



Practice
Tests



Flash
Cards



Review
Exercises



Study
Planner

Cert Guide

Advance your IT career with hands-on learning

AWS Certified Solutions Architect – Associate

(SAA-C03)



MARK WILKINS

FREE SAMPLE CHAPTER |



AWS Certified Solutions Architect – Associate (SAA-C03) Cert Guide

Access interactive study tools on this book's companion website, including practice test software, review questions, Key Term flash card application, a study planner, and more!

To access the companion website, simply follow these steps:

1. Go to www.pearsonITcertification.com/register.
2. Enter the print book ISBN: **9780137941582**.
3. Answer the security question to validate your purchase.
4. Go to your account page.
5. Click on the **Registered Products** tab.
6. Under the book listing, click on the Access Bonus Content link.

If you have any issues accessing the companion website, you can contact our support team by going to <http://pearsonitp.echelp.org>.

Pearson Test Prep online system requirements:

Browsers: Chrome version 73 and above; Safari version 12 and above; Microsoft Edge 44 and above.

Devices: Desktop and laptop computers, tablets running Android v8.0 and above or iPadOS v13 and above, smartphones running Android v8.0 and above or iOS v13 and above with a minimum screen size of 4.7".

Internet access required.

Pearson Test Prep offline system requirements:

Windows 10, Windows 8.1; Microsoft .NET Framework 4.5 Client; Pentium-class 1 GHz processor (or equivalent); 512 MB RAM; 650 MB disk space plus 50 MB for each downloaded practice exam; access to the Internet to register and download exam databases.

AWS Certified Solutions Architect – Associate (SAA-C03) Cert Guide

Mark Wilkins



AWS Certified Solutions Architect – Associate (SAA-C03) Cert Guide

Copyright © 2023 by Pearson Education, Inc.

All rights reserved. No part of this book shall be reproduced, stored in a retrieval system, or transmitted by any means, electronic, mechanical, photocopying, recording, or otherwise, without written permission from the publisher. 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: 978-0-13-794158-2

ISBN-10: 0-13-794158-7

Library of Congress Control Number: 2023930964

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 author 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 or (800) 382-3419.

For government sales inquiries, please contact governmentsales@pearsoned.com.

For questions about sales outside the U.S., please contact intlcs@pearson.com.

**Vice President,
IT Professional**
Mark Taub

**Director, ITP Product
Management**
Brett Bartow

Executive Editor
Nancy Davis

Development Editor
Christopher Cleveland

Managing Editor
Sandra Schroeder

Senior Project Editor
Tonya Simpson

Copy Editor
Bill McManus

Indexer
Jen Hinchliffe

Proofreader
Jen Hinchliffe

Technical Editor
Ralph Parisi

Publishing Coordinator
Cindy Teeters

Cover 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>.

This page intentionally left blank

Contents at a Glance

| | | |
|-------------------------|---|-----|
| | Introduction | xxx |
| CHAPTER 1 | Understanding the Foundations of AWS Architecture | 3 |
| CHAPTER 2 | The AWS Well-Architected Framework | 39 |
| CHAPTER 3 | Designing Secure Access to AWS Resources | 75 |
| CHAPTER 4 | Designing Secure Workloads and Applications | 145 |
| CHAPTER 5 | Determining Appropriate Data Security Controls | 203 |
| CHAPTER 6 | Designing Resilient Architecture | 233 |
| CHAPTER 7 | Designing Highly Available and Fault-Tolerant Architecture | 287 |
| CHAPTER 8 | High-Performing and Scalable Storage Solutions | 357 |
| CHAPTER 9 | Designing High-Performing and Elastic Compute Solutions | 421 |
| CHAPTER 10 | Determining High-Performing Database Solutions | 477 |
| CHAPTER 11 | High-Performing and Scalable Networking Architecture | 523 |
| CHAPTER 12 | Designing Cost-Optimized Storage Solutions | 593 |
| CHAPTER 13 | Designing Cost-Effective Compute Solutions | 631 |
| CHAPTER 14 | Designing Cost-Effective Database Solutions | 665 |
| CHAPTER 15 | Designing Cost-Effective Network Architectures | 693 |
| CHAPTER 16 | Final Preparation | 721 |
| APPENDIX A | Answers to the “Do I Know This Already?” Quizzes and Q&A Sections | 733 |
| APPENDIX B | <i>AWS Certified Solutions Architect – Associate (SAA-C03) Cert Guide</i> Exam Updates | 749 |
| | Glossary of Key Terms | 751 |
| | Index | 761 |
| Online Elements: | | |
| APPENDIX C | Study Planner | |
| | Glossary of Key Terms | |

Table of Contents

Introduction xxx

Chapter 1 Understanding the Foundations of AWS Architecture 3

Essential Characteristics of AWS Cloud Computing 6

AWS Cloud Computing and NIST 8

On-Demand Self-Service 9

Broad Network Access 10

Resource Pooling 10

Rapid Elasticity 11

Measured Service 12

Moving to AWS 13

Infrastructure as a Service (IaaS) 14

Platform as a Service (PaaS) 17

Operational Benefits of AWS 19

Cloud Provider Responsibilities 20

Security at AWS 21

Network Security at AWS 22

Application Security at AWS 23

Migrating Applications 24

Applications That Can Be Moved to AWS and Hosted on an EC2 Instance with No Changes 26

Applications with Many Local Dependencies That Cause Problems When Being Moved to the Cloud 27

Replacing an Existing Application with a SaaS Application Hosted by a Public Cloud Provider 28

Applications That Should Remain On Premises and Eventually Be Deprecated 28

The AWS Well-Architected Framework 28

The Well-Architected Tool 30

AWS Services Cheat Sheet 31

In Conclusion 36

| | | |
|------------------|---|-----------|
| Chapter 2 | The AWS Well-Architected Framework | 39 |
| | “Do I Know This Already?” | 40 |
| | Foundation Topics | 42 |
| | The Well-Architected Framework | 42 |
| | Operational Excellence Pillar | 44 |
| | Security Pillar | 45 |
| | <i>Defense in Depth</i> | 45 |
| | Reliability Pillar | 47 |
| | Performance Efficiency Pillar | 49 |
| | Cost Optimization Pillar | 51 |
| | Sustainability Pillar | 51 |
| | Designing a Workload SLA | 52 |
| | Reliability and Performance Are Linked | 54 |
| | Disaster Recovery | 54 |
| | Placing Cloud Services | 55 |
| | <i>Data Residency and Compute Locations</i> | 55 |
| | <i>Caching Data with CDNs</i> | 56 |
| | <i>Data Replication</i> | 57 |
| | <i>Load Balancing Within and Between Regions</i> | 58 |
| | <i>Failover Architecture</i> | 60 |
| | Deployment Methodologies | 60 |
| | Factor 1: Use One Codebase That Is Tracked with Version Control to Allow Many Deployments | 63 |
| | <i>AWS CodeCommit</i> | 64 |
| | Factor 2: Explicitly Declare and Isolate Dependencies | 65 |
| | Factor 3: Store Configuration in the Environment | 66 |
| | Factor 4: Treat Backing Services as Attached Resources | 66 |
| | Factor 5: Separate Build and Run Stages | 67 |
| | Factor 6: Execute an App as One or More Stateless Processes | 67 |
| | Factor 7: Export Services via Port Binding | 69 |
| | Factor 8: Scale Out via the Process Model | 69 |
| | Factor 9: Maximize Robustness with Fast Startup and Graceful Shutdown | 69 |

| | | |
|------------------|---|-----------|
| | Factor 10: Keep Development, Staging, and Production as Similar as Possible | 70 |
| | Factor 11: Treat Logs as Event Streams | 70 |
| | Factor 12: Run Admin/Management Tasks as One-Off Processes | 71 |
| | Exam Preparation Tasks | 71 |
| | Review All Key Topics | 71 |
| | Define Key Terms | 72 |
| | Q&A | 72 |
| Chapter 3 | Designing Secure Access to AWS Resources | 75 |
| | “Do I Know This Already?” | 75 |
| | Foundation Topics | 79 |
| | Identity and Access Management (IAM) | 79 |
| | IAM Policy Definitions | 81 |
| | IAM Authentication | 82 |
| | Requesting Access to AWS Resources | 84 |
| | The Authorization Process | 85 |
| | Actions | 87 |
| | IAM Users and Groups | 88 |
| | The Root User | 88 |
| | The IAM User | 90 |
| | <i>Creating an IAM User</i> | 91 |
| | <i>IAM User Access Keys</i> | 92 |
| | IAM Groups | 94 |
| | Signing In as an IAM User | 94 |
| | IAM Account Details | 95 |
| | Creating a Password Policy | 96 |
| | Rotating Access Keys | 97 |
| | Using Multi-Factor Authentication | 99 |
| | Creating IAM Policies | 99 |
| | IAM Policy Types | 100 |
| | <i>Identity-Based Policies</i> | 100 |
| | <i>Resource-Based Policies</i> | 102 |

| | |
|--|-----|
| <i>Inline Policies</i> | 104 |
| IAM Policy Creation | 105 |
| <i>Policy Elements</i> | 106 |
| <i>Reading a Simple JSON Policy</i> | 107 |
| <i>Policy Actions</i> | 109 |
| <i>Additional Policy Control Options</i> | 110 |
| <i>Reviewing Policy Permissions</i> | 114 |
| <i>IAM Policy Versions</i> | 115 |
| <i>Using Conditional Elements</i> | 116 |
| <i>Using Tags with IAM Identities</i> | 116 |
| IAM Roles | 118 |
| When to Use IAM Roles | 119 |
| <i>AWS Services Perform Actions on Your Behalf</i> | 119 |
| <i>EC2 Instances Hosting Applications Need Access to AWS Resources</i> | 119 |
| <i>Access to AWS Accounts by Third Parties</i> | 121 |
| <i>Web Identity Federation</i> | 121 |
| <i>SAML 2.0 Federation</i> | 122 |
| <i>Cross-Account Access</i> | 124 |
| AWS Security Token Service | 126 |
| IAM Best Practices | 128 |
| IAM Security Tools | 130 |
| IAM Cheat Sheet | 132 |
| AWS Identity Center | 132 |
| AWS Organizations | 134 |
| AWS Organizations Cheat Sheet | 136 |
| AWS Resource Access Manager | 136 |
| AWS Control Tower | 138 |
| Exam Preparation Tasks | 140 |
| Review All Key Topics | 140 |
| Define Key Terms | 141 |
| Q&A | 142 |

Chapter 4 Designing Secure Workloads and Applications 145

| | |
|--|-----|
| “Do I Know This Already?” | 145 |
| Foundation Topics | 149 |
| Securing Network Infrastructure | 149 |
| Networking Services Located at Edge Locations | 150 |
| <i>AWS Shield (Standard and Advanced)</i> | 151 |
| <i>AWS Web Application Firewall (WAF)</i> | 152 |
| VPC Networking Services for Securing Workloads | 154 |
| <i>Route Tables</i> | 154 |
| <i>The Main Route Table</i> | 155 |
| <i>Security Groups</i> | 158 |
| <i>Security Groups Cheat Sheet</i> | 161 |
| <i>Web Server Inbound Ports</i> | 163 |
| <i>Database Server Inbound Ports</i> | 163 |
| <i>Administration Access</i> | 164 |
| <i>Understanding Ephemeral Ports</i> | 165 |
| <i>Security Group Planning</i> | 167 |
| <i>Network ACLs</i> | 168 |
| <i>Network ACL Implementation Details</i> | 169 |
| Network ACL Cheat Sheet | 169 |
| <i>Network ACL Rule Processing</i> | 170 |
| VPC Flow Logs | 172 |
| NAT Services | 174 |
| <i>NAT Gateway Service</i> | 174 |
| <i>NAT Instance</i> | 175 |
| <i>AWS NAT Gateway Service Cheat Sheet</i> | 176 |
| Amazon Cognito | 176 |
| User Pool | 177 |
| Federated Identity Provider | 179 |
| External Connections | 180 |
| Virtual Private Gateway | 181 |
| Customer Gateway | 182 |
| AWS Managed VPN Connection Options | 183 |

| | | |
|--|---|------------|
| Understanding Route Propagation | 184 | |
| AWS Direct Connect | 185 | |
| <i>AWS Direct Connect Gateway</i> | 186 | |
| AWS Direct Connect Cheat Sheet | 187 | |
| Amazon GuardDuty | 187 | |
| Amazon GuardDuty Cheat Sheet | 189 | |
| Amazon Macie | 189 | |
| Amazon Macie Cheat Sheet | 190 | |
| Security Services for Securing Workloads | 191 | |
| AWS CloudTrail | 191 | |
| <i>Creating an AWS CloudWatch Trail</i> | 192 | |
| <i>AWS CloudTrail Cheat Sheet</i> | 194 | |
| AWS Secrets Manager | 194 | |
| Amazon Inspector | 195 | |
| AWS Trusted Advisor | 196 | |
| AWS Config | 198 | |
| Exam Preparation Tasks | 199 | |
| Review All Key Topics | 199 | |
| Define Key Terms | 200 | |
| Q&A | 201 | |
| Chapter 5 | Determining Appropriate Data Security Controls | 203 |
| “Do I Know This Already?” | 204 | |
| Foundation Topics | 207 | |
| Data Access and Governance | 207 | |
| Data Retention and Classification | 207 | |
| Infrastructure Security | 209 | |
| IAM Controls | 210 | |
| Detective Controls | 210 | |
| Amazon EBS Encryption | 212 | |
| Amazon S3 Bucket Security | 216 | |
| S3 Storage at Rest | 220 | |
| Amazon S3 Object Lock Policies | 221 | |
| Legal Hold | 222 | |

| | | |
|--|---|------------|
| Amazon S3 Glacier Storage at Rest | 222 | |
| Data Backup and Replication | 223 | |
| AWS Key Management Service | 224 | |
| Envelope Encryption | 225 | |
| AWS KMS Cheat Sheet | 226 | |
| AWS CloudHSM | 227 | |
| AWS Certificate Manager | 227 | |
| Encryption in Transit | 228 | |
| Exam Preparation Tasks | 229 | |
| Review All Key Topics | 229 | |
| Define Key Terms | 230 | |
| Q&A | 230 | |
| Chapter 6 | Designing Resilient Architecture | 233 |
| “Do I Know This Already?” | 233 | |
| Foundation Topics | 237 | |
| Scalable and Resilient Architecture | 237 | |
| Scalable Delivery from Edge Locations | 238 | |
| Stateful Versus Stateless Application Design | 239 | |
| Changing User State Location | 241 | |
| User Session Management | 243 | |
| Container Orchestration | 244 | |
| Migrating Applications to Containers | 246 | |
| Resilient Storage Options | 246 | |
| Application Integration Services | 247 | |
| Amazon Simple Notification Service | 248 | |
| <i>Amazon SNS Cheat Sheet</i> | 250 | |
| Amazon Simple Queue Service | 250 | |
| <i>SQS Components</i> | 251 | |
| <i>Amazon SQS Cheat Sheet</i> | 253 | |
| AWS Step Functions | 254 | |
| Amazon EventBridge | 256 | |
| Amazon API Gateway | 258 | |
| API Gateway Cheat Sheet | 261 | |
| Building a Serverless Web App | 262 | |

| | | |
|------------------|---|------------|
| | <i>Step 1: Create a Static Website</i> | 263 |
| | <i>Step 2: User Authentication</i> | 263 |
| | <i>Step 3: Create the Serverless Backend Components</i> | 264 |
| | <i>Step 4: Set Up the API Gateway</i> | 265 |
| | <i>Step 5: Register for the Conference</i> | 266 |
| | Automating AWS Infrastructure | 266 |
| | AWS CloudFormation | 268 |
| | <i>CloudFormation Components</i> | 269 |
| | <i>CloudFormation Templates</i> | 270 |
| | <i>CloudFormation Stacks</i> | 272 |
| | <i>CloudFormation Stack Sets</i> | 276 |
| | <i>Third-Party Solutions</i> | 277 |
| | AWS Service Catalog | 277 |
| | AWS Elastic Beanstalk | 279 |
| | Updating Elastic Beanstalk Applications | 282 |
| | Exam Preparation Tasks | 284 |
| | Review All Key Topics | 284 |
| | Define Key Terms | 285 |
| | Q&A | 285 |
| Chapter 7 | Designing Highly Available and Fault-Tolerant Architecture | 287 |
| | “Do I Know This Already?” | 289 |
| | Foundation Topics | 293 |
| | High Availability and Fault Tolerance | 293 |
| | High Availability in the Cloud | 294 |
| | Reliability | 295 |
| | AWS Regions and Availability Zones | 296 |
| | Availability Zones | 300 |
| | <i>Availability Zone Distribution</i> | 301 |
| | <i>Planning Network Topology</i> | 303 |
| | <i>Local Zones</i> | 306 |
| | <i>Wavelength Zones</i> | 308 |
| | AWS Services Use Cases | 308 |
| | Choosing an AWS Region | 310 |
| | Compliance Rules | 311 |

| | |
|--|-----|
| <i>Understanding Compliance Rules at AWS: Use Case</i> | 312 |
| <i>AWS Compliance Standards</i> | 315 |
| <i>HIPAA</i> | 316 |
| <i>NIST</i> | 316 |
| <i>AWS GovCloud</i> | 318 |
| Latency Concerns | 319 |
| Services Offered in Each AWS Region | 320 |
| Calculating Costs | 321 |
| Distributed Design Patterns | 321 |
| Designing for High Availability and Fault Tolerance | 322 |
| Removing Single Points of Failure | 325 |
| Immutable Infrastructure | 327 |
| Storage Options and Characteristics | 329 |
| Failover Strategies | 330 |
| Backup and Restore | 332 |
| Pilot Light | 333 |
| Warm Standby | 337 |
| Multi-Region Scenarios | 339 |
| <i>Warm Standby with Amazon Aurora</i> | 340 |
| <i>Active-Active</i> | 340 |
| Single and Multi-Region Recovery Cheat Sheet | 343 |
| Disaster Recovery Cheat Sheet | 344 |
| AWS Service Quotas | 345 |
| AWS Service Quotas Cheat Sheet | 347 |
| Amazon Route 53 | 348 |
| Route 53 Health Checks | 349 |
| Route 53 Routing Policies | 350 |
| Route 53 Traffic Flow Policies | 351 |
| Alias Records | 352 |
| Route 53 Resolver | 352 |
| Exam Preparation Tasks | 354 |
| Review All Key Topics | 354 |
| Define Key Terms | 355 |
| Q&A | 355 |

Chapter 8 High-Performing and Scalable Storage Solutions 357

| | |
|--|-----|
| “Do I Know This Already?” | 358 |
| Foundation Topics | 362 |
| AWS Storage Options | 362 |
| Workload Storage Requirements | 363 |
| Amazon Elastic Block Store | 365 |
| EBS Volume Types | 367 |
| General Purpose SSD (gp2/gp3) | 369 |
| Elastic EBS Volumes | 370 |
| Attaching an EBS Volume | 371 |
| Amazon EBS Cheat Sheet | 372 |
| EBS Snapshots | 373 |
| <i>Taking a Snapshot from a Linux Instance</i> | 373 |
| <i>Taking a Snapshot from a Windows Instance</i> | 374 |
| <i>Fast Snapshot Restore</i> | 374 |
| <i>Snapshot Administration</i> | 375 |
| <i>EBS Recycle Bin</i> | 376 |
| <i>Snapshot Cheat Sheet</i> | 376 |
| Local EC2 Instance Storage Volumes | 377 |
| Amazon Elastic File System | 379 |
| EFS Performance Modes | 380 |
| EFS Throughput Modes | 381 |
| EFS Security | 382 |
| EFS Storage Classes | 382 |
| EFS Lifecycle Management | 383 |
| Amazon EFS Cheat Sheet | 383 |
| AWS DataSync | 384 |
| Amazon FSx for Windows File Server | 386 |
| Amazon FSx for Windows File Server Cheat Sheet | 388 |
| Amazon Simple Storage Service | 388 |
| Amazon S3 Bucket Concepts | 390 |
| Amazon S3 Data Consistency | 393 |
| Amazon S3 Storage Classes | 394 |
| Amazon S3 Management | 396 |
| S3 Bucket Versioning | 400 |

| | | |
|------------------|--|------------|
| | Amazon S3 Access Points | 401 |
| | Multi-Region Access Points | 402 |
| | Preselected URLs for S3 Objects | 403 |
| | S3 Cheat Sheet | 403 |
| | Amazon S3 Glacier | 404 |
| | Vaults and Archives | 405 |
| | S3 Glacier Retrieval Policies | 405 |
| | S3 Glacier Deep Archive | 406 |
| | Amazon S3 Glacier Cheat Sheet | 406 |
| | AWS Data Lake | 407 |
| | AWS Lake Formation | 409 |
| | Structured and Unstructured Data | 411 |
| | Analytical Tools and Datasets | 412 |
| | AWS Glue | 413 |
| | Analytic Services | 415 |
| | Amazon Kinesis Data Streams | 417 |
| | Exam Preparation Tasks | 418 |
| | Review All Key Topics | 418 |
| | Define Key Terms | 419 |
| | Q&A | 419 |
| Chapter 9 | Designing High-Performing and Elastic Compute Solutions | 421 |
| | “Do I Know This Already?” | 421 |
| | Foundation Topics | 425 |
| | AWS Compute Services | 425 |
| | AWS EC2 Instances | 427 |
| | Amazon Machine Images | 429 |
| | <i>AWS AMIs</i> | 431 |
| | <i>Creating a Custom AMI</i> | 432 |
| | <i>AMI Build Considerations</i> | 434 |
| | <i>Amazon EC2 Image Builder</i> | 435 |
| | AWS Lambda | 436 |
| | AWS Lambda Integration | 438 |
| | <i>Lambda Settings</i> | 439 |

| | |
|---|-----|
| AWS Lambda Cheat Sheet | 441 |
| Amazon Container Services | 441 |
| Amazon Elastic Container Service | 443 |
| AWS ECS Task Definition Choices | 443 |
| Amazon Elastic Kubernetes Service | 446 |
| Monitoring with AWS CloudWatch | 447 |
| CloudWatch Basic Monitoring | 448 |
| CloudWatch Logs | 449 |
| Collecting Data with the CloudWatch Agent | 451 |
| Planning for Monitoring | 452 |
| Amazon CloudWatch Integration | 453 |
| Amazon CloudWatch Terminology | 455 |
| Creating a CloudWatch Alarm | 459 |
| Additional Alarm and Action Settings | 460 |
| Amazon CloudWatch Cheat Sheet | 461 |
| Auto Scaling Options at AWS | 461 |
| EC2 Auto Scaling | 463 |
| EC2 Auto Scaling Operation | 463 |
| <i>Launch Configuration</i> | 464 |
| <i>Launch Templates</i> | 464 |
| <i>Auto Scaling Groups</i> | 465 |
| <i>Scaling Options for Auto Scaling Groups</i> | 466 |
| <i>Management Options for Auto Scaling Groups</i> | 470 |
| Cooldown Period | 471 |
| Termination Policy | 471 |
| Lifecycle Hooks | 472 |
| EC2 Auto Scaling Cheat Sheet | 473 |
| AWS Auto Scaling | 473 |
| Exam Preparation Tasks | 474 |
| Review All Key Topics | 474 |
| Define Key Terms | 475 |
| Q&A | 475 |

Chapter 10 Determining High-Performing Database Solutions 477

| | |
|--|-----|
| “Do I Know This Already?” | 477 |
| Foundation Topics | 481 |
| AWS Cloud Databases | 481 |
| Amazon Relational Database Service | 481 |
| Amazon RDS Database Instances | 483 |
| Database Instance Class Types | 485 |
| High-Availability Design for RDS | 485 |
| Multi-AZ RDS Deployments | 488 |
| Big-Picture RDS Installation Steps | 488 |
| Monitoring Database Performance | 490 |
| Best Practices for RDS | 491 |
| Amazon Relational Database Service Proxy | 492 |
| Amazon RDS Cheat Sheet | 493 |
| Amazon Aurora | 493 |
| Amazon Aurora Storage | 496 |
| Amazon Aurora Replication | 498 |
| Communicating with Amazon Aurora | 499 |
| Amazon Aurora Cheat Sheet | 500 |
| Amazon DynamoDB | 501 |
| Amazon DynamoDB Tables | 503 |
| <i>Provisioning Table Capacity</i> | 504 |
| <i>Adaptive Capacity</i> | 506 |
| <i>Data Consistency</i> | 507 |
| <i>ACID and Amazon DynamoDB</i> | 509 |
| <i>Global Tables</i> | 510 |
| Amazon DynamoDB Accelerator | 511 |
| Backup and Restoration | 511 |
| Amazon DynamoDB Cheat Sheet | 512 |
| Amazon ElastiCache | 512 |
| Amazon ElastiCache for Memcached | 513 |
| Amazon ElastiCache for Memcached Cheat Sheet | 514 |
| Amazon ElastiCache for Redis | 514 |
| Amazon ElastiCache for Redis Cheat Sheet | 516 |

| | |
|--|------------|
| ElastiCache for Redis: Global Datastore | 516 |
| Amazon Redshift | 517 |
| Amazon Redshift Cheat Sheet | 519 |
| Exam Preparation Tasks | 520 |
| Review All Key Topics | 520 |
| Define Key Terms | 521 |
| Q&A | 521 |
| Chapter 11 High-Performing and Scalable Networking Architecture | 523 |
| “Do I Know This Already?” | 523 |
| Foundation Topics | 527 |
| Amazon CloudFront | 527 |
| How Amazon CloudFront Works | 527 |
| Regional Edge Caches | 528 |
| CloudFront Use Cases | 529 |
| HTTPS Access | 529 |
| Serving Private Content | 530 |
| <i>Using Signed URLs</i> | 530 |
| <i>Using an Origin Access Identifier</i> | 531 |
| <i>Restricting Distribution of Content</i> | 532 |
| CloudFront Origin Failover | 532 |
| Video-on-Demand and Live Streaming Support | 533 |
| Edge Functions | 534 |
| <i>CloudFront Functions</i> | 534 |
| <i>Lambda@Edge Functions</i> | 535 |
| <i>Lambda@Edge Use Cases</i> | 535 |
| CloudFront Cheat Sheet | 536 |
| AWS Global Accelerator | 536 |
| Elastic Load Balancing Service | 539 |
| Application Load Balancer Features | 540 |
| <i>Application Load Balancer Deployment</i> | 541 |
| Health Checks | 548 |
| <i>Target Group Attributes</i> | 550 |
| <i>Sticky Session Support</i> | 551 |

| | |
|--|-----|
| <i>Access Logs</i> | 553 |
| <i>ALB Cheat Sheet</i> | 553 |
| Network Load Balancer | 554 |
| <i>NLB Cheat Sheet</i> | 554 |
| Multi-Region Failover | 555 |
| <i>CloudWatch Metrics</i> | 555 |
| AWS VPC Networking | 556 |
| The Shared Security Model | 557 |
| AWS Networking Terminology | 558 |
| VPC Cheat Sheet | 560 |
| Creating a VPC | 561 |
| <i>Using the Create VPC Wizard</i> | 561 |
| <i>Using the AWS CLI to Create a VPC</i> | 563 |
| How Many VPCs Does Your Organization Need? | 564 |
| <i>Creating the VPC CIDR Block</i> | 565 |
| Subnets | 570 |
| Subnet Cheat Sheet | 572 |
| IP Address Types | 573 |
| Private IPv4 Addresses | 573 |
| Private IPv4 Address Summary | 574 |
| Public IPv4 Addresses | 574 |
| <i>Elastic IP Addresses</i> | 575 |
| <i>Public IPv4 Address Cheat Sheet</i> | 577 |
| Inbound and Outbound Traffic Charges | 578 |
| Bring-Your-Own IP | 579 |
| <i>The BYOIP Process</i> | 580 |
| IPv6 Addresses | 580 |
| VPC Flow Logs | 581 |
| Connectivity Options | 583 |
| VPC Peering | 583 |
| Establishing a Peering Connection | 584 |
| VPC Endpoints | 585 |

| | |
|--|------------|
| <i>VPC Gateway Endpoints</i> | 585 |
| <i>VPC Interface Endpoints</i> | 586 |
| <i>Endpoint Services</i> | 588 |
| Exam Preparation Tasks | 590 |
| Review All Key Topics | 590 |
| Define Key Terms | 591 |
| Q&A | 591 |
| Chapter 12 Designing Cost-Optimized Storage Solutions | 593 |
| “Do I Know This Already?” | 593 |
| Foundation Topics | 597 |
| Calculating AWS Costs | 597 |
| Cloud Service Costs | 598 |
| Tiered Pricing at AWS | 599 |
| Management Tool Pricing Example: AWS Config | 600 |
| <i>AWS Config Results</i> | 601 |
| Cost Management Tools | 602 |
| AWS Cost Explorer | 604 |
| AWS Budgets | 607 |
| AWS Cost and Usage Reports | 609 |
| Managing Costs Cheat Sheet | 610 |
| Tagging AWS Resources | 611 |
| Using Cost Allocation Tags | 612 |
| Storage Types and Costs | 613 |
| AWS Backup | 618 |
| Lifecycle Rules | 619 |
| AWS Backup Cheat Sheet | 620 |
| Data Transfer Costs | 621 |
| AWS Storage Gateway | 625 |
| AWS Storage Gateway Cheat Sheet | 627 |
| Exam Preparation Tasks | 627 |
| Review All Key Topics | 628 |
| Define Key Terms | 628 |
| Q&A | 629 |

Chapter 13 Designing Cost-Effective Compute Solutions 631

| | |
|--|-----|
| “Do I Know This Already?” | 631 |
| Foundation Topics | 633 |
| EC2 Instance Types | 633 |
| What Is a vCPU? | 634 |
| EC2 Instance Choices | 634 |
| Dedicated Host | 636 |
| <i>Dedicated Hosts Cheat Sheet</i> | 637 |
| Dedicated Instances | 638 |
| Placement Groups | 638 |
| EC2 Instance Purchasing Options | 638 |
| EC2 Pricing—On-demand | 640 |
| On-demand Instance Service Quotas | 641 |
| Reserved Instances | 644 |
| Term Commitment | 645 |
| Payment Options | 646 |
| EC2 Reserved Instance Types | 646 |
| Scheduled Reserved EC2 Instances | 646 |
| Regional and Zonal Reserved Instances | 647 |
| Savings Plans | 649 |
| Spot Instances | 650 |
| Spot Fleet Optimization Strategies | 653 |
| Spot Capacity Pools | 653 |
| EC2 Pricing Cheat Sheet | 655 |
| Compute Tools and Utilities | 655 |
| Strategies for Optimizing Compute | 656 |
| Matching Compute Utilization with Requirements | 659 |
| Compute Scaling Strategies | 661 |
| Exam Preparation Tasks | 662 |
| Review All Key Topics | 662 |
| Define Key Terms | 662 |
| Q&A | 663 |

Chapter 14 Designing Cost-Effective Database Solutions 665

- “Do I Know This Already?” 665
- Foundation Topics 668
- Database Design Choices 668
 - RDS Deployments 668
 - RDS Costs Cheat Sheet* 671
 - RDS Database Design Solutions* 672
 - NoSQL Deployments 675
 - NoSQL Costs Cheat Sheet* 676
 - Migrating Databases 680
 - AWS Schema Conversion Tool* 681
- Database Data Transfer Costs 681
 - Data Transfer Costs and RDS 682
 - Data Transfer Costs with DynamoDB 683
 - Data Transfer Costs with Amazon Redshift 685
 - Data Transfer Costs with DocumentDB 686
 - Data Transfer Costs Cheat Sheet 686
- Database Retention Policies 687
 - Database Backup Policies Cheat Sheet 688
- Exam Preparation Tasks 689
- Review All Key Topics 690
- Define Key Terms 690
- Q&A 690

Chapter 15 Designing Cost-Effective Network Architectures 693

- “Do I Know This Already?” 693
- Foundation Topics 695
- Networking Services and Connectivity Costs 695
 - Elastic Load Balancing Deployments 695
 - NAT Devices 696
 - AWS CloudFront 698
 - CloudFront Pricing Cheat Sheet* 699
 - VPC Endpoints 701
 - Network Services from On-Premises Locations 703
- Data Transfer Costs 706
 - Accessing AWS Services in the Same Region 707

| | | |
|-------------------------|---|------------|
| | Workload Components in the Same Region | 709 |
| | Accessing AWS Services in Different Regions | 710 |
| | Data Transfer at Edge Locations | 713 |
| | Network Data Transfer | 714 |
| | Public Versus Private Traffic Charges | 714 |
| | Data Transfer Costs Cheat Sheet | 716 |
| | Exam Preparation Tasks | 717 |
| | Review All Key Topics | 717 |
| | Define Key Terms | 718 |
| | Q&A | 718 |
| Chapter 16 | Final Preparation | 721 |
| | Exam Information | 721 |
| | Tips for Getting Ready for the Exam | 724 |
| | Scheduling Your Exam | 725 |
| | Tools for Final Preparation | 726 |
| | Pearson Test Prep Practice Test Software and Questions on the Website | 727 |
| | <i>Accessing the Pearson Test Prep Software Online</i> | 727 |
| | <i>Accessing the Pearson Test Prep Software Offline</i> | 727 |
| | <i>Customizing Your Exams</i> | 728 |
| | Updating Your Exams | 729 |
| | <i>Premium Edition</i> | 730 |
| | Chapter-Ending Review Tools | 730 |
| | Suggested Plan for Final Review/Study | 730 |
| | Summary | 731 |
| Appendix A | Answers to the “Do I Know This Already?” Quizzes and Q&A Sections | 733 |
| Appendix B | <i>AWS Certified Solutions Architect – Associate (SAA-C03) Cert Guide Exam Updates</i> | 749 |
| | Glossary of Key Terms | 751 |
| | Index | 761 |
| Online Elements: | | |
| Appendix C | Study Planner | |
| | Glossary of Key Terms | |

About the Author

Mark Wilkins is an electronics 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 as an AWS Certified Solutions Architect – Associate. 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 worked as a technical evangelist for IBM SoftLayer from 2013 through 2016 and taught both SoftLayer fundamentals and SoftLayer design classes to many Fortune 500 companies in Canada, the United States, Europe, and Australia. As former course director for Global Knowledge, Mark developed and taught many technical seminars, including Configuring Active Directory Services, Configuring Group Policy, and Cloud and Virtualization Essentials. Mark currently develops AWS curriculum on AWS cloud services and certification for O'Reilly Media and LinkedIn Learning. To learn more about what Mark finds interesting about the cloud, visit The Cloud Thingy, at <https://thecloudthingy.substack.com/>. To learn more about the AWS cloud and AWS certification, check out Mark's YouTube channel at <http://www.youtube.com/@SAA-C03>.

Mark's published books include *Windows 2003 Registry for Dummies*, *Administering SMS 3.0*, *Administering Active Directory*, and *Learning Amazon Web Services (AWS): A Hands-On Guide to the Fundamentals of AWS Cloud*.

Dedication

I would like to dedicate this book to my grandson, Silas, a future nerd. And to Bruce, one of our cats, for making me take breaks when he wanted.

Acknowledgments

This manuscript was made truly great by the incredible project management of Tonya Simpson, who went above and beyond! Thanks so much.

I would also like to express my gratitude to Chris Cleveland, the development editor of this book. I was lucky to work with him on this text. Chris helped make this book several cuts above the rest.

Finally, thanks so much to Nancy Davis, my tireless acquisitions editor. Nancy very patiently made this book a reality.

About the Technical Reviewer

Ralph Parisi is a certified Champion Authorized Amazon instructor and has been teaching AWS courses for 6 years. Ralph has been an instructor for more than 20 years and has taught technical classes for Microsoft Exchange Server, Microsoft Windows Server, Active Directory, Group Policy, Citrix XenDesktop, and XenApp. Ralph has worked as a consultant to large corporations architecting Exchange Server and Active Directory solutions and migrations. Ralph has also worked with various companies as a technical writer. Ralph lives in North Carolina with his wife and Saluki, Dillon.

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

Reader Services

Register your copy of *AWS Certified Solutions Architect – Associate (SAA-C03) Cert Guide* at www.pearsonitcertification.com for convenient access to downloads, updates, and corrections as they become available. To start the registration process, go to www.pearsonitcertification.com/register and log in or create an account*. Enter the product ISBN 9780137941582 and click Submit. When the process is complete, you will find any available bonus content under Registered Products.

*Be sure to check the box that you would like to hear from us to receive exclusive discounts on future editions of this product.

Introduction

There are many reasons to get certified in AWS technology. First of all, AWS certifications validate your AWS cloud knowledge. To fully understand the AWS cloud, preparing for the AWS Certified Solutions Architect – Associate (SAA-C03) exam is a great place to start. There are other AWS certifications that may be a better fit, depending on your technical level, your current knowledge of cloud concepts, and your current and future jobs with AWS technologies and services. Certifications are broken down into Foundational, Associate, Professional, and Specialty certifications. Full details can be found at <https://aws.amazon.com/certification/>. AWS frequently adds new certification tracks, but the following are the certifications that are currently available:

- **Foundational:** There is one Foundational certification: AWS Certified Cloud Practitioner. The recommendation is to have at least 6 months of fundamental AWS cloud knowledge before attempting this certification exam. You might be closer to this certification than you think, depending on your current level of technical skills. One advantage of taking the AWS Certified Cloud Practitioner exam first is that it helps you to get used to answering multiple-choice test questions and to learn about the foundational AWS cloud services.
- **Associate:** There are several Associate certifications:
 - **AWS Certified Solutions Architect – Associate:** For individuals working as solutions architects, designing AWS solutions using AWS services
 - **AWS Certified SysOps Administrator – Associate:** For individuals working as systems administrators, managing and operating AWS services
 - **AWS Certified Developer – Associate:** For individuals working as developers, deploying and debugging cloud-based applications hosted at AWS

Each certification exam expects that you know how the AWS service that you are being tested on works. Each Associate certification has a specific focus:

- **Architect:** The best design possible, based on the question and scenario
- **SysOps:** The administration steps required to carry out a particular task
- **Developer:** How to best use the service for the hosted application you are writing

For example, the three Associate exams would test different aspects of CloudWatch logs:

- **Architect:** The main focus of this exam is on how CloudWatch logs work and the main design features to consider based on specific needs—that is, design knowledge related to using CloudWatch logs for a variety of solutions.
- **SysOps:** The main focus of this exam is on how to configure CloudWatch logs based on specific needs—that is, configuration and deployment of CloudWatch logs using operational knowledge.
- **Developer:** The main focus of this exam is on what CloudWatch logs are useful for when developing applications for tracking performance of an application hosted on an EC2 instance—that is, knowledge of how a particular AWS service can help in the development and testing process with applications.

Before you attempt one of the Associate certifications, AWS recommends that you have at least 1 year of experience solving problems and implementing solutions using AWS services. AWS really wants to ensure that you have hands-on experience solving problems.

- **Professional:** These certifications include the AWS Certified Solutions Architect Professional and the AWS Certified DevOps Engineer Professional. Professional certifications are not where you normally start your certification journey. AWS recommends that you have at least 2 years of hands-on experience before taking a Professional exam.
- **Specialty:** The Specialty certifications for Advanced Networking, Security, Machine Learning, Data Analytics, SAP on AWS, and Database require advanced knowledge of the subject matter. AWS recommends that you have an Associate certification before you attempt one of these certifications.

NOTE The AWS Certified Solutions Architect – Associate (SAA-C03) certification is globally recognized and does an excellent job of demonstrating that the holder has knowledge and skills across a broad range of AWS topics.

The Goals of the AWS Certified Solutions Architect – Associate Certification

The AWS Certified Solutions Architect – Associate certification is intended for individuals who perform in a solutions architect role. This exam validates a candidate's ability to effectively demonstrate knowledge of how to architect and deploy secure and robust applications on AWS technologies. It validates a candidate's ability to

- Have knowledge and skills in the following AWS services: compute, networking, storage, and database and deployment and management services

- Have knowledge and skills in deploying, managing, and operating AWS workloads and implementing security controls and compliance requirements
- Identify which AWS service meets technical requirements
- Define technical requirements for AWS-based applications
- Identify which AWS services meet a given technical requirement

Recommended Prerequisite Skills

While this book provides you with the information required to pass the Certified Solutions Architect – Associate (SAA-C03) exam, Amazon considers ideal candidates to be those who possess the following:

- Experience in AWS technology
- Strong on-premises IT experience
- Understanding of mapping on-premises technology to the cloud
- Experience with other cloud services

The Exam Domains

The AWS Certified Solutions Architect – Associate (SAA-C03) exam is broken down into four major domains. This book covers each of the domains and the task statements.

- **Domain 1: Design Secure Architectures 30%**
 - Task Statement 1: Design secure access to AWS resources
 - Task Statement 2: Design secure workloads and applications
 - Task Statement 3: Determine appropriate data security controls
- **Domain 2: Design Resilient Architectures 26%**
 - Task Statement 1: Design scalable and loosely coupled architectures
 - Task Statement 2: Design highly available and/or fault-tolerant architectures
- **Domain 3: Design High-Performing Architectures 24%**
 - Task Statement 1: Determine high-performing and/or scalable storage solutions
 - Task Statement 2: Design high-performing and elastic compute solutions
 - Task Statement 3: Determine high-performing database solutions

- Task Statement 4: Determine high-performing and/or scalable network architectures
- Task Statement 5: Determine high-performing data ingestion and transformation solutions
- **Domain 4: Design Cost-Optimized Architectures 20%**
 - Task Statement 1: Design cost-optimized storage solutions
 - Task Statement 2: Design cost-optimized compute solutions
 - Task Statement 3: Design cost-optimized database solutions
 - Task Statement 4: Design cost-optimized network architectures

Steps to Becoming an AWS Certified Solutions Architect – Associate

To become an AWS Certified Solutions Architect – Associate, an exam candidate must meet certain prerequisites and follow specific procedures. Exam candidates must ensure that they have the necessary background and technical experience for the exam and then sign up for the exam.

Signing Up for the Exam

The steps required to sign up for the AWS Certified Solutions Architect – Associate exam are as follows:

- Step 1.** Create an AWS Certification account at <https://www.aws.training/Certification> and schedule your exam from the home page by clicking Schedule New Exam.
- Step 2.** Select a testing provider, either Pearson VUE or PSI, and select whether you want to take the exam at a local testing center or online from your home or office. If you choose to take an online exam, you will have to agree to the online testing policies.
- Step 3.** Complete the examination signup by selecting the preferred language and the date of your exam.
- Step 4.** Submit the examination fee.

TIP Refer to the AWS Certification site at <https://aws.amazon.com/certification/> for more information regarding this and other AWS certifications.

How to Use This Book

This book maps directly to the domains of the AWS Certified Solutions Architect – Associate (SAA-C03) exam and includes a number of features that help you understand the topics and prepare for the exam.

Objectives and Methods

This book uses several key methodologies to help you discover the exam topics on which you need more review, to help you fully understand and remember those details, and to help you ensure that you have retained your knowledge of those topics. This book does not try to help you pass the exam only by memorization; it seeks to help you truly learn and understand the topics. This book is designed to help you pass the AWS Certified Solutions Architect – Associate (SAA-C03) exam by using the following methods:

- Helping you discover which exam topics you have not mastered
- Providing explanations and information to fill in your knowledge gaps
- Supplying exercises that enhance your ability to recall and deduce the answers to test questions
- Providing practice exercises on the topics and the testing process via test questions on the companion website

Book Features

To help you customize your study time using this book, the core chapters have several features that help you make the best use of your time:

- **Foundation Topics:** The sections under “Foundation Topics” describe the core topics of each chapter.
- **Exam Preparation Tasks:** The “Exam Preparation Tasks” section lists a series of study activities that you should do at the end of each chapter:
 - **Review All Key Topics:** The Key Topic icon appears next to the most important items in the “Foundation Topics” section of the chapter. The “Review All Key Topics” activity lists the key topics from the chapter, along with the number of the page where you can find more information about each one. Although the contents of the entire chapter could be tested on the exam, you should definitely know the information listed in each key topic, so you should review these.
 - **Define Key Terms:** Although the AWS Certified Solutions Architect – Associate (SAA-C03) exam may be unlikely to word a question “Define

this term,” the exam does require that you learn and know a lot of terminology. This section lists the most important terms from the chapter and asks you to write a short definition and compare your answer to the glossary at the end of the book.

- **Q&A:** Confirm that you understand the content that you just covered by answering these questions and reading the answer explanations.

- **Web-based practice exam:** The companion website includes the Pearson Test Prep practice test engine, which enables you to take practice exam questions. Use it to prepare with a sample exam and to pinpoint topics where you need more study.

How This Book Is Organized

This book contains 14 core chapters—Chapters 2 through 15. Chapter 1 introduces the foundations of AWS, and Chapter 16 provides preparation tips and suggestions for how to approach the exam. Each core chapter covers a specific task statement or multiple task statements of the domains for the AWS Certified Solutions Architect – Associate (SAA-C03) exam.

Companion Website

Register this book to get access to the Pearson Test Prep practice test software and other study materials plus additional bonus content. Check this site regularly for new and updated postings written by the author that provide further insight into the more troublesome topics on the exam. Be sure to check the box indicating that you would like to hear from us to receive updates and exclusive discounts on future editions of this product or related products.

To access this companion website, follow these steps:

- Step 1.** Go to <https://www.pearsonitcertification.com/register> and log in or create a new account.
- Step 2.** Enter the ISBN 9780137941582.
- Step 3.** Answer the challenge question as proof of purchase.
- Step 4.** Click the Access Bonus Content link in the Registered Products section of your account page to be taken to the page where your downloadable content is available.

Please note that many of our companion content files can be very large, especially image and video files.

If you are unable to locate the files for this title by following these steps, please visit <https://www.pearsonITcertification.com/contact> and select the Site Problems/Comments option from the Select a Topic drop-down list. Our customer service representatives will assist you.

Pearson Test Prep Practice Test Software

As noted earlier, the Pearson Test Prep practice test software comes with two full practice exams. These practice exams are available to you either online or as an offline Windows application. To access the practice exams that were developed with this book, see the instructions in the card inserted in the sleeve at the back of the book. This card includes a unique access code that enables you to activate your exams in the Pearson Test Prep practice test software. For more information about the practice exams and more tools for exam preparation, see Chapter 16.

Figure Credits

Cover: Yurchanka Siarhei/Shutterstock

Chapter opener: Charlie Edwards/Getty Images

Figures 1.1, 1.3 through 1.6, 1.10, 1.12 through 1.4, 2.1 through 2.4, 2.6 through 2.8, 2.13, 2.14, 3.1 through 3.4, 3.7 through 3.9, 3.11 through 3.24, 3.27 through 3.37, 3.39 through 3.48, 4.3, 4.4, 4.6 through 4.8, 4.11 through 4.14, 4.22 through 4.34, 5.2, 5.6 through 5.11, 5.14 through 5.16, 5.18, 6.7, 6.11 through 6.15, 6.17 through 6.20, 6.22, 6.23, 6.26 through 6.30, 7.5, 7.11 through 7.14, 7.33, 7.34, 8.1 through 8.13, 8.15, 8.17 through 8.23, 9.2 through 9.5, 9.7, 9.9, 9.10, 9.12, 9.13 through 9.29, 10.1, 10.4, 10.10 through 10.12, 10.17, 10.18, 11.3 through 11.7, 11.10 through 11.21, 11.23, 11.24, 11.27 through 11.31, 11.33, 11.34, 12.1 through 12.10, 12.12 through 12.17, 13.3 through 13.12, 14.1, 14.3 through 14.6, 14.13, 15.4, 16.1, 16.2: Amazon Web Services, Inc

Figure 2.11: Adam Wiggins

Figures 2.9a, 7.1: Andrei Minsk/Shutterstock

Figures 3.10, 3.38, 11.25: Microsoft Corporation

This page intentionally left blank

Determining Appropriate Data Security Controls

Organizations have workloads and associated cloud services fail while operating at AWS. Amazon Elastic Compute Cloud (EC2) instances fail, Amazon Elastic Block Store (EBS) volumes crash, and cloud services can stop working. However, you shouldn't have to go to your boss and announce, "We've lost some data." Fortunately, all data can be securely and redundantly stored at AWS.

All data stored at AWS using any storage service can be encrypted; organizations make the decision about whether encryption is required. However, Amazon S3 objects and S3 Glacier archive storage *is* automatically encrypted at rest. All other storage services at AWS store data records in an unencrypted state to start. For example, Amazon S3 buckets are encrypted using server-side encryption using Amazon S3, the AWS Key Management Service (KMS) with customer master keys (CMK) and data keys, or encryption keys supplied by each organization. Amazon EBS volumes—both boot and data volumes—can be encrypted at rest and in transit using CMKs provided by AWS KMS. Shared storage services such as Amazon EFS and Amazon FSx for Windows File Server can also be encrypted at rest, as can Amazon DynamoDB tables, Amazon Relational Database Service (RDS) deployments, and Amazon Simple Queue Service (SQS) queues.

NOTE Use of master/slave terms in the following chapter is ONLY in association with the official terminology used in industry specifications and/or standards, and in no way diminishes Pearson's commitment to promoting diversity, equity, and inclusion, and challenging, countering, and/or combating bias and stereotyping in the global population of the learners we serve.

AWS does not have single-tenant persistent data storage for individual organizations; all storage services offered at AWS are multi-tenant by design. AWS has the responsibility to ensure that each organization's stored data records are isolated to the AWS account in which they are first created. Organizations can secure data at rest by choosing to encrypt all data records; protecting data in transit can be achieved using Transport Layer Security (TLS).

Each organization is in control of the storage and retrieval of its data records that are stored at AWS. It's the organization's responsibility to define the security and accessibility of all data records stored at AWS. All data storage at AWS starts as private storage only accessible across the AWS private network. Organizations can choose to make select Amazon S3 buckets public, but all other storage services offered by AWS remain private and are not publicly accessible across the Internet. AWS VPN and AWS Direct Connect connections from on-premises locations can directly access AWS storage services; however, EBS volumes can only be accessed through the attached EC2 instance. Figure 5-1 illustrates the options for data encryption at AWS that are discussed in this chapter.

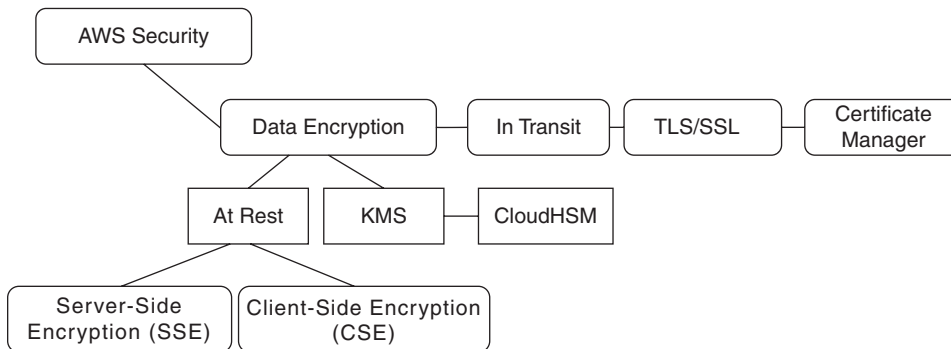
Key Topic


Figure 5-1 Encryption Choices at AWS

“Do I Know This Already?”

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 5-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and Q&A Sections.”

Table 5-1 “Do I Know This Already?” Section-to-Question Mapping

| Foundation Topics Section | Questions |
|----------------------------|-----------|
| Data Access and Governance | 1, 2 |
| Amazon EBS Encryption | 3, 4 |
| Amazon S3 Bucket Security | 5, 6 |
| AWS Key Management Service | 7, 8 |
| AWS Certificate Manager | 9, 10 |

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. What AWS service assists in protecting access to AWS?
 - a. AWS Shield
 - b. Amazon Macie
 - c. Amazon EBS volumes
 - d. Amazon DynamoDB databases
2. What is the purpose of using detective controls?
 - a. To enable and enforce multifactor access
 - b. To detect and alert when security controls change
 - c. To manage AWS Organizations backups
 - d. To analyze compliance levels
3. Which of the following determines whether an attached Amazon EBS volume can be encrypted?
 - a. The type of Amazon EC2 instance
 - b. The size of the Amazon EBS volume
 - c. The type of the Amazon EBS volume
 - d. The IOPS assigned to the Amazon EBS volume
4. Where are data keys stored when they are delivered to an Amazon EC2 instance for safekeeping?
 - a. The associated Amazon EBS volume
 - b. Unsecured RAM
 - c. Secured RAM
 - d. AWS Key Management Service
5. What security policy allows multiple AWS accounts to access the same Amazon S3 bucket?
 - a. Amazon IAM policy
 - b. AWS IAM server control policy

- c. Amazon S3 Bucket policy
 - d. Amazon IAM policy
- 6. What type of encryption can be carried out before uploading objects to Amazon S3 to ensure absolute encryption outside AWS control?
 - a. RSA encryption
 - b. AES 128-bit encryption
 - c. Client-side encryption
 - d. Server-side encryption
- 7. What is the advantage of importing your organization's symmetric keys into AWS KMS?
 - a. High level of compliance
 - b. Faster encryption and decryption
 - c. Absolute control of encryption keys
 - d. None
- 8. What additional AWS service can work with AWS KMS as a custom key store?
 - a. Encrypted EBS volume
 - b. Encrypted Amazon S3 bucket
 - c. AWS CloudHSM
 - d. Encrypted AWS SQS queue
- 9. How does AWS charge for provisioning SSL/TLS certificates for AWS services using AWS Certificate Manager?
 - a. It charges per certificate per year.
 - b. It charges for private TLS certificates only.
 - c. It does not charge for AWS services.
 - d. It charges per certificate check.
- 10. Where are the security certificates for the AWS Application Load Balancer stored?
 - a. Amazon S3 bucket
 - b. Amazon EBS volume
 - c. AWS Certificate Manager
 - d. AWS KMS service

Foundation Topics

Data Access and Governance

Many on-premises and AWS-hosted workloads store their associated data records in the AWS cloud. Personal data stored in the public cloud is sometimes defined as personally identifiable information (PII). Sensitive data types, such as PII, must be protected to comply with privacy regulations such as the General Data Protection Regulation (GDPR), laws such as the Health Insurance Portability and Accountability Act (HIPAA), and industry standards such as the Payment Card Industry Data Security Standard (PCI DSS). More than 13 billion data records have been stolen since 2013, according to the *2022 Thales Data Threat Report* (<https://cpl.thalesgroup.com/data-threat-report>). AWS Artifact, located in the AWS Management console, provides on-demand access to all current AWS compliance and security reports, including Service Organization Control (SOC) and Payment Card Industry (PCI) reports and certifications from accreditation bodies validating the implementation and operating effectiveness of AWS security controls (see Figure 5-2).

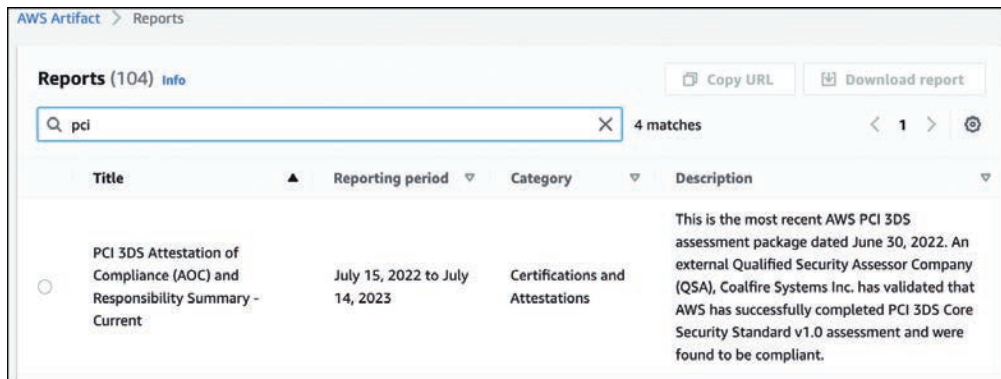


Figure 5-2 AWS Artifact PCI Report



Data Retention and Classification

When classifying data, it's important for each organization to implement data retention policies for each class of stored data. Organizations should design security policies using security zones for all data records, and data classification requirements based on how data is stored and who has access to it (see Figure 5-3). Defined security zones for data records range from highly protected to publicly accessible.

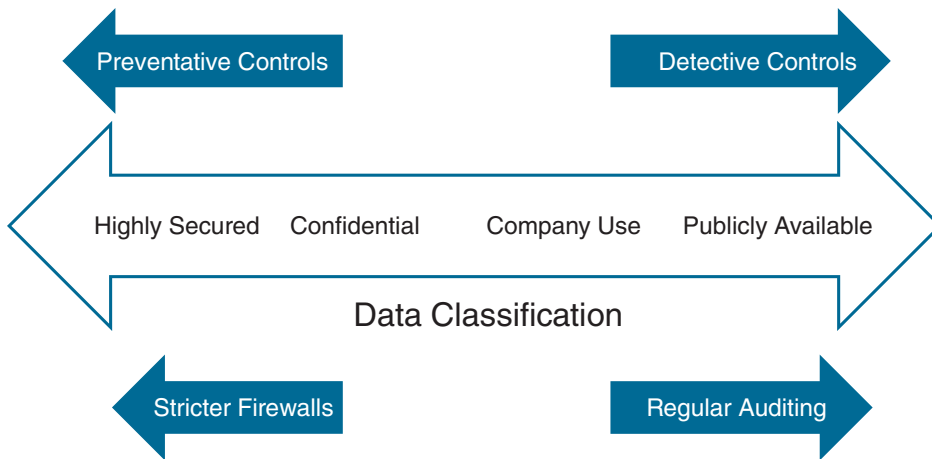


Figure 5-3 Classification of Data Records

Security zones are typically used to segregate different types of organizational data assets based on their sensitivity or importance, with the most sensitive or valuable data being placed in the highest security zone. This segregation enables organizations to implement different levels of security controls and access restrictions based on the sensitivity of the data, ensuring that only authorized users with the appropriate level of clearance can access and view sensitive data records.

Additionally, the creation of relevant security zones can help organizations prevent the spread of security breaches by limiting the potential impact to a specific area of the organization. Organizations also should create a network perimeter with defined network flow and access policies for data records defining where and how data can be accessed. Defense-in-depth security at AWS is applied using infrastructure security controls, AWS IAM security policies, and AWS detective controls (see Figure 5-4).

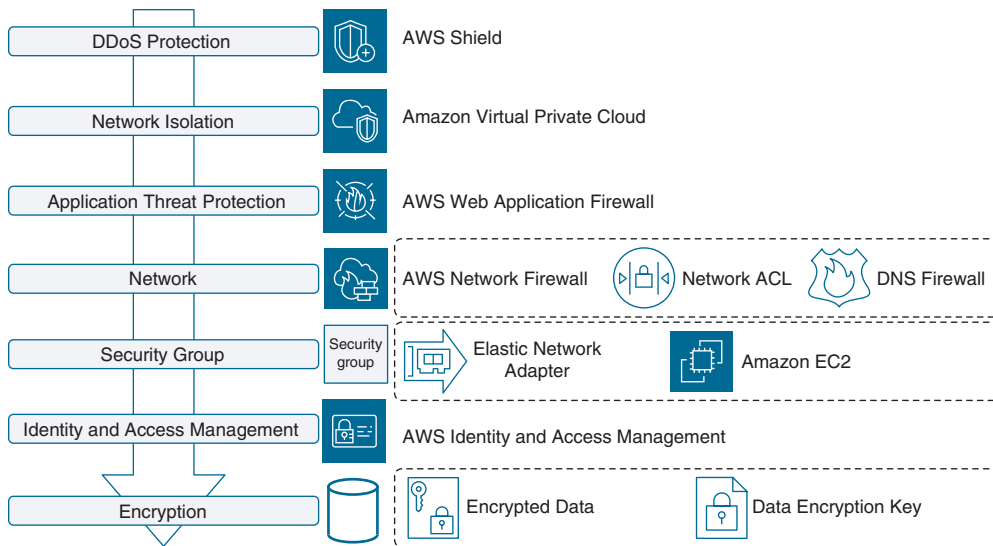


Figure 5-4 Preventative Controls

**Key
Topic**

Infrastructure Security

Infrastructure security requires deploying the following protections:

- **DDoS Protection:** Amazon deploys AWS WAF and Shield to protect the AWS cloud from DDoS attacks.
- **Network isolation:** EC2 instances must be hosted in a virtual private cloud (VPC). Many AWS services can be accessed from a VPC with private VPC endpoints (Interface and Gateway endpoints), ensuring workload traffic remains on the private AWS network.
- **Application-layer threat protection:** The AWS Web Application Firewall (WAF) allows organizations to create rules and filters to accept or reject incoming requests to Amazon CloudFront distributions, Amazon API Gateway deployments, and Application Load Balancers, and HTTP/HTTPS traffic to web servers.
- **Security groups:** Security groups must be designed to allow ingress traffic from associated security groups.
- **Network ACL:** Design network ACLs to implement zone-based models for your workload (web/app servers/database), allowing only legitimate traffic to reach each subnet.

IAM Controls

AWS Identity and Access Management (IAM) policies are useful for controlling access to the data layer (database, queue, AWS EBS volumes, shared data [AWS EFS and AWS FSx for Windows File Server], and Amazon S3 storage) and managing IAM user and federated user activity and infrastructure security. Separate administrative tasks should be created for Amazon RDS with IAM policies (see Example 5-1) that control access to database data records. For authentication and authorization to any workload or organizational data records, enable multifactor authentication (MFA) for all administrators and end users.

Example 5-1 Administrative Access to Amazon RDS

```

"Version": "2012-10-17",
  "Statement": [
    {
      "Sid": " Controlled Admin Tasks",
      "Effect": "Allow",
      "Action": [
        "rds:CreateDBSnapshot",
        "rds:StopDBInstance",
        "rds:StartDBInstance"
      ],
      "Resource": [
        "arn:aws:rds:[AWS_region]:[_AWS_account_
id]:snapshot:*",
        "arn:aws:rds:[AWS_region]:[_AWS_account_
id]:db:demoDB"
      ]
    },
    {
      "Sid": "DescribeInstances",
      "Effect": "Allow",
      "Action": "rds:DescribeDBInstances",
      "Resource": "*"
    }
  ]
}

```



Detective Controls

Detective controls are a type of security control designed to detect and alert when potential security incidents or breaches occur. Detective controls typically are used

with preventive and corrective controls forming a comprehensive security strategy. Examples of detective controls at AWS include intrusion detection systems, and auditing or logging systems that monitor user activity and alert on suspicious behavior. The goal of detective controls is to identify potential security threats or vulnerabilities before they can cause harm, allowing organizations to take appropriate action to prevent or mitigate the impact of a security incident.

Detective controls are an important part of a defense-in-depth security strategy as they provide an additional layer of protection by detecting and responding to potential security threats. Detective controls at AWS include the following security services:

- **VPC Flow Logs:** A feature of Amazon VPC that monitors network traffic at the elastic network interface, subnet, or entire VPC. Captured network traffic can be used for troubleshooting connectivity issues and to check current network access rules.
- **AWS CloudTrail:** Continuously monitor and record API usage and user activity across AWS infrastructure.
- **AWS CloudWatch:** Monitors AWS cloud services such as Amazon RDS databases, EC2 instances, and DynamoDB tables and hosted applications by collecting and tracking metric data, application and operating system log files, and using automated responses to defined alarms.
- **Amazon GuardDuty:** Provides continuous threat detection and analysis of VPC Flow Logs, Amazon Route 53 DNS query logs, and AWS CloudTrail S3 data event logs, and protecting AWS accounts and data stored in Amazon S3 from malicious activity. AWS GuardDuty malware protection can help detect malicious files stored on EBS volumes, protecting attached EC2 instances and Amazon Elastic Kubernetes Service (EKS) clusters.
- **AWS Config:** Detects configuration changes in RDS AWS infrastructure including Amazon RDS, EC2 instances, VPC and database architecture, including security groups, database instances, snapshots, and subnet groups.
- **Amazon Macie:** Uses machine learning and pattern matching to protect Amazon S3 objects and sensitive data types.
- **Access Analyzer for S3:** Monitors Amazon S3 buckets and details public or cross-account access.
- **Amazon Detective:** Graphically analyzes AWS CloudTrail management events, VPC Flow Logs, AWS GuardDuty findings, and Amazon EKS audit logs to help identify the cause of potential security issues.



Amazon EBS Encryption

Amazon Elastic Block Storage (EBS) volumes provide persistent block-level storage volumes for EC2 instances. They can be used to store a wide variety of data, including operating system files, application data, and database records. EBS volumes are automatically replicated within their availability zone to protect against data loss due to failure, and support a range of performance levels and storage options to meet the needs of different workloads.

Amazon Elastic Block Store (EBS) provides the option to encrypt EBS volumes to protect the data records. Encrypting EBS volumes ensures that the data cannot be read or accessed by unauthorized parties, even if the underlying storage volume is compromised. Encryption is performed using a customer master key and data key managed by the AWS Key Management Service (KMS), which provides a secure and auditable encryption service for managing data encryption at AWS using encryption keys. EBS volumes can be encrypted when first created, or volumes can be encrypted after they have been created. EBS also provides the option to encrypt snapshots of EBS volumes, enabling you to create encrypted backups of your EBS volumes.

Both EBS boot and data volumes can be encrypted. Most EC2 instances support EBS volumes' encryption, including the C4, I2, I3, M3, M4, R3, and R4 families. AWS has made the encryption process incredibly easy to deploy; when creating an EBS volume, merely checking off the option to enable encryption starts the encryption process (see Figure 5-5), which is managed by AWS Key Management Service (KMS). More details on AWS KMS are provided throughout this chapter.

Encryption [Info](#)

Use Amazon EBS encryption as an encryption solution for your EBS resources associated with your EC2 instances.

Encrypt this volume

KMS key [Info](#)

(default) aws/ebs ↕ ↻

KMS key description

Default master key that protects my EBS volumes when no other key is defined

Figure 5-5 Enabling EBS Encryption

NOTE Data encrypted using the EBS encryption process is encrypted before it crosses the AWS private network. Data also remains encrypted in-flight and at rest and remains encrypted when a snapshot is created of an encrypted volume.

The CMK protects all the other keys issued for data encryption and decryption of your EBS volumes within your AWS account. All AWS KMS-issued CMKs are protected using envelope encryption, which means AWS is responsible for creating and wrapping the “envelope” that contains the CMKs of the respective AWS account. Envelope encryption encrypts the plaintext data with a data key, and then encrypts the data key using a key that is managed by the AWS Key Management Service (KMS). KMS keys are created inside AWS KMS and never leave AWS KMS unencrypted. AWS cryptographic tools and services support the Advanced Encryption Standard (AES) with 128-, 192-, or 256-bit keys. AES is combined with Galois/Counter Mode (GCM), which provides high-performance *symmetric key* operation using a block size of 128 bits and is used by AWS KMS. AES and GCM are documented as AES-GCM.

After enabling your customer key using KMS for your AWS account, for additional security, it’s a good idea to add another key administrator and to allow key rotation of your Customer Master Keys. Administrators can use the KMS master key provided to create additional AWS KMS administrators, and to optionally enable key rotation of the CMK (see Figure 5-6).

**Key
Topic**

General configuration

| | | |
|---|---|---|
| Alias cloudtrail | Status Enabled | Creation date Oct 06, 2020 23:03 EDT |
| ARN arn:aws:kms:us-east-1:3138:58614000:key/bfe3c811-0430-4348-8432-8f4984293d78 | Description The key created by CloudTrail to encrypt log files. Created Wed Oct 07 03:03:38 UTC 2020 | Regionality Single Region |

Key policy | Cryptographic configuration | Tags | **Key rotation** | Aliases

Key rotation Save

Automatically rotate this KMS key every year. [Learn more](#)

Figure 5-6 Enabling Key Rotation

To encrypt an EBS volume using the AWS Key Management Service, a CMK can be created by AWS and stored in AWS KMS. Optionally, organizations can choose to specify the key material for the CMK, which can be generated by KMS or imported from your own key management infrastructure. After a CMK has been created, you

can create an encrypted EBS volume using the EC2 dashboard and specifying the ID of the CMK when creating the volume (see Figure 5-7). The EBS volume will be encrypted using the specified CMK, and the data on the EBS volume will be encrypted at rest on the underlying storage.

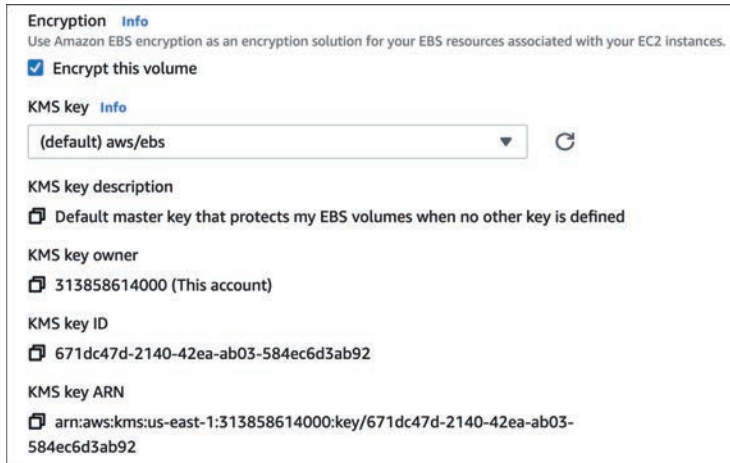


Figure 5-7 Select KMS Key

When you attach the encrypted EBS volume to an EC2 instance, the instance will automatically download and install the necessary encryption and decryption components, including the appropriate version of the AWS Encryption SDK and the public key portion of the CMK. The instance will then use the CMK to encrypt and decrypt data as it is written to and read from the EBS volume. The private key portion of the CMK remains securely stored in AWS KMS, and is never made available to the EC2 instance.

When an EBS volume has been encrypted and attached to an EC2 instance, the following data types are encrypted:

- Data at rest inside the EBS volume
- All data that moves between the attached EBS volume and the EC2 instance
- All snapshots created from the EBS volume
- All volumes created from the encrypted snapshots

AWS KMS performs the following steps, as illustrated in Figure 5-8, to encrypt and decrypt the EBS volume:

- Step 1.** AWS EBS sends a request to KMS, specifying the CMK to use for the AWS EBS volume encryption.

- Step 2.** AWS KMS generates a new data key, encrypts it using the specified CMK, and sends the encrypted key to AWS EBS to be stored with the volume metadata.
- Step 3.** The Amazon EC2 service sends a decrypt request to KMS.
- Step 4.** EBS sends a request to KMS to decrypt the data key.
- Step 5.** KMS uses the CMK to decrypt the encrypted data key and sends the decrypted key to the EC2 service.
- Step 6.** EC2 stores the plaintext decrypted key in protected hypervisor memory on the bare-metal server where the EC2 instance is hosted and uses the key when required to perform decryption for the EBS volume.

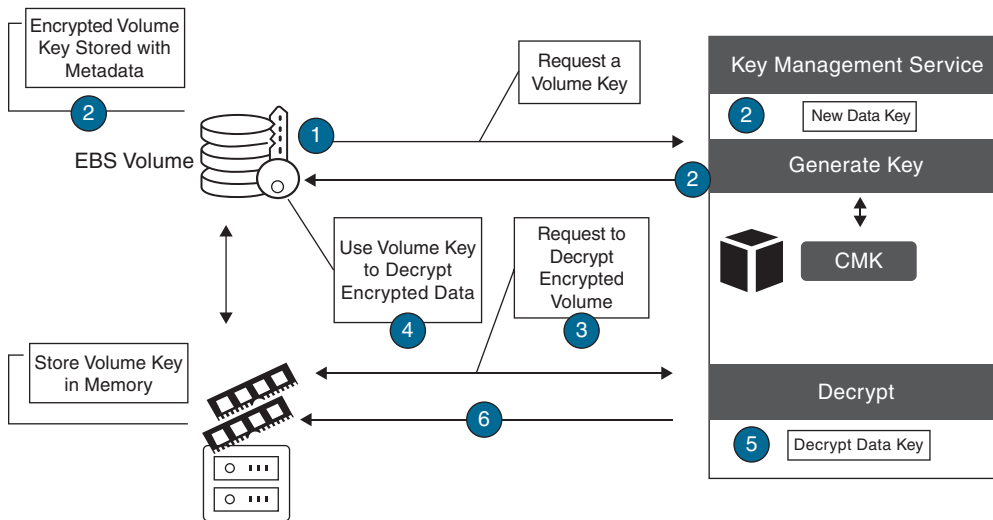


Figure 5-8 EBS Encryption Steps

NOTE The default setting for each AWS region is that EBS encryption is not enabled. To enable EBS encryption in the AWS region, open the EC2 dashboard, and in the upper-right corner under Account Attributes click EBS Encryption. Click Manage and choose the desired AWS-managed CMK or another CMK. Next, click Enable and then click Update EBS encryption. Once encryption is enabled for the AWS region, all new EBS volumes and snapshots will be encrypted at creation.

Amazon S3 Bucket Security

By default, only the owner who created an S3 bucket has access to the objects stored in the bucket. There are several methods for controlling security for an S3 bucket (see Figure 5-9):

- **ACLs:** You can use *access control lists (ACLs)* to control primary access from other AWS accounts for list and write objects and read and write bucket permissions, public access, and access to S3 logging information. ACLs are available for purposes of backward compatibility and are the weakest type of S3 security (and therefore not recommended).

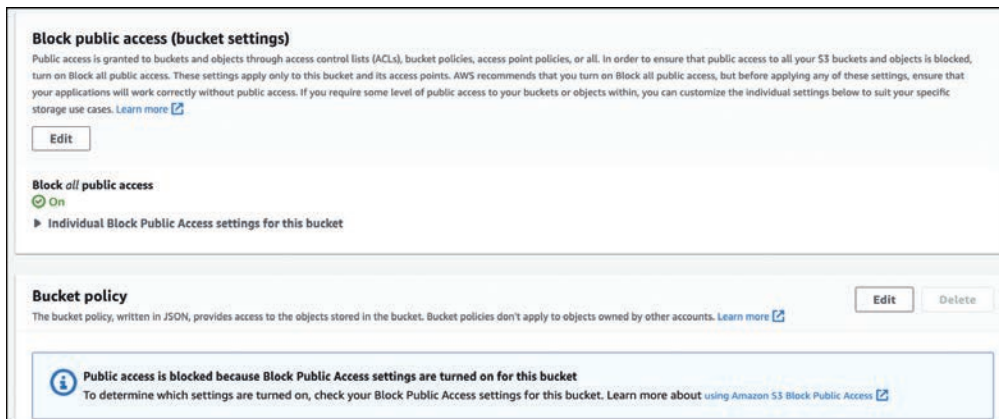


Figure 5-9 S3 Permission Settings

- **IAM policy:** You can grant access to other AWS users and groups of IAM users by using IAM permission policies in partnership with resource policies.
- **S3 Bucket policy:** You can control direct access to an S3 bucket, as shown in Example 5-2, by creating a *bucket policy* assigned directly to the S3 bucket. An S3 bucket policy is a JSON-formatted document that defines which actions are allowed or denied on an S3 bucket and its contents. A bucket policy is attached directly to the bucket it is protecting, and the policy settings list who has access to the bucket and what they can do with the objects in the bucket. An S3 bucket policy might allow a specific IAM user to read and write objects in the bucket, while denying access to all other users. Or, the policy might allow any user to read objects in the bucket but allow only authenticated users to write objects.

S3 bucket policies are defined using the AWS Policy Language, which provides a set of keywords and operations that you can use to specify the conditions under which a policy takes effect. A bucket policy can also allow access from multiple AWS accounts to a single S3 bucket.

Example 5-2 S3 Bucket Policy

```
{
  "Version": "2012-10-17",
  "Id": "S3PolicyId1",
  "Statement": [
    {
      "Sid": "IPAllow",
      "Effect": "Deny",
      "Principal": "*",
      "Action": "s3:*",
      "Resource": [
        "arn:aws:s3:::2021232reports",
        "arn:aws:s3:::2021232reports/*"
      ],
      "Condition": {
        "NotIpAddress": {"aws:SourceIp": "54.242.144.0/24"}
      }
    }
  ]
}
```

- **Query string authentication:** Query string authentication is a method to authenticate requests to an Amazon S3 bucket allowing organizations to generate a URL (see Figure 5-10) that can be shared with end users. When an end user clicks the URL, they are granted access to the specified S3 bucket and its contents.

Share "Drone Note E.mp3" with a presigned URL ✕

Presigned URLs are used to grant access to an object for a limited time. [Learn more](#)

ⓘ Anyone can access the object with this presigned URL until it expires, even if the bucket, and object are private.

Time interval until the presigned URL expires
 Using the S3 console, you can share an object with a presigned URL for up to 12 hours or until your session expires. To create a presigned URL with a longer time interval, use the AWS CLI or AWS SDK. Time intervals for presigned URLs can be restricted by your IAM policy.

Minutes
 Hours

Number of minutes

Must be a whole number between 1 and 720.

After you create the presigned URL, it's automatically copied to your clipboard.

Cancel Create presigned URL

Figure 5-10 Presigned URL for S3 Object Access

The URL includes a set of parameters that specify the credentials that grant access to the bucket. These parameters include the access key ID, an expiration time for the URL, and a signature that is calculated using the access key secret.

When someone attempts to access the URL, the Amazon S3 service checks the signature to verify that it matches the expected value. If the signature is valid, the user is granted access to the bucket; otherwise, the request is denied.

The use case for using query string authentication is useful for granting temporary access to an S3 bucket without having to create an IAM user or provide AWS access keys. However, query string authentication is not as secure as IAM policies or bucket policies because the URL and its parameters are included in each request; therefore, anyone who has access to the URL can potentially gain access to the bucket.

NOTE If you require public access to objects in an S3 bucket, it's recommended that you create a separate AWS account specifically for hosting the S3 buckets that will have public S3 object access.

- **Blocking S3 public access:** S3 Buckets always start as private, with no default public access (see Figure 5-11). When the Block Public Access (Bucket Settings) setting is enabled, attempts at changing security settings to allow public access to objects in the S3 bucket are denied. You can block public access on an individual S3 bucket or on all S3 buckets in your AWS account by editing the public access settings for your account using the S3 console. Choices for blocking S3 public access include the following:
 - **Public:** Everyone has access to list objects, write objects, and read and write permissions.
 - **Objects Can Be Public:** The bucket is not public; however, public access can be granted to individual objects by users with permissions.
 - **Buckets and Objects Not Public:** No public access is allowed to the bucket or the objects within the bucket.

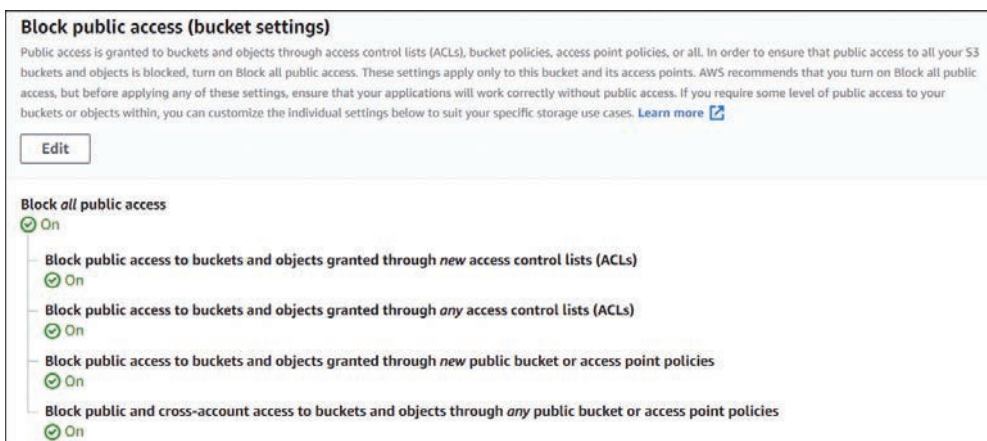


Figure 5-11 Blocking Public Access on an S3 Bucket by Default

NOTE Amazon Macie is a powerful AWS security service that uses artificial intelligence (AI) and machine learning (ML) technology to analyze your S3 objects and access patterns. Amazon S3 data can be classified based on many file formats, such as Personally Identifiable Information (PII) and other file types. AWS SNS notifications can be generated by Amazon Macie when Amazon S3 objects are discovered to be compromised.

**Key
Topic**

S3 Storage at Rest

For the AWS Certified Solutions Architect – Associate (SAA-C03) exam, the key topics to know about S3 storage at rest are as follows:

- SSE-S3:** With SSE-S3, Amazon S3 manages the encryption and decryption of the data in the bucket. Organizations that select this option don't manage the encryption keys but can access the data in the bucket without having to manage the keys. SSE-S3 uses the Advanced Encryption Standard (AES) algorithm with a 256-bit key to encrypt the data in the bucket. The key is automatically generated by Amazon S3 and is regularly rotated to ensure the security of the encrypted data (see Figure 5-12). Note that SSE encrypts the object data but the optional tag object metadata remains unencrypted.

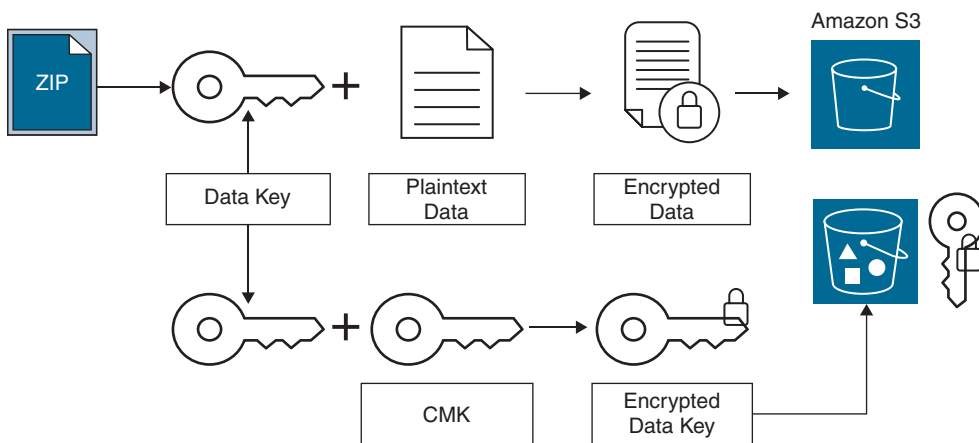


Figure 5-12 SSE-S3 Encryption Process

- SSE-KMS:** Organizations can select AWS KMS to manage their encryption keys. Select the default CMK or choose a CMK that was already created in AWS KMS before starting an S3 encryption process. Accessing encrypted objects managed by KMS can be expensive: If you have an exceptionally large number of encrypted objects, a large volume of decryption requests will be made to KMS. You can configure SSE-KMS to significantly reduce the cost of the encryption and decryption process. When an S3 Bucket Key is configured for SSE-KMS server-side encryption, a short-lived encryption key is created and stored and used to encrypt objects internally inside AWS S3 rather than utilize AWS KMS encryption processes. The S3 Bucket Key creates unique data keys for encrypting objects in the specific S3 bucket that has enabled the S3 Bucket Key option. The encryption process reduces AWS KMS requests

for external encryption keys and can reduce encryption costs by 99%. The S3 Bucket Key is a worker process within the S3 bucket that enables you to perform encryption services without constant communication with KMS.

- **SSE-C:** You can use SSE with a customer-provided encryption key. With each request, the encryption key is provided to AWS, and Amazon S3 manages the encryption and decryption of S3 objects by using the supplied key. The same encryption key that was used to encrypt the object must be provided before the object can be decrypted (see Figure 5-13). After the encryption process is complete, the supplied encryption key is deleted from memory. To upload an object with an organization-provided encryption key (SSE-C), the AWS CLI, AWS SDK, or Amazon S3 REST API must be used.

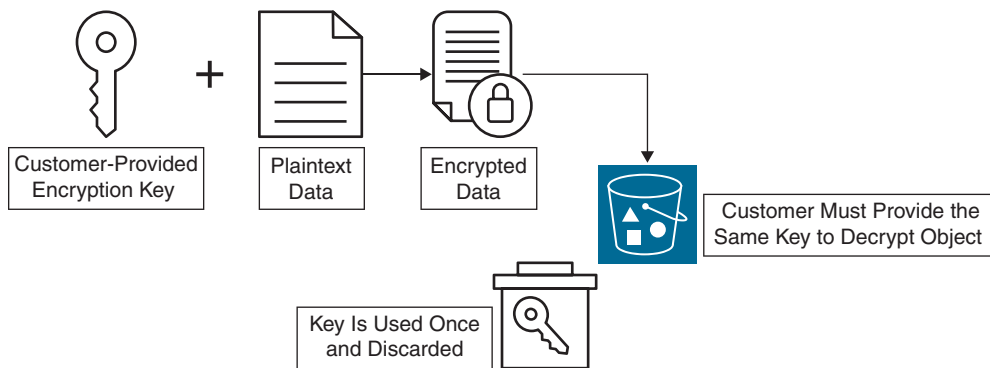


Figure 5-13 SSE-C Encryption Process

Key Topic

Amazon S3 Object Lock Policies

Amazon S3 buckets and Amazon S3 Glacier have data policies that can lock objects so they cannot be deleted or changed. Amazon S3 objects can be locked using a *write-once/read-many (WORM)* policy. Object lock policies enable you to set rules that restrict certain actions on objects, such as deleting or overwriting them, in order to protect objects and ensure they remain available and unaltered. Object lock policies are set at the S3 bucket level and apply to all objects in the bucket, or set on individual objects. This can be useful for complying with legal or regulatory requirements or protecting important or sensitive data. Apply a WORM policy, as shown in Figure 5-14, to stop an Amazon S3 object from being overwritten, or deleted for a fixed time period, or indefinitely. There are several options to WORM policies to understand. First is the *retention period*, which refers to a set number of days or years during which an object will remain

locked, protected, and unable to be overwritten or deleted. There are two retention modes:

- **Governance mode:** An S3 object cannot have its lock settings overwritten and cannot itself be overwritten or deleted unless the user has unique permissions. To override governance mode retention settings, an IAM user must have the **s3: BypassGovernanceRetention** permission and **x-amz-bypass-governance-retention: true** applied.
- **Compliance mode:** A protected object in your AWS account cannot be overwritten or deleted by anyone, including the root user, for the entire retention period.

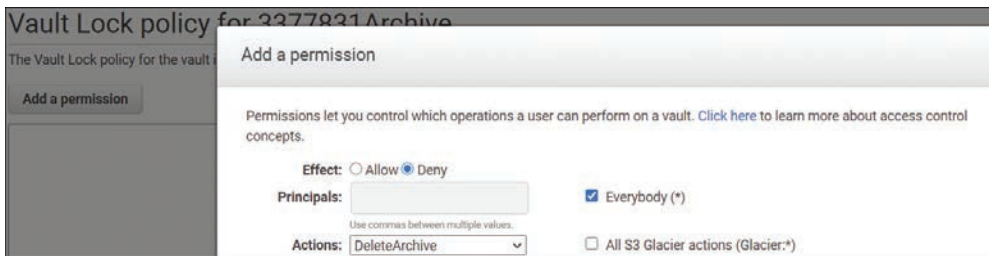


Figure 5-14 WORM Policy Settings

Legal Hold

An object lock allows you to place a legal hold on an S3 object. Legal hold provides the same protection as a previously discussed retention period but does not have an expiration date. Once in force, a legal hold remains in place until it is removed. An object lock works on S3 buckets that have versioning already enabled. Legal hold can be applied to a single S3 object. A legal hold can be placed and removed by any user with the **s3:PutObjectLegalHold** permission applied to their IAM user or group account they are a member of.

NOTE Object lock can only be enabled for new buckets when they are being created.



Amazon S3 Glacier Storage at Rest

Objects stored in Amazon S3 Glacier are automatically encrypted using SSE and AES-256 encryption. Amazon S3 Glacier Vault Lock enables you to deploy and

enforce regulatory and required compliance controls by applying a Vault Lock policy on an Amazon S3 Glacier vault. Once a WORM policy has been applied to an S3 Glacier vault, the policy cannot be changed.

NOTE Both EFS and FSx use AES-256 encryption to encrypt EFS data and metadata at rest. When your file system is mounted, you can also encrypt your EFS data in transit with TLS. FSx also supports the encryption of data in transit on file shares mapped on a computer instance that supports SMB Version 3.0 or newer. Encryption of data records at rest is automatically enabled when an FSx file system is created.

Data Backup and Replication

Amazon S3 object backups can be carried out with the services and utilities listed in Table 5-2. AWS Backup and AWS DataSync can back up additional AWS storage service data records.

Table 5-2 Data Backup and Replication Options

| AWS Service | Use | Data Types |
|--|---|--|
| AWS Backup | Back up all AWS storage services | EBS volumes and snapshots, S3 buckets, EFS, FSx for Windows File Server, RDS, DynamoDB |
| Amazon S3 Same-Region Replication (SRR) | Replicate objects to an S3 bucket in the same AWS region | Objects and versioned objects |
| Amazon S3 Cross-Region Replication (CRR) | Replicate objects to an S3 bucket in a different AWS region | Objects and versioned objects |
| Amazon S3 Multi-Region Access Points | Replicate data sets across multiple AWS regions | Objects and versioned objects |
| AWS DataSync | Copy data to and from AWS storage services | Network File System (NFS) or Server Message Block (SMB) shares, Hadoop Distributed File Systems (HDFS), AWS Snowcone, S3 buckets, EFS, FSx for Windows File Server |



AWS Key Management Service

AWS Key Management Service (KMS) lets organizations create, manage, and control cryptographic keys used to protect data records. AWS KMS integrates with AWS services that can encrypt data records (see Figure 5-15).

The screenshot shows the AWS KMS console interface. On the left is a navigation sidebar with 'Key Management Service (KMS)' selected. The main area displays 'AWS managed keys (13)' with a search bar and pagination controls. Below is a table listing several keys with their aliases, key IDs, and status.

| Aliases | Key ID | Status |
|-----------------------|--------------------------------------|---------|
| aws/lambda | 4e348669-5704-4079-922c-0e6559a47794 | Enabled |
| aws/acm | 5e734f45-b808-4279-a782-948455960f32 | Enabled |
| aws/ebs | 671dc47d-2140-42ea-ab03-584ec6d3ab92 | Enabled |
| aws/elasticfilesystem | 763b4b16-998c-4a54-ae8-eca63bd53cee | Enabled |
| aws/cloud9 | a12a8290-9390-4770-b537-b89fd6ecd52d | Enabled |

Figure 5-15 KMS Console

Organizations do not have to directly interface with AWS KMS to enable data encryption; instead, they can use AWS KMS services through more than 100 integrated AWS services, such as Amazon EBS storage, Amazon RDS, Amazon S3, Amazon EFS, Amazon FSx for Windows File Server, Amazon Aurora, and Amazon DynamoDB. When you enable encryption services using AWS KMS, a CMK is automatically generated in your AWS account for data encryption and decryption services. Organizations can choose to create one or more CMKs and use them to match their security requirements. A custom CMK allows you to control each key's access control and usage policy; you can also grant permissions to other AWS accounts and services to use a specific custom CMK.

You can also choose to create symmetric CMKs, which use the same key to encrypt and decrypt data, or asymmetric CMKs, which use a public/private key pair (one for encrypting and one for decrypting).

The most common way to use KMS is to choose which AWS service will encrypt your data and select the CMK from within the AWS service itself; for example, you can encrypt an RDS database volume, as shown in Figure 5-16.



Figure 5-16 Generating CMKs with KMS for an RDS Instance

Envelope Encryption

KMS uses a process called *envelope encryption* to encrypt data at rest. It involves two layers of encryption: the first layer encrypts the data using a key generated by the organization, and the second layer encrypts the customer-generated key using a key that is managed by the AWS Key Management Service (KMS). This process enables each organization to retain control over their encryption keys and also enables them to rotate and manage the keys as needed, while still benefitting from the security and reliability of using the KMS for encryption key management. When you need to encrypt data, KMS generates a data key that is used to encrypt the data locally within the AWS service or application. The data keys are also encrypted under the organization's CMK. When it's time to decrypt your data, a request is sent to KMS to decrypt the data key (that is, the data key copy that was stored with the encrypted data) using your CMK. The entire encryption or decryption process is logged in AWS CloudTrail for auditing purposes.

NOTE You can create up to 10,000 CMKs per AWS account per AWS region. Keys generated by AWS KMS can be enabled to be automatically rotated on an annual basis. However, automatic key rotation is not supported for external cryptographic keys imported into AWS KMS.

Organizations that choose to import 256-bit symmetric keys into AWS KMS for compliance requirements are responsible for managing the imported keys' expiration dates.

In addition to encrypting your data, AWS KMS provides other security features to help protect your encryption keys:

- **Key management:** As an administrator, you can create, rotate, disable, and delete the CMKs that are used to encrypt your data. You can also view the key policy for a CMK, which specifies who has access to the CMK and what actions they can perform with it.
- **Access control:** Organizations can use AWS IAM policies to control who has access to their CMKs and what actions can be performed with them. For example, users can be granted the ability to encrypt data using a specific CMK, but not to decrypt it or change the key policy.
- **Auditing:** AWS KMS logs all API calls to AWS CloudTrail so organizations can track who is using each CMK and for what purpose. Auditing can help ensure that encryption keys are being used securely and in accordance with an organization's security policies.
- **Key material:** KMS stores the key material for your CMKs in secure hardware devices called hardware security modules (HSMs). This helps protect the security of each organization's keys and ensures that they are only accessible to authorized users.
- **Key rotation:** CMKS can be configured to automatically be rotated on an annual basis, to help prevent security breaches.



AWS KMS Cheat Sheet

For the AWS Certified Solutions Architect – Associate (SAA-C03) exam, you need to understand the following critical aspects of AWS KMS:

- AWS KMS can be used to create symmetric keys within a custom key store such as AWS CloudHSM.
- An organization's symmetric keys can be imported for use with AWS KMS.
- AWS KMS can create symmetric and asymmetric data key pairs for application use.
- CMKs can be automatically rotated annually.
- CMKs can be disabled and re-enabled.
- AWS KMS keys can be audited with AWS CloudTrail.



AWS CloudHSM

Instead of using the default AWS KMS store, you can create a custom key store using a VPC-hosted AWS CloudHSM cluster and authorize KMS to use it as its dedicated key store. AWS CloudHSM clusters are created using multiple single-tenant hardware devices (see Figure 5-17). Amazon maintains the AWS CloudHSM hardware and backs up its contents but never enters an AWS CloudHSM device. Organizations might use an AWS CloudHSM deployment if compliance rules explicitly require that encryption keys are protected in a single-tenant hardware device. AWS CloudHSM can operate as a complete stand-alone hardware device for your synchronous and asynchronous keys and provide you with Federal Information Processing Standard (FIPS) 140-2 Level 3 compliance.




Figure 5-17 CloudHSM Design

AWS Certificate Manager

AWS Certificate Manager (ACM) is a managed service that allows you to provision, manage, and deploy public and private SSL/TLS certificates that can be used with your AWS services and AWS-hosted websites and applications. Certificates can also be deployed on ELB load balancers, CloudFront distributions, Elastic Beanstalk, and APIs hosted on Amazon API Gateway. There is no additional charge for provisioning public or private SSL/TLS certificates for use with AWS services. However, organizations will pay a fee for creating and operating a private *certificate authority (CA)* and for the private certificates that are issued by the private CA that is used by your internally hosted resources, such as application servers or appliances.

ACM can generate the following certificate types (see Figure 5-18):



- Public certificates:** ELB port 443 traffic, CloudFront distributions, and public-facing APIs hosted by Amazon API Gateway all use public certificates. Use AWS Certificate Manager to request a public certificate for a domain name for your site. AWS Certificate Manager validates that you own or control the domain name in your certificate request. Validation options include DNS validation and email validation.

- **Private certificates:** Delegated private certificates are managed by an AWS Certificate Manager–hosted private CA, which can automatically renew and deploy certificates for private-facing Amazon ELB and Amazon API Gateway deployments. Private certificates can also secure Amazon EC2 instances, Amazon ECS containers, and IoT devices.
- **Imported certificates:** Third-party certificates can be imported into AWS Certificate Manager.
- **CA certificates:** Certificates can be issued for creating a private CA up to five levels deep, including a root CA, three levels of subordinate CAs, and a single issuing CA.

Choose **Import a certificate** to import an existing certificate instead of requesting a new one. [Learn more.](#) [Import a certificate](#)

Request a certificate

Choose the type of certificate for ACM to provide.

Request a public certificate - Request a public certificate from Amazon. By default, public certificates are trusted by browsers and operating systems.

Request a private certificate - No Private CAs available for issuance. [Learn more.](#)

Figure 5-18 Certificate Choices in AWS Certificate Manager

Encryption in Transit

AWS uses HTTPS endpoints communication, providing encryption in transit for communicating with AWS APIs. AWS service endpoints can also be accessed using TLS version 1.2. Some AWS services offer endpoints that support the Federal Processing Standard (FIPS) 140-2 in some regions. Each endpoint is the URL of the entry point for each AWS service. AWS SDKs and the AWS Command Line Interface (AWS CLI) automatically use the default endpoint for each service per AWS Region, but an alternative endpoint can be specified for API requests. Most AWS services have regional endpoints that can be used to make requests. The format for a regional endpoint is *protocol://service-code.region-code.amazonaws.com*. AWS endpoints can be referenced here: <https://docs.aws.amazon.com/general/latest/gr/aws-service-information.html>.

Global endpoints are used for global services and services located in edge locations. The global AWS services are

- Amazon CloudFront
- AWS Global Accelerator

- AWS Identity and Access Management (IAM)
- AWS Organizations
- Amazon Route 53
- AWS Shield Advanced
- AWS WAF Classic

HTTP endpoints for domains and hosted workloads hosted at AWS can be blocked with Security Groups and Network ACLs and can automatically be redirected to HTTPS endpoints when using Amazon CloudFront or an Amazon ELB.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 16, “Final Preparation,” and the exam simulation questions in the Pearson Test Prep software online.

Review All Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the margin of the page. Table 5-3 lists these key topics and the page number on which each is found.



Table 5-3 Chapter 5 Key Topics

| Key Topic Element | Description | Page Number |
|-------------------|-----------------------------------|-------------|
| Figure 5-1 | Encryption Choices at AWS | 204 |
| Section | Data Retention and Classification | 207 |
| Section | Infrastructure Security | 209 |
| Section | Detective Controls | 210 |
| Section | Amazon EBS Encryption | 212 |
| Figure 5-6 | Enabling Key Rotation | 213 |
| Section | S3 Storage at Rest | 220 |
| Section | Amazon S3 Object Lock Policies | 221 |
| Section | Amazon S3 Glacier Storage at Rest | 222 |

| Key Topic Element | Description | Page Number |
|-------------------|---|-------------|
| Section | AWS Key Management Service | 224 |
| Section | AWS KMS Cheat Sheet | 226 |
| Section | AWS CloudHSM | 227 |
| List | AWS Certificate Manager certificate types | 227 |

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

Amazon Elastic Block Storage (EBS), symmetric key, access control list (ACL), bucket policy, write-once/read-many (WORM), AWS Key Management Service (KMS), certificate authority (CA)

Q&A

The answers to these questions appear in Appendix A. Use the Pearson Test Prep Software Online for more practice with exam format questions.

1. Which AWS storage service is available with AWS as a single-tenant storage design?
2. What is the default state of an S3 bucket regarding public access when the bucket is first created?
3. What is the security advantage of using SSE-C encryption with Amazon S3 buckets?
4. Describe the concept of envelope encryption that KMS uses.
5. What type of data stored at AWS is always automatically encrypted by default?
6. Why is AWS CloudHSM chosen by companies that must adhere to a high compliance standard?
7. How does AWS KMS carry out automatic key rotation for imported keys?
8. Where can private CAs created by AWS Certificate Manager be deployed?

This page intentionally left blank

Index

A

- Access Advisor, 131
- access key/s, 82, 751
 - IAM user, 92–94
 - rotating, 97–99
- access logs, 553
- access point, S3 (Simple Storage Service), 401–402
- account
 - access, 124–126
 - IAM (Identity and Access Management), 95–96
 - user, 88. *See also* user
- ACID (atomicity, consistency, isolation, durability), 509, 751
- ACL (access control list), 23, 113, 752
- actions, 87–88, 109–110, 546–547
- active-active failover, 340–343, 751
- adaptive capacity, 506–507
- AES (Advanced Encryption Standard), 21
- Agile, 62, 267
- alarm, 460–461, 751
- ALB (Application Load Balancer), 541–543
 - access logs, 553
 - cheat sheet, 553
 - health checks, 548–550
 - listeners and routing, 543–545
 - rules, conditions, and actions, 545–547
 - sticky session support, 551–552
 - target group attributes, 550–551
 - target groups, 547–548
- alias records, 352
- Amazon
 - A2C (App2 Container), 246
 - Amazon EventBridge, 256–258
 - API Gateway, 258–259
 - cheat sheet, 261–262
 - choosing the API protocol to use, 260–261
 - communication options, 259–260
 - selecting an authorizer, 261
 - Aurora, 340, 493–495
 - cheat sheet, 500–501
 - communicating with, 499–500
 - deployment options, 494–496
 - replication, 498–499
 - serverless, 674–675
 - storage, 496–498
 - CloudFront, 151, 238–239, 527, 751
 - cheat sheet, 536
 - edge functions, 534–536
 - how it works, 527–528
 - HTTPS access, 529–530
 - origin failover, 532–533
 - regional edge caches, 528–529
 - restricting distribution of
 - content, 532
 - serving private content, 530–532
 - use cases, 529
 - video-on-demand and live streaming, 533–534
 - Cognito, 176–177
 - federated identity pool, 179–180
 - user pool, 177–179
 - DynamoDB, 238, 299, 501–503

- Accelerator, 511
- ACID, 509
- adaptive capacity, 506–507
- backup and restore, 511–512
- cheat sheet, 512
- data consistency, 507–509
- data transfer costs, 683–685
- global tables, 510–511
- provisioning table capacity, 504–506
- tables, 503–504
- EBS (Elastic Block Storage), 212, 753
 - enabling, 212–213
 - enabling for each AWS region, 215
 - key rotation, 213–214
 - select KMS key, 214–215
- EC2 Image Builder, 435–436
- ECS (Elastic Container Service), 244, 443
 - task definition choices, 443–446
- EFS (Elastic File System), 379–380
 - cheat sheet, 383–384
 - lifecycle management, 383
 - performance modes, 380–381
 - security, 382
 - storage classes, 382
 - throughput modes, 381
- EKS (Elastic Kubernetes Service), 244, 446–447
- ElastiCache, 512–513, 751
 - for Memcached, 513–514
 - for Redis, 514–517
- ELB (Elastic Load Balancer), 539
 - access logs, 553
 - application load balancer
 - deployment, 541–545
 - costs, 695–696
 - features, 540–541
 - health checks, 548–550
 - rules, conditions, and actions, 545–547
 - sticky session support, 551–552
 - target group, 547–548
 - target group attributes, 550–551
- FSx for Windows File Server, 386–388
- Global Accelerator, 536–538
- Glue, 413
 - components, 413–414
 - ETL job flow, 414
- GuardDuty, 187–188
 - cheat sheet, 189
 - types of security analysis, 188–189
- Inspector, 195–196
- Kinesis, 417
- Lambda, 436–438
 - cheat sheet, 441
 - integration, 438–439
 - settings, 439–440
- Macie, 189–191, 219
- RDS (Relational Database Service). *See* RDS (Relational Database Service)
- Redshift, 517–520, 685–686
- Route 53, 59, 150, 348–349
 - alias records, 352
 - health checks, 349–350
 - resolver, 352–353
 - routing policies, 350
 - traffic flow policies, 351
- S3 (Simple Storage Service), 9, 216, 388–390
 - access points, 401–402
 - bucket concepts, 390–393
 - bucket policy, 217–218
 - cheat sheet, 403–404
 - data consistency, 393
 - Glacier storage at rest, 222–223
 - management, 396–400
 - multi-region access points, 402
 - object lock policies, 221–222
 - permission settings, 216–217
 - preselected URLs for objects, 403
 - presigned URL, 218–219
 - storage at rest, 220–221
 - storage classes, 394–396
- S3 Glacier, 404–405

- cheat sheet, 406
 - Deep Archive, 406
 - retrieval policies, 405–406
 - vault, 405
- SNS (Simple Notification Service), 248–249
 - cheat sheet, 250
 - creating a notification topic, 250
 - publisher and subscriber options, 249–250
- SQS (Simple Queue Service), 250–254
- VPC (Virtual Private Cloud), 15
- AMI (Amazon Machine Image), 428–430
- AWS Marketplace, 431–432
 - build considerations, 434–435
 - choosing, 430–431
 - custom, 432–434
 - custom instance store, 434
 - golden pipeline, 436
 - prebuilt, 430
 - Windows, 431
- analytical tools, 412–413, 415–416
- API (application programming interface), 18, 751
- API Gateway, 258–259
 - cheat sheet, 261–262
 - communication options, 259–261
 - selecting an authorizer, 261
- application/s
 - dependency, 27–28
 - deprecation, 28
 - integration services, 247–248
 - Amazon EventBridge, 256–258
 - Amazon SNS (Simple Notification Service), 248–250
 - Amazon SQS (Simple Queue Service), 250–254
 - AWS Step Functions, 254–256
 - load balancer, 23, 58. *See also* load balancer
 - migrating, 24
 - allow access to on-premises data records, 26
 - define a value proposition, 24–25
 - lift and shift, 26–27
 - solve a single problem, 26
 - start with low value/low risk, 25–26
 - replacing, 28
 - security, 23–24, 46
 - stateful, 239
 - stateless, 239–243
- archive, 405, 751
- ASG (Auto Scaling group), 465–466
 - lifecycle hooks, 472–473
 - management options, 470–471
 - scaling, 466–470
- asymmetric key, 751
- authentication
 - Amazon Cognito
 - federated identity pool, 179–180
 - user pool, 177–179
 - external, 81
 - IAM (Identity and Access Management), 82–84
 - multifactor, 80–81, 99
- authorization, IAM (Identity and Access Management), 85–87
- auto scaling, 461–463, 751
 - AWS, 473–474
 - cheat sheet, 473
 - EC2 (Elastic Compute Cloud), 463
 - ASG (Auto Scaling group), 465–471
 - launch configuration, 464
 - launch template, 464
 - lifecycle hooks, 472–473
 - termination policy, 471–472
- automatic failover, 60
- automation, 16
 - cooldown period, 471
 - Elastic Beanstalk, 279–281
 - modifying the capacity of the application infrastructure, 281
 - updating applications, 282–283

- Service Catalog, 277–279
- tools, 266–277
- availability, 293–295. *See also* high
 - availability; reliability
 - outages and, 306
 - workload, 48
- availability zone, 300–301, 752
 - distribution, 301–303
 - RDS (Relational Database Service), 488
 - storage, 329
- AWS. *See also* cloud computing;
 - Well-Architected Framework
 - analytical services, 415–416
 - Application Discovery Service, 26
 - Application Migration Service, 26
 - Architecture Center, 6
 - Artifact, 752
 - Artifact utility, 311
 - auto scaling, 473–474
 - availability zone, 300–303
 - Backup, 337, 618–619
 - cheat sheet, 620–621
 - lifecycle rules, 619–620
 - Budgets, 607–609
 - CDN, placement, 56–57
 - Certificate Manager, 227–228
 - cloud provider responsibilities, 20–21
 - cloud services, 15–16
 - Cloud9, 24–25
 - CloudFormation, 16, 268–269
 - components, 269
 - creating an EC2 instance, 273–274
 - stack sets, 276–277
 - stacks, 272–273
 - templates, 270–272
 - updating with change sets, 275
 - CloudHSM, 227
 - CloudTrail, 16, 191–192
 - cheat sheet, 194
 - creating a custom trail, 192–194
 - CloudWatch, 16
 - CodeCommit, 64
 - Cognito, 83
 - compute, 15, 55–56. *See also* compute
 - Config, 199–200, 600–602
 - Control Tower, 138–139
 - Cost and Usage Reports, 609–610
 - costs, calculating, 597–598
 - data backup and replication, 223–224
 - Data Lake, 407–409, 412–413
 - data replication, placement, 57–58
 - database/s, 481
 - DataSync, 384–385
 - Direct Connect, 149, 185–186, 752
 - cheat sheet, 187
 - gateway, 186–187
 - edge locations. *See also* edge locations
 - AWS Shield, 151–152
 - network services, 150–151
 - Elastic Beanstalk, 18, 67
 - essential characteristics, 6–8
 - failover, architecture, 60
 - GovCloud, 317, 318–319
 - IAS (Identity and Access Management), 14
 - Identity Center, 132–133
 - infrastructure, 16. *See also*
 - infrastructure
 - KMS (Key Management Service), 224
 - cheat sheet, 226–227
 - console, 224–225
 - envelope encryption, 225–226
 - Lake Formation, 409–411
 - Lambda, 238
 - load balancer, placement, 57–58
 - managed service, 19–20, 293, 308–310
 - management console, 9
 - Marketplace, 431–432
 - NIST compliance, 316–317
 - operational benefits, 19–20
 - Organizations, 134–136
 - outages, 306
 - Outposts, 7, 357
 - PaaS (platform as a service), 17–18

- RDS (Relational Database Service), 481–482
 - best practices, 491
 - cheat sheet, 493
 - engines, 482
 - failover, 487–488
 - high-availability design, 485–488
 - installation, 488–490
 - instance class types, 485
 - instances, 483–484
 - Multi-AZ deployment, 488
 - performance monitoring, 490–491
 - Proxy, 492–493
 - standby, 487
 - region, selection criteria, 310.
 - See also* region/s
 - compliance rules, 311–319
 - latency concerns, 319–320
 - pricing, 321
 - services, 320
 - regulatory compliance
 - rules, 311–314
 - standard/s, 315
 - Resource Access Manager, 136–138
 - Schema Conversion tool, 681
 - Secrets Manager, 194–195
 - security, 21
 - application, 23–24
 - data, 21–22
 - network, 22–23
 - self-service, 9
 - servers, 19
 - Service Catalog, 277–279
 - services, cheat sheet, 31–36. *See also* service/s
 - Shield, 150, 151–152
 - SLA (service-level agreement), 47–48
 - SP 800–145, “The NIST Definition of Cloud Computing”, 8–9
 - broad network access, 10–11
 - on-demand self-service, 9
 - measured service, 12–13
 - rapid elasticity, 11–12
 - resource pooling, 10
 - stateless processes, 68
 - Step Functions, 254–256
 - storage, 19, 362, 363–365. *See also* storage
 - Storage Gateway, 625–627
 - STS (Security Token Service), 120, 126–128
 - tiered pricing, 599–600
 - Trusted Advisor, 196–198
 - uptime, 331
 - user, 88–90
 - VMWare, 16
 - VPN (virtual private network)
 - route propagation, 184–185
 - solutions, 183–184
 - WAF (Web Application Firewall), 151
 - Well-Architected Framework, 4, 28–30, 39–40, 752
 - cost optimization, 51
 - operational excellence, 44–45
 - performance efficiency, 49–51
 - reliability, 47–49
 - security, 45–47
 - sustainability, 51–52
 - AZ (availability zone), 155
- ## B
- BAA (Business Associate Addendum), 316
 - backing services, 66–67
 - backup and restore, 223–224, 332–333
 - Amazon DynamoDB, 511–512
 - AWS Backup, 618–619
 - cheat sheet, 620–621
 - lifecycle rules, 619–620
 - database retention policies, 687–689
 - fast snapshot, 374–375
 - snapshot, 295, 362
 - warm standby, 337–339
 - Amazon Aurora, 340
 - multi-region, 339

- bastion host, 164–165
 - best practices
 - IAM (Identity and Access Management), 128–130
 - RDS (Relational Database Service), 491
 - Big Bang, 62
 - billing. *See also* cost/s; pricing
 - measured service, 12–13
 - traffic, 578–579
 - block storage, 362, 752
 - born-in-the-cloud mentality, 14
 - broad network access, 10–11
 - bucket, 752
 - policy, 217–218
 - S3 (Simple Storage Service), 390–393
 - versioning, 400–401
 - budget, 607–609
 - build stage, Elastic Beanstalk, 67
 - building, serverless web app, 262
 - create a static website, 263
 - create the backend components, 264
 - register for the conference, 266
 - set up the API gateway, 265
 - user authentication, 263–264
 - burst capacity, 752
 - burst credit, 369–370, 752
 - bursting mode, EFS, 381
 - business continuity, 60
 - BYOIP (Bring Your Own IP), 579–580
- C**
- CA (certificate authority), 752
 - canary deployment, 327
 - capacity units, 752
 - CDN
 - placement, 56–57
 - POP (point of presence), 56
 - change sets, 275
 - cheat sheet
 - ALB (Application Load Balancer), 553
 - Amazon API Gateway, 261–262
 - Amazon Aurora, 500–501
 - Amazon DynamoDB, 512
 - Amazon EFS (Elastic File System), 383–384
 - Amazon ElastiCache, 514, 516
 - Amazon Macie, 190–191
 - Amazon Redshift, 519–520
 - Amazon S3 Glacier, 406
 - Amazon SNS (Simple Notification Service), 250
 - Amazon SQS (Simple Queue Service), 253–254, 403–404
 - auto scaling, 473
 - AWS Backup, 620–621
 - AWS CloudTrail, 192–194
 - AWS Lambda, 441
 - AWS Storage Gateway, 625–627
 - CloudFront, 536
 - CloudWatch, 461
 - cost management, 610–611
 - data transfer costs, 686–687, 716–717
 - dedicated host, 637
 - disaster recovery, 344–345
 - EBS (Elastic Block Store), 372–373
 - FSx for Windows File Server, 388
 - IAM (Identity and Access Management), 132
 - IP address, 577–578
 - KMS (Key Management Service), 226–227
 - NACL (network access control list), 169–170
 - NAT (network address translation), 176
 - NoSQL costs, 676–680
 - RDS (Relational Database Service), 493, 671
 - route table, 158
 - service quota, 347–348
 - SG (security group), 161–162
 - single and multi-region recovery, 343–344
 - snapshot, 376–377
 - subnet, 572–573

- VPC (Virtual Private Cloud), 560–561
- CIDR block
 - primary, 566–568
 - secondary, 568–569
- Cloud CoE (Cloud Center of Excellence), 44
- cloud computing. *See also* Well-Architected Framework
 - availability, 293–295
 - AWS, essential characteristics, 6–8
 - failover, architecture, 60
 - IaaS (infrastructure as a service), 14–16
 - load balancer, placement, 57–58.
 - See also* load balancer providers, 39
 - public cloud, 6–7
 - reliability, 295–296
 - SaaS (software as a service), 13
 - service/s. *See also* service/s
 - CDN, 56–57
 - costs, 598–599
 - data replication, 57–58
 - data residency and compute locations, 55–56
 - placing, 55
 - shared responsibility model, 79–80
 - SP 800–145, “The NIST Definition of Cloud Computing”
 - broad network access, 10–11
 - on-demand self-service, 9
 - measured service, 12–13
 - rapid elasticity, 11–12
 - resource pooling, 10
- Cloud Foundry, 17
- Cloud9, 24–25
- CloudFormation, 16, 268–269
 - components, 269
 - creating an EC2 instance, 273–274
 - stack sets, 276–277
 - stacks, 272–273
 - templates, 270–272
 - updating with change sets, 275
- CloudFront, 151, 527
 - cheat sheet, 536
 - costs, 698–701
 - edge functions, 534–536
 - how it works, 527–528
 - origin failover, 532–533
 - regional edge caches, 528–529
 - restricting distribution of
 - content, 532
 - serving private content
 - HTTPS access, 530
 - using an origin access identity, 531–532
 - using signed URLs, 530–531
 - use cases, 529–530
 - video-on-demand and live streaming, 533–534
- CloudTrail, 16
- CloudWatch, 16, 53, 421, 447–448
 - alarm and action settings, 460–461
 - basic monitoring, 448–449
 - cheat sheet, 461
 - collecting data, 451–452
 - creating an alarm, 459–460
 - integration, 453–455
 - log group, 752
 - logs, 449–451
 - metrics, 555–556
 - planning for monitoring, 452–453
 - terminology, 455–459
- code repository, 63
- codebase, 63–64, 752
- CodeCommit, 64
- cold storage, 753
- command/s
 - create-policy, 105
 - iostat, 370
 - list-policies, 105
- compliance
 - NIST, 316–317
 - regulatory. *See* regulatory compliance

components

- Amazon SQS (Simple Queue Service), 251–253
- AWS CloudFormation, 269
- Composer, 65
- compute, 15, 425–427. *See also* EC2 (Elastic Compute Cloud)
- Amazon Lambda, 436–438
- EC2 (Elastic Compute Cloud), 427–428
 - AMI (Amazon Machine Image), 428–435
 - dedicated host, 636–637
 - dedicated instance, 638
 - on-demand instance service quotas, 641–643
 - on-demand pricing, 640–641
 - Fleet, 655
 - Image Builder, 435–436
 - instance choices, 634–636
 - instance purchasing options, 638–640
 - instance types, 633
 - placement groups, 638
 - pricing, 655
 - Reserved instance, 644–647
 - Savings Plans, 649–650
 - vCPU, 634
- matching utilization with requirements, 659–660
- optimizing, 656–659
- scaling, 661
- selecting a location, 55–56
- Spot Fleet, 651–653
- spot instance, 650–651
- tools and utilities, 655–656
- conditional policy, 86, 116
- configuration files, 66
- connection draining, 753
- connection tracking, 161
- connectivity options, VPC, 583
- containers and container management, 441–443
 - Amazon ECS (Elastic Container Service), 443–446
 - Amazon EKS (Elastic Kubernetes Service), 446–447
 - migrating applications to, 246
 - orchestration, 244–245
 - Control Tower, 138–139
 - controlled storage, 373
 - controls
 - detective, 210–212
 - IAM (Identity and Access Management), 210
 - cooldown period, 471, 753
 - corporate mindset, 13
 - Cost Explorer, 604–607
 - cost/s
 - allocation tags, 612–613, 753
 - AWS, 321, 597–598
 - cheat sheet, 610–611
 - cloud service, 598–599
 - CloudFront, 698–701
 - data transfer, 681–682, 706–707
 - accessing AWS services in different regions, 710–713
 - accessing AWS services in the same region, 707–709
 - cheat sheet, 686–687
 - DocumentDB, 686–687
 - DynamoDB, 683–685
 - edge locations, 713
 - network, 714
 - public versus private traffic charges, 714
 - RDS, 682–683
 - Redshift, 685–686
 - workload components in the same region, 709–710
 - ELB (Elastic Load Balancer), 695–696
 - management tools, 602–604
 - NAT (network address translation), 696–697

- network services from on-premises
 - locations, 703–705
 - optimization, 51
 - protection, 152
 - reliability, 295
 - storage, 613–617
 - createdBy tag, 612
 - create-policy command, 105
 - creating
 - CloudTrail trail, 192–194
 - IAM policy, 105–106
 - IAM user, 91–92
 - VPC (Virtual Private Cloud), 561–564
 - Credential Report, 130
 - CRM (customer relationship management), 156
 - cross-account access, 124–126
 - CUR (Cost and Usage Report), 753
 - custom AMI, 432–434
 - custom policy, 102
 - custom route table, 155–158
 - custom SG (security group), 162–163
 - customer gateway, 182–183
- D**
- dashboard, IAM, 79
 - data
 - access, governance, 207
 - classification, 207–209
 - consistency, 507–509, 753
 - lake, 407–409
 - replication, 57–58, 223–224
 - security, 21–22
 - stateful, 239–243
 - stateless, 239–243
 - storage. *See* storage
 - structured, 411–412
 - transfer, 621–625, 753
 - accessing AWS services in different regions, 710–713
 - accessing AWS services in the same region, 707–709
 - costs, 706–707
 - costs cheat sheet, 716–717
 - edge locations, 713
 - public versus private traffic
 - charges, 714
 - workload components in the same region, 709–710
 - unstructured, 412
 - database/s, 299
 - Amazon Aurora, 493–495
 - cheat sheet, 500–501
 - communicating with, 499–500
 - deployment options, 494–496
 - replication, 498–499
 - serverless, 674–675
 - storage, 496–498
 - Amazon DynamoDB, 501–503
 - Accelerator, 511
 - ACID, 509
 - adaptive capacity, 506–507
 - backup and restore, 511–512
 - cheat sheet, 512
 - data consistency, 507–509
 - global tables, 510–511
 - provisioning table capacity, 504–506
 - tables, 503–504
 - Amazon Redshift, 517–520
 - AWS, 15, 481
 - data transfer costs, 681–682
 - DocumentDB, 686–687
 - DynamoDB, 683–685
 - RDS, 682–683
 - Redshift, 685–686
 - design choices, 668
 - migration, 680–681
 - NoSQL, 675–677
 - costs cheat sheet, 676–680
 - service comparisons, 676–677
 - RDS (Relational Database Service), 481–482, 668–670
 - best practices, 491
 - cheat sheet, 493

- costs cheat sheet, 671
- design solutions, 672–675
- engines, 482
- failover, 487–488
- high-availability design, 485–488
- installation, 488–490
- instance class types, 485
- instances, 483–484
- Multi-AZ deployment, 488
- performance monitoring, 490–491
- Proxy, 492–493
- read replica, 673
- standby, 487
- retention policies, 687–689
- schema conversion, 681
- SQL (Structured Query Language), 503
- declare and isolate dependencies, 65
- dedicated host, 636–637
- dedicated instance, 638
- default VPC, 569–570
- Defense in Depth, 45–47, 753
- on-demand instance service quotas, 641–643
- on-demand self-service, 9
- dependency/ies, 753
 - application, 27–28
 - declare and isolate, 65
 - infrastructure-level, 63–64
 - manager, 65
 - workload, 48, 54
- deployment
 - Amazon Aurora, 494–496
 - canary, 327
 - Multi-AZ, 488
 - pilot light, 333–337
- detective controls, 210–212
- development, 70
 - Agile, 61–62
 - Big Bang, 62
 - frameworks, 66
 - Waterfall, 61–62
- DevOps, 267

- Direct Connect, 149, 185–186, 753
 - cheat sheet, 187
 - gateway, 186–187
- disaster recovery, 54
- distributed design, 321–322
 - high availability and fault tolerance, 322–325
 - removing single points of failure, 325–327
- distributed session management, 243–247, 753
- DocumentDB, data transfer costs, 686–687
- DR (disaster recovery), 330, 331
 - backup and restore, 332–333
 - cheat sheet, 344–345
 - pilot light, 333–337
 - warm standby, 337–339

E

- EBS (Elastic Block Storage), 50, 365–366, 751, 753
 - attaching a volume, 371–372
 - cheat sheet, 372–373
 - encryption, 212, 294–295
 - enabling, 212–213
 - enabling for each AWS region, 215
 - key rotation, 213–214
 - select KMS key, 214–215
 - multi-attach, 366
 - recycle bin, 376
 - snapshot, 373
 - administration, 375–376
 - cheat sheet, 376–377
 - fast restore, 374–375
 - taking from a Linux instance, 373–374
 - taking from a Windows instance, 374
 - volume types, 367–369
 - elastic EBS, 370–371
 - General Purpose SSD, 369–370
- EC2 (Elastic Compute Cloud), 80, 751.
 - See also* Reserved instance
 - access to AWS resources, 119–121
 - auto scaling

- ASG (Auto Scaling group), 465–471
 - launch configuration, 464
 - launch template, 464
 - bastion host, 164–165
 - dedicated host, 636–637
 - on-demand pricing, 640–641
 - Fleet, 655
 - Image Builder, 435–436
 - immutable infrastructure, 327
 - instance
 - choices, 634–636
 - dedicated, 638
 - on-demand service quotas, 641–643
 - purchasing options, 638–640
 - Reserved, 644–647
 - storage volume, 377–378
 - types, 633
 - placement groups, 638
 - pricing, 655
 - Savings Plans, 649–650
 - spot capacity pool, 653–655
 - task definition choices, 443–446
 - vCPU, 634
 - ECS (Elastic Container Service), 443
 - edge locations, 303
 - AWS Shield, 151–152
 - data transfer costs, 713
 - scalable delivery, 238–239
 - WAF (Web Application Firewall), 152–153, 154–167
 - EFS (Elastic File System), 379–380
 - cheat sheet, 383–384
 - lifecycle management, 383
 - performance modes, 380–381
 - security, 382
 - storage classes, 382
 - throughput modes, 381
 - EKS (Elastic Kubernetes Service), 244, 446–447
 - elastic*, 305
 - Elastic Beanstalk, 18, 279–281
 - build stage, 67
 - modifying the capacity of the application infrastructure, 281
 - updating applications, 282–283
 - elastic EBS volumes, 370–371
 - elastic IP address, 575–577, 753
 - elasticity, 12, 462
 - encryption, 21
 - Amazon EBS (Elastic Block Storage), 212
 - enabling, 212–213
 - enabling for each AWS region, 215
 - key rotation, 213–214
 - select KMS key, 214–215
 - envelope, 225–226
 - field-level, 238–239
 - in transit, 228–229
 - endpoint, 753
 - services, 588–589
 - VPC (Virtual Private Cloud), 585
 - costs, 701–703
 - gateway, 585–586
 - interface, 586–588
 - entity, IAM (Identity and Access Management), 82
 - envelope encryption, 225–226
 - EOIG (egress-only Internet gateway), 753
 - ephemeral ports, 159, 165–167
 - ephemeral storage, 362–363, 754
 - event notification, 754
 - explicit allow permission, 94
 - external authentication, 81, 83
 - external connections, VPC (Virtual Private Cloud), 180–181
 - customer gateway, 182–183
 - VPG (virtual private gateway), 181–182
 - externally authenticated user, 754
- F**
- failover, 21, 330–331. *See also* DR (disaster recovery); high availability
 - active-active, 340–343
 - architecture, 60
 - multi-region, 555

- origin, 532–533
- RDS (Relational Database Service), 487–488
- FAQs, 4
- fast snapshot restore, 374–375
- fast startup, 69–70
- fault tolerance, 288, 293, 322–325
- federated identity pool, Amazon Cognito, 179–180
- federation
 - SAML 2.0, 122–124
 - web identity, 121–122
- FedRAMP (Federal Risk and Authorization Management Program), 317, 754
- field-level encryption, 238–239
- Firecracker, 437–438
- firewall
 - NACL (network access control list), 168–169
 - cheat sheet, 169–170
 - implementation, 169
 - rule processing, 170–172
 - Web Application, 152
 - behaviors, 152–153
 - rules, 154–167
- FISMA, 317
- flow log, VPC, 172–174, 581–582
- FSx for Windows File Server, 386–388

G

- gateway endpoint, 585–586
- Gateway Load Balancer, 695
- gateway service, NAT (network address translation), 174–175
- General Purpose SSD, 369–370
- GitHub, 63
- Glacier, storage at rest, 222–223
- global service, 303
- global tables, Amazon DynamoDB, 510–511
- golden AMI pipeline, 436

- GovCloud, 317, 318–319
- governance, data access, 207
- graceful shutdown, 69–70
- group, IAM (Identity and Access Management), 82, 94
- GuardDuty, 187–188
 - cheat sheet, 189
 - types of security analysis, 188–189
- guardrails, 139

H

- health check, 754
 - ELB, 466, 548–550
 - Route 53, 349–350
- Heroku, 17–18, 60
- high availability, 21, 287, 288, 293, 754
 - distributed design, 322–325
 - endpoints, 304–305
 - failover strategies, 330–331
 - infrastructure, third-party solutions, 277
 - RDS (Relational Database Service), 485–488
- HIPAA (Health Insurance Portability and Accountability Act), regulatory compliance, 316
- horizontal scaling, 12, 51
- hosting, re-, 26–27
- hyperthreading, 634

I

- IaaS (infrastructure as a service), 6, 14–16
- IAM (Identity and Access Management), 14, 46, 79, 752
 - account, options, 95–96
 - actions, 87–88
 - authentication, 82–84
 - external, 83
 - multifactor, 99
 - authorization, 85–87
 - best practices, 128–130
 - cheat sheet, 132
 - controls, 210

- dashboard, 79
- entity, 82
- features, 80–81
- group, 82, 94, 754
- permission, explicit allow, 94
- policy, 81–82, 99–100
 - ACL (access control list), 113
 - actions, 109–110
 - conditional, 86, 116
 - creating, 105–106
 - elements, 106–107
 - identity-based, 100–102
 - inline, 104–105
 - managed, 100–101
 - password, 96
 - permission boundaries, 110–112
 - permissions, 114–115
 - resource-based, 102–104
 - rules, 107–109
 - service control, 112
 - session, 113–114
 - statement, 82, 107
 - trust, 118
 - version, 106, 115
- principal, 82
- requesting access to AWS resources, 84–85
- resource, 82
- role/s, 82, 118–119, 754
 - attaching to EC2 instance, 119–121
 - cross-account access, 124–126
 - SAML 2.0 federation, 122–124
 - service-linked, 119
 - for third-party access, 121
 - when to use, 119
- rotating access keys, 97–99
- security tools, 130–132
- service-linked roles, 80–81
- tags, 116–117
- user, 81–82, 88, 90–91
 - access keys, 92–94
 - creating, 91–92
 - signing in as, 94
- ID key, 83
- identity, 82. *See also* web identity
 - federation
 - based policy, 100–102
 - origin access, 531–532
- Identity Center, 132–133
- IG (Internet gateway), 569, 754
- Image Builder, 435–436
- immutable infrastructure, 327–329
- implementation, NACL (network access control list), 169
- infrastructure. *See also* network;
Twelve-Factor App Methodology
 - authentication, 266
 - automation, 277. *See also* AWS, Service Catalog; CloudFormation
 - AWS, 16
 - as code, 267
 - dependencies, 63–64
 - distributed design, 321–322
 - high availability and fault tolerance, 322–325
 - removing single points of failure, 325–327
 - immutable, 327–329
 - security, 209
 - zone
 - Local, 306–307
 - Wavelength, 308
- inline policy, 104–105
- installation, RDS (Relational Database Service), 488–490
- instance
 - Amazon RDS (Relational Database Service), 483–484
 - NAT (network address translation), 175–176
 - storage volume, 377–378
- integration and integration services
 - Amazon EventBridge, 256–258

- Amazon SNS (Simple Notification Service), 248–249
 - cheat sheet, 250
 - creating a notification topic, 250
 - publisher and subscriber options, 249–250
- Amazon SQS (Simple Queue Service), 250–251
 - cheat sheet, 253–254
 - compatibility with AWS services, 253
 - components, 251–253
 - triggered Lambda function, 251
- AWS Step Functions, 254–256
- CloudWatch, 453–455
- interface endpoint, 586–588
- intra-AZ connections, 302
- IOPS (input/output operations per second), 365, 754
- iostat command, 370
- IP address. *See also* BYOIP (Bring Your Own IP)
 - cheat sheet, 577–578
 - elastic, 575–577
 - private, 573–574
 - public, 574–575
- IPv6, 580–581
- ISO/IEC 27001 security standard, 80
- ITIL (Information Technology Infrastructure Library), 267

J-K

- Jassy, A., 6
- key, 390
- key rotation, EBS, 213–214
- key-value, 754
- KMS (Key Management Service), 224, 752, 754
 - cheat sheet, 226–227
 - console, 224–225
 - envelope encryption, 225–226

L

- labs, AWS Well-Architected, 4–5
- Lambda@Edge, 535–536, 754
- latency, region selection and, 319–320
- launch configuration, 464
- launch template, 464, 754
- LCU (Load Balancer Capacity Unit), 695, 755
- least privilege, 46
- lifecycle
 - hook, 472–473, 755
 - management, EFS (Elastic File System), 383
 - policy, 755
 - rules, 619–620, 755
- lift and shift, 26–27
- Linux, taking an EBS snapshot, 373–374
- listener, 543–545, 755
- list-policies command, 105
- live streaming, 238, 533–534
- load balancer, 240–241
 - Amazon ELB, 539
 - application load balancer
 - deployment, 541–545
 - costs, 695–696
 - features, 540–541
 - sticky session support, 551–552
 - application, 23, 541–543
 - access logs, 553
 - cheat sheet, 553
 - health checks, 548–550
 - listeners and routing, 543–545
 - rules, conditions, and actions, 545–547
 - sticky session support, 551–552
 - target group attributes, 550–551
 - target groups, 547–548
 - network, 554
 - cheat sheet, 554–555
 - multi-region failover, 555
 - placement, 57–58
- local instance storage, 377–378

Local Zone, 306–307, 755
 logs and logging, 70
 access, 553
 CloudWatch, 449–451
 flow, 172–174, 581–582

M

main route table, 155
 managed policy, 100–101
 managed service, 293
 AWS, 19–20
 Lambda@Edge, 535–536
 use cases, 308–310
 management console, AWS, 9
 management options, ASG (Auto Scaling group), 470–471
 measured service, 12–13
 Memcached, Amazon ElastiCache, 513–514
 metadata
 object, 391
 XML, 123
 metrics, CloudWatch, 53, 447, 455, 555–556
 MFA (multifactor authentication), 22, 80–81, 99, 755
 Microsoft Azure, 6, 39
 migration
 application, 24
 allow access to on-premises data records, 26
 define a value proposition, 24–25
 lift and shift, 26–27
 with many local dependencies, 27–28
 solve a single problem, 26
 start with low value/low risk, 25–26
 applications that should remain on premises, 28
 to containers, 246
 data transfer options, 621–625
 database, 680–681
 mindset
 born-in-the-cloud, 14

 corporate, 13
 startup, 14
 modular design, 237
 monitoring, 16, 490–491. *See also*
 CloudWatch
 multipart upload, 755
 multi-region warm standby, 339

N

NACL (network access control list), 168–169, 244, 755
 cheat sheet, 169–170
 implementation, 169
 rule processing, 170–172
 NAT (network address translation), 174
 cheat sheet, 176
 costs, 696–697
 gateway service, 174–175, 755
 instance, 175–176
 network, 51. *See also* edge locations
 access control list, 168–169, 244
 cheat sheet, 169–170
 implementation, 169
 rule processing, 170–172
 address translation, 174
 cheat sheet, 176
 gateway service, 174–175
 instance, 175–176
 BYOIP (Bring Your Own IP), 579–580
 data transfer costs, 714
 IP address
 elastic, 575–577
 private, 573–574
 public, 574–575
 load balancer, 58, 554
 cheat sheet, 554–555
 multi-region failover, 555
 resiliency, 304
 security, 22–23, 149–150, 151–152.
 See also security
 shared security model, 557–558
 terminology, 558–559

- topology, planning, 303–306
 - traffic charges, 578–579
 - VPC (Virtual Private Cloud), 154, 556–557. *See also* VPC (Virtual Private Cloud)
 - calculating number required, 564–565
 - cheat sheet, 560–561
 - connectivity options, 583
 - creating, 561–564
 - creating the CIDR block, 565–569
 - default, 569–570
 - endpoints, 585–589
 - flow log, 581–582
 - peering, 583–585
 - route table, 154–158
 - SG (security group), 158–168
 - subnet, 570–573
 - NIST (National Institute of Standards and Technology)
 - compliance, 316–317
 - SP 800–145, “The NIST Definition of Cloud Computing”, 8–9
 - broad network access, 10–11
 - on-demand self-service, 9
 - measured service, 12–13
 - rapid elasticity, 11–12
 - resource pooling, 10
 - Nitro, 755
 - non-persistent data store, 378
 - NoSQL, 675–677, 755
 - costs cheat sheet, 676–680
 - service comparisons, 676–677
 - NVMe (Non-Volatile Memory Express), 755
- O**
- OAI (origin access identity), 531–532, 756
 - object
 - metadata, 391
 - S3, 390
 - storage, 362, 756
 - object lock policy, Amazon S3, 221–222
 - operational benefits, AWS, 19–20
 - operational excellence, 43, 44–45
 - origin failover, 756
 - outages, 306
 - Outposts, 7, 357
- P**
- PaaS (platform as a service), 17–18
 - Cloud Foundry, 17
 - Elastic Beanstalk, 18
 - Heroku, 17–18
 - password policy, 96, 756
 - Paxos, 508
 - PCI DSS (Payment Card Industry Data Security Standard), compliance
 - checklist, 313–314
 - peering, 583–585, 756
 - performance
 - efficiency, 49–51
 - modes, EFS, 380–381
 - RDS (Relational Database Service), 490–491
 - and reliability, 54
 - Well-Architected Framework, 29
 - permission/s, 105, 114
 - Amazon S3, 216–217
 - boundaries, 110–112
 - explicit allow, 94
 - summary table, 114–115
 - PII (personally identifiable information), 207
 - pilot light, 333–337, 756
 - PIOPS (provisioned input/output operations per second), 365
 - placement group, EC2 (Elastic Compute Cloud), 638
 - placing cloud services, 55
 - CDN, 56–57
 - data replication, 57–58
 - data residency and compute locations, 55–56
 - load balancer, 57–58

- planning
 - network topology, 303–306
 - security group, 167–168
 - policy/ies
 - ACL (access control list), 113
 - bucket, 217–218
 - conditional, 86, 116
 - database retention, 687–689
 - IAM (Identity and Access Management), 81–82, 99–100
 - actions, 109–110
 - creating, 105–106
 - elements, 106–107
 - identity-based, 100–102
 - inline, 104–105
 - permission boundaries, 110–112
 - resource-based, 102–104
 - rules, 107–109
 - session, 113–114
 - statement, 82
 - version, 106
 - identity-based
 - custom, 102
 - managed, 100–101
 - lifecycle, 755
 - object lock, 221–222
 - password, 96
 - permissions, 105, 114–115
 - retrieval, 405–406
 - routing, 350
 - scaling, cooldown period, 471
 - service control, 112
 - stickiness, 552
 - termination, 471–472
 - traffic flow, 351
 - trust, 118
 - version, 115
 - WORM (write-once/read-many), 221
 - POP (point of presence), 56
 - port binding, 69
 - pricing
 - AWS, 321
 - CloudFront, 700–701
 - EC2 (Elastic Compute Cloud), 655
 - Reserved instance, 648–649
 - tiered, 599–600
 - primary CIDR block, 566–568
 - primary database, 756
 - principal, IAM (Identity and Access Management), 82
 - private IP address, 573–574
 - product, Service Catalog, 277–279
 - production, 70
 - providers, 39
 - provisioned mode, EFS, 381
 - public cloud, 6–7
 - public IP address, 574–575
- ## Q
- queue, 756
 - quotas
 - on-demand service, 641–643
 - service, 345–348, 391
- ## R
- rapid elasticity, 11–12
 - RDS (Relational Database Service),
 - 481–482, 668–670
 - best practices, 491
 - cheat sheet, 493
 - costs cheat sheet, 671
 - data transfer costs, 682–683
 - design solutions, 672–675
 - engines, 482
 - failover, 487–488
 - high-availability design, 485–488
 - installation, 488–490
 - instance class types, 485
 - instances, 483–484
 - Multi-AZ deployment, 488
 - performance monitoring, 490–491
 - Proxy, 492–493
 - read replica, 673
 - standby, 487

- read capacity unit, 756
- read replica, 673, 756
- recycle bin, EBS (Elastic Block Store), 376
- Redis, Amazon ElastiCache, 514–517
- redundancy, 21, 48, 54
- regional edge cache, 756
- regional Reserved instance, 647
- region/s, 296–299, 756
 - cheat sheet, 343–344
 - DR (disaster recovery)
 - backup and restore, 332–333
 - pilot light, 333–337
 - edge cache, 528–529
 - GovCloud, 318–319
 - selection criteria, 310
 - compliance rules, 311–319
 - latency concerns, 319–320
 - pricing, 321
 - services, 320
 - warm standby, 337–339, 340
- regulatory compliance, 207
 - HIPAA (Health Insurance Portability and Accountability Act), 316
 - rules, 311–314
 - standards, 315
- re-hosting, 26–27
- reliability, 287, 295–296, 757
 - and performance, 54
 - Well-Architected Framework, 29, 47–49
- replacing, applications, 28
- replication, 223–224
 - Amazon Aurora, 498–499
 - S3 (Simple Storage Service), 397–398
- Reserved instance, 644–645, 757
 - payment options, 646
 - pricing, 648–649
 - regional versus zonal, 647
 - reviewing monthly charges, 648
 - scheduled reservation, 646–647
 - scope, 647
 - term commitment, 645
 - types, 646
- resiliency, 237, 246–247, 288, 304
- resolver, Route 53, 352–353
- Resource Access Manager, 136–138
- resource pooling, 10
- resource/s
 - actions, 87–88
 - based policy, 102–104
 - IAM (Identity and Access Management), 82
 - requesting access, 84–85
- responsibilities, AWS cloud provider, 20–21
- retrieval policy, Amazon S3 Glacier, 405–406
- role/s
 - IAM (Identity and Access Management), 82, 118–119
 - attaching to EC2 instance, 119–121
 - cross-account access, 124–126
 - service-linked, 119
 - when to use, 119
 - SAML 2.0 federation, 122–124
 - for third-party access, 121
 - web identity federation, 121–122
- root user, 88–90
- Route 53, 59, 150, 348–349
 - alias records, 352
 - health checks, 349–350
 - resolver, 352–353
 - routing policies, 350
 - traffic flow policies, 351
- route propagation, 184–185
- route table
 - cheat sheet, 158
 - custom, 155–158
 - main, 155
- routing, ALB (Application Load Balancer), 543–545
- RPO (recovery point objective), 54, 331, 756
- RTO (recovery time objective), 54, 331, 756

- rules
 - actions, 546–547
 - Amazon ELB, 545–547
 - Amazon Inspector, 195–196
 - AWS Service Catalog, 278
 - compliance, 311–319
 - IAM policy, 107–109
 - lifecycle, 619–620, 755
 - NACL (network access control list), 170–172
 - regulatory compliance, 311–314
 - SG (security group), 162
 - WAF (Web Application Firewall), 152, 154–167
- S**
- S3 (Simple Storage Service)
 - batch operations, 396
 - bucket versioning, 400–401
 - inventory, 399
 - object lock, 396
 - replication, 397–398
- SAA-CO3 exam, 721–724
 - preparation tools, 726–731
 - sample questions, 5–6
 - scaled scoring, 4
 - scheduling, 725–726
 - tips, 724–725
 - updates, 749–750
- SaaS (software as a service), 13
- SAML 2.0 federation, 122–124
- sample questions, SAA-CO3 exam, 5–6
- Savings Plans, 649–650
- scale out, 757
- scaled scoring, 4
- scaling. *See also* auto scaling
 - ASG (Auto Scaling group), 466–470
 - auto, 461–463
 - AWS, 473–474
 - cheat sheet, 473
 - EC2 (Elastic Compute Cloud), 463–473
 - compute, 661
 - cooldown period, 471
 - horizontal, 12, 51
 - policy, 757
 - termination policy, 471–472
 - scope, Reserved instance, 647
 - Scrum, 267
 - SDN (software-defined network), 14
 - secondary CIDR block, 568–569
 - secret access key, 83
 - security. *See also* authentication; encryption
 - Amazon Macie, 189–191
 - AWS, 21
 - application, 23–24
 - data, 21–22
 - network, 22–23
 - controls
 - detective, 210–212
 - IAM, 210
 - Defense in Depth, 45–47
 - edge location, 150–151
 - AWS Shield, 151–152
 - WAF (Web Application Firewall), 152–154
 - EFS, 382
 - group, 23, 757
 - IAM (Identity and Access Management), 79. *See also* IAM (Identity and Access Management)
 - access keys, 92–94
 - account options, 95–96
 - ACL (access control list), 113
 - actions, 87–88
 - authorization, 85–87
 - best practices, 128–130
 - cheat sheet, 132
 - conditional policy, 116
 - custom policy, 102
 - dashboard, 79
 - entity, 82
 - explicit allow permission, 94

- external authentication, 83
- features, 80–81
- group, 82, 94
- identity-based policy, 100–102
- inline policy, 104–105
- managed policy, 100–101
- MFA (multifactor authentication), 99
- password policy, 96
- permission boundaries, 110–112
- policy, 81–82, 99–100
- policy, creating, 105–106
- policy actions, 109–110
- policy elements, 106–107
- policy rules, 107–109
- policy statement, 107
- policy version, 106, 115
- principal, 82
- requesting access to AWS resources, 84–85
- resource, 82
- resource-based policy, 102–104
- role, 82, 118–121
- rotating access keys, 97–99
- service-linked roles, 80–81
- session policy, 113–114
- signing in as a user, 94
- tags, 116–117
- tools, 130–132
- trust policy, 118
- user, 81–82, 90–91
- user, creating, 91–92
- infrastructure, 209
- network, 149–150
 - AWS Shield, 151–152
 - VPC, 154–176
- Well-Architected Framework, 29, 43–44, 45–47
- workshops, 5
- self-service, AWS, 9
- server
 - AWS, 19
 - immutable, 327–328
 - serverless, 237–238, 757
 - Amazon Aurora, 674–675
 - web app, building, 262
 - create a static website, 263
 - create the backend components, 264
 - register for the conference, 266
 - set up the API gateway, 265
 - user authentication, 263–264
 - service control policy, 112
- service/s
 - analytical, 415–416
 - AWS, 15–16
 - cheat sheet, 31–36
 - compute, 15
 - database, 15
 - monitoring, 16
 - PaaS, 17–18
 - storage, 15
 - VMWare, 16
 - AWS CloudTrail, 191–192
 - cheat sheet, 194
 - creating a custom trail, 192–194
 - AWS Config, 199–200
 - AWS Trusted Advisor, 196–198
 - backing, 66–67
 - backup and restore, 332–333
 - CDN, placement, 56–57
 - compute, 425–427
 - Amazon Lambda, 436–441
 - EC2, 427–436. *See also* EC2 (Elastic Compute Cloud)
 - container, 441–443. *See also* containers and container management
 - Amazon ECS (Elastic Container Service), 443–446
 - Amazon EKS (Elastic Kubernetes Service), 446–447
 - costs, 598–599
 - data replication, placement, 57–58
 - detective control, 211–212
 - endpoint, 588–589
 - global, 303

- IaaS, 14–16
- immutable infrastructure, 328–329
- linked roles, 80–81
- load balancer, placement, 57–58
- managed, 293, 308–310
- PaaS
 - Cloud Foundry, 17
 - Heroku, 17–18
- placing, 55–56
- quota, 3–4, 345–348, 757
- serverless, 237–238
- storage, 329–330
- tiered pricing, 599–600
- session policy, 113–114
- SG (security group), 158–161
 - administration access, 164–165
 - cheat sheet, 161–162
 - custom, 162–163
 - database server inbound ports, 163–164
 - ephemeral ports, 159, 165–167
 - planning, 167–168
 - rules, 162
 - web server inbound ports, 163
- shared memory segment, 371
- shared responsibility model, 79–80
- shared security model, 557–558
- signing in as user, IAM (Identity and Access Management), 94
- simple scaling, 757
- single points of failure, removing, 325–327
- SLA (service-level agreement), 14, 20–21, 47–48, 52–53, 294, 757
- SLI (service-level indicator), 52–53, 757
- SLO (service-level objective), 52–53, 757
- SMB (Server Message Block), 757
- snapshot, 295, 362, 757
 - cheat sheet, 376–377
 - EBS (Elastic Block Store), 373
 - administration, 375–376
 - taking from a Linux instance, 373–374
 - taking from a Windows instance, 374
 - fast restore, 374–375
 - Snow device, 757
 - SP 800–145, “The NIST Definition of Cloud Computing”, 8–9
 - broad network access, 10–11
 - on-demand self-service, 9
 - measured service, 12–13
 - rapid elasticity, 11–12
 - spot capacity pool, 653–655
 - Spot Fleet, 651–653
 - spot instance, 650–651
 - SQL (Structured Query Language), 503, 758
 - SSE (server-side encryption), 757
 - SSO (single sign-on), 83, 132–133
 - stack sets, 276–277
 - stacks, AWS CloudFormation, 272–273
 - staging, 70
 - standard/s
 - ISO/IEC 27001, 80
 - regulatory compliance, 311–312, 315
 - standby database, 757
 - startup, mentality, 14
 - stateful, 161, 239–243, 758
 - stateless, 239–243, 758
 - statement, policy, 82, 107
 - Step Functions, 254–256
 - step scaling, 758
 - sticky session, 243, 551–552, 758
 - storage, 329–330, 362
 - Amazon Aurora, 496–498
 - Amazon S3, 15
 - AWS, 19
 - block, 362
 - classes
 - EFS, 382
 - S3, 394–396
 - cold, 753
 - controlled, 373
 - costs, 613–617
 - EBS (Elastic Block Store), 365–366. *See also* EBS (Elastic Block Store)
 - administration, 375–376

- attaching a volume, 371–372
- cheat sheet, 372–373
- elastic EBS, 370–371
- fast snapshot restore, 374–375
- General Purpose SSD, 369–370
- multi-attach, 366
- recycle bin, 376
- snapshot, 373–376
- volume types, 367–369
- ephemeral, 362–363
- instance, 377–378
- object, 362
- resiliency, 246–247
- at rest, 220–221, 222–223
- workload requirements, 363–365
- streaming, 238
- structured data, 411–412
- STS (Security Token Service), 120, 126–128
- subnet, 570–573, 758
- sustainability, Well-Architected Framework, 30, 51–52
- symmetric key, 758
- syntax, IAM policy, 107–109

T

- T instance, 758
- tag
 - cost allocation, 612–613
 - createdBy, 612
 - IAM (Identity and Access Management), 116–117
- target group, 465–466, 547–548, 550–551, 758
- task definition, 443–446, 758
- template
 - AWS CloudFormation, 270–272
 - launch, 464
- termination policy, 471–472
- throughput modes, EFS, 381
- tiered pricing, 599–600, 758
- time

- availability, 293–295. *See also* availability up, 323, 331
- tool/s
 - analytical, 412–413
 - Artifact, 311
 - automation, 266–277
 - AWS Config, 600–602
 - AWS Schema Conversion, 681
 - cost management, 602–604
 - Budgets, 607–609
 - Cost and Usage Reports, 609–610
 - Cost Explorer, 604–607
 - exam preparation, 726–731
 - IAM (Identity and Access Management), 130–132
 - Well-Architected Framework, 5, 30–31
- traffic
 - billing, 578–579
 - flow policy, 351
- trust policy, 118
- TTL (time to live), 758
- Twelve-Factor App Methodology, 60–61, 62
 - declare and isolate dependencies, 65
 - execute an app as one or more stateless processes, 67–68
 - export services via port binding, 69
 - keep development, staging, and production similar, 70
 - maximize robustness with fast startup and graceful shutdown, 69–70
 - run admin/management tasks as on-off processes, 71
 - scale out via the process model, 69
 - separate build and run stages, 67
 - store configuration in the environment, 66
 - treat backing services as attached resources, 66–67
 - treat logs as event streams, 70
 - use one codebase, 63–64

U

- unstructured data, 412
- uptime, 323, 331, 758
- use cases
 - Amazon CloudFront, 529
 - Lambda@Edge, 535–536
 - managed service, 308–310
- user, 88
 - IAM (Identity and Access Management), 81–82, 90–91
 - creating, 91–92
 - signing in as, 94
 - root, 88–90
 - session management
 - distributed, 243–247
 - sticky sessions, 243
 - state, 758
- user pool, Amazon Cognito, 177–179

V

- value proposition, define, 24–25
- vault, Amazon S3 Glacier, 405
- vCPU (virtual CPU), 634
- versioning, 758
 - bucket, 400–401
 - IAM policy, 106, 115
- virtual machine/s, 8
- VM (virtual machine), 442
- VMWare, on AWS, 16
- Vogels, W., 288
- volume
 - EBS (Elastic Block Store)
 - attaching, 371–372
 - snapshot, 373
 - types, 369–371
 - instance storage, 377–378
- VPC (Virtual Private Cloud), 15, 22–23, 154, 556–557, 758
 - calculating number required, 564–565
 - cheat sheet, 560–561
 - connectivity options, 583
 - creating, 561–564

- creating the CIDR block, 565–566
 - primary, 566–568
 - secondary, 568–569
- default, 569–570
- endpoint/s, 585
 - costs, 701–703
 - gateway, 585–586
 - interface, 586–588
 - services, 588–589
- external connections, 180–181
 - customer gateway, 182–183
 - route propagation, 184–185
 - VPG (virtual private gateway), 181–182
- flow log, 172–174, 581–582
- NACL (network access control list), 168–169
 - cheat sheet, 169–170
 - implementation, 169
 - rule processing, 170–172
- NAT (network address translation), 174
 - cheat sheet, 176
 - gateway service, 174–175
 - instance, 175–176
- network terminology, 558–559
- peering, 583–585
- route table, 154
 - cheat sheet, 158
 - custom, 155–158
 - main, 155
- SG (security group), 158–161
 - administration access, 164–165
 - cheat sheet, 161–162
 - custom, 162–163
 - database server inbound ports, 163–164
 - ephemeral ports, 159, 165–167
 - planning, 167–168
 - rules, 162
 - web server inbound ports, 163
- shared security model, 557–558
- subnet, 570–573

VPG (virtual private gateway),
181–182, 758
VPN (virtual private network), 10, 149
AWS solutions, 183–184
route propagation, 184–185

W

WAF (Web Application Firewall), 23,
151, 152–153, 154–167
warm standby, 759
Amazon Aurora, 340
multi-region, 339
Waterfall, 61–62
Wavelength Zone, 308
WCU (write capacity unit), 759
web identity federation, 121–122
web server inbound ports, security
group, 163
Well-Architected Framework, 28–30,
39–40, 42,
287–288, 752
best practices, 42
cost optimization, 51
Microsoft Azure, 39
operational excellence, 43, 44–45
performance efficiency, 49–51

reliability, 47–49
security, 43–44, 45–47
sustainability, 51–52
tool, 30–31

Wiggins, A., 60

Windows

AMI (Amazon Machine Image), 431
taking an EBS snapshot, 374

workload, 293

availability, 48, 294
dependencies, 48, 54
reliability, 295

SLA (service-level agreement), 52–53
storage requirements, 363–365

workshop, AWS security, 5

WORM (write-once/read-many) policy,
221, 759

X-Y-Z

XML, metadata, 123

zonal Reserved instance, 647

zone. *See also* availability zone

Local, 306–307

Wavelength, 308