

AWS[®] Certified Database

STUDY GUIDE

SPECIALTY (DBS-C01) EXAM

Includes one year of FREE access after activation to the interactive
online learning environment and study tools:

Custom practice exam

100 electronic flashcards

Searchable key term glossary

**MATHEUS ARRAIS, RENE MARTINEZ BRAVET, LEONARDO CICCONE,
ANGIE NOBRE COCHARERO, ÉRIKA KURAUCHI, AND HUGO ROZESTRATEN**

 **SYBEX[®]**
A Wiley Brand

AWS®

Certified Database

Study Guide

Specialty (DBS-C01) Exam



Matheus Arrais

Rene Martinez Bravet

Leonardo Ciccone

Angie Nobre Cocharero

Érika Kurauchi

Hugo Rozestraten

 **SYBEX**
A Wiley Brand

Copyright © 2023 by John Wiley & Sons, Inc. All rights reserved.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey.

Published simultaneously in Canada and the United Kingdom.

ISBN: 978-1-119-77895-0

ISBN: 978-1-119-77896-7 (ebk.)

ISBN: 978-1-119-77897-4 (ebk.)

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 750-4470, or on the web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at www.wiley.com/go/permission.

Trademarks: WILEY, the Wiley logo, and the Sybex logo are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates, in the United States and other countries, and may not be used without written permission. AWS is a registered trademark of Amazon Technologies, Inc. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc. is not associated with any product or vendor mentioned in this book.

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Further, readers should be aware that websites listed in this work may have changed or disappeared between when this work was written and when it is read. Neither the publisher nor authors shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic formats. For more information about Wiley products, visit our web site at www.wiley.com.

Library of Congress Control Number: 2023932626

Cover image: © Jeremy Woodhouse/Getty Images

Cover design: Wiley

I dedicate this book to my two funny and witty daughters and my lovely husband, who were always supportive throughout the journey.

—Angie Nobre

I dedicate this book to my family and friends who supported and inspired me in this fantastic journey.

—Érika Kurauchi

I dedicate this book to my wife and son, who support me in all adventures.

—Hugo Rozestraten

To my family: Erika, Antonio, Clelia, and Leandro, your love and support made this book a reality.

—Leonardo Ciccone

I dedicate this book to my wife, my daughter Helena, my parents, and my sister, who are the foundation of my life, for supporting me since I started to be curious about technology.

—Matheus Arrais

To my wife, my son, my daughters, and my parents. Specially to my dad who also wrote a database book but couldn't get it published.

—Rene Martinez

Acknowledgments

First, we would like to thank our wonderful families for supporting us and encouraging us to start this amazing journey. We needed to focus hours on writing and reviewing, and their understanding was key for the team to work together and to succeed on finishing this study guide.

Also, we would like to show our appreciation for Amazon Web Services (AWS) for providing cloud-computing platforms, APIs, and the Specialty (DBS-C01) exam to a world-wide audience.

We'd also like to thank acquisitions editor Ken Brown for entrusting us to create this study guide for Wiley. We also appreciate the assistance of John Sleeva, the project editor, and the entire team at Wiley for their guidance in creating this book.

A special thanks to Murilo Nascimento for the detailed reviews and guidance, who agreed to work with our team as the technical editor, after having been a product manager of the AWS Database Blog for years.

Our sincere thanks to Gus Santana, who wrote the *AWS Certified Security Study Guide* years ago, for bringing the team together and showing us by example that the challenge could be met.

Finally, thanks to our readers for trusting us to guide and prepare you for this incredible learning journey.

—The Authors

About the Authors

Matheus Arrais is a World Wide Cross Service Specialist Solutions Architect at Amazon Web Services focusing in cloud foundational and operations services. Prior to his tenure at AWS, he worked for a consulting company as a cloud engineer and networking administrator, accumulating more than 14 years of experience. He has worked in infrastructure, experienced networking, databases, and migrations processes. Matheus holds a dozen AWS certifications and is a frequent speaker at summits, webinars, and re:invent.

Rene Martinez Bravet is principal solutions architect at AWS, focused on enterprise and financial services customers. Before AWS, he worked as an Authorized Google Cloud Platform trainer and as a solutions architect at Ericsson. Rene has more than 15 years of experience in IT and has worked in different roles, including as a software architect, DevOps engineer, and database administrator.

Leonardo Ciccone is a senior database migration specialist at Amazon Web Services, with a focus on database migration and modernization. Leonardo started his career working with embedded hardware, networks, and Linux administration. In 2006, he switched his focus to databases and has never stopped since. He is a certified professional with eight AWS certifications, Google Cloud Architect, Oracle Database, and Oracle Exadata. Leonardo's technical interests include databases, data engineering, cloud, and security.

Angie Nobre Cocharero has been a database specialist solution architect at Amazon Web Services for Latin America since 2020. She has worked for more than 20 years in database implementations, supporting multiple customers in architecting database environments. She has experience with high availability and disaster recovery and has supported migrations and provided best practices for deploying databases.

Érika Kurauchi is the specialist solution architecture manager at Amazon Web Services Latin America, has 20 years of experience in the IT industry. Prior to her tenure at AWS, she worked for Itaú bank for almost 12 years. She has extensive experience in migrations and mergers of large institutions in the financial market and Itaú. She collaborated with Itaú's data-driven journey, leading the enterprise Data Governance Foundation, and became an SME in data management, quality, and protection. Érika holds a degree in data processing, an MBA in business management from FVG with an advanced certificate from UC Irvine international extension, certification in negotiation and leadership from Harvard Law School, and two AWS certifications.

Hugo Rozestraten is a principal data strategy solutions architect and author of several AWS blog posts. Hugo has worked for Google, IBM, Oracle, and ZUP as a data expert. Starting as a developer in 2001, Hugo quickly moved his career toward databases, becoming a certified DBA (Oracle 8i OCP) in 2003 and diving deep into databases ever since. Hugo has a set of content creation and is a frequent presenter in summits and webinars.

About the Technical Editor

Murilo Cerone Nascimento has experience with a variety of relational and nonrelational databases. He has been working with databases for the past 13 years, most of the time as a consultant in private companies. Murilo holds a bachelor's degree in engineering with a specialization in computer engineering from Escola Politécnica de São Paulo. He currently works as a product manager at Amazon Web Services Brazil.

Contents at a Glance

<i>Introduction</i>		<i>xxv</i>
<i>Assessment Test</i>		<i>xxxv</i>
Part I	Workload-Specific Database Design	1
Chapter 1	Databases—from Your Server to AWS Cloud	3
Chapter 2	Basic AWS Concepts	13
Chapter 3	Purpose-Built Databases	35
Part II	Management and Operations, Database Security, Monitoring and Troubleshooting per Workload	51
Chapter 4	Relational Databases on AWS	53
Chapter 5	Low-Latency Response Time for Your Apps and APIs	109
Chapter 6	Document Databases in the Cloud	165
Chapter 7	Better Places Than Databases to Store Large Objects	195
Chapter 8	Deliver Valuable Information at the Speed Your Business Needs	225
Chapter 9	Discovering Relationships Using Graph Databases	275
Chapter 10	Immutable Database and Traceable Transactions	295
Chapter 11	Caching Data with In-Memory Databases	313
Part III	Deployment and Migration	337
Chapter 12	Migrating Your Data to AWS	339
Chapter 13	Disaster Recovery	369
Chapter 14	Save Time and Reduce Errors Automating Your Infrastructure	389
Appendix	Answers to Review Questions	421
<i>Index</i>		<i>439</i>

Contents

<i>Introduction</i>		<i>xxv</i>
<i>Assessment Test</i>		<i>xxxv</i>
Part I	Workload-Specific Database Design	1
Chapter 1	Databases—from Your Server to AWS Cloud	3
	Databases from the Beginning	4
	1960s	5
	1970s	5
	1980s	6
	1990s	7
	2000–2010	7
	2010–Today	8
	Databases on Premises	9
	Basic Infrastructure	9
	Complex Infrastructure and Resiliency Management	10
	Databases in the Cloud	10
	Data Remains Data	11
	DBAs Are the Initial DevOps	11
	DBA Career—Now What?	11
	Summary	12
Chapter 2	Basic AWS Concepts	13
	AWS Global Infrastructure	14
	AWS Regions	14
	AWS Availability Zones	14
	AWS Local Zones	15
	AWS Wavelength Zones	15
	Points of Presence	16
	Networking	16
	IP Addressing	17
	Subnets	17
	VPC Routing	17
	Gateways and Connectivity Options	18
	Security	21
	Network Security	21
	AWS Shared Responsibility Model	22
	Identity and Access Management	24

	Data Encryption	25
	Storage at AWS	26
	Amazon Elastic Block Storage	27
	Amazon Elastic File System	28
	Amazon Simple Storage Service	29
	Operations	31
	Monitoring and Logging	32
	Amazon CloudWatch	32
	AWS CloudTrail	33
	Summary	34
Chapter 3	Purpose-Built Databases	35
	Data store Concepts	36
	Data Access Patterns	37
	Latency	37
	Scaling	37
	Transaction Support	38
	Consistency	38
	Volume	38
	Durability	39
	Availability	39
	Security/Compliance	39
	Business Logic	39
	Cost	40
	Purpose-Built Databases on AWS	40
	Relational Databases	40
	Nonrelational Databases	42
	Summary	46
	Exam Essentials	46
	Review Questions	48
Part II	Management and Operations, Database Security, Monitoring and Troubleshooting per Workload	51
Chapter 4	Relational Databases on AWS	53
	Relational Databases	54
	Structured Query Language	55
	Install and Manage Databases Yourself	57
	I/O Requirement	58
	Managing Databases on EC2	60
	Monitoring Databases on EC2	60
	Scaling Databases	61
	Upgrading Databases	61

Managed Services for Relational Databases	62
Launching an RDS Instance	63
Managing High Availability and Scalability	72
Configuring RDS Parameter Groups	74
Configuring RDS Option Groups	76
Deletion Protection	78
RDS Pricing Model	79
Amazon Aurora Cloud-Native Relational Database	79
Amazon Aurora Storage	80
Amazon Aurora DB Clusters	81
Amazon Aurora High Availability	82
Amazon Aurora Global Database	82
Amazon Aurora Read Replica Across Regions	84
Amazon Aurora Serverless	84
Amazon Aurora Multi-master	85
Patch Management and Upgrade	86
Monitoring and Performance Management	86
Backup and Restore	90
Backtrack	94
Cloning an Amazon Aurora DB Cluster Volume	94
Auditing	97
Database Activity Stream in Amazon Aurora	98
Security	98
Migrating Databases	101
Summary	101
Exam Essentials	101
Exercises	102
Review Questions	106
Chapter 5	
Low-Latency Response Time for Your Apps and APIs	109
Getting Started with Modern Applications and NoSQL Databases	110
Amazon DynamoDB	112
Design Considerations	112
Migrating Your Data into DynamoDB	114
Query Considerations	116
Amazon DynamoDB Index Options	120
DynamoDB Capacity Modes	122
Other Features of DynamoDB	124
Backup/Restore	134
Scalability	136
Security	139
Monitoring	143

	Amazon Keyspaces	148
	Design Considerations	148
	Migrating Your Data into Keyspaces	149
	Query Considerations	150
	Keyspaces Capacity Modes	151
	Consistency Models	152
	Uses Cases	153
	Best Practices	153
	Backups	154
	Scalability	155
	Security	155
	Monitoring	157
	Summary	158
	Exam Essentials	159
	Exercises	159
	Review Questions	162
Chapter 6	Document Databases in the Cloud	165
	Introducing Document Databases	166
	Getting Started with Amazon DocumentDB	170
	Creating an Amazon DocumentDB	
	Cluster	174
	Amazon DocumentDB Architecture	175
	Security	177
	Access Control	177
	Data Protection	178
	Other Features	179
	Backup and Restore	180
	Performance and Scaling	182
	Compatibility between Amazon DocumentDB and MongoDB	182
	Migrating from MongoDB to Amazon DocumentDB	183
	Amazon DocumentDB Monitoring	184
	Developing with Amazon DocumentDB	187
	When to Use DynamoDB vs. DocumentDB	187
	Amazon DocumentDB Pricing	188
	Summary	189
	Exam Essentials	189
	Exercises	190
	Review Questions	192
Chapter 7	Better Places Than Databases to Store Large Objects	195
	Databases and Large Objects	196
	Introducing Amazon S3	198

Costs of Amazon S3 vs. Elastic Block Storage	200
Moving LOBs to Amazon S3	204
Creating an S3 Bucket	204
Putting or Uploading Objects	206
Indexing LOBs in Amazon S3	208
Additional S3 Features	211
Backup and Dump Files	211
Other Use Cases	211
Pay per Usage and Scalability	212
Availability and Durability	213
Security	214
Access Control	214
Data Encryption	215
Summary	216
Exam Essentials	216
Exercises	217
Review Questions	219

Chapter 8 Deliver Valuable Information at the Speed Your Business Needs 225

Information Latency	226
Data Warehouses	229
Database Engines for Data Warehouses	231
Migrating Data Warehouses to Amazon Redshift Using AWS SCT	233
Amazon Redshift	235
Redshift Cluster Architecture	235
Table Design in Redshift	235
Loading Data into Redshift	240
Data Lakes in AWS	241
Redshift Spectrum	242
Redshift Federated Queries	242
Data Lakehouse	242
Redshift Cluster Node Types	243
Redshift Monitoring	244
Redshift Scalability	245
Redshift Security	246
Redshift Data Resilience and Backup	247
Time-Series Databases	249
Amazon Timestream	251
Amazon Timestream Architecture	252
Loading Data into Amazon Timestream	253
Querying Data from Amazon Timestream	254
Amazon Timestream Monitoring	255

	Amazon Timestream Scalability	255
	Amazon Timestream Security	256
	Amazon Timestream Data Resilience and Backup	257
	Amazon OpenSearch Service	257
	Amazon OpenSearch Service Domain Architecture	257
	Loading Data to Amazon OpenSearch Service	257
	Searching for Data in Amazon OpenSearch Service	258
	Amazon OpenSearch Service Monitoring	259
	Amazon OpenSearch Service Scalability	259
	Amazon OpenSearch Service Security	259
	Amazon OpenSearch Service Data Resilience and Backup	259
	Summary	260
	Exam Essentials	261
	Exercises	261
	Review Questions	271
Chapter 9	Discovering Relationships Using Graph Databases	275
	Graph Databases	276
	When to Use Graph Databases	277
	Common Use Cases	278
	Amazon Neptune	279
	High-Level Architecture	279
	Graph Models and Query Languages	280
	Using and Extracting Data from Amazon Neptune	282
	Storage Architecture	284
	Data Resilience	284
	Read Replicas	284
	Scalability	286
	Availability	287
	Failover Policy	287
	Security	288
	Automatic Backup and Restore	289
	Monitoring	290
	Summary	291
	Exam Essentials	291
	Review Questions	292
Chapter 10	Immutable Database and Traceable Transactions	295
	Amazon Quantum Ledger Database	296
	Amazon QLDB Components	297
	Working with Amazon QLDB	300
	Backup and Durability	303
	Performance and Scalability	304

	Security	305
	Monitoring	306
	Best Practices	307
	Summary	307
	Exam Essentials	308
	Exercises	308
	Review Questions	311
Chapter 11	Caching Data with In-Memory Databases	313
	Built-in Database Cache	314
	Local Application Cache	315
	In-Memory Databases	315
	Caching Use Cases	315
	Remote Cache Strategies	316
	Caching Data in AWS	317
	Caching Data with ElastiCache	319
	Memcached or Redis?	319
	Memcached Architecture on Amazon	
	ElastiCache Service	320
	Redis Architecture on the Amazon	
	ElastiCache Service	321
	Backup and Restore	324
	Security	324
	Monitoring	325
	Amazon MemoryDB for Redis	325
	Summary	326
	Exam Essentials	326
	Exercises	327
	Review Questions	334
Part III	Deployment and Migration	337
Chapter 12	Migrating Your Data to AWS	339
	Network Communication and Data Migration	340
	Optimizing Bandwidth Usage	343
	Database Migration to AWS	343
	AWS Schema Conversion Tool	344
	AWS Data Migration Service	347
	Setting Up AWS DMS	350
	AWS DMS Continuous Replication	352
	AWS DMS Best Practices	353
	AWS DMS Security	356
	AWS DMS Resilience	356

	Other AWS Services for Data Migration	357
	AWS DataSync	357
	AWS Snow Family	357
	AWS Storage Gateway	358
	Choosing the Migration Path	358
	One or Many Target Databases	358
	Small, Noncritical Databases	359
	Very Large Databases	359
	Summary	360
	Exam Essentials	360
	Exercises	361
	Review Questions	365
Chapter 13	Disaster Recovery	369
	Understanding and Planning Disaster Recovery Requirements	370
	Recovery Point Objective and Recovery Time Objective	371
	Challenges in Disaster Recovery	371
	Managing Disaster Recovery Strategies	373
	Backup and Recovery Strategy	373
	Database Replication Methods: Multi-AZ and	
	Cross-Region Replication	373
	Databases Deployed on EC2	374
	Amazon RDS	375
	Amazon Aurora	378
	Amazon Neptune	383
	Summary	384
	Exam Essentials	384
	Review Questions	386
Chapter 14	Save Time and Reduce Errors Automating Your Infrastructure	389
	AWS CloudFormation	390
	Components	393
	Important Concepts	400
	Updating AWS CloudFormation Stacks	403
	Deleting AWS CloudFormation Stacks	408
	AWS Systems Manager Parameter Store	409
	AWS Secrets Manager	411
	Summary	415
	Exam Essentials	415
	Review Questions	416

Appendix	Answers to Review Questions	421
	Chapter 3: Purpose-Built Databases	422
	Chapter 4: Relational Databases on AWS	422
	Chapter 5: Low-Latency Response Time for Your Apps and APIs	424
	Chapter 6: Document Databases in the Cloud	425
	Chapter 7: Better Places Other Than Databases to Store Large Objects	426
	Chapter 8: Deliver Valuable Information at the Speed Your Business Needs	427
	Chapter 9: Discovering Relationships Using Graph Databases	429
	Chapter 10: Immutable Database and Traceable Transactions	431
	Chapter 11: Caching Data with In-Memory Databases	432
	Chapter 12: Migrating Your Data to AWS	433
	Chapter 13: Disaster Recovery	434
	Chapter 14: Save Time and Reduce Errors Automating Your Infrastructure	436
<i>Index</i>		439

Table of Exercises

Exercise	4.1	Create an Amazon RDS MySQL and enable multi-AZ	103
Exercise	4.2	Create an Amazon Aurora cluster	104
Exercise	4.3	Add an Amazon Aurora replica and modify the priority	105
Exercise	5.1	Create a DynamoDB Table and Populate It	160
Exercise	5.2	Create a Global Secondary Index.	160
Exercise	5.3	Query the Data on Your Table	161
Exercise	5.4	Clean Up the Created Resources	161
Exercise	6.1	Create a Secure Amazon DocumentDB Cluster.	190
Exercise	6.2	Apply a Custom Parameter Group in Your Cluster	190
Exercise	7.1	Create a Secure S3 Bucket for Your LOBs	217
Exercise	7.2	Create a Lifecycle Rule for Your S3 Bucket.	218
Exercise	7.3	Enable Versioning for S3.	218
Exercise	8.1	Create an Amazon Redshift Cluster	262
Exercise	8.2	Use the Query Editor to Create and Interact with a Table.	262
Exercise	8.3	Use the Query Editor to Insert Data and Export Resultset	264
Exercise	8.4	Take a Manual Snapshot of Your Cluster	265
Exercise	8.5	Explore the Resize Options for Your Cluster	266
Exercise	8.6	Destroy Your Cluster	266
Exercise	8.7	Create an Amazon Timestream Database	267
Exercise	8.8	Run Sample Queries on Amazon Timestream.	267
Exercise	8.9	Delete Your Amazon Timestream Database.	268
Exercise	8.10	Create an Amazon OpenSearch Service Domain	268
Exercise	8.11	Log In to Your Amazon OpenSearch Service Domain.	269
Exercise	8.12	Delete Your Amazon OpenSearch Service Domain.	270
Exercise	10.1	Create a New Amazon QLDB Ledger Database.	309
Exercise	10.2	Load the Sample Data	309
Exercise	10.3	Query the Data	309
Exercise	10.4	Request a Digest and Perform a Document Revision.	310
Exercise	11.1	Create an Amazon Linux EC2 Machine to Be Your Bastion Host	327
Exercise	11.2	Create a Memcached Cluster	328
Exercise	11.3	Cache Data with the pymemcache Python Library	330
Exercise	11.4	Create a Redis Cluster	331
Exercise	11.5	Cache Data with the Redis Library for the Python Library	332

Exercise	11.6	Delete the Resources to Avoid Unnecessary Costs	333
Exercise	12.1	Create a MariaDB RDS	361
Exercise	12.2	Create an Aurora PostgreSQL Target Database	361
Exercise	12.3	Create an EC2 Bastion Host to Install the AWS SCT	362
Exercise	12.4	Set Up JDBC Drivers and Schema in EC2	362
Exercise	12.5	Convert the Schema	363
Exercise	12.6	Migrate the Data	363
Exercise	12.7	Delete the AWS Resources	364

Introduction

In 2022, 97 zettabytes of data were created, copied, captured, and consumed, and 181 zettabytes are projected by 2025. Databases play a crucial role in this data world. Over the last decade, database technologies matured: from time series to graph, from ledger to key-value, from relational to nonrelational. It's undeniable: the database field is flourishing. Understanding and applying the concepts covered in this book will improve your contributions and impact as a DBA, developer, architect, or analyst.

The purpose of this book is to help you pass the AWS Certified Database – Specialty exam. As of the writing of this book, the exam is composed of 65 questions, either multiple response or multiple choice, with a total length of 180 minutes.



Don't just study the questions and answers! The questions on the actual exam will be different from the practice questions included in this book. The exam is designed to test your knowledge of a concept or objective, so use this book to learn the objectives behind the questions.



Non-native English speakers can request a 30-minute exam extension when taking the exam in English. For more information, check <https://aws.amazon.com/certification/policies/before-testing>.



When reading exam questions, try to outline the key requirements, and work backwards from the responses. This will help you eliminate wrong responses from the start, saving time.

What Does This Book Cover?

This book covers the topics outlined in the AWS Certified Database – Specialty (DBS-C01) exam guide available here:

https://d1.awsstatic.com/training-and-certification/docs-database-specialty/AWS-Certified-Database-Specialty_Exam-Guide.pdf

Chapters follow a logical evolution of the workload-specific types of databases and their use cases, and they cover monitoring, troubleshooting and security features.

Chapter 1: Databases—From Your Server to AWS Cloud This chapter presents an evolution of databases and the differences of managing databases on premises and in the

AWS cloud. It presents important concepts and features of a self-managed database in AWS, configuring a database using Amazon EC2, and understanding the limitations. It also introduces database options in AWS.

Chapter 2: Basic AWS Concepts This chapter introduces the basics of AWS infrastructure, networking, security, storage, and operations, which are important to understand when designing database solutions in AWS. If you are very experienced with AWS components, you may want to skip this chapter; otherwise, take your time to learn more about AWS regions, availability zones, edge locations, VPC, VPN, direct connect, network connectivity, network security, security model, identity and access management, data encryption, block storage, file storage, object storage, monitoring, logging, and auditing.

Chapter 3: Purpose-Built Databases This chapter discusses workload-specific database design. It introduces how to evaluate application requirements and datastore characteristics, taking into consideration several important aspects: data access patterns, latency, scaling, transaction support, consistency, volume, durability, availability, security and compliance, business logic, and cost. It concludes by presenting a comprehensive way to compare the requirements with AWS managed services for databases.

Chapter 4: Relational Databases on AWS This chapter presents transactional databases in AWS with a focus on relational databases, which are mostly used to support transactional environments using concepts of ACID transactions. It presents options for managing and installing relational databases on Amazon EC2, storage requirements and options on AWS, and monitoring, updating, and scale options for this type of installation. A large part of this chapter is dedicated to Amazon RDS, including the engine options and a deep dive into Amazon Aurora, covering all the topics related to RDS: deployment, migration options, management, operations, monitoring, troubleshooting, and database security.

Chapter 5: Low-Latency Response Time for Your Apps and APIs This chapter introduces the concepts of low-latency and NoSQL databases, presenting two low-latency database services on AWS: Amazon DynamoDB and Amazon Keyspaces (for Apache Cassandra). It details Amazon DynamoDB table design tenets as well as partition and sort key definitions, indexing options, queries, scans, transactions, caching, and global options. For Amazon Keyspaces (for Apache Cassandra) it presents partitions, clustering keys, and static columns concepts. It also covers all the topics related to DynamoDB and Keyspaces: capacity, deployment, scalability, migration options, management, operations, monitoring, troubleshooting, and database security.

Chapter 6: Document Databases in the Cloud This chapter discusses document database concepts and options in AWS. It presents Amazon DocumentDB (with MongoDB compatibility) with its resilient storage layer and the key architecture components such as cluster, instance, and reader endpoints. The chapter covers details of the topics related to Amazon DocumentDB: capacity, deployment, scalability, migration options, management, operations, monitoring, troubleshooting, and database security.

Chapter 7: Better Places Than Databases to Store Large Objects This chapter presents the concept of large objects and how they were handled before AWS and how they are handled with the availability of AWS options. It introduces Amazon S3, compares it to other AWS storage options, and goes into the details of S3 storage classes, data load and retrieval, life-cycle management, operations, monitoring, access control, and security.

Chapter 8: Deliver Valuable Information at the Speed Your Business Needs This chapter describes how to store and use analytics and time series and operational data, with purpose-built databases in AWS to extract value for business. It introduces the information latency concept, presents Amazon Redshift and Amazon Timestream in detail, and also mentions Amazon OpenSearch, a topic that is not covered on the exam. For Amazon Redshift, the chapter goes into the details of table design, loading data, cluster architecture, and cluster options. For Amazon Timestream, it also covers the architecture and loading data options. For Amazon Redshift and Amazon Timestream, it goes into detail on deployment, scalability, migration options, management, operations, monitoring, troubleshooting, and database security.

Chapter 9: Discovering Relationships Using Graph Databases This chapter discusses the concepts of graph databases and their usage to handle relationships between data items. It presents the basic concepts of graph databases and then goes deep into Amazon Neptune data loading and query options, cluster architecture, scalability, migration options, management, operations, monitoring, troubleshooting, and database security.

Chapter 10: Immutable Database and Traceable Transactions This chapter introduces the concept of immutable databases and the AWS service Amazon Quantum Ledger Database (QLDB). It goes into the details of Amazon QLDB and how it ensures no data can be modified or deleted after it is inserted, using an immutable transaction log. It discusses the Amazon QLDB components: ledger, table, document, journal, query engine, and cryptographic verification. It also presents how to load data to and query data from Amazon QLDB, as well as service scalability, management, operations, monitoring, troubleshooting, and database security.

Chapter 11: Caching Data with In-Memory Databases This chapter presents caching strategies and how AWS services can help to implement them. It introduces Amazon ElastiCache service for Redis and Memcached and also Amazon MemoryDB, discussing their features and how they can handle caching requirements and strategies. It also discusses in detail how to distribute data in each engine and how to address deployment, scalability, migration options, management, operations, monitoring, troubleshooting, and database security.

Chapter 12: Migrating Your Data to AWS This chapter presents the options to migrate data to AWS database services. It discusses network communication and specialized database migration services like AWS SCT and AWS DMS, which are native database tools that help in data migration for each engine, including AWS DataSync, AWS Snow Family, and AWS Storage Gateway. It also details database migration strategy and downtime minimization along with security and resilience options and requirements.

Chapter 13: Disaster Recovery This chapter discusses disaster recovery requirements and strategies for database applications in AWS. It presents the concepts of RTO and RPO and how they are affected by the deployment and replication strategies. It also explores how AWS managed services for databases handle data resilience and replication using Multi-AZ and global database options.

Chapter 14: Save Time and Reduce Errors Automating Your Infrastructure This chapter introduces the concept of infrastructure as code (IaC) and how this relates to database deployment and management. It presents AWS CloudFormation and goes into the details of how to simplify infrastructure management using templates, stacks, and change sets. It also discusses best practices related to database security and access management, using AWS System Manager Parameter and AWS Secrets Manager.

Interactive Online Learning Environment and Test Bank

Studying the material in the AWS® *Certified Database Study Guide: Specialty (DBS-C01) Exam* is an important part of preparing for the AWS Certified Database Specialty (DBS-C01) certification exam, but we provide additional tools to help you prepare. The online test bank helps you understand the types of questions that appear on the certification exam. The online test bank runs on multiple devices.

- **Sample tests:** The sample tests in the test bank include all the questions at the end of each chapter as well as the questions from the assessment test. In addition, there are two practice exams with 50 questions each. You can use these tests to evaluate your understanding and identify areas that may require additional study.
- **Flashcards:** The flashcards in the test bank will push the limits of what you should know for the certification exam. There are 100 questions that are provided in digital format. Each flashcard has one question and one correct answer.
- **Glossary:** The online glossary is a searchable list of key terms introduced in this exam guide that you should know for the AWS Certified Database Specialty (DBS-C01) certification exam.

Visit www.wiley.com/go/sybextestprep to register and gain access to this interactive online learning environment and test bank with study tools. To start using these tools to study for the AWS Certified Database Specialty (DBS-C01) exam, go to www.wiley.com/go/sybextestprep, register your book to receive your unique PIN, and once you have the