

ANTHONY SEQUEIRA , CCIE NO. 15626

Cert Guide

Learn, prepare, and practice for exam success



AWS Certified

SysOps Administrator Associate

(SOA-C01)

PEARSON IT
CERTIFICATION

FREE SAMPLE CHAPTER

SHARE WITH OTHERS



AWS Certified SysOps Administrator–Associate (SOA-C01) Cert Guide

Anthony Sequeira, CCIE No. 15626



Pearson

221 River St
Hoboken, NJ 07030

AWS Certified SysOps Administrator–Associate (SOA-C01) Cert Guide

Copyright © 2020 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.

For information regarding permissions, request forms, and the appropriate contacts within the Pearson Education Global Rights & Permissions Department, please visit www.pearson.com/permissions.

ISBN-10: 0-13-585325-7

ISBN-13: 978-0-13-585325-2

Library of Congress Control Number: 2019912877

ScoutAutomatedPrintCode

Editor-in-Chief

Mark Taub

Director, Product Management

Brett Bartow

Acquisitions Editor

Paul Carlstroem

Managing Editor

Sandra Schroeder

Development Editor

Christopher Cleveland

Project Editor

Mandie Frank

Copy Editor

Bart Reed

Technical Editor

Ryan Dymek

Editorial Assistant

Cindy Teeters

Designer

Chuti Prasertsith

Composition

codeMantra

Indexer

Tim Wright

Proofreader

Karen Davis

Figure 5-7	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure 5-8	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure 5-9	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure 5-10	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure 6-2	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure 6-3	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure 6-4	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure 6-5	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure 6-6	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure 6-7	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure 6-8	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure 6-9	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure 7-1	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure 7-2	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure 7-3	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure 7-4	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure 7-5	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure 7-6	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure 7-7	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure 7-8	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure 7-9	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure 7-10	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure 7-11	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure C-1	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure C-2	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure C-3	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure C-4	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure C-5	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure C-6	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Figure C-7	Screenshot of console.aws.amazon.com©Amazon Web Services, Inc.
Cover	ProStockStudio/Shutterstock

Contents at a Glance

	Introduction	xix
CHAPTER 1	Monitoring and Reporting	3
CHAPTER 2	High Availability	39
CHAPTER 3	Deployment and Provisioning	71
CHAPTER 4	Storage and Data Management	99
CHAPTER 5	Security and Compliance	129
CHAPTER 6	Networking	155
CHAPTER 7	Automation and Optimization	193
CHAPTER 8	Final Preparation	225
	Glossary of Key Terms	235
APPENDIX A	Answers to the “Do I Know This Already?” Quizzes and Q&A Sections	241
APPENDIX B	AWS Certified SysOps Administrator Associate (SOA-C01) Certification Guide Exam Updates	247
APPENDIX C	Select Frequently Asked Questions (FAQs)	249
	Index	271
Online Only Elements		
APPENDIX D	Study Planner	

This page intentionally left blank

Table of Contents

Introduction	xix
Chapter 1 Monitoring and Reporting	3
“Do I Know This Already?” Quiz	4
Performance and Availability Metrics	7
Accessing CloudWatch in AWS	7
<i>Amazon CloudWatch Console</i>	7
<i>AWS CLI</i>	7
<i>CloudWatch Query API</i>	11
<i>AWS SDKs</i>	13
Services Related to CloudWatch	13
Viewing Key CloudWatch Metrics for Various Services	15
Create and Maintain Metrics and Alarms	16
Using CloudWatch Dashboards	17
Using CloudWatch Metrics	21
Publishing Your Own Metrics	24
Using CloudWatch Alarms	26
Remediation Based on Metrics	31
Services That Publish Metrics to CloudWatch	31
Authentication and Access Control	34
Remediation of Issues Using CloudWatch: An Example	35
Review All Key Topics	37
Define Key Terms	37
Q&A	37
Chapter 2 High Availability	39
“Do I Know This Already?” Quiz	39
Implement Scalability and Elasticity	42
AWS Auto Scaling	42
Highly Available Versus Reliable and Resilient Environments	46
Limit Management	48
Networking	48
High Availability for Applications	50
SQS	52

SNS	58
RDS	60
ElastiCache	63
Multi-Region HA	65
Common Disaster Recovery (DR) Approaches	65
An HA Example Solution	67
Review All Key Topics	67
Define Key Terms	68
Q&A	68

Chapter 3 Deployment and Provisioning 71

“Do I Know This Already?” Quiz	71
Tools and Best Practices	74
The Importance of Automation	74
Deployment Strategies	75
<i>Provisioning Infrastructure</i>	75
<i>Deploying Applications</i>	75
<i>Configuration Management</i>	75
<i>Tagging</i>	76
<i>Custom Variables</i>	76
<i>Baking Amazon Machine Images (AMI)</i>	77
<i>Logging</i>	78
<i>Instance Profiles</i>	78
<i>Scalability Capabilities</i>	79
<i>Monitoring</i>	80
<i>Continuous Deployment</i>	80
<i>Elastic Beanstalk</i>	81
<i>Elastic Container Service</i>	83
<i>OpsWorks Stacks</i>	84
CloudFormation	86
AWS CLI	87
AWS Systems Manager	87
Deploying a REST API in API Gateway	88
Deploying Lambda Applications	91
Elastic Load Balancers	92

Troubleshoot and Remediate	93
EC2 Launch Issues	93
ELB Error Messages	94
ELB CloudWatch Metrics	95
CloudFormation Issues	96
Review All Key Topics	96
Define Key Terms	97
Q&A	97
Chapter 4 Storage and Data Management	99
“Do I Know This Already?” Quiz	99
Object and Block Storage	102
S3	102
<i>S3 Storage Classes</i>	<i>105</i>
<i>S3 Versioning</i>	<i>106</i>
MFA Delete	107
Lifecycle Policies	107
EBS	111
Other Storage Technologies	112
EFS	112
AMIs	113
AWS Storage Gateway	115
Snowball	117
<i>Snowball Edge</i>	<i>118</i>
Athena	119
Storage Encryption	120
AWS KMS	121
CloudHSM	122
S3 Client-Side Encryption	122
S3 Server-Side Encryption	122
EBS Volume Encryption	123
<i>Snapshots</i>	<i>125</i>
Review All Key Topics	126
Define Key Terms	126
Q&A	126

Chapter 5 Security and Compliance 129

“Do I Know This Already?” Quiz	129
The Shared Responsibility Model	132
Amazon Responsibilities	133
Client Responsibilities	134
Security Policies in AWS	135
DDoS Mitigation	135
<i>AWS Shield Standard</i>	137
<i>AWS Shield Advanced</i>	137
Data Encryption	138
Inventory and Configuration	139
Monitoring and Logging	139
Penetration Testing	140
Access Controls	140
Infrastructure Security	141
Identity and Access Management	141
Best Practices with IAM	148
Review All Key Topics	152
Define Key Terms	152
Q&A	153

Chapter 6 Networking 155

“Do I Know This Already?” Quiz	155
AWS Networking Features	157
AWS Global Infrastructure	157
<i>Regions</i>	157
<i>Availability Zones</i>	159
Edge Locations and CloudFront	160
Virtual Private Cloud	163
<i>The Default VPC</i>	165
<i>Network Interfaces</i>	166
<i>Route Tables</i>	168
<i>Internet Gateways</i>	170
<i>Egress-Only Internet Gateways</i>	171

<i>DHCP Option Sets</i>	172
<i>DNS</i>	174
<i>Elastic IP Addresses</i>	174
<i>VPC Endpoints</i>	175
<i>Interface Endpoints (Powered by AWS PrivateLink)</i>	176
<i>Gateway Endpoints</i>	176
<i>NAT</i>	177
AWS CLI	177
AWS Connectivity Services	178
Network to Amazon VPC	178
<i>Hardware VPN</i>	178
<i>Direct Connect</i>	180
<i>Direct Connect and VPN</i>	181
<i>VPN CloudHub</i>	182
<i>Software VPN</i>	183
Amazon VPC to Amazon VPC	184
<i>VPC Peering</i>	185
<i>Software VPN</i>	186
<i>Software-to-Hardware VPN</i>	186
<i>Hardware VPN</i>	186
<i>Direct Connect</i>	187
Internal User to Amazon VPC	187
Network Troubleshooting	187
Network Troubleshooting Tools	188
<i>VPC Flow Logs</i>	188
<i>Route 53 Record Routing Policies</i>	189
Review All Key Topics	190
Complete Tables and Lists from Memory	190
Define Key Terms	190
Q&A	191
Chapter 7 Automation and Optimization	193
“Do I Know This Already?” Quiz	193
Managing Resource Utilization	196

Prepare for Operational Excellence	197
Operate to Achieve Operational Excellence	200
Evolve for Operational Excellence	202
Best Practices	205
<i>Compute</i>	206
<i>Storage</i>	207
<i>Database</i>	208
<i>Network</i>	209
Trade-Offs	210
Key AWS Services	211
Monitoring	212
Cost Optimization Strategies	213
Best Practices	213
Cost Monitoring	215
Deploy Automation	218
Automation Tools and Techniques	218
<i>CodePipeline</i>	218
<i>CodeBuild</i>	218
<i>CodeDeploy</i>	218
<i>CodeStar</i>	219
<i>Elastic Container Service</i>	219
<i>Lambda</i>	219
<i>CloudFormation</i>	220
<i>OpsWorks</i>	220
<i>Systems Manager</i>	220
<i>AWS Config</i>	221
<i>CloudWatch</i>	221
<i>X-Ray</i>	221
<i>CloudTrail</i>	222
<i>Elastic Beanstalk</i>	222
<i>CodeCommit</i>	222
Automation Best Practices	222

Review All Key Topics	223
Define Key Terms	224
Q&A	224
Chapter 8 Final Preparation	225
Exam Information	225
Getting Ready	228
Tools for Final Preparation	229
Pearson Cert Practice Test Engine and Questions on the Website	229
<i>Accessing the Pearson Test Prep Software Online</i>	229
<i>Accessing the Pearson Test Prep Software Offline</i>	230
Customizing Your Exams	231
Updating Your Exams	232
<i>Premium Edition</i>	232
Chapter-Ending Review Tools	233
Suggested Plan for Final Review/Study	233
Summary	233
Glossary of Key Terms	235
APPENDIX A Answers to the “Do I Know This Already?” Quizzes and Q&A Sections	241
APPENDIX B AWS Certified SysOps Administrator Associate (SOA-C01) Certification Guide Exam Updates	247
APPENDIX C Select Frequently Asked Questions (FAQs)	249
Index	271

About the Author

Anthony Sequeira, CCIE No. 15626, is a seasoned trainer and author regarding various levels and tracks of Cisco, Microsoft, and AWS certifications. Anthony formally began his career in the information technology industry in 1994 with IBM in Tampa, Florida. He quickly formed his own computer consultancy, Computer Solutions, and then discovered his true passion—teaching and writing about information technologies.

Anthony joined Mastering Computers in 1996 and lectured to massive audiences around the world about the latest in computer technologies. Mastering Computers became the revolutionary online training company KnowledgeNet, and Anthony trained there for many years.

Anthony is currently pursuing his second CCIE in the area of Cisco Data Center! Anthony is happier than he has ever been in his career as a freelance author and trainer. Keep up with his latest projects at AJSnetworking.com.

Dedication

This book is dedicated to my best friend, Pierre Smith. Pierre, thanks for the lifetime of laughs mixed with great advice, and the occasional brilliant football bet.

Acknowledgments

This manuscript was made truly great by the incredible technical review of Ryan Dymek. Sometimes I think he might have invented AWS.

I would also like to express my gratitude to Chris Cleveland, the development editor of this book. I was so incredibly lucky to work with him again on this text. Like Ryan, he made this book several cuts above the rest.

Finally, thanks you so much to Paul Carlstroem. Paul very patiently made this book a reality.

About the Technical Reviewer

Ryan Dymek has been working with Amazon Web Services (AWS) for more than 9 years and holds all nine AWS certifications as well as various Google Cloud Platform (GCP) certifications. Ryan trains and advises some of the largest companies in the world on sound architectural practices in cloud strategy and DevOps principles. While working with business leaders, developers, and engineers, Ryan bridges the gap between business and technology, maintaining the understanding and skills required to be able to perform at a deep technical level. Ryan runs his own cloud consulting practice, advising more than 20 companies on the Fortune 500 list, and has helped many startups find their way in the cloud.

In addition to cloud and technical acumen, Ryan is a certified business coach personally trained by John Maxwell. He uses these professional skills not only to advise companies on best cloud practices but also on how to align with a business's needs and culture, making confident business and technical decisions and cultivating a transformation into DevOps.

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

Introduction

The AWS Certified SysOps Administrator–Associate is a cloud-related certification that tests a candidate’s ability to operate effective solutions by calling upon the most popular aspects of Amazon Web Services. The SysOps Administrator candidates must demonstrate their skills on how to effectively implement a sophisticated design that saves costs, is secure, and, perhaps most importantly, operates with excellence. Candidates are also required to know the most important facts regarding various services and their capabilities.

The AWS Certified SysOps Administrator–Associate is an Associate-level cloud career certification. This certification is an excellent second step after the achievement of the AWS Certified Solutions Architect–Associate certification. For some students, this certification might actually be their third step. This is due to the fact they may have started with the AWS Certified Cloud Practitioner exam, which is an entry-level exam considered by those who arrive to the study of AWS with little to no prior experience.

Following the SysOps Associate certification, AWS offers a Professional level of certification for the SysOps Administrator.

AWS also offers certifications you might be interested in for different tracks. For example, there is a Developer track for AWS that also includes Associate and Professional levels. There are also Specialty certifications that Amazon will use to deep-dive into many different areas such as security and advanced networking.

NOTE The AWS Certified SysOps Administrator–Associate 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 SysOps Administrator–Associate Certification

The AWS Certified SysOps Administrator–Associate certification is intended for individuals who have technical expertise in deployment, management, and operations on AWS. It seeks to validate that the candidate can do the following:

- Deploy, manage, and operate scalable, highly available, and fault tolerant systems on AWS.
- Implement and control the flow of data to and from AWS.
- Select the appropriate AWS service based on compute, data, or security requirements.

- Identify appropriate use of AWS operational best practices.
- Estimate AWS usage costs and identify operational cost control mechanisms.
- Migrate on-premises workloads to AWS.

Recommended Prerequisite Skills

While this text provides you with the information required to pass this exam, Amazon considers ideal candidates to be those that possess the following:

- Minimum of one year of hands-on experience with AWS
- Experience managing/operating systems on AWS
- Understanding of the AWS tenets—architecting for the cloud
- Hands-on experience with the AWS CLI and SDKs/API tools
- Understanding of network technologies as they relate to AWS
- Understanding of security concepts with hands-on experience in implementing security controls and compliance requirements

The Exam Objectives (Domains)

The AWS Certified SysOps Administrator–Associate exam is broken down into five major domains. The contents of this book cover each of the domains and the sub-topics included in them, as illustrated in the following descriptions.

The following table breaks down each of the domains represented in the exam.

Domain	Percentage of Representation in Exam
1: Monitoring and Reporting	22%
2: High Availability	8%
3: Deployment and Provisioning	14%
4: Storage and Data Management	12%
5: Security and Compliance	18%
6: Networking	14%
7: Automation and Optimization	12%
	Total 100%

Here are the details of each domain:

Domain 1: Monitoring and Reporting: This domain is covered primarily in Chapter 1.

- 1.1 Create and maintain metrics and alarms utilizing AWS monitoring services
- 1.2 Recognize and differentiate performance and availability metrics
- 1.3 Perform the steps necessary to remediate based on performance and availability metrics

Domain 2: High Availability: This domain is covered primarily in Chapter 2.

- 2.1 Implement scalability and elasticity based on use case
- 2.2 Recognize and differentiate highly available and resilient environments on AWS

Domain 3: Deployment and Provisioning: This domain is covered primarily in Chapter 3.

- 3.1 Identify and execute steps required to provision cloud resources
- 3.2 Identify and remediate deployment issues

Domain 4: Storage and Data Management: This domain is covered primarily in Chapter 4.

- 4.1 Create and manage data retention
- 4.2 Identify and implement data protection, encryption, and capacity planning needs

Domain 5: Security and Compliance: This domain is covered primarily in Chapter 5.

- 5.1 Implement and manage security policies on AWS
- 5.2 Implement access controls when using AWS
- 5.3 Differentiate between the roles and responsibility within the shared responsibility model

Domain 6: Networking: This domain is covered primarily in Chapter 6.

- 6.1 Apply AWS networking features
- 6.2 Implement connectivity services of AWS
- 6.3 Gather and interpret relevant information for network troubleshooting

Domain 7: Automation and Optimization: This domain is covered primarily in Chapter 7.

- 7.1 Use AWS services and features to manage and assess resource utilization
- 7.2 Employ cost-optimization strategies for efficient resource utilization
- 7.3 Automate manual or repeatable process to minimize management overhead

Steps to Becoming an AWS Certified SysOps Administrator–Associate

To become an AWS Certified SysOps Administrator–Associate, a test candidate must meet certain prerequisites and follow specific procedures. Test candidates must qualify for the exam and sign up for the exam.

Signing Up for the Exam

The steps required to sign up for the AWS Certified SysOps Administrator–Associate are as follows:

1. Create an AWS Certification account at <https://www.aws.training/Certification> and schedule your exam.
2. Complete the examination agreement, attesting to the truth of your assertions regarding professional experience and legally committing to the adherence of the testing policies.
3. Submit the examination fee.

Facts About the Exam

The exam is a computer-based test. The exam consists of multiple-choice questions only. You must bring a government-issued identification card. No other forms of ID will be accepted.

TIP Refer to the AWS Certification site at <https://aws.amazon.com/certification/> for more information regarding this, and other, AWS certifications. I am also in the process of building a simple hub site for everything AWS certification related at awscerthub.com. This site is made up of 100 percent AWS solutions, of course!

About the AWS Certified SysOps Administrator–Associate Certification Guide

This book maps directly to the topic areas of the exam and uses a number of features to 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 prove to yourself 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 to truly learn and understand the topics. This book is designed to help you pass the AWS Certified SysOps Administrator–Associate (SOA-C01) 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:** These are the core sections of each chapter. They explain the concepts for the topics in that chapter.
- **Exam Preparation Tasks:** After the “Foundation Topics” section of each chapter, the “Exam Preparation Tasks” section lists a series of study activities that you should do at the end of the 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 their page numbers. Although the contents of the entire chapter could be on the exam, you should definitely know the information listed in each key topic, so you should review these.

- **Define Key Terms:** Although the SysOps - Associate exam may be unlikely to ask a question such as “Define this term,” the exam does require that you learn and know a lot of AWS-related cloud terminology. This section lists the most important terms from the chapter, asking you to write a short definition and compare your answer to the glossary at the end of the book.
- **Review Questions:** 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 Cert Practice Test engine that allows 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 seven core chapters—Chapters 1 through 7. Chapter 8 includes preparation tips and suggestions for how to approach the exam. Each core chapter covers a subset of the topics on the AWS Certified SysOps Administrator–Associate (SOA-C01) exam. The core chapters map to the AWS Certified SysOps Administrator–Associate (SOA-C01) exam topic areas and cover the concepts and technologies that you will encounter on the exam.

Security and Compliance

It is amazing just how many engineers are often scared to move to the cloud due to security reasons. In all actuality, there are many reasons to move there that might encourage a more secure infrastructure. Just think, because Amazon can afford the latest in physical security measures at their data centers, you will enjoy a level of physical security that might not be possible in your own enterprise environment.

This chapter focuses on important security topics you should know and know well for AWS. This includes a look at the Shared Responsibility Model as well as an exploration of key security policies and access controls available to you.

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess if you should read the entire chapter. Table 5-1 lists the major headings in this chapter and the “Do I Know This Already?” quiz questions covering the material in those headings so you can assess your knowledge of these specific areas. The answers to the “Do I Know This Already?” quiz appear in Appendix A.

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

Foundations Topics Section	Questions
The Shared Responsibility Model	1–2
Security Policies in AWS	3–4
Access Controls	5–6

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. Who is responsible for creating users, groups, and roles in IAM for use in an AWS architecture?
 - a. The AWS customer
 - b. AWS staff
 - c. The managed service provider
 - d. There are no users, roles, or groups in IAM

2. Who is responsible for securing the hypervisor in use in AWS?
 - a. AWS staff
 - b. The client of AWS
 - c. The managed service provider
 - d. There is no hypervisor in use in AWS

3. You would like to add DDoS protection against your EC2 instances and your Elastic Load Balancing services. What service should you use?
 - a. AWS CloudIPS
 - b. AWS Shield Advanced
 - c. AWS Cognito
 - d. AWS Shield Standard

4. What credentials would you require in order to submit a penetration testing request?
 - a. AWSFullAdmin
 - b. Root account
 - c. AWSIAMAdmin
 - d. AWS Region Admin

5. What is the IAM component that is often ideal for allowing EC2 instances to other AWS services and resources?
 - a. Groups
 - b. Users
 - c. Clusters
 - d. Roles

- 6.** When creating a user account in AWS IAM, what are the options for access type? (Choose two.)
- a.** AWS Management Console access
 - b.** Restore
 - c.** Programmatic access
 - d.** CLI only

Foundation Topics

The Shared Responsibility Model

**Key
Topic**

The AWS Shared Responsibility Model is very simple. It divides the security responsibilities between two parties—the AWS customer (you) and Amazon (AWS). The fact that you are no longer responsible for a massive portion of the security required for scalable data centers is a huge advantage. You can leverage the massive budgets of Amazon and their intense expertise.

The next two sections of this chapter provide many examples of responsibilities in each part of the model. But for now, realize the Amazon responsibilities include the host operating system and virtualization layer down. From there, Amazon is also responsible for the physical security of the facilities in which the service operates. It is your (the customer's) responsibility to secure the guest operating system (including updates and security patches), application software, and the AWS network security group firewall. Be aware that the client responsibilities will vary depending on which services the client chooses to use. The client responsibilities further vary based on the level of integration of AWS services consumed and their IT infrastructure. Laws and regulations that must be followed will also vary.

As shown in Figure 5-1, AWS is considered “Security of the Cloud”, and the customer's responsibility is considered “Security in the Cloud.”

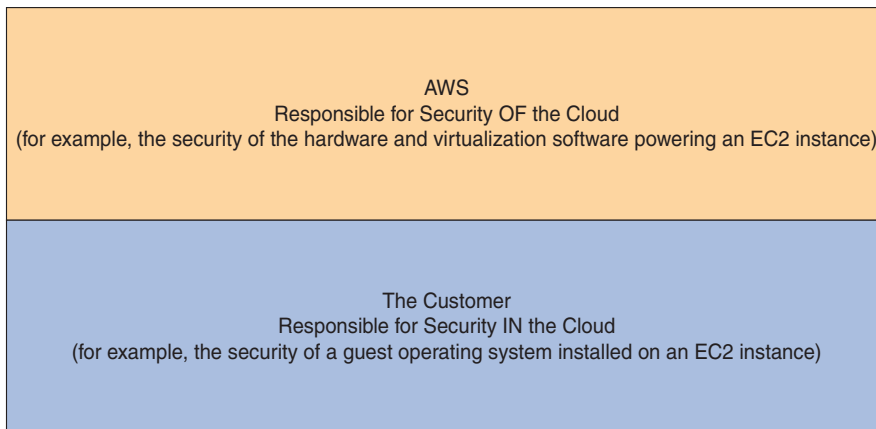


FIGURE 5-1 The AWS Shared Responsibility Model

In addition to partitioning the operational security concerns between the AWS client and AWS themselves, the Shared Responsibility Model applies to IT controls that are in use. Amazon categorizes these controls into three categories:

- **Inherited controls:** These are security controls that the customer fully inherits from AWS. Perfect examples are the physical and environmental security controls used by Amazon.
- **Shared controls:** These refer to controls that apply to both the infrastructure layer of Amazon and the customer responsibilities. Note that these shared controls apply to each domain in completely separate contexts or perspectives. For example, AWS provides the requirements (through controls) for the infrastructure. Then clients provide their own control implementation within their use of the services. Consider Identity and Access Management (IAM). The IAM service must be secured, meet regulatory compliance, and function as intended, while the customer should create well-crafted policies.
- **Customer-specific controls:** These are security controls that the customer is solely responsible for. This varies based on the services they selected, of course. A great example would be when you apply specific patches to one of your operating systems on an EC2 instance.

Amazon Responsibilities

Remember, Amazon is considered responsible for security *of* the cloud. This means that AWS is responsible for protecting the infrastructure that runs the services that customers select. This encompasses the hardware and software required to power the AWS service, including the networking and facilities used.

Specific Amazon responsibilities would include the following:

- Cloud software, including compute, storage, networking, and database software
- Hardware
- AWS Global Infrastructure (Regions, Availability Zones, Edge Locations)



Client Responsibilities

Remember, we consider the client responsible for security *in* the cloud. The specific services selected will cause variations in the client responsibilities. For example, if you are relying heavily on S3 for storage, you will be responsible for knowledge and proper configuration of the security permissions for your resources. Another example would be if the client chooses to use EC2 and run an operating system like Windows Server 2016. The client will be required to keep the operating system updated and patched. The client is also responsible for the application software required on this guest operating system. In addition, the client is responsible for the appropriate security group configuration for the EC2 instance.

Key Topic

Specific examples of client responsibilities would include the following:

- Customer data
- Platform, applications, Identity and Access Management (IAM)
- Guest operating systems
- Network and firewall configurations
- Client-side data encryption
- Server-side encryption (file system and or data)
- Networking traffic protection (encryption, integrity, and identity)

Figure 5-2 shows an example of a customer checking the security groups settings that would apply to an EC2 instance. This is a perfect example of client responsibilities. AWS is responsible for making sure the security group functions as intended, but it is the client's responsibility to configure it correctly.

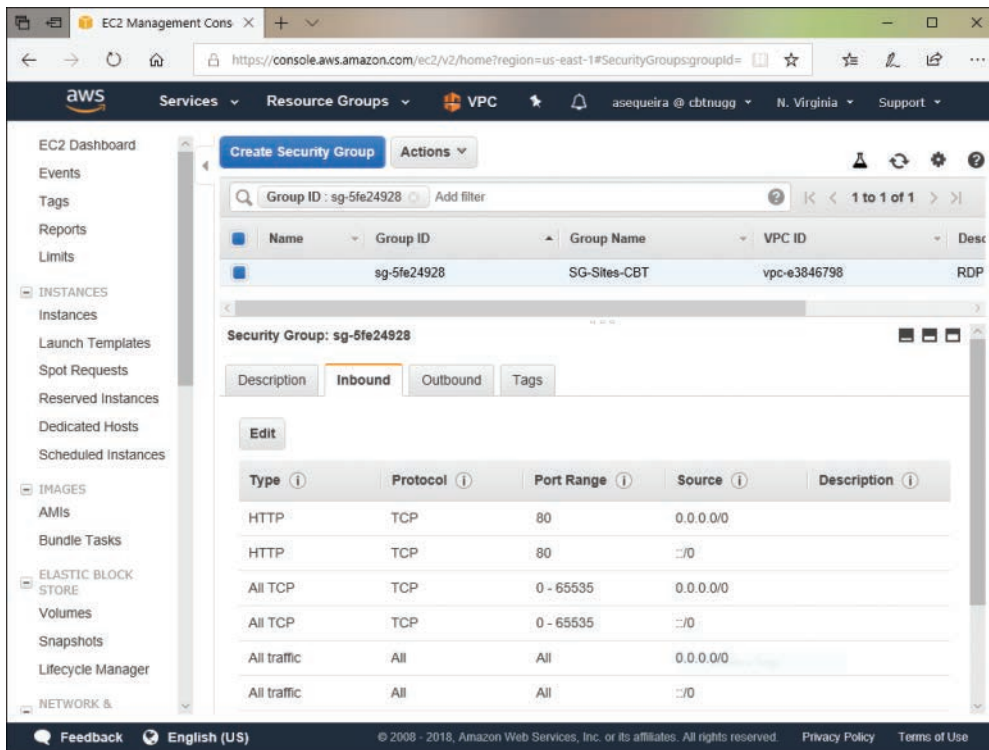


FIGURE 5-2 Checking the Security Groups Settings for an EC2 Instance

Security Policies in AWS

There are common security policies and practices that you should be aware of when operating AWS solutions. This section of the chapter covers some of the more important ones.

DDoS Mitigation

The distributed denial of service (DDoS) attack is one to be feared. Famous examples of this attack include stories about how huge chunks of the entire Internet itself were made unavailable for relatively long periods of time. Just like with a regular old denial of service (DoS) attack, the goal is resource exhaustion so that disruption is in place for legitimate traffic that is attempting to flow or access a service or resource. Having many systems (potentially) participate in the attack

(DDoS) can make the attack that much more effective due to the increase in frequency of the communications.

It is worth restating for clarity—there are two main and related objectives behind DDoS (and DoS):

- Exhaust resources on the server side of the computing model.
- Once exhaustion occurs, disrupt desired traffic flows or requests.

NOTE At the most precise level, DDoS (and DoS) attacks can be tricky to detect. That is because they might be made up of “normal” requests that would be transpiring against your AWS system anyway. So, it is often imperative to analyze the frequency of such requests in order to correlate the data properly and recognize that an attack is actually taking place. In fact, some of the best DDoS attacks may not be possible to detect at all beyond the simple increase in traffic flows. Thus, if all other anti-DDoS measures have been implemented, your last measure might be to simply “out-scale” the attack. This is a unique advantage to the cloud since it is often not possible in the traditional on-prem data center.

We often use the Open Systems Interconnection (OSI) model in order to help us think about and mitigate DDoS attacks. Figure 5-3 shows the OSI model.




FIGURE 5-3 The OSI Model

DDoS attacks that tend to focus on the lower layers (1 through 4) of the OSI model are often called *infrastructure attacks*, whereas upper layers that come under attack are referred to as *application-layer attacks*. An example of a Layer 4 attack might be a

SYN flood or an amplified UDP reflection attack. An attack at Layer 7 (Application) might be an HTTP flood.

Let's examine one of these in more detail. In an amplified UDP reflection attack, the attacker uses the connectionless UDP protocol to ask a server for some piece of information. The attacker forges the packet header so that it contains a different sender address. The machine that receives these "spoofed" packets will send a response back to the forged source address.

ICMP, NTP, DNS, DHCP, TFTP, and many more are all examples of UDP services that, if left unchecked, can be abused. Depending on the command sent and data requested, the amplification ratio can range from 2× to over 200×. This is to say that the attacker sends a small request to the vulnerable server, and the server sends a much larger response to the target system.

Fortunately, AWS knows of these many potentially devastating DDoS attacks and includes some powerful protections for us for free, as well as ensures these protections are in an always-on state.

AWS Shield Standard

If you are using the AWS services of Route 53 (DNS) and CloudFront (CDN), you are already taking advantage of the free DDoS prevention methods of AWS Shield Standard. AWS engages in powerful protection methods for these services that include powerful network flow monitoring as well as protection mechanisms against Layer 3 and Layer 4 attacks. For example, the amplified UDP reflection attack described previously should be blocked thanks to the default behaviors of AWS Shield Standard.



NOTE These protections do require that you have configured your DNS in Route 53 and your CloudFront services correctly. Incorrectly configuring these services might render the security protections useless, of course.

AWS Shield Advanced

While it is not free like the AWS Shield Standard's functionality, you might be compelled to take advantage of the more advanced version, AWS Shield Advanced. This is most commonly acquired through an Enterprise-level support agreement with AWS.

As you might guess, AWS Shield Advanced has the ability to protect a wider range of services than the standard version can. Here are some of the services that are provided protection by the suite of features:

- EC2
- Elastic Load Balancing
- Elastic IP Addressing
- CloudFront
- Route 53
- AWS Global Accelerator

Not only do you enjoy a wider range of services that are protected, your features expand as well, including the following:

- Advanced analysis
- Resource baselining and trending
- Protection against Application (Layer 7) attacks
- AWS DDoS Response Team (DRT)
- DDoS Cost Protection
- Real-time Threat Dashboard access

As if this was not enough, if you use AWS Shield Advanced to protect your EC2 instances, during an attack AWS Shield Advanced automatically deploys your VPC network ACLs to the border of the AWS network. This allows the security suite to provide protection against larger DDoS events.

Data Encryption

It is well known that encrypting your data at rest is often necessary to obtain the level of security you require. Fortunately, AWS not only supports this, but provides many tools to allow you a variety of protections in a variety of configurations. Data encryption capabilities include the following:

- Data encryption capabilities available in AWS storage and database services, such as EBS, S3, Glacier, Oracle RDS, SQL Server RDS, and Redshift
- Flexible key management options, including AWS Key Management Service; allowing you to choose whether to have AWS manage the encryption keys or to have you keep complete control over your keys

- Encrypted message queues for the transmission of sensitive data using server-side encryption (SSE) for Amazon SQS
- Dedicated, hardware-based cryptographic key storage using AWS CloudHSM, allowing you to satisfy compliance requirements

In addition, AWS provides APIs for you to integrate encryption and data protection with any of the services you develop or deploy. For more information on data encryption, see Chapter 4, “Storage and Data Management.”

Inventory and Configuration

One of the legitimate concerns when moving to a cloud service like AWS is the flexibility and ease of resource creation getting out of hand. You can have inventory and the configuration of devices become unmanageable. AWS has tools such as the following to assist with this potential problem:



- Amazon Inspector is a security assessment service that automatically checks applications for vulnerabilities or deviations from best practices. This inspection includes impacted networks, OS, and attached storage.
- Deployment tools to manage the creation and decommissioning of AWS resources according to organization standards.
- Inventory and configuration management tools, including AWS Config, that identify AWS resources and then track and manage changes to those resources over time.
- Template definition and management tools, including AWS CloudFormation to create standard, preconfigured environments; for more information on CloudFormation, see Chapter 7, “Automation and Optimization.”

Monitoring and Logging

“Track everything” is the war cry for many AWS engineers with concerns about cloud security. AWS provides tools for monitoring and logging that include the following:

- Deep visibility into API calls through CloudTrail, including details on the calls themselves
- Log aggregation options, streamlining investigations, and compliance reporting

- Alert notifications through CloudWatch when specific events occur or thresholds are exceeded

Consistent use of these tools can improve the security posture, and reduce the risk profile, of your AWS solutions.

Penetration Testing

Key Topic

In order to perform penetration testing to or originating from any AWS resources, you must complete a request form to obtain permissions from Amazon.

NOTE AWS does now permit penetration testing within many services without the formal request process. For purposes of your exam, this recent fact might not be indicated.

There are several important things to note about penetration testing requests. As previously mentioned, there have been modifications to some of these parameters, but your exam might not reflect the current changes:

- To request permission, you must be logged in to the AWS portal using the root credentials associated with the instances you wish to test; otherwise, the form will not pre-populate correctly. If you have hired a third party to conduct your testing, Amazon suggests that you complete the form and then notify your third party when approvals are granted.
- You are only permitted testing of EC2 and RDS instances that you own. Tests against any other AWS services or AWS-owned resources are prohibited.
- Amazon does not permit testing small or micro RDS instance types; testing of m1.small or t1.micro EC2 instance types is not permitted.

Access Controls

AWS solutions must provide secure access by clients and providers of the technologies. This is accomplished using a robust set of technologies.

Infrastructure Security

Amazon provides security capabilities and services to increase privacy and control network access. These include the following:

- Network firewalls built into Virtual Private Cloud (VPC), and web application firewall capabilities in AWS WAF, let you create private networks and control access to your instances and applications.
- Encryption in transit with TLS across all services.
- Connectivity options that enable private, or dedicated, connections from your office or on-premises environment.

Identity and Access Management

IAM is a cloud service that helps you securely control access to AWS resources. You use IAM to control who is authenticated and authorized to use resources.

Upon AWS account creation, you begin with a single sign-in that has complete access to all AWS services in the account. This sign-in is called the AWS account root user. You access AWS with the account by signing in with the email address and password you used at sign-up.

Amazon strongly recommends that you do not use the root account for your everyday tasks, even the administrative ones. Instead, follow the best practice of using the root account only to create your first IAM user. Then securely lock away the root account credentials and use them to perform only a few account and service management tasks.

IAM permits extremely fine-grained permissions. For example, you might grant someone read access to only a single bucket of objects in S3. Or you might use IAM to control specific calls (`GetObject`) against a single object stored in S3. Perhaps you examine a particular time/date range or the source IP address of the call.

Other features of IAM include the following:

- **Access from service to resource in AWS:** For example, you can have an application running on an EC2 instance access an S3 bucket. As you will learn later in this chapter, we often use roles for such access.
- **Multi-factor authentication (MFA):** Permitting access through a password and a code from an approved device, thus strengthening security greatly. Figure 5-4 shows the configuration area for MFA in the IAM Management Console.

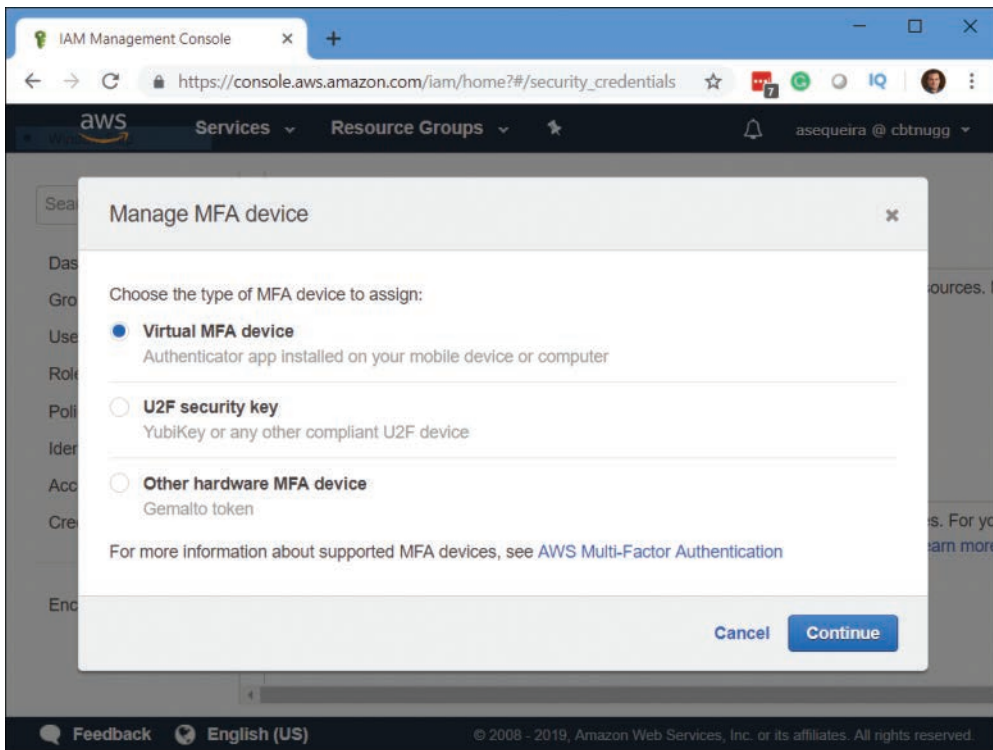


FIGURE 5-4 Configuring MFA for an Account

- **Identity federation:** Users who have already authenticated with another service can gain temporary access to resources and services in your account.
- **Identity information for assurance:** CloudTrail can trace and log all API activity against every service and resource in your account. Figure 5-5 shows the CloudTrail Dashboard in AWS.

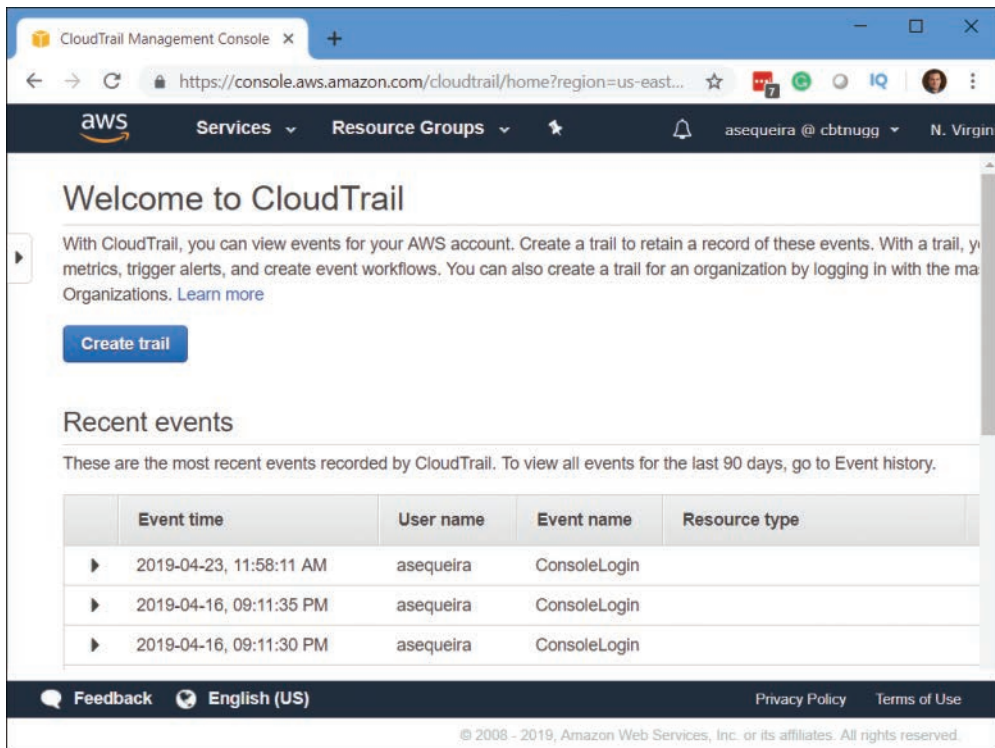


FIGURE 5-5 The CloudTrail Dashboard

- **PCI DSS compliance:** IAM supports the processing, storage, and transmission of credit card data by a merchant or service provider, and it has been validated as being compliant with the Payment Card Industry (PCI) Data Security Standard (DSS).
- **Integration:** IAM integrates with every major service of AWS.
- **Eventually consistent:** Amazon replicates important data around the world with their Global Infrastructure to help ensure high availability (HA). As a result, data in some locations might lag others. Therefore, with IAM, consider implementing your changes for IAM first, and then verify full replication before working with dependent service deployments.
- **Always free:** Whereas some services of AWS can be used for one year free (using the Free Tier account), IAM services remain free for the life of your account.
- **Accessibility options:** You can access the components of IAM in a variety of ways, including the AWS Management Console, AWS command-line tools, AWS SDKs, and IAM HTTPS API.

Key Topic

It is critical that you understand the main identities you'll use in IAM. Realize that there is much more to IAM than these identities, but at this point in your AWS education, we are covering the main foundational components.

Remember, an account that supersedes the IAM service is root. As stated earlier in this chapter, this account should rarely be used.

Identities in IAM consist of the following:

- Users:** These are the entities you create in AWS to represent the people or services that use the IAM user to interact with AWS. When you create an IAM user, you grant it permissions by making it a member of a group. You assign appropriate permission policies to the group. This is the recommended approach from Amazon. Note that you could directly attaching policies to the user, but this is not recommended because it is not a scalable approach and could make security management more difficult. You can also clone the permissions of an existing IAM user. This approach automatically makes the new user a member of the same groups and attaches all the same policies. Figure 5-6 shows a user in AWS.

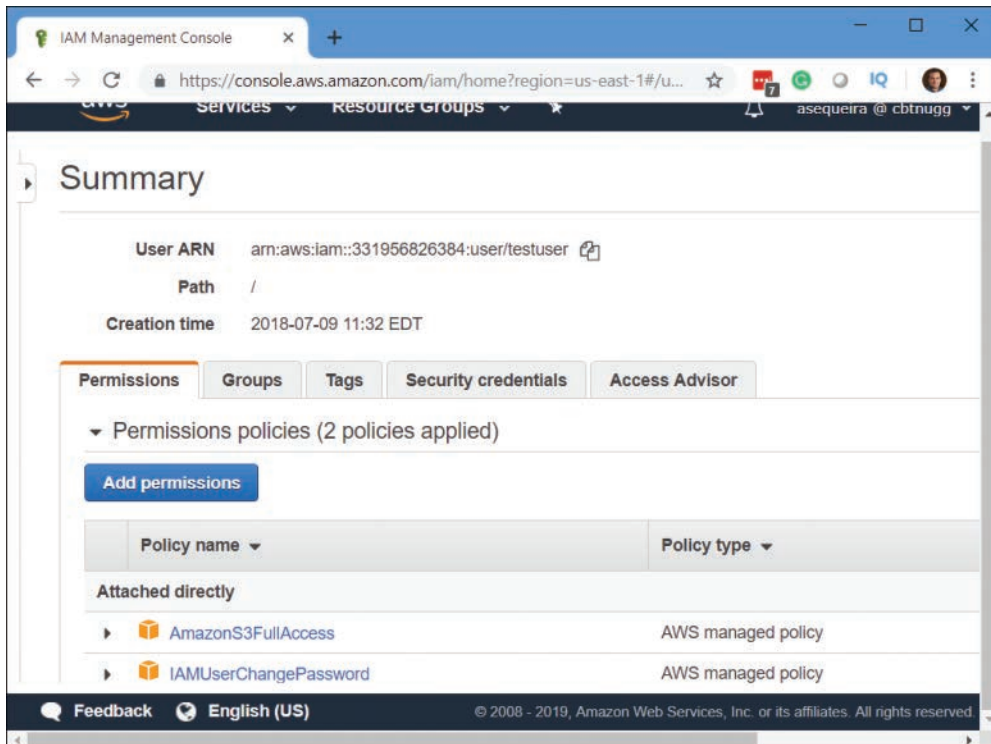


FIGURE 5-6 A User in AWS IAM

- **Groups:** A collection of IAM users. You can use groups to specify permissions for a collection of users, which can make those permissions easier to manage for those users.
- **Roles:** These are similar to user accounts, but they do not have credentials (password or access keys) associated with them.

In the following steps, we create a group that provides full access to S3 in AWS and then create a user, adding it to this group:

Key Topic

- Step 1.** Navigate to the AWS Management Console and then search for the IAM service.
- Step 2.** Select **Groups** in the left navigation pane. Figure 5-7 shows the Groups console.

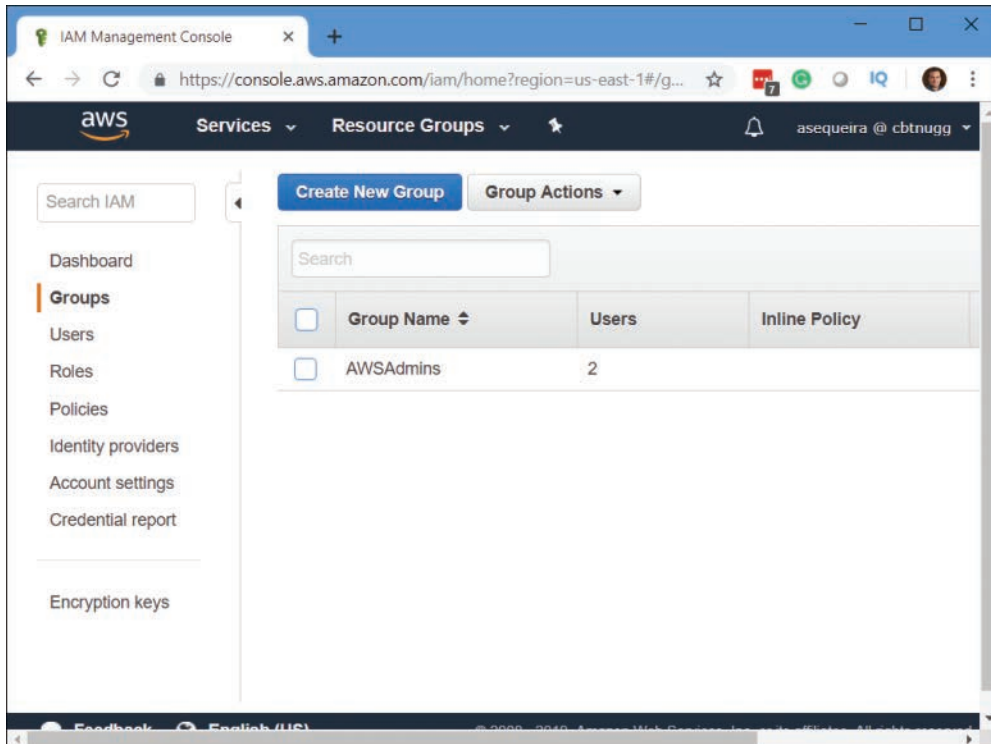


FIGURE 5-7 Groups in IAM

- Step 3.** Click the **Create New Group** button.
- Step 4.** Set the Group Name and click **Next Step**.
- Step 5.** In the Attach Policy page (shown in Figure 5-8), enter **S3** in the Filter option.

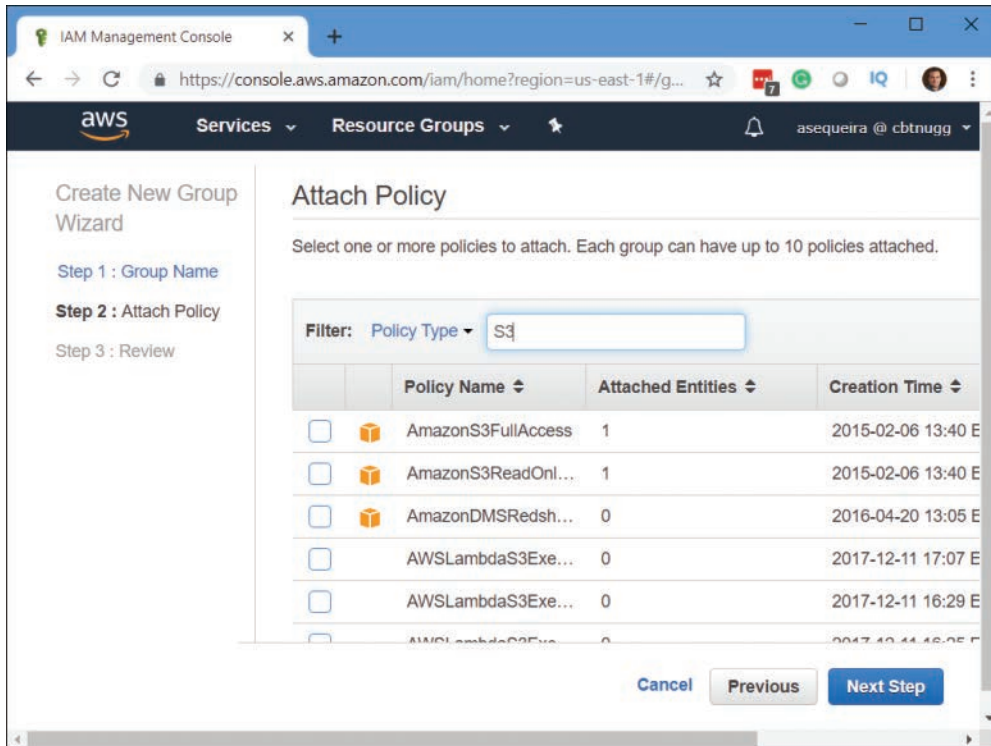


FIGURE 5-8 The Attach Policy Page

- Step 6.** Check **AmazonS3FullAccess** and click **Next Step**.
- Step 7.** Review the configuration and click **Create Group**.
- Step 8.** Click the **Users** option in the left navigation pane.
- Step 9.** Click the **Add User** button.
- Step 10.** Provide the username and then allow both types of access to the accounts. Leave the defaults in place regarding the password. Figure 5-9 shows this page.

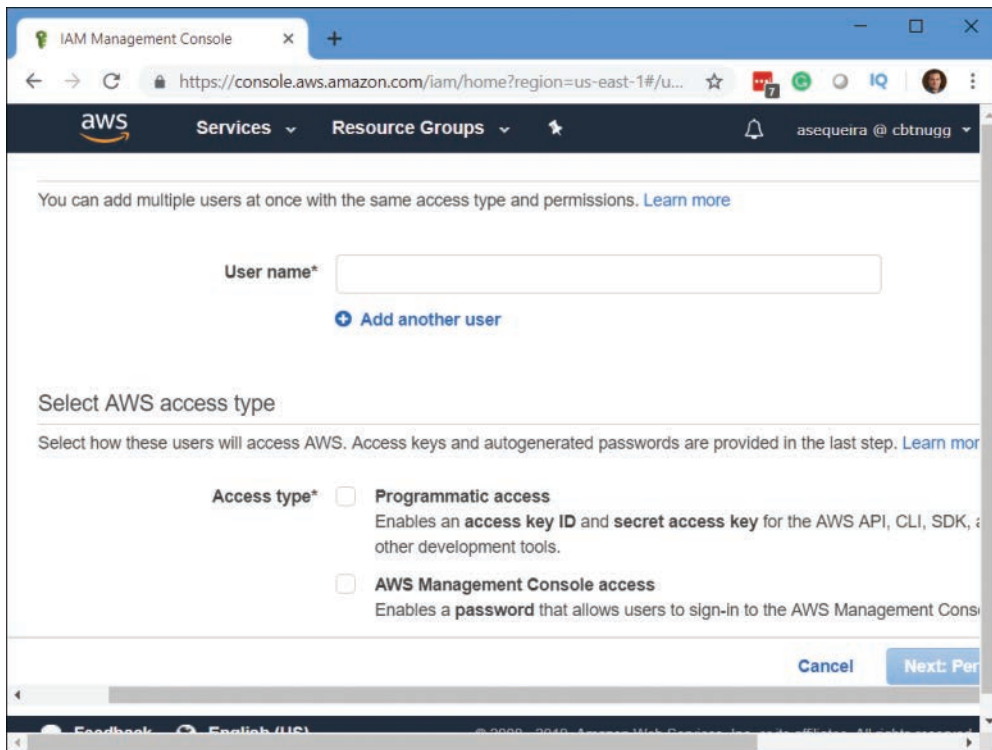


FIGURE 5-9 Naming the User Account

- Step 11.** Click **Next: Permissions**.
- Step 12.** Add the user to the group you created earlier in these steps. Click **Next: Tags**.
- Step 13.** Click **Next: Review**. Remember, adding tags is an excellent idea to indicate various identifiers, but here in a lab environment, we will just skip it.
- Step 14.** Review your settings and click **Create User**.

Best Practices with IAM

While IAM in AWS provides many exciting capabilities, its complexity can cause organizations to make fatal flaws when working with the service. This is why following best practices is critical.

Key Topic

You should consider following most (if not all) of these recommendations.

- **Care of the root account:** The root account for your AWS implementation should be used infrequently. The current best practice is to delete any access keys associated with root. Root should never have automation keys. You should never automate against root, and the only reason to have keys is for automation. Root should have only a login (email address and password) and physical MFA. Physical MFA is the best practice because you do not want a single person with root access on the phone; it should be a separate hardware device locked up and not used except in an emergency. As you no doubt realize, MFA for root on a phone, which could be lost, could be obtained easily. Some companies have one team manage the password for root, while another team manages the physical MFA device. This ensures checks and balances to gain access to root. Exceptions to these best practices may be in the case of organizations where new AWS accounts are managed via automation.
- **Create individual IAM users:** Because you do not want to use root for your AWS implementation, it is critical that you create additional users. This would include for yourself so that you are not required to use root. In larger organizations, you will have a large team working on AWS. You must create multiple users for your staff to ensure that everyone is authenticating and being authorized for only those resources and permissions that are required for members to do their jobs. You will most likely have one user in IAM for every person who requires administrative access.

NOTE This recommendation is assuming no federation is in place. Some companies larger than just a few IT staff will typically use Active Directory federation and may actually have no IAM users at all, but rather simply a SAML trust and use of roles.

- **Use groups to assign permissions to IAM users:** Even though it might seem silly, if you are the sole administrator of your AWS implementation, you will want to create a group and assign permissions to this group. Why? If you do need to grow and hire another administrator, you can just add that user account to the group you created. We always want our AWS implementations

to scale, and using groups helps ensure this. It should also be noted that applying permissions to groups instead of individual user accounts will help eliminate assignment errors, as we are minimizing the number of permissions we must grant.

- **Use AWS-defined policies for permissions:** Amazon was very kind to us. They defined a ton of policies we can easily leverage when working with IAM. What's more, AWS maintains and updates these policies as they introduce new services and API operations. The policies that AWS created for us are defined around the most common tasks we need to perform. These make up an excellent starting place for your own policies. You can copy a given policy and customize it to make it even more secure. Oftentimes, you will find the default defined policies are too broad with access.
- **Grant least privilege:** Create the IAM user identity for your AWS user that provides the least privileges they require. That way, if an attacker does manage to capture security credentials and begins acting as that user in the AWS architecture, he can do a limited amount of damage.
- **Review IAM permissions:** You should not use a “set and forget” policy when it comes to your permissions in IAM. You should consistently review the permissions level assigned to ensure that you are following least privilege concepts and that you are still granting those permissions to the groups that require them. There is even a policy summary option within IAM to facilitate this.
- **Always configure a strong password policy for your users:** It is a sad fact of human nature: Your users will tend to be lazy about setting (and changing) their passwords. They will tend to use simple passwords that are easy for them to remember. Unfortunately, these simple passwords are also easy to crack. Help your security by setting a strong password policy that your users must adhere to. Figure 5-10 shows the configuration of a password policy for user accounts in the IAM Management Console.

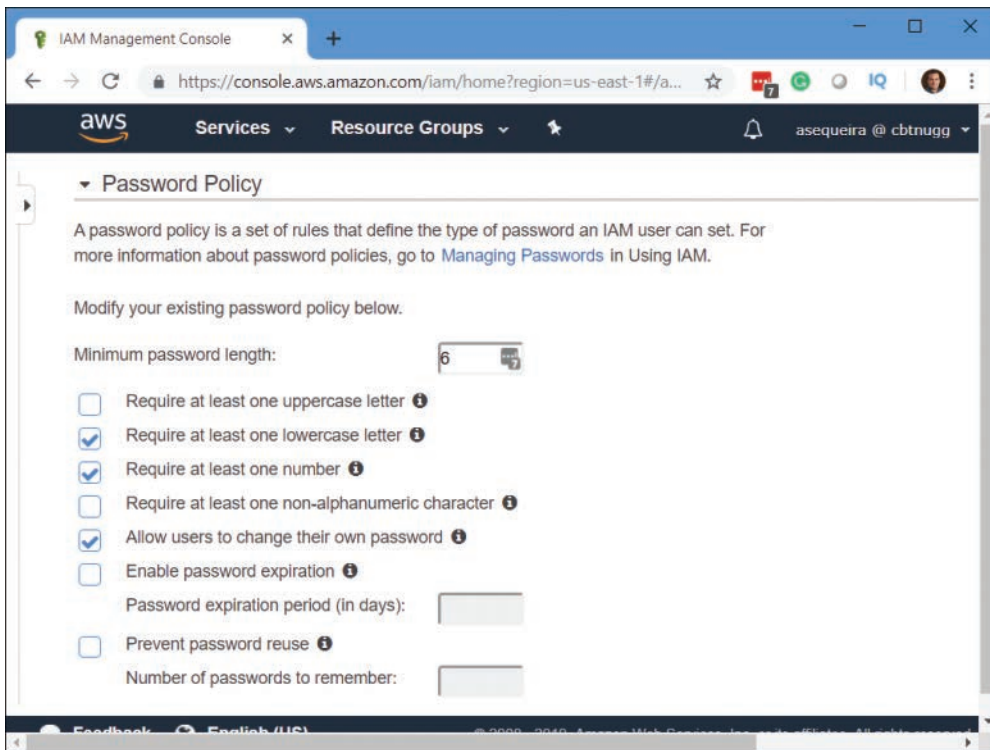


FIGURE 5-10 Configuring a Password Policy

- **Enable multi-factor authentication for privileged user accounts:** Of course, you do this for the seldom-used AWS root account, but you should also protect key admin accounts you have created in AWS. Using multi-factor authentication (MFA) ensures the user knows something (like a password) and also possesses something (like a smartphone). With most AWS environments today, MFA is considered mandatory.
- **Use roles:** You should consider the use of roles in AWS when you have applications or services running on EC2 instances that need to access other services or resources.
- **Use roles to delegate permissions:** Roles can also prove valuable when you need to permit one AWS account to access resources in another AWS account. This is a much more secure option to providing the other AWS account with username and password information for your account. And remember, the use of roles is always recommended within an AWS account.

- **Do not share access keys:** It might be tempting to take the access keys that permit programmatic access to a service or resource and just share those with another account that needs the same access. Resist this temptation. Remember, you can always create a role that encompasses the required access.
- **Rotate credentials:** Be sure to change passwords and access keys regularly in AWS. The reason for this, of course, is the fact that if these credentials are compromised, you will have minimized the damage that can be done when the stolen credentials no longer function. Roles rotate credentials automatically for you many times per day. This is a huge security advantage and makes their use desirable, especially at scale.
- **Remove unnecessary credentials:** Because it is so easy to learn and test new features in AWS, it can get messy as far as IAM components you leave in place that are no longer needed are concerned. Be sure to routinely audit your resources for any “droppings” that are no longer needed. AWS even assists in this regard with structuring reports around credentials that have not been recently used. Again, roles provide another built-in advantage in this regard.
- **Use policy conditions:** Always consider building conditions into your security policies. For example, access might have to come from a select range of IP addresses, or MFA might be required.
- **Monitor, monitor, monitor:** AWS services provide the option for an intense amount of logging. Here are just some of the services where careful logging and analysis can dramatically improve security:
 - CloudFront
 - CloudTrail
 - CloudWatch
 - AWS Config
 - S3

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 8, “Final Preparation,” and the exam simulation questions in the Pearson Test Prep Software Online.

Review All Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. Table 5-2 lists these key topics and the page numbers on which each is found.



Table 5-2 Key Topics for Chapter 5

Key Topic Element	Description	Page Number
Overview	The AWS Shared Responsibility Model	132
List	Amazon responsibilities	133
List	Client responsibilities	134
Overview	AWS Shield Standard and AWS Shield Advanced	137
Overview	Inventory and Configuration	139
Overview	Penetration Testing	140
List	AWS IAM identities	144
Steps	Creating users and groups in IAM	145
List	IAM best practices	148

Define Key Terms

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

The AWS Shared Responsibility Model, Security of the Cloud, Security in the Cloud, DDoS, AWS Shield, AWS Account Root User, IAM Users, IAM Groups, IAM Roles

Q&A

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

1. Provide at least three examples each for client and AWS responsibilities in the Shared Security Model.
2. Name two services protected by AWS Shield Standard.
3. What service would you use to log API calls in AWS?

Index

A

- access control, 34
- accessing
 - APIs, 258
 - CloudWatch Console, 7
 - EFS (Elastic File System), 264
- accounts, aliases, 265
- ALARM state, 27
- alarms, 251
 - characteristics, 28
 - CloudWatch, 26–27
 - CPU utilization, 35–36
 - creating, 28–31
 - false positives, 27
 - M out of N, 30
 - missing data points, 28
 - parameters, 27
 - states, 27
- Amazon Athena, 119–120
- Amazon VPC Endpoints, 260
- Amazon-to-Amazon VPC, 184–185
 - VPC peering, 185
- AMIs (Amazon Machine Images), 113–114
 - baking, 77–78
- API Gateway, FAQs, 256–258
- APIs
 - CloudWatch Query API, 11–12
 - configuring, 89–90
- Application Load Balancer, 92, 255
 - error messages, 94–95
- application-layer attacks, 136
- applications
 - deployment, 75
 - in Elastic Beanstalk, 81
- asynchronous decoupling, 53
 - SQS (Simple Queue Services), 52–53
- Athena, 119–120
- Aurora, 43
- authentication, 34
 - MFA Delete, 107
- Auto Scaling, 42
 - best practices, 44
 - configuring, 45
 - FAQs, 254–255
 - predictive scaling, 44
 - pre-implementation considerations, 44
 - troubleshooting areas, 45–46
- automation, 47, 57, 74, 197. *See also*
 - automation tools
 - best practices, 222–223
 - of deployment provisioning, 75
- automation tools
 - AWS Config, 221
 - CloudFormation, 220
 - CloudTrail, 222
 - CloudWatch, 221
 - CodeBuild, 218
 - CodeCommit, 222
 - CodeDeploy, 218
 - CodePipeline, 218
 - CodeStar, 219
 - Elastic Container Service, 219

- Lambda, 219
- OpsWorks, 220
- Systems Manager, 220–221
- X-Ray, 221
- availability, S3 (Simple Storage Service), 104
- AWS (Amazon Web Services), 3, 74.
 - See also* CloudWatch
- accounts, 34
- achieving operational excellence in, 199–204
- alarms, creating, 28–31
- AMIs (Amazon Machine Images), 113–114
 - baking, 77–78
- authentication, 34
- Auto Scaling, 42
 - best practices, 44
 - predictive scaling, 44
 - pre-implementation considerations, 44
 - troubleshooting areas, 45–46
- BYOIP (Bring Your Own IP), 267
- ClassicLink, 267
- CLI, 177–178
- CLI (command line interface), 87
 - aws configure command, 8–9
- CloudFormation, 86–87
- CloudFront, 160–163
- CloudHSM, 122
- CloudWatch
 - alarms, 26–27
 - console, accessing, 7
 - dashboards, 17
 - installing on Windows, 10, 8
 - metrics, 21–22
- configuration management approaches, 75–76
- Connectivity services
 - Amazon-to-Amazon VPC, 184–187
 - internal user-to-Amazon VPC, 187
 - Network-to-Amazon VPC, 178–184
- connectivity services, 178
- continuous deployment, 80
- Cost Management dashboard, 216–218
- Cross Service Dashboard, 15–16
 - removing services from, 16
- dashboards
 - alarms, 26
 - creating, 17–18
 - creating with JSON, 19–20
 - editing, 19
- DataSync, 264
- deployment
 - of applications, 75
 - blue-green, 74
 - canary, 74
 - custom variables, 76
 - provisioning infrastructure, 75
- design goals, 47
- Detailed Monitoring, 21
- EBS (Elastic Block Storage), 111–112
 - volume encryption, 123–125
- ECS (Elastic Container Service), 83–84
- Edge Locations, 160–163
- EFS (Elastic File System), 112–113
- Elastic Beanstalk, 81–83
- Elastic Load Balancers, 92–93
- example HA solution, 67
- Fargate, 83
- Global Infrastructure
 - AZs, 159–160
 - regions, 157–159
- HA for applications, 50–51
- IAM (Identity and Access Management), 34
- identity-based policies, 35
- key services, 211–212
- KMS (Key Management Service), 121
- Lambda applications, deploying, 91–92
- limit management, 48
- logging, 78
- metrics, 3

- publishing, 24–25
 - viewing, 21–24
 - navigation pane, pinning dashboard
 - to, 19
 - networking
 - redundancy, 50
 - reliability, 49
 - resiliency, 49
 - services, 49–50
 - VPC, 48
 - VPG (Virtual Private Gateway), 50
 - object storage, 102
 - OpsWorks, 84–86
 - PrivateLink, 267
 - provisioning infrastructure, 75
 - root user account, 34
 - S3 (Simple Storage Service), 102–103
 - advantages of, 104–105
 - client-side encryption, 122
 - configuring versioning, 107
 - Glacier Deep Archive storage
 - class, 106
 - Glacier storage class, 105–106
 - Intelligent-Archiving storage class, 105
 - lifecycle policies, 107–108
 - MFA Delete, 107
 - One Zone-Infrequent Access storage
 - class, 105
 - server-side encryption, 122–123
 - snapshots, 125
 - standard storage class, 105
 - standard-IA storage class, 105
 - storage buckets, 102–103
 - uses, 103
 - versioning, 106
 - scalability capabilities, 79
 - SDKs, 13
 - security policies
 - data encryption, 138–139
 - DDoS mitigation, 135–137
 - IAM, 141–151
 - infrastructure security, 141
 - inventory and configuration, 139
 - monitoring and logging, 139–140
 - penetration testing, 140
 - Serverless Application Model (SAM), 92
 - services, 13–14
 - Shield Advanced, 137–138
 - Shield Standard, 137
 - Snowball, 117–118
 - storage, 207
 - storage encryption, 120–121
 - Storage Gateway, 115–117
 - System Manager, 87–88
 - Systems Manager Parameter Store, 76
 - VPC
 - components, 163–165
 - default, 165–166
 - DHCP option sets, 172–173
 - DNS, 174
 - egress-only Internet gateways, 171–172
 - elastic IP addresses, 174–175
 - endpoints, 175
 - gateway endpoints, 176
 - interface endpoints, 176
 - Internet gateways, 170–171
 - NAT, 177
 - network interfaces, 166–167
 - route tables, 168–170
 - aws cloudwatch get-metric-statistics
 - command, 9–10
 - AWS Config, 221
 - aws configure command, 8–9
 - AWS Shield, 50
 - AZs (Availability Zones), 157, 159–160
- ## B
- backups, 63, 66–67
 - EFS (Elastic File System), 264
 - snapshots, 125

baking AMIs (Amazon Machine Images),
77–78

Batch Operations (S3), 262

benefits, of RDS, 61

best practices

for achieving operational excellence,
199–204

for automation, 222–223

AWS Auto Scaling, 44

cost optimization, 213–214

HA (high availability), 52

IAM (Identity and Access
Management), 148–151

for managing resource utilization,
205–206

RDS (Relational Database Service), 61

block storage, 102

blue-green deployments, 74

BYOIP (Bring Your Own IP), 267

C

calling, deployed APIs, 89

canary deployments, 74

categories of disruptions, 51

change deployment, 67

characteristics

of alarms, 28

of SNS, 59

CIDR (Classless Inter-Domain
Routing), 48

Classic Load Balancer, 93

error messages, 94–95

ClassicLink, 164, 267

CLI (command line interface), 87,
177–178

SQS management, 57

clients, SNS (Simple Notification
Service), 58

client-side encryption, 122

cloud computing

AWS (Amazon Web Services)

design goals, 47

limit management, 48

PaaS (Platform as a Service), 81

VPC (Virtual Private Cloud), 48

VPC, components, 163–165

CloudFormation, 86–87, 91–92, 220

FAQs, 268–269

templates, 86–87, 268–269

troubleshooting, 96

CloudFront, 160–163

FAQs, 268

supported content, 268

CloudHSM, 122

CloudTrail, 14, 222

FAQs, 252–254

CloudWatch, 3, 221

alarms, 26–27, 251

characteristics, 28

creating, 28–31

false positives, 27

M out of N, 30

missing data points, 28

parameters, 27

states, 27

CLI commands, 10–11

dashboards, 17

creating, 17–18

editing, 19

dimensions, 25

FAQs, 249–252

home page, 4

identity-based policies, 35

installing on Windows, 10, 8

logging, 249–251

metrics, 21–24

custom, 26

publishing, 24–25

publishing services, 31–33

remediation of issues, 35–36

services related to, 13–15

CloudWatch Console, accessing from
AWS, 7

CloudWatch Query API, 11–12

- HTTP requests, 12
 - CodeBuild, 218
 - CodeCommit, 222
 - CodeDeploy, 218
 - CodePipeline, 218
 - CodeStar, 219
 - commands
 - aws cloudwatch get-metric-statistics, 9–10
 - aws configure, 8–9
 - CloudWatch CLI, 10–11
 - get-metric-statistics, 24
 - compute solutions, 206–207
 - configuring
 - APIs, 89–90
 - AWS Auto Scaling, 45
 - lifecycle policies, 108–111
 - SNS (Simple Notification Service), 59–60
 - SQS (Simple Queue Services), 54–56
 - System Manager, 88
 - versioning in S3, 107
 - connectivity, categories of disruptions, 51
 - connectivity services, 178
 - containers, 83, 206
 - continuous deployment, 80
 - control plane, 52
 - controls, Shared Responsibility Model, 133
 - Cost Management dashboard, 216–218
 - cost optimization
 - best practices, 213–214
 - design principles, 214–215
 - pricing models, 214
 - strategies, 213
 - CPU utilization, monitoring, 35–36
 - creating
 - alarms, 28–31
 - dashboards
 - CloudWatch, 17–18
 - using JSON, 19–20
 - Elastic Beanstalk environments, 81–83
 - encrypted root volume, 123–124
 - OpsWorks stacks, 85–86
 - queues, 54–56
 - Cross Service Dashboard, 15–16
 - services, removing, 16
 - cross-zone load balancing, 255–256
 - CRR (Cross-Region Replication), 263
 - custom metrics, 26
 - custom variables, 76
 - customer-specific controls, Shared Responsibility Model, 133
 - CWE (CloudWatch Events), 251–252
- ## D
- dashboards, 80
 - alarms, 26
 - CloudWatch, 17
 - creating, 17–18
 - editing, 19
 - creating, using JSON, 19–20
 - pinning to navigation pane, 19
 - widget type, choosing, 18
 - data encryption, 138–139
 - data plane, 52
 - data points, 28
 - data transfers
 - Snowball, 117–118
 - Snowball Edge, 118–119
 - database solutions, 208
 - DataSync, 264
 - DDoS mitigation, 135–137
 - default VPC (Virtual Private Cloud), 165–166
 - deployment
 - of applications, 75
 - blue-green, 74
 - canary, 74
 - configuration management approaches, 75–76
 - continuous, 80
 - custom variables, 76

- feature toggles, 75
 - of Lambda applications, 91–92
 - monitoring, 80
 - OpsWorks, 84–86
 - of REST API, 88–91
 - scalability capabilities, 79
 - tagging, 76
 - design goals, for AWS, 47
 - design principles
 - for cost optimization, 214–215
 - for performance efficiency, 204–205
 - Detailed Monitoring, 21
 - DHCP option sets, 163, 172–173
 - dimensions parameter, 25
 - Direct Connect, 180–182
 - disruptions, 51
 - DNS (Domain Naming Service), 163, 174
 - DR (disaster recovery), 46
 - backup and restore approach, 66
 - backups, 63
 - multi-site solution method, 66
 - pilot light method, 66
 - RPO (recovery point objective), 65
 - RTO (recovery time objective), 65
 - warm-standby method, 66
 - durability, 103
 - DynamoDB, 43, 65, 208
- E**
- EBS (Elastic Block Storage), 102, 111–112
 - FAQs, 263
 - performance, 263
 - snapshots, 263
 - volume encryption, 123–125
 - EC2 (Elastic Compute Cloud), 3, 21, 49
 - Auto Scaling, 14, 254
 - Detailed Monitoring, 21
 - instance profiles, 78
 - ECS (Elastic Container Service), 43, 83–84
 - launch types, 83–84
 - troubleshooting, 93–94
 - edge computing, Snowball Edge, 118–119
 - Edge Locations, 160–163
 - editing, CloudWatch dashboards, 19
 - EFS (Elastic File System), 112–113
 - accessing, 264
 - backups, 264
 - business cases, 264
 - FAQs, 264
 - egress-only Internet gateways, 163, 171–172
 - Elastic Beanstalk, 81–83, 222
 - applications, 81
 - environment tiers, 81
 - Elastic Container Service, 219
 - elastic IP addresses, 163, 174–175
 - Elastic Load Balancers, 92–93
 - CloudWatch metrics, 95–96
 - error messages, 94–95
 - FAQs, 255–256
 - ElastiCache, 63
 - HA aspects, 64–65
 - in-memory caching engines, 64
 - elasticity, 38, 67
 - AWS Auto Scaling, 42
 - best practices, 44
 - configuring, 45
 - predictive scaling, 44
 - pre-implementation considerations, 44
 - troubleshooting areas, 45–46
 - load balancing, 49
 - encryption, 120–121
 - client-side, 122
 - server-side, 122–123
 - volume, 123–125
 - endpoints
 - gateway, 176
 - interface, 176
 - VPC, 164, 175
 - endpoints, VPC (Virtual Private Cloud), 265–267

environments, Elastic Beanstalk, 81–83
 error messages, Elastic Load Balancing, 94–95
 exam, 225–227
 objectives, 227–228
 preparing for, 228–229
 tools for final preparation, 229–233
 exam information, 222–223

F

failure scenarios, testing, 47
 false positives, 27
 FAQs
 API Gateway, 256–258
 CloudFormation, 268–269
 CloudFront, 268
 CloudTrail, 252–254
 CloudWatch, 249–252
 EBS (Elastic Block Storage), 263
 EFS (Elastic File System), 264
 Elastic Load Balancers, 255–256
 IAM (Identity and Access Management), 264–265
 Lambda, 258–260
 S3 (Simple Storage Service), 260–263
 VPC (Virtual Private Cloud), 265–267
 Fargate, 83
 feature toggles, 75
 file gateways, 115–116
 filter policy, SNS (Simple Notification Service), 59
 flat storage, 102
 Flow Logs, 188–189
 FT (fault tolerance), 46
 functions, 207
 functions, Lambda, 258, 259

G

gateway endpoints, 176
 get-metric-statistics command, 24
 Glacier, 102

Glacier Deep Archive storage class, 106
 Glacier storage class, 105–106
 Global Accelerator, 49
 Global Infrastructure
 AZs, 159–160
 regions, 157–159

H

HA (high availability), 39. *See also* SQS (Simple Queue Services)
 for applications, 50–51
 backups, 63
 best practices, 52
 categories of disruptions, 51
 EFS (Elastic File System), 113
 Elastic Load Balancers, 92–93
 ElastiCache, 63–65
 in-memory caching engines, 64
 example solution in AWS, 67
 FT (fault tolerance), 46
 multi-region, 65
 versus reliability, 46
 versus resiliency, 46
 and RTO (recovery time objective), 46
 high resolution metrics, 24–25
 horizontal scalability, 47
 HSMs (Hardware Security Modules), 121
 HTTP requests, CloudWatch Query API, 11–12

I

IAM (Identity and Access Management), 14, 34, 102, 141–151
 best practices, 148–151
 FAQs, 264–265
 identities, 144–145
 identity federation, 265
 policy simulator, 265
 role account, 34
 temporary security credentials, 265
 user account, 34

- identity federation, 265
- identity-based policies, 35
- infrastructure attacks, 136
- inherited controls, Shared Responsibility Model, 133
- InsufficientInstanceCapacity error, troubleshooting, 94
- installing, CloudWatch, 8
- instance profiles, 78
- InstanceLimitExceeded errors, troubleshooting, 94
- instances, 206
- INSUFFICIENT_DATA state, 27
- Intelligent-Archiving storage class, 105
- interactive query services, Athena, 119–120
- interface endpoints, 176
- internal user-to-Amazon VPC, 187
- Internet gateways, 163, 170–171
- inventory and configuration, 139

J

- JSON, creating dashboards, 19–20

K

- key services, 211–212
- KMS (Key Management Service), 121
- KPIs (key performance indicators), 47

L

- Lambda applications, 219
 - deploying, 91–92
 - event source, 259
 - FAQs, 258–260
- Lambda@Edge, 259
- launch types, ECS (Elastic Container Service), 83–84
- LCU (Load Balancer Capacity Unit), 255
- lifecycle policies, 107–108
 - configuring, 108–111
- limit management, 48
- load balancing, 49

- Elastic Load Balancers, 92–93, 255–256
 - CloudWatch metrics, 95–96
- logging, 78, 139–140, 250–251, 260
 - Flow Logs, 188–189

M

- M out of N alarms, 30
- Management Console
 - alarms, creating, 28–31
 - encrypted root volume, creating, 123–124
 - SQS (Simple Queue Services), configuring, 54–56
- managing resource utilization, 91–111
 - best practices, 205–206
- master keys, 121
- Memcached, 64
- in-memory caching engines, ElastiCache, 64
- metrics, 3, 17
 - CloudWatch, 21–22
 - publishing, 24–25
 - publishing services, 31–33
 - custom, 26
 - high resolution, 24–25
 - standard resolution, 24–25
 - viewing, 21–24
- MFA (multi-factor authentication), 141
- MFA Delete, 107
- monitoring, 67, 139–140, 212. *See also*
 - alarms
 - CPU utilization, 35–36
 - deployments, 80
 - Detailed Monitoring, 21
- Multi-AZ RDS, 61–62
- multi-region HA, 65
- multi-site solution approach, 66

N

- NAT (Network Address Translation), 164, 177
- NAT Gateway, 50

- network interfaces, VPC, 166–167
 - Network Load Balancer, 93, 256
 - network resources, 209–210
 - networking
 - AWS Global Infrastructure
 - AZs, 159–160
 - regions, 157–159
 - CIDR (Classless Inter-Domain Routing), 48
 - CloudFront, 160–163
 - Edge Locations, 160–163
 - redundancy, 50
 - reliability, 49
 - resiliency, 49
 - services, 49–50
 - troubleshooting, 187–189
 - VPC, 48
 - components, 163–165
 - default, 165–166
 - DHCP option sets, 172–173
 - DNS, 174
 - egress-only Internet gateways, 171–172
 - elastic IP addresses, 174–175
 - endpoints, 175
 - gateway endpoints, 176
 - interface endpoints, 176
 - Internet gateways, 170–171
 - NAT, 177
 - network interfaces, 166–167
 - route tables, 168–170
 - network-to-Amazon VPC, 178
 - Direct Connect, 180–182
 - hardware VPN, 178–180
 - software VPN, 183–184
 - VPN CloudHub, 182–183
 - North American regions, 158
- O**
- Object Lock (S3), 263
 - object storage, 102
 - OK state, 27
 - One Zone-Infrequent Access storage class, 105
 - operating solutions, tradeoffs, 210
 - operational excellence
 - best practices, 199–204
 - preparing for, 197
 - OpsWorks, 84–86, 220
 - OpsWorks Stacks, 84
- P**
- PaaS (Platform as a Service), 81
 - parameters
 - for alarms, 27
 - dimensions, 25
 - Partner Networks, 50
 - PCI DSS (Payment Card Industry Data Security Standard), 143
 - penetration testing, 140
 - performance
 - of EBS, 263
 - monitoring, 212
 - performance efficiency, design principles, 204–205
 - permissions, 34
 - physical limits, of AWS, 48
 - pilot light approach, 66
 - pinning dashboard to navigation pane, 19
 - in-place upgrades, 80
 - policies, identity-based, 35
 - PowerShell, 8
 - predictive scaling, 44, 255
 - preparing for operational excellence, 197
 - pricing models, 214
 - PrivateLink, 267
 - programmatic access to CloudWatch data
 - AWS SDKs, 13
 - CloudWatch Query API, 11–12
 - provisioning infrastructure, 75
 - publishing, CloudWatch metrics, 24–25

Q

queues, creating, 54–56

R

RDS (Relational Database Service), 60

- automated failover, 62
- benefits of, 61
- best practices, 61
- database engine version, 63
- Multi-AZ, 61–62
- Read Replicas, 62–63
- snapshots, 63

Read Replicas, 62–63

Redis, 64

RedShift, 208

redundancy, AWS networking, 50

regions, 157–159

reliability

- of AWS networking, 49
- versus HA, 46

removing, services from Cross Service

Dashboard, 16

requirements, for AWS database

solutions, 208

resiliency, 67

- of AWS networking, 49
- versus HA, 46

SQS (Simple Queue Services), 53

resource types

- compute, 206–207
- database, 208
- network, 209–210
- storage, 207

REST API, deploying in API Gateway,

88–91

restores, 66

root user account (AWS), 34

Route, 53, 49, 65

record routing policies, 189

route tables, 163

VPC, 168–170

RPO (recovery point objective), 65

RTO (recovery time objective), 46, 65

S

S3 (Simple Storage Service), 102–103

- advantages of, 104–105
- Amazon VPC Endpoints, 260
- Batch Operations, 262
- client-side encryption, 122
- configuring versioning, 107
- controlling access to data, 260
- CRR (Cross-Region Replication), 263
- FAQs, 260–263
- Intelligent-Tiering, 260
- Inventory reports, 261
- lifecycle policies, 107–108
 - configuring, 108–111
- logging, 260
- Object Lock, 263
- object tags, 261
- server-side encryption, 122–123
- snapshots, 125
- standard storage class, 105
- storage buckets, 102–103
- Storage Class Analysis, 261
- Transfer Acceleration, 261
- uses, 103
- versioning, 106
 - MFA Delete, 107

scalability, 38

- AWS Auto Scaling, 42
 - best practices, 44
 - configuring, 45
 - predictive scaling, 44
 - pre-implementation considerations, 44
 - troubleshooting areas, 45–46
- capabilities in AWS, 79
- horizontal, 47
- RDS (Relational Database Service), 61
- S3 (Simple Storage Service), 104

- SDKs (software development kits),
 - AWS, 13
- security. *See also* encryption
 - S3 (Simple Storage Service), 104
 - storage encryption, 120–121
- security policies
 - data encryption, 138–139
 - DDoS mitigation, 135–137
 - IAM, 141–151
 - infrastructure security, 141
 - inventory and configuration, 139
 - monitoring and logging, 139–140
 - penetration testing, 140
- Serverless Application Model (SAM), 92
- Serverless Application Repository, 92
- server-side encryption, 122–123
- services
 - API Gateway, deploying REST API in, 88–91
 - Cloudwatch, 13–15
 - connectivity, 178
 - ECS (Elastic Container Service), 83–84
 - Elastic Beanstalk, 81–83
 - ElastiCache, 63–65
 - key, 211–212
 - networking, 49–50
 - publishing, 31–33
 - RDS (Relational Database Service), 60
 - automated failover, 62
 - benefits of, 61
 - best practices, 61
 - Multi-AZ, 61–62
 - Read Replicas, 62–63
 - snapshots, 63
 - removing from Cross Service
 - Dashboard, 16
 - S3 (Simple Storage Service), 102–103
 - advantages of, 104–105
 - configuring lifecycle policies, 108–111
 - configuring versioning, 107
 - Glacier Deep Archive storage class, 106
 - Glacier storage class, 105–106
 - Intelligent-Archiving storage class, 105
 - lifecycle policies, 107–108
 - MFA Delete, 107
 - One Zone-Infrequent Access storage class, 105
 - snapshots, 125
 - standard storage class, 105
 - standard-IA storage class, 105
 - storage buckets, 102–103
 - uses, 103
 - versioning, 106
 - SNS (Simple Notification Service), 58
 - characteristics, 59
 - clients, 58
 - filter policy, 59
 - SQS (Simple Queue Services), 52–53
 - asynchronous decoupling, 53
 - configuring, 54–56
 - managing from CLI, 57
 - resiliency, 53
 - shared controls, Shared Responsibility Model, 133
 - Shared Responsibility Model
 - Amazon responsibilities, 133
 - client responsibilities, 134–135
 - controls, 133
 - Shield Advanced, 137–138
 - Shield Standard, 137
 - snapshots, 63, 125, 263
 - Snowball, 117–118
 - Snowball Edge, 118–119
 - SNS (Simple Notification Service), 13, 58
 - characteristics, 59
 - clients, 58
 - configuring, 59–60
 - filter policy, 59
 - software-to-hardware VPN, 186–187

SQS (Simple Queue Services), 52–53
 asynchronous decoupling, 53
 configuring, 54–56
 managing from CLI, 57
 resiliency, 53
 stacks, 84
 CloudFormation, 86
 stages, 89
 standard resolution metrics, 24–25
 standard storage class, 105
 standard-IA storage class, 105
 states, of alarms, 27
 storage buckets, 102–103
 storage classes
 Glacier, 105–106
 Glacier Deep Archive, 106
 Intelligent-Archiving, 105
 One Zone-Infrequent Access s, 105
 standard, 105
 standard-IA, 105
 storage encryption, 120–121
 Storage Gateway, 115–117
 storage solutions, 207
 sustained load testing, 52
 Systems Manager, 220–221
 configuring, 88

T

tagging, 76
 tape gateways, 117
 templates, CloudFormation, 86–87,
 268–269
 temporary security credentials, 265
 testing, 67
 failure scenarios, 47
 TPM (Trusted Platform Module), 118
 tradeoffs, 210
 trails, 253
 troubleshooting
 AWS Auto Scaling, 45–46
 CloudFormation, 96

ECS (Elastic Container Service), 93–94
 ELB CloudWatch metrics, 95–96
 networking, 187–189

U

upgrades, in-place, 80
 user accounts, AWS, 34

V

versioning
 MFA Delete, 107
 S3 (Simple Storage Service), 106
 viewing
 Cross Service Dashboard, 15–16
 metrics, 21–24
 volume encryption, 123–125
 volume gateways, 116–117
 VPC (Virtual Private Cloud), 48
 Amazon-to-Amazon, 184–185
 software VPN, 186
 VPC peering, 185
 CIDR (Classless Inter-Domain
 Routing), 48
 components, 163–165
 cost of using, 266
 default, 165–166, 266
 DHCP option sets, 172–173
 DNS, 174
 egress-only Internet gateways,
 171–172
 elastic IP addresses, 174–175
 endpoints, 175, 265–267
 FAQs, 265–267
 Flow Logs, 188–189
 gateway endpoints, 176
 interface endpoints, 176
 internal user-to-Amazon, 187
 Internet gateways, 170–171
 NAT, 177
 network interfaces, 166–167
 network-to-Amazon

Direct Connect, 180–182
 hardware VPN, 178
 software VPN, 183–184
 VPN CloudHub,
 182–183
 peering connections, 267
 route tables, 168–170
 security groups, 266
VPG (Virtual Private Gateway), 50

W

warm-standby approach, 66
web services, ElastiCache, 63

widget type, choosing for your
 dashboard, 18
Windows, 10, installing CloudWatch
 on, 8

X

X-Ray, 221

Z

zero RTO (recovery time objective), 46