


ALIBABA CLOUD CERTIFIED ASSOCIATE (ACA) PRACTICE TEST (SAMPLE)

STUDY GUIDE



**SAMPLE STUDY GUIDE
WITH LIMITED QUESTIONS**

**VISIT US ONLINE FOR
THE FULL VERSION STUDY GUIDE**

<https://aca.examzify.com>



Copyright © 2026 by Examzify - A Kaluba Technologies Inc. product.

ALL RIGHTS RESERVED.

No part of this book may be reproduced or transferred in any form or by any means, graphic, electronic, or mechanical, including photocopying, recording, web distribution, taping, or by any information storage retrieval system, without the written permission of the author.

Notice: Examzify makes every reasonable effort to obtain accurate, complete, and timely information about this product from reliable sources.

SAMPLE

Table of Contents

Copyright	1
Table of Contents	2
Introduction	3
How to Use This Guide	4
Questions	5
Answers	8
Explanations	10
Next Steps	16

SAMPLE

Introduction

Preparing for a certification exam can feel overwhelming, but with the right tools, it becomes an opportunity to build confidence, sharpen your skills, and move one step closer to your goals. At Examzify, we believe that effective exam preparation isn't just about memorization, it's about understanding the material, identifying knowledge gaps, and building the test-taking strategies that lead to success.

This guide was designed to help you do exactly that.

Whether you're preparing for a licensing exam, professional certification, or entry-level qualification, this book offers structured practice to reinforce key concepts. You'll find a wide range of multiple-choice questions, each followed by clear explanations to help you understand not just the right answer, but why it's correct.

The content in this guide is based on real-world exam objectives and aligned with the types of questions and topics commonly found on official tests. It's ideal for learners who want to:

- Practice answering questions under realistic conditions,
- Improve accuracy and speed,
- Review explanations to strengthen weak areas, and
- Approach the exam with greater confidence.

We recommend using this book not as a stand-alone study tool, but alongside other resources like flashcards, textbooks, or hands-on training. For best results, we recommend working through each question, reflecting on the explanation provided, and revisiting the topics that challenge you most.

Remember: successful test preparation isn't about getting every question right the first time, it's about learning from your mistakes and improving over time. Stay focused, trust the process, and know that every page you turn brings you closer to success.

Let's begin.

How to Use This Guide

This guide is designed to help you study more effectively and approach your exam with confidence. Whether you're reviewing for the first time or doing a final refresh, here's how to get the most out of your Examzify study guide:

1. Start with a Diagnostic Review

Skim through the questions to get a sense of what you know and what you need to focus on. Your goal is to identify knowledge gaps early.

2. Study in Short, Focused Sessions

Break your study time into manageable blocks (e.g. 30 – 45 minutes). Review a handful of questions, reflect on the explanations.

3. Learn from the Explanations

After answering a question, always read the explanation, even if you got it right. It reinforces key points, corrects misunderstandings, and teaches subtle distinctions between similar answers.

4. Track Your Progress

Use bookmarks or notes (if reading digitally) to mark difficult questions. Revisit these regularly and track improvements over time.

5. Simulate the Real Exam

Once you're comfortable, try taking a full set of questions without pausing. Set a timer and simulate test-day conditions to build confidence and time management skills.

6. Repeat and Review

Don't just study once, repetition builds retention. Re-attempt questions after a few days and revisit explanations to reinforce learning. Pair this guide with other Examzify tools like flashcards, and digital practice tests to strengthen your preparation across formats.

There's no single right way to study, but consistent, thoughtful effort always wins. Use this guide flexibly, adapt the tips above to fit your pace and learning style. You've got this!

Questions

SAMPLE

- 1. Which of the following is a consequence of performing a full table scan on a large database?**
 - A. Increased query speed
 - B. Significant resource consumption
 - C. Reduction of storage costs
 - D. Improved data integrity
- 2. What is one of the core functions of the Alibaba Cloud CDN?**
 - A. Direct data compression
 - B. Content caching at edge nodes
 - C. SEO performance enhancement
 - D. Cost tracking of data usage
- 3. What kind of analytics does Alibaba Cloud's Log Service provide?**
 - A. Real-time log analysis and data visualization
 - B. Historical performance reviews
 - C. User behavior tracking
 - D. Financial forecasting
- 4. In a microservice architecture, how many spans would a typical call chain entail if it includes a load balancer, authentication service, billing service, resource request, and returning a result?**
 - A. 3
 - B. 4
 - C. 5
 - D. 6
- 5. Which statement is correct regarding Data Transmission Service (DTS)?**
 - A. DTS supports data migration between self-built Oracle databases.
 - B. DTS can only be used for data backup tasks.
 - C. DTS requires manual intervention for all migration processes.
 - D. DTS offers real-time data synchronization only.

- 6. What type of billing does Alibaba Cloud Log Service follow?**
- A. Monthly subscriptions
 - B. Pay-As-You-Go
 - C. Flat rate billing
 - D. Credit-based system
- 7. Which of the following options is part of the relation between spans in the OpenTracing model?**
- A. Connection
 - B. References
 - C. Connections
 - D. Branches
- 8. When deploying applications, why might ECI be preferred over ECS?**
- A. Higher resource commitment is needed
 - B. Lower cost of operations and quick startup
 - C. Simplistic management of physical servers
 - D. Complex networking requirements
- 9. In the context of ApsaraDB for PolarDB, what is the best practice for creating a database index when querying based on multiple fields?**
- A. Create a single index for all fields combined.
 - B. Create a separate index for each field.
 - C. Create a gin index for multiple fields.
 - D. Create an index only on the most frequently queried field.
- 10. Which statement about the Message Service (MS) is correct?**
- A. The event notification message is permanently stored.
 - B. Event messages can be retained according to policy.
 - C. The consumer type is limited to Queue only.
 - D. Event notifications require manual sending.

Answers

SAMPLE

1. B
2. B
3. A
4. C
5. A
6. B
7. B
8. B
9. C
10. B

SAMPLE

Explanations

SAMPLE

1. Which of the following is a consequence of performing a full table scan on a large database?

- A. Increased query speed
- B. Significant resource consumption**
- C. Reduction of storage costs
- D. Improved data integrity

A full table scan involves reading every row in a database table to fulfill a query. This can lead to significant resource consumption because it requires scanning through potentially millions or billions of records, consuming CPU, memory, and I/O bandwidth. The larger the table, the more resources are needed to perform the operation, which can affect overall system performance and response times for other queries. Options indicating increased query speed and improved data integrity do not align with the nature of a full table scan, as this process typically does not enhance the speed of data retrieval; it tends to slow it down due to the volume of data being processed. Additionally, while a full table scan might not directly impact storage costs, it does not lead to a reduction in those costs either, as storage usage remains unchanged regardless of the scanning process. Therefore, the correct understanding focuses on the significant resource consumption that occurs when executing a full table scan on large databases.

2. What is one of the core functions of the Alibaba Cloud CDN?

- A. Direct data compression
- B. Content caching at edge nodes**
- C. SEO performance enhancement
- D. Cost tracking of data usage

One of the core functions of Alibaba Cloud CDN (Content Delivery Network) is content caching at edge nodes. CDNs are designed to improve the delivery of web content to users by storing cached copies of the content at strategically located servers, known as edge nodes, which are closer to the end-users. This process reduces latency, as content can be served more quickly from a node that is geographically closer to the requestor rather than from a central server. By caching content at these nodes, the CDN minimizes the distance that data must travel, which helps in speeding up website load times and improving the overall user experience. Additionally, it reduces the load on the origin server, allowing it to handle more requests efficiently and maintain better performance under high traffic conditions. This caching mechanism is essential for high-demand scenarios, such as when delivering multimedia content or handling sudden spikes in user activity. While direct data compression, SEO performance enhancement, and cost tracking of data usage can be beneficial in various contexts, they do not represent fundamental functions of the CDN itself. Therefore, the emphasis on content caching at edge nodes highlights its primary role in functioning effectively as a content delivery network.

3. What kind of analytics does Alibaba Cloud's Log Service provide?

A. Real-time log analysis and data visualization

B. Historical performance reviews

C. User behavior tracking

D. Financial forecasting

Alibaba Cloud's Log Service is designed to handle vast amounts of log data in real-time, which allows for immediate analysis and insight generation. This functionality is particularly beneficial for organizations seeking to monitor system behavior and application performance continuously. By providing real-time log analysis, it enables users to react swiftly to any anomalies or performance issues that may arise, contributing to improved operational efficiency. Additionally, the data visualization capabilities enhance the analysis process by allowing users to create informative dashboards and charts that represent the log data visually. This empowering aspect of log analytics helps stakeholders easily comprehend data trends and insights. In contrast, the other options, while they may represent important aspects of data analysis, do not specifically align with the core functionalities of Alibaba Cloud's Log Service. Historical performance reviews, user behavior tracking, and financial forecasting typically require specialized tools and methodologies that may not encompass the real-time processing and visualization strengths of the Log Service. Thus, the focus on real-time log analysis and data visualization makes the first choice the most accurate representation of what Alibaba Cloud's Log Service offers.

4. In a microservice architecture, how many spans would a typical call chain entail if it includes a load balancer, authentication service, billing service, resource request, and returning a result?

A. 3

B. 4

C. 5

D. 6

In a microservice architecture, each distinct component or operation involved in processing a request is represented as a "span" in the context of monitoring and tracing. When a typical call chain is analyzed that includes a load balancer, an authentication service, a billing service, a resource request, and then the returning of a result, it is essential to account for the individual spans created by each of these components. The load balancer is the entry point for the request, which would be the first span. Next, once the load balancer forwards the request, the authentication service processes the request, constituting the second span. Following successful authentication, the billing service processes the transaction, adding the third span. The resource request that follows the billing service can be considered as the fourth span, as it involves another distinct operation. Lastly, the act of returning the result wraps up the entire processing and would represent the fifth span in this call chain. Thus, each distinct operation or service in this architecture contributes to the overall span count, leading to a total of five spans in this scenario. This comprehensive counting is crucial for understanding performance bottlenecks, latency, and the overall behavior of the microservices involved.

5. Which statement is correct regarding Data Transmission Service (DTS)?

A. DTS supports data migration between self-built Oracle databases.

B. DTS can only be used for data backup tasks.

C. DTS requires manual intervention for all migration processes.

D. DTS offers real-time data synchronization only.

The correct statement about Data Transmission Service (DTS) is that it supports data migration between self-built Oracle databases. DTS is designed to facilitate seamless data migration, synchronization, and backup across various database systems. This includes the capability to handle different types of databases, including both self-built and cloud-based Oracle databases, making it a versatile tool for organizations needing to manage their data more effectively. DTS offers a range of functionalities that extend beyond mere data backup tasks, enabling tasks like real-time data synchronization and data migration across different databases and environments. This capability makes it invaluable for businesses that require flexibility in their data handling strategies. While the other options suggest limitations or inaccuracies regarding DTS, such as its use being confined to backup tasks or the necessity of manual intervention for all processes, the focus on supporting migrations between self-built Oracle databases showcases the true capacity of DTS to operate efficiently in a multi-database environment. It exemplifies DTS's robust features that cater to diverse data management needs, which is essential for organizations looking to optimize their database operations.

6. What type of billing does Alibaba Cloud Log Service follow?

A. Monthly subscriptions

B. Pay-As-You-Go

C. Flat rate billing

D. Credit-based system

Alibaba Cloud Log Service operates on a Pay-As-You-Go billing model, which is designed to provide flexibility and cost-effectiveness for users. This billing approach allows organizations to pay only for the resources and services they actually use, making it an appealing option for businesses that may want to scale their usage up or down as needed without committing to a fixed monthly fee. Under the Pay-As-You-Go model, users are charged based on specific metrics such as the amount of data ingested, stored, and queried within the Log Service. This allows for more granular control over costs, as users can align their spending with their actual logging needs. In contrast, other billing models like monthly subscriptions or flat rate billing require upfront commitments and may not be as flexible in adapting to varying usage patterns. The credit-based system, while applicable in other contexts, does not specifically reflect the operational mechanics of the Log Service's billing approach.

7. Which of the following options is part of the relation between spans in the OpenTracing model?

- A. Connection
- B. References**
- C. Connections
- D. Branches

In the OpenTracing model, spans represent individual operations within a trace, and the relationships between these spans are crucial for understanding the flow of operations. The term "References" is specifically used to denote how spans relate to one another. References allow spans to link to one another, indicating how they are connected in relation to the overall operation being traced. For instance, a span may reference another span to show that it is a child span, meaning it is a sub-operation of the parent span. This connection helps in creating a full picture of the request lifecycle across distributed systems, showing dependencies and the order in which operations occur. Understanding the concept of References in the OpenTracing model is essential for accurately tracing requests through various services, making it a key aspect of distributed tracing. Other terms, like "Connection," "Connections," and "Branches," while they may seem related, do not specifically define the nature of the relationships among spans in the context of OpenTracing, hence making References the correct choice in this scenario.

8. When deploying applications, why might ECI be preferred over ECS?

- A. Higher resource commitment is needed
- B. Lower cost of operations and quick startup**
- C. Simplistic management of physical servers
- D. Complex networking requirements

Choosing ECI (Elastic Container Instance) over ECS (Elastic Compute Service) is often preferred when factors like cost and quick deployment speed are crucial. ECI is designed for running containers without the overhead of managing the underlying virtual machines. This allows for a more efficient use of resources, as you pay only for the resources you consume during the time your applications are running. The quick startup time is particularly beneficial for applications with unpredictable workloads or those requiring rapid scaling, as ECI enables faster deployment of containerized applications compared to the more traditional approach with ECS. This can lead to significant operational cost savings, especially in scenarios where applications have varying usage patterns and need to be spun up or down frequently. Additionally, ECI abstracts the complexity of underlying infrastructure management, allowing developers to focus on application development and deployment rather than worrying about configurations and maintenance associated with virtual machines. This makes ECI an attractive choice for scenarios where agility and cost efficiency are priorities.

9. In the context of ApsaraDB for PolarDB, what is the best practice for creating a database index when querying based on multiple fields?

- A. Create a single index for all fields combined.
- B. Create a separate index for each field.
- C. Create a gin index for multiple fields.**
- D. Create an index only on the most frequently queried field.

Creating a gin index for multiple fields is considered a best practice in the context of ApsaraDB for PolarDB when querying based on multiple fields. This type of index is designed to optimize searches across multiple columns, particularly when dealing with queries that involve more complex data types or full-text search scenarios. Gin (Generalized Inverted Index) indexes help improve the performance of queries that involve array values or nested data structures, making them particularly suitable for handling scenarios where multiple fields must be evaluated together. By using a gin index, the database can efficiently search through multiple attributes, reducing the need for scanning large volumes of data and consequently speeding up query response times. This can lead to significantly enhanced performance in applications that require rapid access to data based on multiple criteria. Moreover, using a combined index approach enables the database engine to better utilize resources, leading to optimal performance for read operations. The other options have their use cases but may not provide the same level of performance or efficiency as a gin index for queries involving multiple fields. For instance, creating a single index for all fields combined may not leverage the specific strengths of the data, while separate indexes for each field could lead to inefficient query processing, as the database may have to evaluate multiple indexes rather than leveraging a

10. Which statement about the Message Service (MS) is correct?

- A. The event notification message is permanently stored.
- B. Event messages can be retained according to policy.**
- C. The consumer type is limited to Queue only.
- D. Event notifications require manual sending.

The statement regarding the retention of event messages according to policy is spot on. In Alibaba Cloud's Message Service (MS), event messages can be configured to follow specific retention policies. This means that you can set rules determining how long messages should be stored before they are automatically deleted. This flexibility allows organizations to manage their data effectively, ensuring they keep relevant information accessible for a predetermined period while also complying with data retention regulations or practices. Understanding the implications of message retention is crucial for architects who design communication systems using Alibaba Cloud services. It enables effective data management strategies and ensures that resources are allocated efficiently. In contrast, other statements do not accurately represent the features of the Message Service. Event notification messages are not stored permanently, allowing for a more dynamic approach to message handling. Also, the consumer type is not limited to Queue only, as various mechanisms support different consumer formats. Lastly, event notifications do not require manual sending; they can be automated based on events, aiding in real-time communication.

Next Steps

Congratulations on reaching the final section of this guide. You've taken a meaningful step toward passing your certification exam and advancing your career.

As you continue preparing, remember that consistent practice, review, and self-reflection are key to success. Make time to revisit difficult topics, simulate exam conditions, and track your progress along the way.

If you need help, have suggestions, or want to share feedback, we'd love to hear from you. Reach out to our team at hello@examzify.com.

Or visit your dedicated course page for more study tools and resources:

<https://aca.examzify.com>

We wish you the very best on your exam journey. You've got this!

SAMPLE