

Database Systems Practice Test (Sample)

Study Guide



Everything you need from our exam experts!

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 15

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 function is commonly used to extract a portion of a string by specifying a starting position and length?**
 - A. Upper**
 - B. Substring**
 - C. Left**
 - D. Right**

- 2. Which statement defines an Index?**
 - A. A mechanism for increasing the speed of data search and data retrieval on relations with a large number of records.**
 - B. The method to ensure the uniqueness of records.**
 - C. A method to enforce referential integrity across all transactions.**
 - D. A tool for compressing data to save storage.**

- 3. What term describes dispersed master data connected via a central master data registry?**
 - A. Centralized Approach**
 - B. Registry**
 - C. Master Data Management**
 - D. NOSQL**

- 4. Which term describes data following the specified structure or format?**
 - A. Completeness**
 - B. Timeliness**
 - C. Consistency**
 - D. Conformity**

- 5. Which function truncates date/time to a specified interval?**
 - A. Date_Trunc**
 - B. Joins**
 - C. Cross Joins**
 - D. Full joins**

- 6. Which operation is NOT performed by the ALTER TABLE statement?**
- A. Add a column**
 - B. Rename a column**
 - C. Alter a column's data type**
 - D. Create a new index**
- 7. What does EXCEPT do?**
- A. It returns rows that appear in both queries.**
 - B. It returns rows that appear only in the first query.**
 - C. It returns rows that are in the second query but not the first.**
 - D. It returns the cross product of rows from both queries.**
- 8. CAST changes the data type of a column in the query**
- A. CAST**
 - B. Trim**
 - C. IN**
 - D. Derived Column**
- 9. Which operator is the case-insensitive version of LIKE?**
- A. LIKE**
 - B. BETWEEN**
 - C. IN**
 - D. ILIKE**
- 10. Which of the following is an example of a Data Control Language statement?**
- A. SELECT**
 - B. INSERT INTO**
 - C. ALTER**
 - D. GRANT**

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. Which function is commonly used to extract a portion of a string by specifying a starting position and length?

- A. Upper
- B. Substring**
- C. Left
- D. Right

Extracting a portion of a string by giving a starting position and a length is done with a substring operation. The function named Substring is designed to return exactly that slice: it takes the original string, a starting index, and a length, and outputs the characters in that range. For example, many languages or SQL dialects use something like `substring('abcdef', 2, 3)` to yield 'bcd' (depending on indexing rules). This is what makes it the best fit here. The other options don't fit the same pattern. Upper changes every character to uppercase, not a slice of the string. Left takes characters from the left end of the string only, while Right takes from the right end, both without specifying an arbitrary middle start. If you need a portion from anywhere inside the string with a chosen length, Substring is the appropriate choice.

2. Which statement defines an Index?

- A. A mechanism for increasing the speed of data search and data retrieval on relations with a large number of records.**
- B. The method to ensure the uniqueness of records.
- C. A method to enforce referential integrity across all transactions.
- D. A tool for compressing data to save storage.

An index is a data structure designed to speed up searching and data retrieval on large tables. It works like a lookup table: it stores the values of one or more columns in a sorted or otherwise organized form and keeps pointers to the actual rows. With the index, the database can quickly narrow down where a matching row is located instead of scanning every record, which dramatically reduces query time for large datasets. Common implementations include B-tree and hash-based indexes, and there are variants like clustered and non-clustered indexes that affect how data is stored and accessed. Keep in mind that while indexes speed reads, they add some overhead to inserts, updates, and deletes because the index itself must be maintained. The idea that an index primarily serves to speed up data search and retrieval on large relations captures the core purpose, while the other statements describe separate features like enforcing uniqueness, enforcing referential integrity, or data compression.

3. What term describes dispersed master data connected via a central master data registry?

- A. Centralized Approach**
- B. Registry**
- C. Master Data Management**
- D. NOSQL**

In master data management, when master data is spread across multiple systems but tied together through a single reference point, you're looking at the registry approach. Here, a central master data registry contains metadata, identifiers, and pointers that indicate where the actual master data resides in each source system. It acts as an index and governance layer that coordinates data across environments without moving all the data into one place. This is different from a centralized approach, which would physically consolidate all master data in a single repository, and it isn't about NOSQL, which is just a type of database technology. Master data management encompasses various models, and the registry model specifically describes dispersed data connected via a central registry.

4. Which term describes data following the specified structure or format?

- A. Completeness**
- B. Timeliness**
- C. Consistency**
- D. Conformity**

Conformity means data adheres to a defined structure or format. When data conforms, it follows the prescribed model—correct field order, data types, formats, and allowed values—so it can be validated automatically and integrated with other data sources. For instance, a date field might require the ISO format YYYY-MM-DD, or a country field might only accept codes from a sanctioned list. If a value breaks these rules, it fails the conformity check even if no fields are missing. Completeness is about having all required fields filled; timeliness is about being up-to-date; consistency is about uniform values across datasets. Conformity is the right term for following a specified structure or format.

5. Which function truncates date/time to a specified interval?

- A. Date_Trunc**
- B. Joins**
- C. Cross Joins**
- D. Full joins**

Truncating date/time values to the start of a chosen interval is what this function does. The `date_trunc` function takes a unit like 'hour', 'day', or 'month' and a timestamp, returning the timestamp rounded down to the start of that interval. For example, `date_trunc('hour', '2023-08-01 12:34:56')` gives `2023-08-01 12:00:00`. This is especially useful when you want to group time-series data into consistent buckets, such as hourly or daily groups, for aggregation. The other options describe different operations that combine rows from tables (joins) and do not modify the date/time value itself, so they don't serve the purpose of truncating timestamps.

6. Which operation is NOT performed by the ALTER TABLE statement?

- A. Add a column**
- B. Rename a column**
- C. Alter a column's data type**
- D. Create a new index**

ALTER TABLE is used to change the structure of an existing table. It can add a new column, rename an existing column, or change a column's data type, as well as adjust constraints, defaults, and nullability. These actions modify the table's schema itself. Creating a new index is handled by a separate operation known as indexing, which builds a data structure to speed up data retrieval. While some systems may support index creation in conjunction with table alterations in certain syntax, the standard, portable understanding is that index creation is not the primary function of ALTER TABLE. Therefore, the operation that isn't performed by ALTER TABLE is creating a new index.

7. What does EXCEPT do?

- A. It returns rows that appear in both queries.**
- B. It returns rows that appear only in the first query.**
- C. It returns rows that are in the second query but not the first.**
- D. It returns the cross product of rows from both queries.**

EXCEPT computes the difference between two query results: it returns the rows that come from the first query but do not appear in the second query. Think of it as subtracting the second result set from the first one. In most SQL systems, this yields distinct rows, removing any duplicates. If you ever need to keep duplicates, some dialects offer EXCEPT ALL. For example, if the first query returns (1, 'apple'), (2, 'banana'), and the second returns (2, 'banana'), the result of EXCEPT is (1, 'apple'). This shows why the correct interpretation is that EXCEPT gives rows that are in the first set but not in the second. This is different from an INTERSECT, which would give rows common to both; the other options would correspond to considering rows in the second but not the first, or to combining every row from both queries (a cross product), neither of which matches EXCEPT.

8. CAST changes the data type of a column in the query

- A. CAST**
- B. Trim**
- C. IN**
- D. Derived Column**

The idea being tested is how to change a value's data type within a query. This is done with CAST, which explicitly converts an expression to a specified datatype using the syntax CAST(expression AS datatype). For example, CAST(price AS INT) turns a numeric or decimal price into an integer, or CAST(name AS VARCHAR(50)) converts a value to a string of up to 50 characters. This explicit conversion is essential when you need to perform operations that require matching types or to control the format of the result. The other options don't perform a type change. Trimming removes whitespace and leaves the type the same (often still a string). IN is a predicate that checks whether a value is in a given list, not a conversion. Derived Column is an ETL/tool-specific concept for creating or transforming columns in a data flow, not a standard SQL datatype conversion within a query.

9. Which operator is the case-insensitive version of LIKE?

- A. LIKE
- B. BETWEEN
- C. IN
- D. ILIKE**

Pattern matching with wildcard characters uses the LIKE operator, but in many databases LIKE is sensitive to case. The way to perform the same kind of matching without regard to case is to use an operator designed for that purpose: it mirrors LIKE's syntax and wildcards while ignoring whether characters are uppercase or lowercase. This means a pattern like 'abc%' will match strings starting with any case variation of "abc," such as "ABC" or "Abc." The other operators shown serve different purposes—BETWEEN checks if a value falls within a range, and IN checks membership in a set—so they aren't used for case-insensitive pattern matching.

10. Which of the following is an example of a Data Control Language statement?

- A. SELECT
- B. INSERT INTO
- C. ALTER
- D. GRANT**

Data Control Language statements manage who can do what with the data. GRANT is the classic example because it assigns permissions (like SELECT, INSERT, UPDATE, or DELETE) to users or roles, controlling access to database objects. The other options belong to different categories: SELECT is used to query data (and is considered DML or DQL in many classifications), INSERT INTO adds data (DML), and ALTER changes the structure of database objects (DDL). So GRANT uniquely represents permission control, which is the essence of Data Control Language.

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://databasesystems.examzify.com>

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

SAMPLE