

ORACLE AUTONOMOUS DATABASE CLOUD SPECIALIST PRACTICE TEST (SAMPLE) STUDY GUIDE



SAMPLE STUDY GUIDE WITH LIMITED QUESTIONS

VISIT US ONLINE FOR
THE FULL VERSION STUDY GUIDE

<https://oracleautonomoudbcloud.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. What value do NUMBER data type columns return when fully redacted?**
 - A. Null
 - B. Zero (0)
 - C. Empty string
 - D. None

- 2. In the default configuration of the Autonomous Data Warehouse, which statement is true?**
 - A. Direct access to the database node and local file system is allowed
 - B. Degree of parallelism is set based on the number of OCPUs in the system
 - C. Adding or modifying tablespaces is allowed
 - D. Compression is enabled by default

- 3. Which of the following best describes the capabilities of Oracle's Data Safe?**
 - A. Only data migration
 - B. Data encryption and monitoring
 - C. Monitoring resource usage
 - D. Data warehousing solutions

- 4. On what infrastructure does the Autonomous Database run?**
 - A. VM on Oracle Cloud Infrastructure
 - B. Exadata on Oracle Cloud Infrastructure
 - C. Any Oracle Engineered system
 - D. Bare-metal on Oracle Cloud Infrastructure

- 5. Which sources can Data Sync connect to?**
 - A. Only Oracle databases
 - B. CSV, XLSX, tables, views, and SQL statements
 - C. Just CSV files
 - D. Only SQL statements

6. What is the maximum capacity of an Autonomous Database in terms of CPU for the high-performance option?
- A. 8 CPUs
 - B. 16 CPUs
 - C. 32 CPUs
 - D. 64 CPUs
7. What happens to the database password by default in the Autonomous Database?
- A. It expires after 30 days
 - B. It never expires
 - C. It expires after 90 days
 - D. It does not expire until 360 days
8. In which scenarios is DBMS_CLOUD.COPY_DATA typically used?
- A. To back up databases
 - B. To copy files between instances
 - C. To load data into existing tables
 - D. To list available data files in cloud storage
9. What feature facilitates serverless deployment in Oracle Autonomous Database?
- A. High availability configuration
 - B. Pay-as-you-go pricing model
 - C. Dedicated resource allocation
 - D. Pre-paid transaction costs
10. How are optimizer statistics managed in ATP?
- A. Gathered manually by users
 - B. Automatically done for all tables during maintenance
 - C. Executed upon direct-path load operations only
 - D. Updated based on user requests

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. What value do NUMBER data type columns return when fully redacted?

- A. Null
- B. Zero (0)**
- C. Empty string
- D. None

When NUMBER data type columns are fully redacted, they return a value of zero (0). This behavior is defined by Oracle's data redaction policy, which ensures that sensitive data, while still accessible in terms of its datatype, does not reveal meaningful information. The rationale behind returning zero instead of null, an empty string, or no value at all is to provide a consistent output that keeps the integrity of the data type. For instance, if the data type is expected to be a number, returning zero keeps the structure intact while masking the actual value. This approach allows applications that rely on numeric data types to continue functioning without encountering unexpected nulls or data type mismatches that could arise from pushing an empty string or no value. This mechanism not only protects sensitive information by providing a meaningless value but also avoids creating scenarios where applications might misinterpret absent data as a failure or error in data retrieval. By returning zero in the context of redacted NUMBER data types, Oracle ensures that systems that depend on this data can handle it without issues, thereby facilitating both security and usability within applications.

2. In the default configuration of the Autonomous Data Warehouse, which statement is true?

- A. Direct access to the database node and local file system is allowed
- B. Degree of parallelism is set based on the number of OCPUs in the system**
- C. Adding or modifying tablespaces is allowed
- D. Compression is enabled by default

In the default configuration of the Autonomous Data Warehouse, the statement regarding the degree of parallelism being set based on the number of OCPUs in the system is accurate. This configuration reflects the optimized performance capabilities of the Autonomous Database, allowing it to handle workloads efficiently by leveraging the available OCPU resources. The degree of parallelism is crucial for managing how many processes can work on the same query concurrently, which ultimately enhances the speed and efficiency of data processing and query execution. By aligning this setting with the number of OCPUs, the Autonomous Data Warehouse ensures that resources are utilized effectively, enabling the database to scale with the workload and maintain optimal performance levels. In contrast, direct access to the database node and local file system, adding or modifying tablespaces, and ensuring that compression is enabled by default do not reflect the operational limitations and features of the Autonomous Data Warehouse's standard setup. Direct access is restricted to ensure security and stability, modifying database structures is not permitted for users due to the system's autonomous nature, and while compression is commonly used for efficient storage management, it is a feature that can vary based on specific configurations or use cases.

3. Which of the following best describes the capabilities of Oracle's Data Safe?

- A. Only data migration
- B. Data encryption and monitoring**
- C. Monitoring resource usage
- D. Data warehousing solutions

Oracle's Data Safe is designed to enhance the security and management of data in Oracle databases. One of its primary capabilities is to provide comprehensive data encryption and monitoring features. Data Safe helps organizations secure sensitive data by offering encryption for data at rest and in transit, ensuring that information is protected from unauthorized access. Additionally, it includes robust monitoring capabilities that track user activities, database configurations, and potential vulnerabilities. This dual focus on data encryption and continuous monitoring strengthens overall database security and compliance with regulatory requirements. Other options like data migration, monitoring resource usage, and data warehousing solutions do not accurately capture the core purpose of Data Safe, which is centered around security and monitoring rather than the functions described in those choices.

4. On what infrastructure does the Autonomous Database run?

- A. VM on Oracle Cloud Infrastructure
- B. Exadata on Oracle Cloud Infrastructure**
- C. Any Oracle Engineered system
- D. Bare-metal on Oracle Cloud Infrastructure

The Autonomous Database runs specifically on Oracle Exadata hardware within the Oracle Cloud Infrastructure. This infrastructure is designed for optimal database performance and includes advanced features such as storage optimization and improved I/O throughput, which significantly enhance the efficiency of database operations. Exadata is a specialized engineered system that combines hardware and software to deliver high-performance database services. By utilizing Exadata, Oracle Autonomous Database can leverage its capabilities, such as smart storage and intelligent caching, to automatically manage and optimize performance, scalability, and reliability for cloud database workloads. This environment enables Autonomous Database to offer its self-scaling, self-securing, and self-repairing capabilities, making it an ideal solution for organizations looking to simplify database management while ensuring high availability and performance.

5. Which sources can Data Sync connect to?

- A. Only Oracle databases
- B. CSV, XLSX, tables, views, and SQL statements**
- C. Just CSV files
- D. Only SQL statements

The correct answer is that Data Sync can connect to CSV, XLSX, tables, views, and SQL statements. This reflects the versatility of Data Sync in handling various data formats and sources, allowing for efficient data synchronization and integration processes. Data Sync is designed to facilitate the movement of data from a variety of sources into the Oracle Autonomous Database. By supporting multiple file formats such as CSV and XLSX, as well as database-specific formats like tables and views, Data Sync enhances its usability across different data environments. The capability to process SQL statements further allows users to directly interact with databases, enabling complex queries to be executed as part of the synchronization process. This comprehensive support for multiple data types ensures that organizations can easily utilize their diverse data assets and streamline their data management workflows without being limited to a narrow range of sources.

6. What is the maximum capacity of an Autonomous Database in terms of CPU for the high-performance option?

- A. 8 CPUs
- B. 16 CPUs
- C. 32 CPUs**
- D. 64 CPUs

The maximum capacity of an Autonomous Database for the high-performance option is indeed 32 CPUs. This specification is important because the high-performance option is designed to handle workloads that demand more computational power, making it suitable for intensive database operations, analytics, and large-scale applications. The 32 CPU limit allows users to leverage significant resources for enhanced performance, especially when dealing with large datasets or complex queries. This capability supports high concurrency and enables efficient data processing, which is essential for enterprise-level applications. The other options propose lower CPU counts, which would not provide the same level of performance and might not support more demanding workload requirements effectively. Therefore, the correct maximum capacity reflects the intended performance benchmarks of the high-performance Autonomous Database offering.

7. What happens to the database password by default in the Autonomous Database?

- A. It expires after 30 days
- B. It never expires
- C. It expires after 90 days
- D. It does not expire until 360 days**

In Oracle Autonomous Database, by default, the database password is set to expire after a period of 360 days. This long expiration period helps ensure that users do not have to frequently change their passwords, which can be a burden, especially in environments that require access to multiple databases or applications. Such a policy can enhance usability while still maintaining a level of security over an extended timeframe. While many databases enforce shorter password expiration timelines to promote security best practices, the Autonomous Database is designed to balance usability with security needs. It allows for password management strategies that accommodate various operational requirements, thereby enabling users to access their databases with less frequent disruption. This feature can also simplify user management by reducing the number of password-related issues that arise from frequent expirations.

8. In which scenarios is DBMS_CLOUD.COPY_DATA typically used?

- A. To back up databases
- B. To copy files between instances
- C. To load data into existing tables**
- D. To list available data files in cloud storage

DBMS_CLOUD.COPY_DATA is specifically designed for loading data from external sources, typically located in cloud storage, directly into existing tables within an Oracle database. This functionality is particularly useful when dealing with large datasets that need to be ingested efficiently without manual intervention or extensive coding. In scenarios where data needs to be imported from cloud object storage or another cloud service into an Oracle Autonomous Database or Oracle Database, DBMS_CLOUD.COPY_DATA automates the process of reading the data files and inserting the records into specified tables. It supports various file formats and provides options for configuring the import process, such as specifying delimiters and handling null values. Other options do not align with the primary use case of DBMS_CLOUD.COPY_DATA. Backing up databases, copying files between instances, and listing available data files in cloud storage are functionalities better served by other tools and commands within Oracle's suite of database services. For example, backup operations typically involve using RMAN or Data Pump, and file copy operations may utilize other file transfer or data movement services.

9. What feature facilitates serverless deployment in Oracle Autonomous Database?

- A. High availability configuration
- B. Pay-as-you-go pricing model**
- C. Dedicated resource allocation
- D. Pre-paid transaction costs

The pay-as-you-go pricing model is a key feature that facilitates serverless deployment in Oracle Autonomous Database. This model allows users to only pay for the resources they consume rather than committing to a fixed capacity or resource baseline. As workloads fluctuate, users can scale their resources dynamically and align costs directly with usage, making it particularly suitable for applications with variable workloads or unpredictable traffic. In a serverless environment, the database resource management fully abstracts the underlying infrastructure, enabling automatic scaling to meet demand without requiring user intervention. This model effectively supports the flexibility and cost-efficiency that are hallmarks of serverless computing, allowing organizations to focus on application development rather than resource management. Other features related to database offerings, such as high availability configuration or dedicated resource allocation, do not directly pertain to the serverless nature of deployment. Similarly, pre-paid transaction costs relate to upfront payment options rather than the on-demand scalability that a pay-as-you-go model provides, which is central to the serverless paradigm.

10. How are optimizer statistics managed in ATP?

- A. Gathered manually by users
- B. Automatically done for all tables during maintenance**
- C. Executed upon direct-path load operations only
- D. Updated based on user requests

In Autonomous Transaction Processing (ATP), optimizer statistics are automatically managed as part of the database maintenance activities. This means that the system regularly collects and updates statistics for all tables without requiring manual intervention from users. This automatic process ensures that the database optimizer has accurate and up-to-date information to make informed decisions about query execution plans, ultimately leading to improved performance. The automatic management of optimizer statistics is designed to optimize system resources and enhance efficiency, as it eliminates the need for database administrators to manually gather statistics. This feature aligns with the autonomous nature of the database, allowing it to adapt and adjust dynamically. This approach also helps maintain the database's overall health and performance, as it ensures statistics are continuously refreshed, accounting for changes in data distribution and volume. In contrast, other options suggest methods that either necessitate user involvement or apply only in specific circumstances, which are not reflective of the inherent capabilities of ATP's automatic maintenance.

SAMPLE

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://oracleautonomousdbcloud.examzify.com>

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

SAMPLE