

Google Cloud Certified

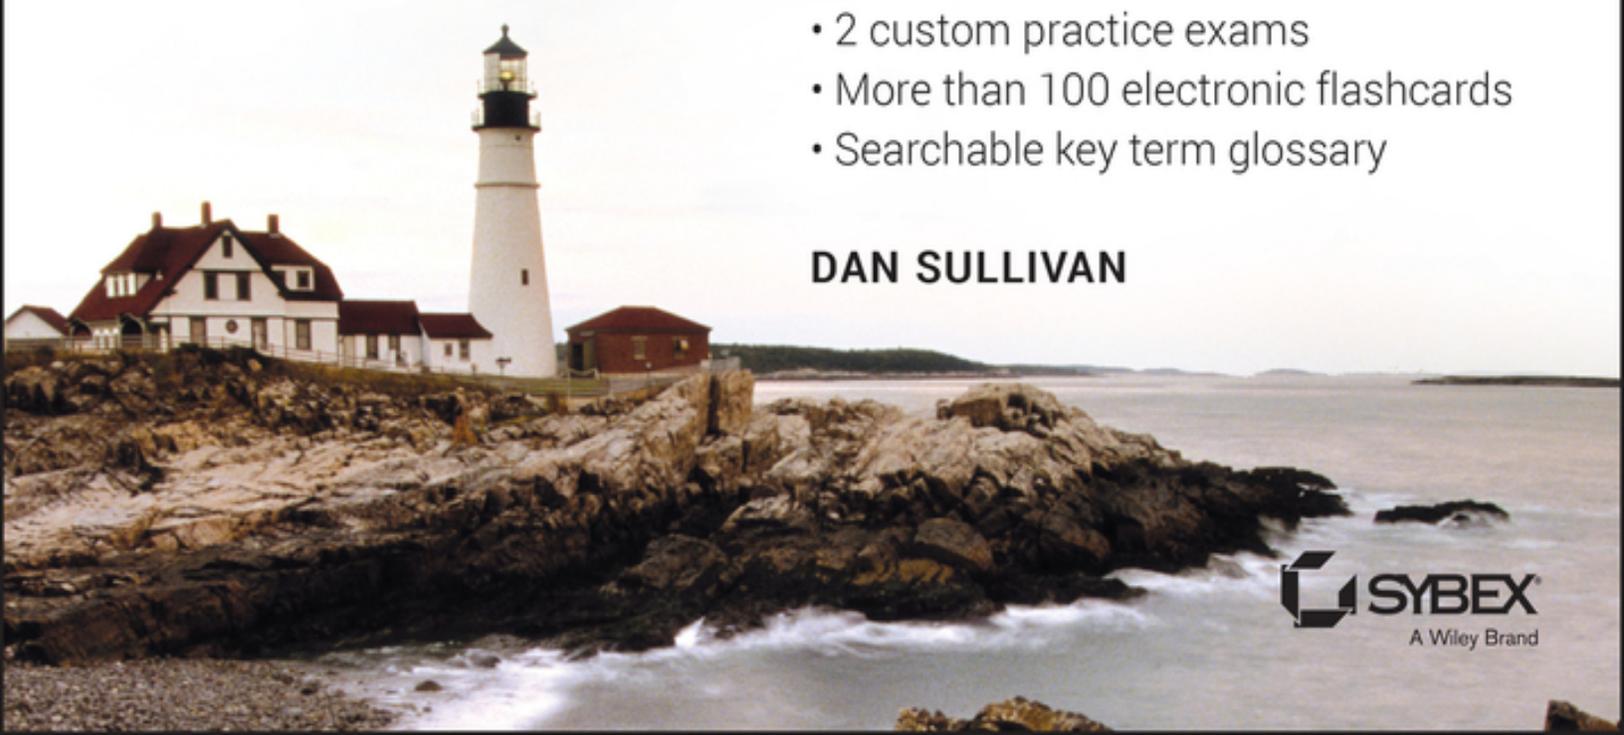
Professional Cloud Architect Study Guide

Second Edition

**Includes one year of FREE access
to interactive online learning
environment and study tools:**

- 2 custom practice exams
- More than 100 electronic flashcards
- Searchable key term glossary

DAN SULLIVAN



SYBEX
A Wiley Brand

Table of Contents

[Cover](#)

[Title Page](#)

[Copyright](#)

[Dedication](#)

[Acknowledgments](#)

[About the Author](#)

[About the Technical Editors](#)

[Introduction](#)

[What Does This Book Cover?](#)

[Interactive Online Learning Environment and Test Bank](#)

[Additional Resources](#)

[Objective Map](#)

[Assessment Test](#)

[Answers to the Assessment Test](#)

[Chapter 1: Introduction to the Google Professional Cloud Architect Exam](#)

[Exam Objectives](#)

[Analyzing Business Requirements](#)

[Analyzing Technical Requirements](#)

[Exam Case Studies](#)

[Summary](#)

[Exam Essentials](#)

[Review Questions](#)

[Chapter 2: Designing Solutions to Meet Business Requirements](#)

[Business Use Cases and Product Strategy](#)
[Application Design and Cost Considerations](#)
[Systems Integration and Data Management](#)
[Compliance and Regulation](#)
[Security](#)
[Success Measures](#)
[Summary](#)
[Exam Essentials](#)
[Review Questions](#)

[Chapter 3: Designing Solutions to Meet Technical Requirements](#)

[High Availability](#)
[Scalability](#)
[Reliability](#)
[Summary](#)
[Exam Essentials](#)
[Review Questions](#)

[Chapter 4: Designing Compute Systems](#)

[Compute Services and Use Cases](#)
[Anthos](#)
[AI and Machine Learning Services](#)
[Data Flows and Pipelines](#)
[Compute System Provisioning](#)
[Additional Design Issues](#)
[Summary](#)
[Exam Essentials](#)
[Review Questions](#)

[Chapter 5: Designing Storage Systems](#)

[Overview of Storage Services](#)

[Object Storage with Google Cloud Storage](#)

[Network-Attached Storage with Google Cloud Filestore](#)

[Databases](#)

[Data Retention and Lifecycle Management](#)

[Networking and Latency](#)

[Summary](#)

[Exam Essentials](#)

[Review Questions](#)

[Chapter 6: Designing Networks](#)

[IP Addressing, Firewall Rules, and Routers](#)

[Virtual Private Clouds](#)

[Hybrid-Cloud Networking](#)

[Service-Centric Networking](#)

[Load Balancing](#)

[Additional Network Services](#)

[Summary](#)

[Exam Essentials](#)

[Review Questions](#)

[Chapter 7: Designing for Security and Legal Compliance](#)

[Identity and Access Management and Related Access Control Services](#)

[Organization Constraints](#)

[Data Security](#)

[Security Evaluation](#)

[Security Design Principles](#)

[Major Regulations](#)

[ITIL Framework](#)

[Summary](#)

[Exam Essentials](#)

[Review Questions](#)

[Chapter 8: Designing for Reliability](#)

[Improving Reliability with Cloud Operations Suite](#)

[Release Management](#)

[Systems Reliability Engineering](#)

[Summary](#)

[Exam Essentials](#)

[Review Questions](#)

[Chapter 9: Analyzing and Defining Technical Processes](#)

[Software Development Lifecycle Plan](#)

[Continuous Integration/Continuous Delivery](#)

[Troubleshooting and Post-Mortem Analysis Culture](#)

[IT Enterprise Processes](#)

[Business Continuity Planning and Disaster Recovery](#)

[Summary](#)

[Exam Essentials](#)

[Review Questions](#)

[Chapter 10: Analyzing and Defining Business Processes](#)

[Stakeholder Management](#)

[Change Management](#)

[Team Skill Management](#)

[Customer Success Management](#)

[Cost Optimization/Resource Optimization](#)

[Summary](#)

[Exam Essentials](#)

[Review Questions](#)

[Chapter 11: Development and Operations](#)

[Application Development Methodologies](#)

[Technical Debt](#)

[API Best Practices](#)

[Testing Frameworks](#)

[Data and System Migration Tooling](#)

[Interacting with Google Cloud Programmatically](#)

[Summary](#)

[Exam Essentials](#)

[Review Questions](#)

[Chapter 12: Migration Planning](#)

[Integrating Cloud Services with Existing Systems](#)

[Migrating Systems and Data to Support a Solution](#)

[Software Licensing Mapping](#)

[Network Planning](#)

[Summary](#)

[Exam Essentials](#)

[Review Questions](#)

[Appendix: Answers to the Review Questions](#)

[Index](#)

[End User License Agreement](#)

List of Tables

Chapter 2

[TABLE 2.1 Examples of Google Cloud Platform managed services](#)

Chapter 3

[TABLE 3.1 Example availability SLAs and corresponding downtimes](#)

Chapter 8

[TABLE 8.1 Example of a CPU utilization time series for a VM instance](#)

List of Illustrations

Chapter 4

[FIGURE 4.1 Kubernetes clusters have a set of worker nodes that are managed b...](#)

[FIGURE 4.2 Pods are deployed on nodes, which may be grouped into multiple no...](#)

[FIGURE 4.3 Kubernetes uses multiple types of IP addresses for diff...](#)

Chapter 7

[FIGURE 7.1 Google Cloud Platform resource hierarchy](#)

Chapter 8

[FIGURE 8.1 Service dashboard showing time-series data](#)

Google Cloud Certified

Professional Cloud Architect Study Guide

Second Edition



Dan Sullivan



Copyright © 2022 by Dan Sullivan. All rights reserved.

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

Published simultaneously in Canada.

ISBN: 978-1-119-87105-7

ISBN: 978-1-119-87106-4 (ebk.)

ISBN: 978-1-119-87107-1 (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.

Limit of Liability/Disclaimer of Warranty: The publisher and the author make no representations or warranties with respect to the accuracy or completeness of the contents of this work and specifically disclaim all warranties, including without limitation warranties of fitness for a particular purpose. No warranty may be created or extended by sales or promotional materials. The advice and strategies contained herein may not be suitable for every situation. This work is sold with the understanding that the publisher is not engaged in rendering legal, accounting, or other professional services. If professional assistance is required, the services of a competent professional person should be sought. Neither the publisher nor the author shall be liable for damages arising herefrom. The fact that an organization or Website is referred to in this work as a citation and/or a potential source of further information does not mean that the author or the publisher endorses the information the organization or Website may provide or recommendations it may make. Further, readers should be aware the Internet Websites listed in this work may have changed or disappeared between when this work was written and when it is read.

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: 2022931858

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. 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.

Cover image: © Jeremy Woodhouse/Getty Images

Cover design: Wiley

for Katherine

Acknowledgments

I have been fortunate to work again with professionals from Waterside Productions and Wiley to create this study guide.

Carole Jelen, vice president of Waterside Productions, and Jim Minatel, associate publisher at John Wiley & Sons, led the effort to continue to create Google Cloud certification guides. It was a pleasure to work with Kristi Bennett, Melissa Burlock, Saravanan Dakshinamurthy, and Pete Gaughan, who managed the revision process and made this project go more smoothly than I expected.

I am especially grateful to Ammett Williams and Mark Grand for their deep knowledge of Google Cloud and the thorough technical review they provided of the second edition of this book. Their input has significantly improved the explanation of multiple topics. Thanks to Valerie Parham-Thompson for her technical review of the first edition of this book.

My sons, James and Nicholas, both technology writers themselves, were my first readers and helped me get the original manuscript across the finish line. Katherine, my wife and partner in so many ventures, was again key to completing yet another project.

—Dan Sullivan

About the Author

Dan Sullivan is a principal engineer specializing in cloud architecture, data architecture, and data analytics. Dan is the author of the *Official Google Cloud Certified Associate Cloud Engineer Study Guide* (Sybex, 2019), *Official Google Cloud Certified Professional Data Engineer Study Guide* (Sybex, 2019), and *NoSQL for Mere Mortals* (Addison-Wesley Professional, 2015). He is an online instructor with numerous Google Cloud training courses on Udemy, including *Google Cloud Professional Architect: Get Certified*, *Google Cloud Professional Data Engineer: Get Certified*, and *Google Cloud Associate Engineer: Get Certified*. He is also the author of several LinkedIn Learning courses on databases, data science, and machine learning.

About the Technical Editors

Ammett Williams is a very simple and sometimes avid daydreamer who has more than 14 years of experience in the IT industry. He has a strong inclination to help others learn and challenge themselves with a lot of experience gained as a team leader. Ammett has started the platform called Start Cloud Now with the aim to inspire others along their IT career path. Ammett holds several IT certifications, including CCIE #43569, CISSP, AWS, and a few Google Cloud professional level certs. Ammett can be found online on LinkedIn at www.linkedin.com/in/ammiett and is also a developer relations engineer at Google.

Mark Grand has over 30 years of experience in software development and architecture. The author of eight books on software architecture and Java, he has deep experience in distributed applications, the Java ecosystem, and database design. He has worked with Java since before 1.0. He can translate English to SQL. Mark is also a GCP Certified Professional Cloud Architect.

The application development areas that Mark has been involved with include social media, analytics, what-if analysis, e-commerce, security, machine learning, blockchain, EDI translation, data warehouse, big data, BPM internals, and database internals.

Mark's areas of domain expertise include retail, travel, shipping, credit card processing, healthcare, facilities management, accounting, advertising, and bioinformatics. Companies that Mark has worked with include JFrog, IBM, HP, InComm, AutoZone, Whole Foods, Home Depot, TSYS, Macy's, Deloitte, Oracle, Young & Rubicam, and Bridge2 Solutions.

In his spare time, Mark enjoys cooking and composing music.

Introduction

The Google Cloud Platform is a diverse and growing set of services. To pass the Google Cloud Professional Cloud Architect exam, you will need to understand how to reason about both business requirements and technical requirements. This is not so much a test of knowledge about how to do specific tasks in GCP, such as attaching a persistent disk to a VM instance, which is the type of question you are more likely to get or see on the Google Cloud Associate Cloud Engineer exam. The Google Cloud Professional Architect exam tests your ability to perform high-level design and architecture tasks related to the following:

- Designing applications
- Planning migrations
- Ensuring feasibility of proposed designs
- Optimizing infrastructure
- Building and deploying code
- Managing data lifecycles

You will be tested on your ability to design solutions using a mix of compute, storage, networking, and managed services. The design must satisfy both business and technical requirements. If you find a question that seems to have two correct technical answers, look closely at the business requirements. There is likely a business consideration that will make one of the options a better choice than the other. For example, you might have a question about implementing a stream processing system, and the options include a solution based on Apache Flink

running in Compute Engine and a solution using Cloud Dataflow. If the business requirements indicate a preference for managed services, then the Cloud Dataflow option is a better choice.

You will be tested on how to plan the execution of work required to implement a cloud solution. Migrations to the cloud are often done in stages. Consider the advantages of starting with low-risk migration tasks, such as setting up a test environment in the cloud before moving production workloads to GCP.

The business and technical requirements may leave you open to proposing two or more different solutions. In these cases, consider the feasibility of the implementation. Will it be scalable and reliable? Even if GCP services have high SLOs, your system may depend on a third-party service that may go down. If that happens, what is the impact on your workflow? Should you plan to buffer work in a Cloud Pub/Sub queue rather than sending it directly to the third-party service? Also consider costs and optimizations, but only after you have a technically viable solution that meets business requirements. As computer science pioneer Donald Knuth realized, “The real problem is that programmers have spent far too much time worrying about efficiency in the wrong places and at the wrong times; premature optimization is the root of all evil (or at least most of it) in programming.”¹ The same can be said for architecture as well—meet business and technical requirements before trying to optimize.

The exam guide states that architects should be familiar with the software development lifecycle and agile practices. These will be important to know when answering questions about developing and releasing code, especially how to release code into production environments without shutting down the service. It is important to understand

topics such as Blue/Green deployments, canary deployments, and continuous integration/continuous delivery.

In this context, managing is largely about security and monitoring. Architects will need to understand authentication and authorization in GCP. The IAM service is used across GCP, and it should be well understood before attempting the exam. Cloud Monitoring and Cloud Logging are the key services for monitoring and logging in GCP.

How Is the Professional Cloud Architect Exam Different from the Associate Cloud Engineer Exam?

There is some overlap between the Professional Cloud Architect and Associate Cloud Engineer exams. Both exams test for an understanding of technical requirements and the ability to build, deploy, and manage cloud resources. In addition, the Professional Cloud Architect exam tests the ability to work with business requirements to design, plan, and optimize cloud solutions.

The questions on the Professional Cloud Architect exam are based on the kinds of work cloud architects do on a day-to-day basis. This includes deciding which of several storage options is best, designing a network to meet industry regulations, or understanding the implications of horizontally scaling a database.

The questions on the Associate Cloud Engineer exam are based on the tasks that cloud engineers perform, such as creating instance groups, assigning roles to identities, or monitoring a set of VMs. The engineering exam is more likely to have detailed questions about `gcloud`, `gsutil`, and `bq` commands. Architects need to be familiar with these commands and their function, but a detailed knowledge of command options and syntax is not frequently needed on the Professional Cloud Architect exam.

This book is designed to help you pass the Professional Cloud Architect certification exam. If you'd like additional preparation, review the *Official Google Cloud*

Certified Associate Cloud Engineer Study Guide (Sybex, 2019).

What Does This Book Cover?

This book covers the topics outlined in the Professional Cloud Architect exam guide available here:

cloud.google.com/certification/guides/professional-cloud-architect

Chapter 1: Introduction to the Google

Professional Cloud Architect Exam This chapter outlines the exam objectives, scope of the exam, and case studies used in the exam. One of the most challenging parts of the exam for many architects is mapping business requirements to technical requirements. This chapter discusses strategies for culling technical requirements and constraints from statements about nontechnical business requirements. The chapter also discusses the need to understand functional requirements around computing, storage, and networking as well as nonfunctional characteristics of services, such as availability and scalability.

Chapter 2: Designing Solutions to Meet Business

Requirements This chapter reviews several key areas where business requirements are important to understand, including business use cases and product strategies, application design and cost considerations, systems integration and data management, compliance and regulations, security, and success measures.

Chapter 3: Designing Solutions to Meet Technical

Requirements This chapter discusses ways to ensure high availability in compute, storage, and applications. It also reviews ways to ensure scalability in compute, storage, and network resources. The chapter also introduces reliability engineering.

Chapter 4: Designing Compute Systems This chapter discusses Compute Engine, App Engine, Kubernetes Engine, Anthos, and Cloud Functions. Topics in this chapter include use cases, configuration, management, and design. Other topics include managing state in distributed systems, data flows and pipelines, and data integrity. Monitoring and alerting are also discussed.

Chapter 5: Designing Storage Systems This chapter focuses on storage and database systems. Storage systems include object storage, network-attached storage, and caching. Several databases are reviewed, including Cloud SQL, Cloud Spanner, BigQuery, Cloud Firestore, and Bigtable. It is important to know how to choose among storage and database options when making architectural choices. Other topics include provisioning, data retention and lifecycle management, and network latency.

Chapter 6: Designing Networks This chapter reviews VPCs, including subnets and IP addressing, hybrid cloud networking, VPNs, peering, Shared VPCs, and direct connections. This chapter also includes a discussion of regional and global load balancing. Hybrid cloud computing and networking topics are important concepts for the exam.

Chapter 7: Designing for Security and Legal Compliance This chapter discusses IAM, data security including encryption at rest and encryption in transit, key management, security evaluation, penetration testing, auditing, and security design principles. Major regulations and ITIL are reviewed.

Chapter 8: Designing for Reliability This chapter begins with a discussion of Cloud Operations (formerly Stackdriver) for monitoring, logging, and alerting.

Next, the chapter reviews continuous deployment and continuous integration. Systems reliability engineering is discussed, including overloads, cascading failures, and testing for reliability. Incident management and post-mortem analysis are also described.

Chapter 9: Analyzing and Defining Technical Processes This chapter focuses on software development lifecycle planning. This includes troubleshooting, testing and validation, business continuity, and disaster recovery.

Chapter 10: Analyzing and Defining Business Processes This chapter includes several business-oriented skills including stakeholder management, change management, team skill management, customer success management, and cost management.

Chapter 11: Development and Operations This chapter reviews application development methodologies, API best practices, and testing frameworks, including load, unit, and integration testing. The chapter also discusses data and systems migration tooling. The chapter concludes with a brief review of using Cloud SDK and programmatically working with GCP.

Chapter 12: Migration Planning This chapter describes how to plan for a cloud migration. Steps include integrating with existing systems, migrating systems and data, license mapping, network management and planning, as well as testing and developing proof-of-concept systems.

Like all exams, the Professional Cloud Architect certification from Google is updated periodically and may eventually be retired or replaced. At some point after Google no longer offers this exam, the old editions of our

books and online tools will be retired. If you have purchased this book after the exam was retired, or are attempting to register in the Sybex online learning environment after the exam was retired, please know that we make no guarantees that this exam's online Sybex tools will be available once the exam is no longer available.

Interactive Online Learning Environment and Test Bank

Studying the material in the *Google Cloud Certified Professional Cloud Architect Study Guide* is an important part of preparing for the Professional Cloud Architect certification exam, but we also provide additional tools to help you prepare. The online Test Bank will help you understand the types of questions that will appear on the certification exam.

The sample tests in the Test Bank include all the questions in 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.

The flashcards in the Test Bank will push the limits of what you should know for the certification exam. There are more than 100 questions that are provided in digital format. Each flashcard has one question and one correct answer.

The online glossary is a searchable list of key terms introduced in this exam guide that you should know for the Professional Cloud Architect certification exam.

Go to www.wiley.com/go/sybextestprep to register and gain access to this interactive online learning environment and test bank with study tools.

Additional Resources

People learn in different ways. For some, a book is an ideal way to study, while auditory learners may find audio and video resources a more efficient way to study. A combination of resources may be the best option for many of us. In addition to this study guide, here are some other resources that can help you prepare for the Google Cloud Professional Cloud Architect exam.

The Professional Cloud Architect Certification Exam Guide:

cloud.google.com/certification/guides/professional-cloud-architect

Exam FAQs:

cloud.google.com/certification/faqs/#0

Google's Sample Questions:

cloud.google.com/certification/cloud-architect

Google Cloud Platform documentation:

cloud.google.com/docs

Online course Google Cloud Professional Architect: Get Certified by Dan Sullivan

www.udemy.com/course/google-cloud-professional-architect-get-certified



Exam objectives are subject to change at any time without prior notice and at Google's sole discretion. Please visit the Google Professional Cloud Architect website (cloud.google.com/certification/cloud-architect) for the most current listing of exam objectives.

Objective Map

Objective	Chapter
Section 1: Designing and planning a cloud solution architecture	
1.1 Designing a solution infrastructure that meets business requirements	1, 2
1.2 Designing a solution infrastructure that meets technical requirements	2, 3
1.3 Designing network, storage, and compute resources	4
1.4 Creating a migration plan (i.e., documents and architectural diagrams)	12
1.5 Envisioning future solution improvements	2
Section 2: Managing and provisioning solutions infrastructure	
2.1 Configuring network topologies	6
2.2 Configuring individual storage systems	5
2.3 Configuring compute systems	4
Section 3: Designing for security and compliance	
3.1 Designing for security	7
3.2 Designing for compliance	7

Objective	Chapter
Section 4: Analyzing and optimizing technical and business processes	
4.1 Analyzing and defining technical processes	9
4.2 Analyzing and defining business processes	10
4.3 Developing procedures to ensure reliability of solutions in production (e.g., chaos engineering, penetration testing)	8
Section 5: Managing implementation	
5.1 Advising development/operations team(s) to ensure successful deployment of the solution	11
5.2 Interacting with Google Cloud programmatically	11
Section 6: Ensuring solutions and operations reliability	
6.1 Monitoring/logging/profiling/alerting solution	8
6.2 Deployment and release management	8
6.3 Assisting with support of deployed solutions	8
6.4 Evaluating quality control measures	8

Assessment Test

1. Building for Builders LLC manufactures equipment used in residential and commercial building. Each of its 500,000 pieces of equipment in use around the globe has IoT devices collecting data about the state of equipment. The IoT data is streamed from each device every 10 seconds. On average, 10 KB of data is sent in each message. The data will be used for predictive maintenance and product development. The company would like to use a managed service in Google Cloud. What would you recommend?

- A. Apache Cassandra
 - B. Cloud Bigtable
 - C. BigQuery
 - D. Cloud SQL
2. You have developed a web application that is becoming widely used. The front end runs in Google App Engine and scales automatically. The backend runs on Compute Engine in a managed instance group. You have set the maximum number of instances in the backend managed instance group to five. You do not want to increase the maximum size of the managed instance group or change the VM instance type, but there are times the front end sends more data than the backend can keep up with and data is lost. What can you do to prevent the loss of data?
- A. Use an unmanaged instance group.
 - B. Store ingested data in Cloud Storage.
 - C. Have the front end write data to a Cloud Pub/Sub topic, and have the backend read from that topic.
 - D. Store ingested data in BigQuery.
3. You are setting up a cloud project and want to assign members of your team different roles that have appropriate permissions for their responsibilities. What GCP service would you use to do that?
- A. Cloud Identity
 - B. Identity and Access Management (IAM)
 - C. Cloud Authorizations
 - D. LDAP
4. You would like to run a custom stateless container in a managed Google Cloud service. What are your three

options?

- A. App Engine Standard, Cloud Run, and Kubernetes Engine
 - B. App Engine Flexible, Cloud Run, and Kubernetes Engine
 - C. Compute Engine, Cloud Functions, and Kubernetes Engine
 - D. Cloud Functions, Cloud Run, and App Engine Flexible
5. PhotosForYouToday prints photographs and ships them to customers. The front-end application uploads photos to Cloud Storage. Currently, the back end runs a cron job that checks Cloud Storage buckets every 10 minutes for new photos. The product manager would like to process the photos as soon as they are uploaded. What would you use to cause processing to start when a photo file is saved to Cloud Storage?
- A. A Cloud Function
 - B. An App Engine Flexible application
 - C. A Kubernetes pod
 - D. A cron job that checks the bucket more frequently
6. The chief financial officer of your company believes that you are spending too much money to run an on-premises data warehouse and wants to migrate to a managed cloud solution. What GCP service would you recommend for implementing a new data warehouse in GCP?
- A. Compute Engine
 - B. BigQuery
 - C. Cloud Dataproc