

1. Which AWS service enables you to easily deploy a horizontally scalable in-memory caching cluster?
  - A. ElastiCache Memcached
  - B. ElastiCache Redis, Cluster mode enabled
  - C. ElastiCache Redis, Cluster mode disabled
  - D. CloudFront
  
2. You have configured a CloudFront distribution to cache static content from an Apache2 web server. The content on the web server is refreshed every 15 minutes when the application is updated. However, the users are complaining that they seem to see updates only every 2 hours or so. What is most likely the problem, and how would you resolve this issue?
  - A. CloudFront TTL is too long. Set the Min TTL to 15 minutes. This will ensure the content is refreshed every 15 minutes.
  - B. Origin TTL is too long. Set the Max TTL to 15 minutes. This will ensure the content is refreshed every 15 minutes.
  - C. Origin TTL does not exist. Set the Default TTL to 15 minutes. This will ensure the content is refreshed every 15 minutes.
  - D. CloudFront TTL does not exist. Set the TTL to enabled and Default TTL to 15 minutes. This will ensure the content is refreshed every 15 minutes.

## Cram Quiz Answers

1. Answer: B is correct. ElastiCache Memcached is a high-performance, distributed, in-memory key-value store that can scale horizontally.
  2. Answer: B is correct. Max TTL defines the longest possible cache TTL and is used to override any cache-control headers defined at the origin server that are likely to be misconfigured.
-

# Read Replicas

This section covers the following official AWS Certified SysOps Administrator - Associate (SOA-C02) exam domain:

- ▶ Domain 2: Reliability and Business Continuity

## CramSaver

If you can correctly answer these questions before going through this section, save time by skimming the Exam Alerts in this section and then completing the Cram Quiz at the end of the section.

1. Your three-tier application has been connected to a business intelligence (BI) forecasting platform. While the forecasts are improving business practices, the users of your application are reporting the performance has decreased. The web and app tier are scaling appropriately, and the caching cluster is at about 40 percent capacity. What could be the cause of the slowdown seen by the users, and how could you resolve it?
2. True or False: Aurora natively supports both MySQL and PostgreSQL.

## Answers

1. Answer: The BI platform has introduced additional load on the database. Because the BI forecasting requires access to most or all of the dataset, the cache cannot be used to offload the required reads. To mitigate, implement a database read replica.
2. Answer: True. The two engines are fully supported at the time of writing. Other database engines might be supported in the future.

There are five general approaches to scaling database performance:

- ▶ **Vertical scaling:** You can add more CPU and RAM to the primary instance.
- ▶ **Horizontal scaling:** You can add more instances to a database cluster to increase the available CPU and RAM, but this approach is not always supported.

- ▶ **Sharding:** When horizontal scaling is not supported, you can distribute the dataset across multiple primary database engines, thus achieving higher write performance.
- ▶ **Caching:** You can add a caching cluster to offload reads, which can be expensive.
- ▶ **Read replicas:** You can add read replicas to offload read traffic, possibly asynchronous.

Read replicas can potentially offer the same benefit of read offloading that caching can have for your application. One big benefit of read replicas is that the whole database is replicated to the read replica, not just frequently read items. This means that read replicas can be used where the reads are frequently distributed across the majority of the data in the database. As you can imagine, this would be very expensive and resource intensive to achieve in the cache. You would need to provision enough in-memory capacity (potentially terabytes), read the entire database, and write the data into the cache before you could perform any complex operation on the data in the cache. Sometimes complex operations can't even be performed on the caching server; for example, joins of multiple tables for analytics, business intelligence, or data mining purposes are just not possible within the cache.

This is where read replicas excel. You can introduce read replicas in most instance-based AWS database services, including RDS, Aurora, DocumentDB, and Neptune.

In the RDS service, up to five read replicas can be deployed in any MySQL, MariaDB, PostgreSQL, or Oracle database. Because the data resides on the volume attached to the instance, built-in database replication tools are used to replicate the data across the network to the read replica. This means that read replicas introduce additional load on the primary instance and that the replication is always asynchronous with a potential lag of a few seconds to potentially a few minutes in extreme cases. Figure 4.12 illustrates RDS read replicas.

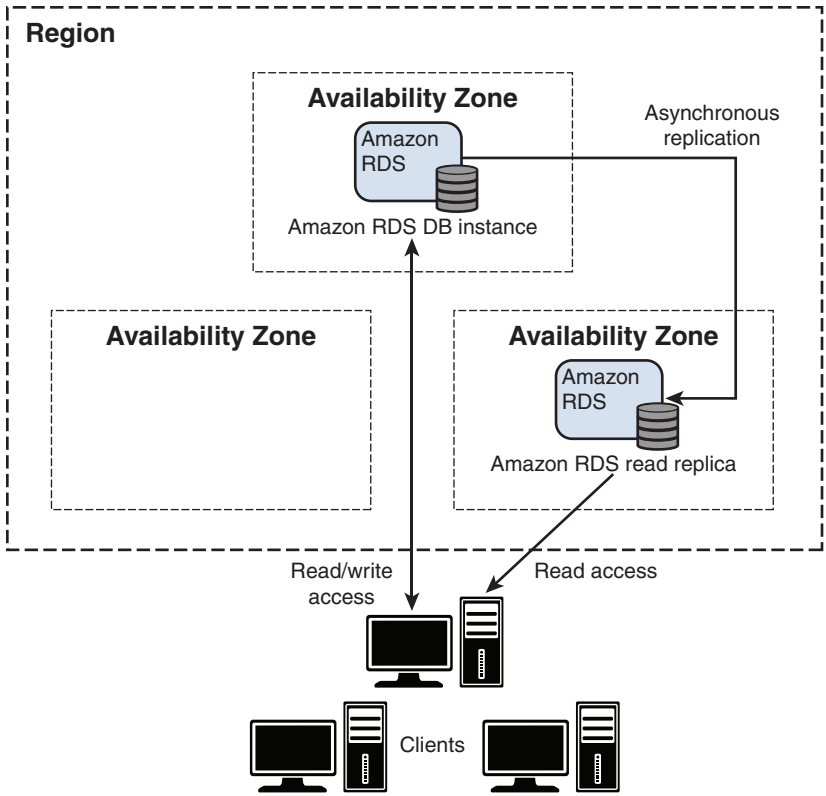


FIGURE 4.12 RDS read replicas

The replica can be placed in any availability zone in the same region as the primary database or can be deployed in a cross-region deployment. A read replica can also be promoted to a primary instance, which means that establishing a read replica can be an easy way to clone a database. After the replica is promoted to primary only, a sync of all the missing data is required.

**ExamAlert**

Traditional RDS database read replicas are a very cost-efficient way to provide data for small (terabyte)-scale analytics and should be selected as a preferred option when the entire dataset needs to be read offloaded. Always evaluate whether the question requires you to offload a certain portion of the data (with caching) or the entire dataset is required (with read replicas).

Other AWS database services employ a decoupled compute–datastore approach. For example, the AWS Aurora service stores all the data on a cluster volume that is replicated to six copies across (at least) three availability zones. The primary instance has read and write access. All writes are sent as changelogs directly to the six nodes of the storage volume, and the commit to the database page is done at the storage cluster. This means that the replication is near synchronous with potentially no more than a few milliseconds of replication lag due to the network distance between the cluster volume nodes in other availability zones and several hundred milliseconds if the replication is done across regions. Additionally, Aurora supports up to 15 read replicas per region and can be deployed to multiple regions with an additional 16 read replicas in other regions. All of the read replicas read from the same cluster volume, meaning they can deliver near synchronous data for each read request. The read replicas can also be seamlessly scaled because there is no requirement to replicate the data on to the replica. When a replica is started, it is simply connected to the cluster volume. This allows the cluster to avoid initial replication lags that you would see in traditional databases. This feature can be used both for elasticity and vertical as well as horizontal scaling of the cluster. Figure 4.13 illustrates an Aurora cluster design.

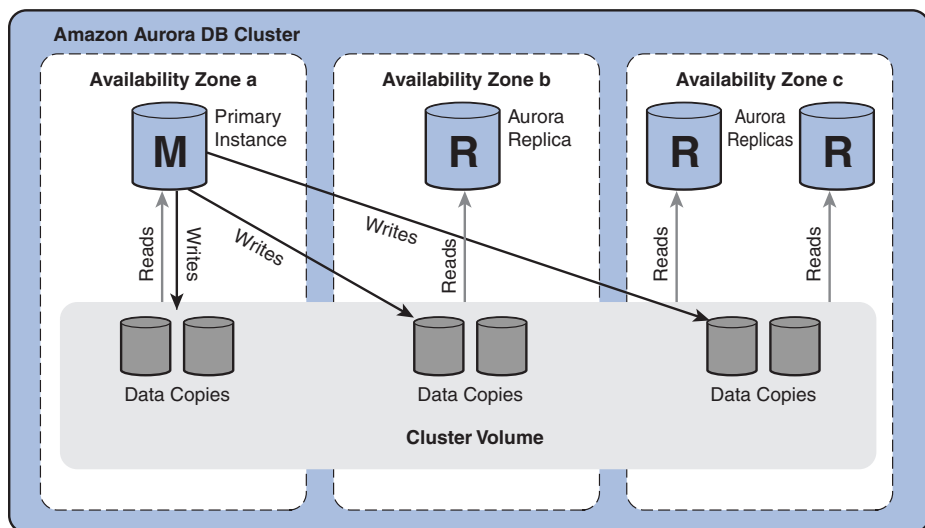


FIGURE 4.13 Aurora cluster design

A similar approach to decoupling the storage volume and the database nodes is used in DocumentDB, Neptune, and so on.



On top of this capability, Aurora can forward some read-heavy analytical operations to the cluster volume nodes and offload the primary from having to perform the reads for any JOIN, GROUP BY, UNION, and such operations across approximately a million or more rows. This capability extends the lightweight analytics capabilities of RDS read replicas to deliver much more power to your analytics queries.

Aurora natively supports MySQL and PostgreSQL databases and is slightly more expensive than RDS MySQL and RDS PostgreSQL.

### ExamAlert

When an exam question indicates that high performance at high scale is required for MySQL or PostgreSQL, always evaluate whether the answer includes Aurora as an option. We recommend you select RDS MySQL or RDS PostgreSQL only if any cost considerations are explicitly stated in the question. Also, consider reducing the number of read replicas by powering them off or terminating them if they are not in use.

---

## Cram Quiz

Answer these questions. The answers follow the last question. If you cannot answer these questions correctly, consider reading this section again until you can.

1. You have been instructed to eliminate any inefficiencies in the following deployment:

**Web tier:** EC2 autoscaling group scaling on CPU usage, 30% floor, 70% ceiling, minimum 1, maximum 10.

**App tier:** EC2 autoscaling group scaling on CPU usage, 20% floor, 80% ceiling, minimum 1, maximum 6.

**Cache tier:** 1 ElastiCache Memcached cluster with 2 partitions

**Database tier:** Multi-AZ MySQL RDS with 3 read replicas

Which of the following would allow you to cost optimize the cluster? (Choose all that apply.)

- A. Evaluate whether all RDS read replicas are required.
- B. Evaluate the maximum on the App and Web EC2 autoscaling groups.
- C. Evaluate the partition configuration of the ElastiCache cluster.
- D. Evaluate the network performance if the app tier matches the network performance of the database tier.

2. You have been asked to propose a solution to scale the read portion of an application database back end in the most cost-effective manner possible. Your application runs Linux LAMP stack with a 30 GB MySQL RDS Multi-AZ back end. The read-heavy operations are very predictable and occur during a particularly heavy four-hour operation each week and require the whole dataset. Which option would you recommend?
- A. Replace the RDS MySQL with Aurora MySQL and let Aurora add read replicas automatically as needed.
  - B. Point the read-heavy operation at the RDS Multi-AZ replica in the other availability zone.
  - C. Deploy a read replica in RDS. Point the application at the RDS read replica. Terminate the read replica after the read-heavy operation is complete. Repeat the process with a script next week.
  - D. Deploy a read replica in RDS. Point the application at the RDS read replica. Create a script that powers off the read replica after the read-heavy operation is complete and powers it on before the next operation with enough time to synchronize the changes.

## Cram Quiz Answers

1. Answer: A and C are correct. You should always evaluate whether the read replicas and caching clusters you deployed are required and properly scaled.
2. Answer: C is correct. Although A and D are possible solutions, deploying a read replica on a weekly basis would be the most cost-effective way to achieve this goal because the operation is sparse in nature and lasts for only four hours. Deploying a read replica is performed from a snapshot and can be just as fast or even faster than replicating a week's worth of data to the powered-off replica. Although the charges of powered-off instances are reduced, they are not zero. Remember also that you can stop a DB instance for up to seven days. If you don't manually start your DB instance after seven days, your DB instance is automatically started so that it doesn't fall behind any required maintenance updates. B is impossible because a Multi-AZ replica is not accessible for any SQL operations.

---

## What Next?

If you want more practice on this chapter's exam objectives before you move on, remember that you can access all of the Cram Quiz questions on the Pearson Test Prep software online. You can also create a custom exam by objective with the Online Practice Test. Note any objective you struggle with and go to that objective's material in this chapter.

*This page intentionally left blank*

# Index

## Numbers

---

- 400 - bad request errors, 45
- 403 - access denied errors, 45
- 404 - page not found errors, 45
- 500 - internal failure errors, 45
- 503 - service unavailable errors, 45

## A

---

### access

- 403 - access denied errors, 45
- access keys, 163–164
- AWS, 11
- broad network access, 4
- ELB access logs, 245–246
- failed access attempts, troubleshooting/  
remediation, 45
- IAM, 19
  - Access Analyzer, 166
  - AWS credentials, 157
  - best practices, 165–167
  - components, 156–157
  - IDP, 161
  - password policies, 162
  - permissions, 157
  - policies, 159–160
  - roles, 158–159
  - tags, 158
  - user accounts, 157–158
  - user groups, 158
- MFA, 164–165
- NACL, VPC, 194–195
- RAM, provisioning, 134–135
- Trusted Advisor, 168

### account management

- access keys, 163–164
- IAM
  - Access Analyzer, 166

## account management

- AWS credentials, 157
- best practices, 165–167
- components, 156–157
- IDP, 161
- password policies, 162
- permissions, 157
- policies, 159–160
- roles, 158–159
- tags, 158
- user accounts, 157–158
- user groups, 158

MFA, 164–165

Trusted Advisor, 168

**ACID requirements, databases, 103–104****ACM (AWS Certificate Manager), 177–178****actions, IAM policies, 160****active-passive upgrades, 142****Aggregation phase, performance monitoring, 275****alarms**

- billing alarms, 267
- CloudWatch, 33
- performance monitoring, 275

**ALB (Application Load Balancing), 95****alerting, 27****alias records, 220****allocation tags, cost, 260–262****Amazon Inspector, 19, 180–181****AMI (Amazon Machine Image), provisioning, 129****analytics**

- performance monitoring, 275
- services, 21

**API Gateway, 21****appliances, virtual, patches, cloud components, 149****applications**

- asynchronous requests, scalability/elasticity, 58
- building, 140
  - canary releases, 141
  - LTS, 141
  - stable builds, 140–141
- database bottlenecks, 61

decoupled components, scalability/elasticity, 58

deployments, 138

- blue/green deployments, 139–140

- development environments, 139

- DR, 140

- OpsWorks, 145–146

- production environments, 139

- QA environments, 139

- staging environments, 139

end-user applications services, 19

lifecycle management

- deprecations (end of life), 144

- migrations, 143–144

- roadmaps, 142–143

- upgrades, 143–144

- versioning, 143

patches, 147

- cloud components, 150

- hot fixes, 148

- scheduled updates, 148

- security patches, 147–148

- signature updates, 148

- virtual patches, 148

persistent layer, scalability/elasticity, 57

scalability/elasticity, compute layer, 57

services, 21

SLA, scalability/elasticity, 58

software vs, 151

troubleshooting/remediation, 44

upgrades

- active-passive upgrades, 142

- rolling upgrades, 142

uptime, defining, 87

**assuming roles from unknown locations, troubleshooting/remediation, 45****asymmetric encryption, 173****asynchronous requests, scalability/elasticity, 58****Athena, 21****Aurora**

- high availability, 105–106

- read replicas, 81–82

**authentication, MFA, 164–165****automation, Systems Manager, 49–50****AutoScaling, 62**

- dynamic scaling, 63
- manual scaling, 63–64
- predictive scaling, 64–65
- scheduled scaling, 64

**availability, high, 86**

- ALB, 95
- Aurora, 105–106
- availability zones, 88–89
- cross-region replication, 89
- databases, 103–104
- datastores, 98–101
- DynamoDB, 106–107
- ELB, 94–97
- GLB, 96
- multi-A high availability, 89
- multiregion high availability, 89
- NLB, 95–96
- RDS, 104–105
- Route 53, 96
- single AZ high availability, 89
- uptime
  - applications, defining in, 87
  - percentage chart, 86–87
- VPC, 89, 91
  - host addresses, 90
  - network addresses, 90
  - private subnets, 91–92
  - public subnets, 91–92
  - subnet ranges, 90–91

**availability zones, 9, 88–89****AWS (Amazon Web Services). See also cloud computing**

- access, 11
- availability zones, 9
- benefits of, 13–15
- CLI
  - command syntax, 128
  - configuring, 127–128
  - installing, 127
  - provisioning, 127–128

- subcommands, 128
- compute services, 18
- datacenters, 9
- development of, 8–9
- edge locations, 10–11
- foundation services, 17–19
- global architecture, 9–11
- management and deployment services, 22
- Management Console, 126
- Marketplace, 198
- Network Firewall, 197–198
- network services, 17
- OpsWorks, 22, 145–146
- Pinpoint, 21
- platform services, 20–22
- RAM, provisioning, 134–135
- regions, 10
- SageMaker, 22
- SDK, 126
- Secrets Manager, 178–179
- security and identity services, 19
- Security Hub, 181
- Shield, 214–215
- Site-to-Site VGW, 17
- Snow Family, 18
- Storage Gateway, 18
- storage services, 18
- Systems Manager, 22
  - automation, 49–50
  - Explorer, 134
  - Incident Manager, 134
  - OpsCenter, 134
  - provisioning, 133–134
- Trusted Advisor, 168
- VPN management, 206–207, 257
- WAF, 19, 212–213, 248–249
- Well-Architected Framework, 272–273
  - Operational Excellence pillar, 273–274
  - Performance Efficiency pillar, 273–274

**B****backups, 109, 115**

- cloud backups, 110–111
- DataSync, 121
- DR, 114
- Glacier, 122–123
- life-cycling rules, 121–122
- RPO, 113–114
- RTO, 113–114
- S3, 119
  - cross-region replication, 123
  - S3 sync, 120–121
- snapshots, 112–113
- stateful services, 111
- strategies
  - comparing, 116–117
  - multisite active-active backup/restore strategy, 116
  - overview, 115
  - pilot light backup/restore strategy, 115–116
  - warm standby backup/restore strategy, 116
  - versioning, 121–122

**benefits of AWS, 13–15****billing alarms, 267****blue/green deployments, 139–140****bottlenecking applications, 61****broad network access, 4****bucket policies, OAI, 236****Budgets, 266–267****bug fixes, 149****building applications, 140**

- canary releases, 141
- LTS, 141
- stable builds, 140–141

**C****CA (Certificate Authority), 176****caching, 67–68, 74, 79**

- client-side caching, 68
- CloudFront, 74–76, 253–254
- database caching, 69–71

- edge caching, 68, 74–76
- ElastiCache, 71–74
- in-line caching, 69–70
- server-side caching, 68–69
- side-loaded caching, 69–71
- write-through caching, 70–71

**canary releases, 141****CapEx (Capital Expenses), 13****certificate management, 175–176**

- ACM, 177–178
- CA, 176
- CRL, 177
- CSR, 177
- digital certificates, 176
- PKI, 176
- private keys, 177
- public keys, 176

**CIDR ranges, VPC, 185****classifying data, 174, 175**

- confidential data, 174
- internal data, 174
- public data, 174
- restricted data, 174

**CLI (Command-Line Interface)**

- configuring, 127–128
- installing, 127
- provisioning, 127–128
- subcommands, 128

**client-side caching, 68****cloud computing. See also AWS**

- backups, 110–111
- broad network access, 4
- community cloud deployments, 4
- defined, 2–7
- EC2, 18
- elasticity, 56–57
  - asynchronous requests, 58
  - compute layer, 57
  - decoupled components, 58
  - disposable resources, 58
  - persistent layer, 57
- features of, 4
- hybrid cloud deployments, 3

- IaaS, 4
- measured service, 4
- on-demand self-service, 4
- PaaS, 5
- patches, 149
  - bug fixes, 149
  - firmware, 150
  - N-1, 151–152
  - networking components, 149–150
  - policies, 151
  - rollbacks, 152
  - software, 151
  - storage components, 150
  - virtual appliances, 149
  - VM, 149
- private cloud deployments, 3
- public cloud deployments, 3
- rapid elasticity, 4
- resource pooling, 4
- responsibility models, 5
- SaaS, 5
- scalability, 56–57
  - asynchronous requests, 58
  - compute layer, 57
  - decoupled components, 58
  - disposable resources, 58
  - persistent layer, 57
- visual model of cloud computing, 5–6
- VPC, 17, 89, 91, 184
  - CIDR ranges, 185
  - connectivity, 189–192
  - EIP addresses, 186
  - ENI, 185–186
  - host addresses, 90
  - IGW, 187–188
  - NAT gateways, 191–192
  - network addresses, 90
  - private subnets, 91–92
  - public subnets, 91–92
  - route tables, 190
  - security, 193–196
  - subnet ranges, 90–91
  - subnets, 189–190
- CloudFormation, 22**
  - Elastic Beanstalk, 130–132
  - provisioning, 129–133
  - StackSets, 132–133
- CloudFront, 17, 74–76, 231**
  - caching
    - hit rates, 253–254
    - origin status codes, 254
  - configuring, 233
  - creating distributions, 231–232
  - HTTP, 233
  - logs, 251–252
  - price classes, 232–233
  - TTL, 232, 233
- CloudHSM, 19**
- CloudTrail, 22, 28, 36–37**
- CloudWatch, 22, 28, 30–31**
  - alarms, 33
  - dimensions, 31–32
  - logs, 34
  - Logs Insights, 34
  - metrics, 32–33
  - namespaces, 31–32
  - percentiles, 33
- clusters, placement groups, 278**
- Cognito, 22**
- cold datacenters, 114**
- commands**
  - CLI
    - command syntax, 128
    - subcommands, 128
  - ec2 commands, 128
- community cloud deployments, 4**
- compliance**
  - account management
    - access keys, 163–164
    - IAM, 156–162, 165–167
    - MFA, 164–165
    - Trusted Advisor, 168
  - monitoring services, 27–28



components (decoupled), scalability/elasticity

**components (decoupled), scalability/elasticity, 58**

**compute layer, scalability/elasticity, 57**

**Compute Optimizer, 263–265**

**compute performance, optimizing, 275–276**

instance stores, 276–277

placement groups, 277–278

**compute services, 18**

**confidential data, 174**

**Config, 22, 51–53**

**configuring**

CLI, 127–128

CloudFront, 233

flow logs, VPC, 240–243

IAM IDP, 161

**connectivity**

dedicated connections, 210

Direct Connect, 17

hosted connections, 210

troubleshooting

AWS-managed VPN, 257

CloudFront, 251–254

Direct Connect, 256–257

ELB access logs, 245–246

flow logs, VPC, 240–243

hybrid links, 256–257

private links, 256–257

WAF ACL logs, 240–243

VPC

Direct Connect, 209–210

subnets, 189–190

**containers, ECS, 18**

**content delivery**

CloudFront, 231

configuring, 233

creating distributions, 231–232

HTTP, 233

price classes, 232–233

TTL, 232, 233

S3

OAI, 235–236

static website hosting, 227–230

**Cost Explorer, 265–266**

**cost optimization, 259**

billing alarms, 267

Budgets, 266–267

Compute Optimizer, 263–265

cost allocation tags, 260–262

Cost Explorer, 265–266

Fargate, 268

managed services, 267–268

operations, 260

RDS, 267–268

Trusted Advisor, 262–263

cost optimization checks, 263

EC2 instances, 263

**credentials (AWS), IAM, 157**

**CRL (Certificate Revocation Lists), 177**

**cross-region replication**

high availability, 89

S3, 123

**CSR (Certificate Signing Requests), 177**

## D

**data at rest encryption, 172**

**data classification, 174, 175**

confidential data, 174

internal data, 174

public data, 174

restricted data, 174

**data in transit encryption, 172**

**data in use encryption, 173**

**data protection**

ACM, 177–178

Amazon Inspector, 180–181

AWS Secrets Manager, 178–179

AWS Security Hub, 181

certificate management, 175–178

data classification, 174–175

encryption, 172–174

GuardDuty, 179–180

**data transfers**

AWS Snow Family, 18

spikes, troubleshooting/remediation, 44

**databases**

- ACID requirements, 103–104
- bottlenecking applications, 61
- caching, 69–71
- high availability, 103–104
- performance, 78–79, 280
- RDS, 104–105
- scalability, 78–79
- services, 20–21

**datacenters, 9, 114****datastores**

- high availability, 98
- optimizing performance, 278–280
- S3, 98–99
  - security, 100–102
  - snapshots, 113
  - storage tiers, 99–100

**DataSync, backups, 121****DAX, in-line caching, 69–70****decoupled components, scalability/elasticity, 58****dedicated connections, 210****deployments**

- applications, 138
  - blue/green deployments, 139–140
  - development environments, 139
  - DR, 140
  - OpsWorks, 145–146
  - production environments, 139
  - QA environments, 139
  - staging environments, 139
- community cloud deployments, 4
- hybrid cloud deployments, 3
- management and deployment services, 22
- Multi-AZ deployments, 61–62
- private cloud deployments, 3
- public cloud deployments, 3

**deprecations (end of life), lifecycle management, 144****development environments, application deployments, 139****Device Farm, 22****digital certificates, 176****dimensions, CloudWatch, 31–32****Direct Connect, 209–210, 256–257****disposable resources, scalability/elasticity, 58****DNS (Domain Name Servers), 217**

- alias records, 220
- Route 53, 218
  - domain registration, 219–220
  - failover routing, 224
  - geolocation routing, 225
  - latency-based routing, 223–224
  - private zones, 218–219
  - public zones, 218–219
  - simple routing, 222–223
  - weighted routing, 223, 225

**DocumentDB, 20****documents, WorkDocs, 19****domains**

- CloudFront, 233
- DNS, Route 53, 218–226
  - registering with Route 53, 219–220

**DR (Disaster Recovery), 114, 140****dynamic scaling, 63****DynamoDB, 20**

- global tables, 107
- high availability, 106–107
- snapshots, 112–113

**E****EBS (Elastic Block Storage), 18, 89, 112****EC2 (Elastic Cloud Computing), 18**

- AMI, 129
- AutoScaling, 62
  - dynamic scaling, 63
  - manual scaling, 63–64
  - predictive scaling, 64–65
  - scheduled scaling, 64
- CLI commands, 128
- compute performance, optimizing, 275–276
  - instance stores, 276–277
  - placement groups, 277–278

**EC2 (Elastic Cloud Computing)**

## instances

- changes to number of instances, 43
- health check failures, 43
- low utilization, cost optimization, 263
- reserved instances, 263- 263
- stores, 276-277

## Launch templates, 62

**ECS (Elastic Container Service), 18****edge caching, 68, 74-76****edge locations, 10-11****EFS (Elastic File System), 18, 112****EIP (Elastic IP) addresses, 186****EKS (Elastic Kubernetes Service), 18****Elastic Beanstalk, 130-132****Elastic Transcoder, 21****ElastiCache, 20, 71-72**

Memcached, 72-73

Redis, 73-74

**elasticity, 56-57, 62**

- asynchronous requests, 58
- compute layer, 57
- decoupled components, 58
- disposable resources, 58
- persistent layer, 57
- rapid elasticity, 4
- scalability, 87

**ELB (Elastic Load Balancing), 17**

- access logs, 245-246
- high availability, 94-97

**email, WorkMail, 19****EMR (Elastic Map Reduce), 21****encryption, 172**

- asymmetric encryption, 173
- data at rest encryption, 172
- data in transit encryption, 172
- data in use encryption, 173
- KMS, 19, 173-174
- symmetric encryption, 173

**end of life (deprecations), lifecycle management, 144****endpoints, VPC, 199-200****end-user applications services, 19****ENI (Elastic Network Interface), 185-186****error handling, 45**

- 400 - bad request errors, 45
- 403 - access denied errors, 45
- 404 - page not found errors, 45
- 500 - internal failure errors, 45
- 503 - service unavailable errors, 45

**EventBridge, 48-49****expenses**

- CapEx, 13
- OpEx, 13

**Explorer, Systems Manager, 134****F****failed access attempts, troubleshooting/ remediation, 45****failed login attempts, troubleshooting/ remediation, 44****failover routing, 96, 224****Fargate, 268****features roadmaps, 143****file systems, EFS, 18, 112****firewalls**

- AWS Network Firewall, 197-198
- WAF, 19, 212-213

**firmware, patches, 150****fixes**

- bug fixes, 149
- hot fixes, 148

**flow logs, 240**

- configuring, 240-243
- VPC, 195

**foundation services, 17-19****frameworks, AWS Well-Architected, 272-274****FSx, snapshots, 112, 113****G****gateway endpoints, VPC, 200-201****gateways**

- AWS Storage Gateway, 18
- IGW, 187-188
- VPC

instance stores, optimizing compute performance

- NAT gateways, 191–192
- transit gateways, 202–205
- General-Purpose SSD, 279**
- Generation phase, performance monitoring, 274**
- geolocation routing, 96, 225**
- geoproximity routing, 96**
- Glacier, 18, 122–123**
- GLB (Gateway Load Balancing), 96**
- global architecture, 9–11**
- global tables, DynamoDB, 107**
- Glue, 21**
- GuardDuty, 179–180**

## H

**hardware, security and CloudHSM, 19**

**high availability, 86**

- ALB, 95
- Aurora, 105–106
- availability zones, 88–89
- cross-region replication, 89
- databases, 103–104
- datastores, 98–101
- DynamoDB, 106–107
- ELB, 94–97
- GLB, 96
- multi-A high availability, 89
- multiregion high availability, 89
- NLB, 95–96
- RDS, 104–105
- Route 53, 96
- single AZ high availability, 89
- uptime
  - applications, defining in, 87
  - percentage chart, 86–87
- VPC, 89, 91
  - host addresses, 90
  - network addresses, 90
  - private subnets, 91–92
  - public subnets, 91–92
  - subnet ranges, 90–91

**hit rates, CloudFront caching, 253–254**

**horizontal scaling, 59–61, 78**

**host addresses, VPC, 90**

**hosted connections, 210**

**hosting websites, S3 static website hosting, 227–230**

**hot fixes, 148**

**HTTP (Hypertext Transfer Protocol)**

- CloudFront, 233

- origin status codes, CloudFront caching, 254

**hybrid cloud deployments, 3**

**hybrid links, troubleshooting, 256–257**

## I

**IaaS (Infrastructure-as-a-Service), 4**

**IAM (Identity and Access Management), 19**

- Access Analyzer, 166

- AWS credentials, 157

- best practices, 165–167

- components, 156–157

- IDP, 161

- password policies, 162

- permissions, 157

- policies, 159–160

- roles, 158–159

- tags, 158

- user accounts, 157–158

- user groups, 158

**identity**

- IAM, 19, 161

- services, 19

**IDP, IAM, 161**

**IGW (Internet Gateways), 187–188**

**Incident Manager, Systems Manager, 134**

**infrastructure issues**

- performance, 43–44

- troubleshooting/remediation, 43–44

**in-line caching, 69–70**

**Inspector, Amazon, 19, 180–181**

**instance stores, optimizing compute performance, 276–277**

instances

### **instances, EC2**

- changes to number of instances, 43
- health check failures, 43

### **interface endpoints, VPC, 199–200**

### **internal data, 174**

### **IoT (Internet of Things) Services, 22**

## **J - K**

---

### **Keyspaces, 20**

### **Kinesis, 21, 248–249**

### **KMS (Key Management Service), 19, 173–174**

### **Kubernetes, EKS, 18**

## **L**

---

### **Lambda, 18**

### **latency-based routing, 96, 223–224**

### **Launch templates, EC2, 62**

### **lazy loading, 70**

### **lease expiration, EC2 reserved instances, 263**

### **lifecycle management**

- backups, 121–122
- deprecations (end of life), 144
- migrations, 143–144
- roadmaps, 142–143
- upgrades, 143–144
- versioning, 143

### **load balancing**

- ALB, 95
- ELB, 17
  - access logs, 245–246
  - high availability, 94–97
- GLB, 96
- NLB, 95–96
- Route 53, 96

### **logins (failed), troubleshooting, 44**

### **logs**

- CloudFront, 251–252
- CloudTrail, 36–37
- CloudWatch, 34
- ELB access logs, 245–246
- Logs Insights, 34

traffic analysis, 25

VPC flow logs, 195, 240–243

WAF ACL logs, 248–249

### **LTS, applications, 141**

## **M**

---

### **magnetic volumes, 279**

### **Management Console, 126**

### **management services, 22, 267–268**

### **managing**

- access, IAM, 19
- accounts
  - access keys, 163–164
  - IAM, 156–162, 165–167
  - MFA, 164–165
  - Trusted Advisor, 168
- applications, 142–144
  - building applications, 140–141
  - deployments, 138–140, 145–146
  - deprecations (end of life), 144
  - migrations, 143–144
  - OpsWorks, 145–146
  - patches, 147–148
  - upgrades, 143–144
- certificates, 175–178
- cloud components, patches, 149–152
- identity
  - access keys, 163–164
  - IAM, 19, 156–162, 165–167
  - MFA, 164–165
  - Trusted Advisor, 168
- keys, KMS, 19
- VPN, 206–207

### **manual scaling, 63–64**

### **Marketplace, AWS, 198**

### **measured service, 4**

### **Memcached, 72–73**

### **metering, 27**

### **metrics, CloudWatch, 32–33**

### **MFA (Multifactor Authentication), 164–165**

### **migrations, lifecycle management, 143–144**

**model of cloud computing, visual, 5–6****monitoring**

- alerting, 27
- CloudTrail, 28, 36–37
- CloudWatch, 28, 30–31
  - alarms, 33
  - dimensions, 31–32
  - logs, 34
  - Logs Insights, 34
  - metrics, 32–33
  - namespaces, 31–32
  - percentiles, 33
- compliance, 27–28
- logs, traffic analysis, 25
- metering, 27
- monitoring, 27
- NOC, 25
- Performance Efficiency, 272–273
  - Aggregation phase, 275
  - alarms, 275
  - analytics, 275
  - Generation phase, 274
  - Operational Excellence pillar Well-Architected Framework, 273–274
  - Performance Efficiency pillar Well-Architected Framework, 273–274
  - real-time processing, 275
  - storage, 275
  - security, 27–28
  - SOC, 25

**multi-A high availability, 89****Multi-AZ deployments, 61–62****multiregion high availability, 89****multisite active-active backup/restore strategy, 116****multivalued answer routing, 96****N****N-1, patches, cloud components, 151–152****NACL, VPC, 194–195****namespaces, CloudWatch, 31–32****NAT gateways, VPC, 191–192****Neptune, 20****Network Firewall, AWS, 197–198****networking**

- access, broad, 4
- addresses, VPC, 90
- AWS Shield, 214–215
- patches, 149–150
- services, 17
- troubleshooting
  - AWS-managed VPN, 257
  - CloudFront, 251–254
  - Direct Connect, 256–257
  - ELB access logs, 245–246
  - hybrid links, 256–257
  - private links, 256–257
  - VPC flow logs, 240–243
  - WAF ACL logs, 240–243
- VPC, 184
  - AWS Network Firewall, 197–198
  - CIDR ranges, 185
  - connectivity, 189–192
  - EIP addresses, 186
  - endpoints, 199–201
  - ENI, 185–186
  - IGW, 187–188
  - NAT gateways, 191–192
  - peering, 202–205
  - route tables, 190
  - security, 193–196
  - subnets, 189–190
  - transit gateways, 202–205
- VPN, 206
  - AWS-managed VPN, 206–207, 257
  - software site-to-site VPN, 207–208
  - WAF, 19, 212–213

**NIST, visual model of cloud computing, 5–6****NLB (Network Load Balancing), 95–96****NOC (Network Operating Centers), 25**

## O

---

**OAI (Origin Access Identity), 235–236**

**on-demand self-service, 4**

**OpenID Connect, IAM IDP, 161**

**Operational Excellence pillar Well-Architected Framework, 273–274**

**operational optimization, 260**

**OpEx (Operating Expenses), 13**

**OpsCenter, Systems Manager, 134**

**OpsWorks, 22, 145–146**

**optimizing**

compute performance, 275–276

instance stores, 276–277

placement groups, 277–278

costs, 259

billing alarms, 267

Budgets, 266–267

Compute Optimizer, 263–265

cost allocation tags, 260–262

Cost Explorer, 265–266

Fargate, 268

managed services, 267–268

operations, 260

RDS, 267–268

Trusted Advisor, 262–263

database performance, 280

datastore performance, 278–280

EC2 reserved instances, 263

performance, 271–275

values, Compute Optimizer, 265

**origin status codes, CloudFront caching, 254**

**over-provisioned values, Compute Optimizer, 265**

## P

---

**PaaS (Platform-as-a-Service), 5**

**partitions, placement groups, 278**

**password policies, IAM, 162**

**patches**

applications, 147

cloud components, 150

hot fixes, 148

scheduled updates, 148

security patches, 147–148

signature updates, 148

virtual patches, 148

cloud components, 149

bug fixes, 149

firmware, 150

N-1, 151–152

networking components, 149–150

policies, 151

rollbacks, 152

software, 151

storage components, 150

virtual appliances, 149

VM, 149

**peering, VPC, 202–205**

**percentiles, CloudWatch, 33**

**performance**

compute performance, optimizing, 275–276

instance stores, 276–277

placement groups, 277–278

databases, 78–79, 280

datastore performance, optimizing, 278–280

infrastructures, 43–44

monitoring efficiency, 272–273

Aggregation phase, 275

alarms, 275

analytics, 275

Generation phase, 274

Operational Excellence pillar Well-Architected Framework, 273–274

Performance Efficiency pillar Well-Architected Framework, 273–274

real-time processing, 275

storage, 275

optimizing, 271–275

Performance Insights RDS, 280

**Performance Efficiency pillar Well-Architected Framework, 273–274**

**permissions, IAM, 157, 158**

- persistent layer, scalability/elasticity, 57**
- pilot light backup/restore strategy, 115–116**
- Pinpoint, 21**
- PKI (Public Key Infrastructure), 176**
- placement groups**
  - clusters, 278
  - compute performance, optimizing, 277–278
  - partitions, 278
  - spreads, 278
- platform services, 20–22**
- policies**
  - IAM, 159–160
  - patches, cloud components, 151
- portfolio roadmaps, 143**
- predictive scaling, 64–65**
- price classes, CloudFront, 232–233**
- private cloud deployments, 3**
- private keys, certificate management, 177**
- private links, troubleshooting, 256–257**
- private subnets, VPC, 91–92, 195**
- private zones, Route 53, 218–219**
- production environments, application deployments, 139**
- protecting data**
  - ACM, 177–178
  - Amazon Inspector, 180–181
  - AWS Secrets Manager, 178–179
  - AWS Security Hub, 181
  - certificate management, 175–178
  - data classification, 174–175
  - encryption, 172–174
  - GuardDuty, 179–180
- Provisioned IOPS SSD, 279**
- provisioning, 125–126**
  - AMI, 129
  - AWS CLI, 127–128
  - CloudFormation, 129–133
  - RAM, 134–135
  - Systems Manager, 133–134
- public cloud deployments, 3**
- public data, 174**

- public keys, certificate management, 176**
- public subnets, VPC, 91–92, 195**
- public zones, Route 53, 218–219**

## Q

---

- QA, application deployments, 139**
- QLDB (Quantum Ledger Database), 20**

## R

---

- RAM, provisioning, 134–135**
- rapid elasticity, 4**
- RDS (Relational Database Service), 20, 267–268**
  - high availability, 104–105
  - Performance Insights, 280
  - RDS Proxy, 280
  - read replicas, 79–80
- read replicas, 78–82**
- real-time processing, performance monitoring, 275**
- recovery**
  - DR, 114
  - RPO, 113–114
  - RTO, 113–114
- Redis, 73–74**
- RedShift, 21**
- regions, 10**
- registering domains, with Route 53, 219–220**
- release roadmaps, 143**
- remediation, 41**
  - application issues, 44
  - data transfer spikes, 44
  - EC2 instances
    - changes to number of instances, 43
    - health check failures, 43
  - error handling, 45
    - 400 - bad request errors, 45
    - 403 - access denied errors, 45
    - 404 - page not found errors, 45
    - 500 - internal failure errors, 45
    - 503 - service unavailable errors, 45



## remediation

- EventBridge, 48–49
- failed access attempts, 45
- failed login attempts, 44
- infrastructure issues, 43–44
- infrastructure performance, 43–44
- role assumption from unknown locations, 45
- security issues, 44–45

**replication, cross-region**

- high availability, 89
- S3, 123

**reports, cost allocation, 262****requests, asynchronous, 58****reserved EC2 instances**

- lease expiration, 263
- optimizing, 263

**resilience, 86**

- ALB, 95
- Aurora, 105–106
- availability zones, 88–89
- cross-region replication, 89
- databases, 103–104
- datastores, 98–101
- DynamoDB, 106–107
- ELB, 94–97
- GLB, 96
- multi-A high availability, 89
- multiregion high availability, 89
- NLB, 95–96
- RDS, 104–105
- Route 53, 96
- single AZ high availability, 89
- uptime
  - applications, defining in, 87
  - percentage chart, 86–87
- VPC, 89, 91
  - host addresses, 90
  - network addresses, 90
  - private subnets, 91–92
  - public subnets, 91–92
  - subnet ranges, 90–91

**resources**

- disposable resources, scalability/elasticity, 58

IAM policies, 160

- pooling, 4

**responsibility models, 5****rest, data at, 172****restores, 109, 115**

- cloud restores, 110–111
- DataSync, 121
- DR, 114
- Glacier, 122–123
- life-cycling rules, 121–122
- RTO, 113–114
- S3, 119
  - cross-region replication, 123
  - S3 sync, 120–121
- snapshots, 112–113
- stateful services, 111
- strategies
  - comparing, 116–117
  - multisite active-active backup/restore strategy, 116
  - overview, 115
  - pilot light backup/restore strategy, 115–116
  - warm standby backup/restore strategy, 116
- versioning, 121–122

**restricted data, 174****roadmaps, lifecycle management, 142–143****roles**

- assumption from unknown locations, troubleshooting/remediation, 45
- IAM, 158–159

**rollbacks, patches, 152****rolling upgrades, 142****Route 53, 17, 96, 218**

- domain registration, 219–220
- failover routing, 224
- geolocation routing, 225
- latency-based routing, 223–224
- private zones, 218–219
- public zones, 218–219
- simple routing, 222–223
- weighted routing, 223, 225

**route tables, VPC, 190****routing**

- failover routing, 96, 224
- geolocation routing, 96, 225
- geoproximity routing, 96
- latency-based routing, 96, 223–224
- multivalued answer routing, 96
- simple routing, 96, 222–223
- weighted routing, 96, 223, 225

**RPO (Recovery-Point Objective), 113–114****RTO (Recovery-Time Objective), 113–114****S****S3 (Simple Storage Service), 18, 98–99**

- backups, 119
- cross-region replication, 123
- Glacier, 122–123
- multipart uploads, 280
- OAI, 235–236
- security, 100–102
- snapshots, 113
- static website hosting, 227–230
- storage tiers, 99–100
- transfer acceleration, 280
- versioning, 121–122
- WAF ACL logs, 248–249

**S3 sync, backups, 120–121****SaaS (Software-as-a-Service), 5****SageMaker, 22****SAML, IAM IDP, 161****scalability, 56–57, 62, 79**

- asynchronous requests, 58
- AutoScaling, 62
  - dynamic scaling, 63
  - manual scaling, 63–64
  - predictive scaling, 64–65
  - scheduled scaling, 64
- bottlenecking applications, 61
- compute layer, 57
- databases, 78–79
- decoupled components, 58

- disposable resources, 58
- dynamic scaling, 63
- elasticity, 87
- horizontal scaling, 59–61, 78
- manual scaling, 63–64
- Multi-AZ deployments, 61–62
- persistent layer, 57
- predictive scaling, 64–65
- read replicas, 78–82
- scheduled scaling, 64
- vertical scaling, 59–60, 78

**scheduled scaling, 64****scheduled updates, 148****SDK (Software Development Kit), 126****Secrets Manager, AWS, 178–179****security**

- access keys, 163–164
- account management
  - access keys, 163–164
  - IAM, 156–162, 165–167
  - MFA, 164–165
  - Trusted Advisor, 168
- AWS Security Hub, 181
- CloudFront, 75–76
- data protection
  - ACM, 177–178
  - Amazon Inspector, 180–181
  - AWS Secrets Manager, 178–179
  - AWS Security Hub, 181
  - certificate management, 175–178
  - data classification, 174–175
  - encryption, 172–174
  - GuardDuty, 179–180
  - KMS, 173–174
- firewalls
  - AWS Network Firewall, 197–198
  - WAF, 19
- hardware, CloudHSM, 19
- IAM
  - Access Analyzer, 166
  - AWS credentials, 157
  - best practices, 165–167

## security

- components, 156–157
- IDP, 161
- password policies, 162
- permissions, 157
- policies, 159–160
- roles, 158–159
- tags, 158
- user accounts, 157–158
- user groups, 158
- MFA, 164–165
- monitoring services, 27–28
- patches, 147–148
  - hot fixes, 148
  - scheduled updates, 148
  - signature updates, 148
  - virtual patches, 148
- S3, 100–102
- services, 19
- troubleshooting/remediation, 44–45
- Trusted Advisor, 168
- VPC, 193
  - AWS Network Firewall, 197–198
  - flow logs, 195
  - NACL, 194–195
  - private subnets, 195
  - public subnets, 195
  - security by default, 193
  - security groups, 194
- Security Hub, AWS, 181**
- self-service (on-demand), 4**
- server-side caching, 68–69**
- services**
  - 503 - service unavailable errors, 45
  - Amazon Inspector, 19
  - analytics services, 21
  - application services, 21
  - AWS
    - compute services, 18
    - end-user applications services, 19
    - foundation services, 17–19
    - network services, 17
    - platform services, 20–22
    - security and identity services, 19
    - storage services, 18
  - database services, 20–21
  - ECS, 18
  - EKS, 18
  - foundation services, 19
  - IaaS, 4
  - managed services, 267–268
  - management and deployment services, 22
  - measured service, 4
  - PaaS, 5
  - SaaS, 5
  - stateful services, backup/restore strategies, 111
- sharding, 79**
- Shield, AWS, 214–215**
- side-loaded caching, 69–71**
- signature updates, 148**
- simple routing, 96, 222–223**
- single AZ high availability, 89**
- Site-to-Site VGW, 17**
- SLA (Service-Level Agreements)**
  - defining, 87
  - scalability/elasticity, 58
  - uptime, 87
- snapshots, 112–113**
- Snow Family, 18**
- SOC (System Operating Centers), 25**
- software**
  - applications vs, 151
  - firmware, patches, cloud components, 150
  - SDK, 126
  - systems software, 151
- software site-to-site VPN, 207–208**
- spikes in data transfers, troubleshooting/remediation, 44**
- spreads, placement groups, 278**
- SSD (Solid-State Drives), 279– 279**
- stable builds, applications, 140–141**
- StackSets, CloudFormation, 132–133**

**staging environments, application deployments, 139**  
**stateful services, backup/restore strategies, 111**  
**statements, IAM policies, 160**  
**static website hosting, S3, 227–230**  
**statistics, CloudWatch metrics, 33**  
**status codes, CloudFront caching, 254**  
**storage**  
     Amazon Glacier, 18  
     Amazon Simple Storage Service (S3), 18  
     AWS Storage Gateway, 18  
     EBS, 18, 89, 112  
     patches, cloud components, 150  
     performance monitoring, 275  
     S3 storage tiers, 99–100  
     services, 18  
     Storage Gateway, AWS, 18  
**strategy roadmaps, 143**  
**subcommands, CLI, 128**  
**subnets, VPC, 189–190**  
     private subnets, 91–92, 195  
     public subnets, 91–92, 195  
     ranges, 90–91  
**support, LTS applications, 141**  
**symmetric encryption, 173**  
**synchronization**  
     DataSync, backups, 121  
     S3 sync, backups, 120–121  
**syntax, CLI commands, 128**  
**Systems Manager, 22**  
     automation, 49–50  
     Explorer, 134  
     Incident Manager, 134  
     OpsCenter, 134  
     provisioning, 133–134  
**systems software, 151**

## T

---

**tags, IAM, 158**  
**templates, Launch, 62**  
**transferring data**

    AWS Snow Family, 18  
     spikes, troubleshooting/remediation, 44  
**transit, data in, 172**  
**transit gateways, VPC peering, 203–205**  
**troubleshooting, 41**  
     application issues, 44  
     data transfer spikes, 44  
     Direct Connect, 256–257  
     EC2 instances  
         changes to number of instances, 43  
         health check failures, 43  
     error handling, 45  
         400 - bad request errors, 45  
         403 - access denied errors, 45  
         404 - page not found errors, 45  
         500 - internal failure errors, 45  
         503 - service unavailable errors, 45  
     EventBridge, 48–49  
     failed access attempts, 45  
     failed login attempts, 44  
     hybrid links, 256–257  
     infrastructure issues, 43–44  
**infrastructure performance, 43–44**  
     network connectivity  
         AWS-managed VPN, 257  
         CloudFront, 251–254  
         Direct Connect, 256–257  
         ELB access logs, 245–246  
         hybrid links, 256–257  
         private links, 256–257  
         VPC flow logs, 240–243  
         WAF ACL logs, 240–243  
     private links, 256–257  
     role assumption from unknown locations, 45  
     security issues, 44–45  
**Trusted Advisor, 168, 262–263**  
     cost optimization checks, 263  
     EC2 instances, 263–263  
**TTL, CloudFront, 232, 233**

under-provisioned values, Compute Optimizer

## U

---

**under-provisioned values, Compute Optimizer, 265**

### updating

- scheduled updates, 148
- signature updates, 148

### upgrades

- active-passive upgrades, 142
- lifecycle management, 143–144
- rolling upgrades, 142

### uptime

- applications, defining in, 87
- percentage chart, 86–87

**use, data in, 173**

**user accounts, IAM, 157–158**

**user groups, IAM, 158**

## V

---

### versioning

- backups, 121–122
- IAM policies, 160
- lifecycle management, 143
- S3, 121–122

**vertical scaling, 59–60, 78**

**VGW (AWS Site-to-Site VPN gateway), 17**

**VIF (Virtual Interfaces), 209–210**

**virtual appliances, patches, 149**

**virtual patches, 148**

**visual model of cloud computing, 5–6**

**VM, patches, 149**

**VPC (Virtual Private Cloud), 17, 89, 91, 184**

- AWS Marketplace, 198
- CIDR ranges, 185
- connectivity, 189–192
- Direct Connect, 209–210
- EIP addresses, 186
- endpoints, 199
  - gateway endpoints, 199–200
  - interface endpoints, 199–200

ENI, 185–186

flow logs, 195, 240–243

host addresses, 90

IGW, 187–188

NAT gateways, 191–192

network addresses, 90

peering, 202–205

route tables, 190

security, 193

AWS Network Firewall, 197–198

by default, 193

flow logs, 195

groups, 194

NACL, 194–195

private subnets, 195

public subnets, 195

subnets, 189–190

private subnets, 91–92, 195

public subnets, 91–92, 195

ranges, 90–91

transit gateways, 202–205

**VPN (Virtual Private Networks), 206**

AWS-managed VPN, 206–207, 257

software site-to-site VPN, 207–208

## W - X - Y - Z

---

**WAF (Web Application Firewall), 19, 212–213, 248–249**

**warm standby backup/restore strategy, 116**

**websites, S3 static website hosting, 227–230**

**weighted routing, 96, 223, 225**

**Well-Architected Framework, 272–273**

Operational Excellence pillar, 273–274

Performance Efficiency pillar, 273–274

**WorkDocs, 19**

**WorkMail, 19**

**WorkSpaces, 19**

**write-through caching, 70–71**