

DATABASE SYSTEMS PRACTICE TEST (SAMPLE)

STUDY GUIDE



SAMPLE STUDY GUIDE WITH LIMITED QUESTIONS

VISIT US ONLINE FOR
THE FULL VERSION STUDY GUIDE

<https://databasesystems.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	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 term reverts the information to its original state?**
 - A. Recovery
 - B. Encryption
 - C. Authentication
 - D. Decryption Key

- 2. Which term identifies a point in the recovery process when updates are written to disk?**
 - A. recovery log
 - B. checkpoint
 - C. commit
 - D. rollback

- 3. Which statement defines an Index in databases?**
 - A. A mechanism to ensure the uniqueness of records.
 - B. A mechanism for increasing the speed of data search and data retrieval on relations with a large number of records.
 - C. A feature to enforce referential integrity across all transactions.
 - D. A technique to normalize data.

- 4. What does a commit operation do in a database?**
 - A. Marks updates as permanent but not necessarily written to disk
 - B. Inserts a new transaction log entry without persisting changes
 - C. Causes all the updates to be recorded on the disk
 - D. Reverts updates since the last commit

- 5. Which concept best describes the cost associated with a query's execution, used by the optimizer to compare plans?**
 - A. recovery log
 - B. query optimization
 - C. query cost
 - D. operational information

- 6. Which function returns the current date or current time?**
- A. current_date/time
 - B. NOW
 - C. make_date/time
 - D. Date_Trunc
- 7. Which statement best describes the data stored in a data warehouse?**
- A. It contains analytically useful information
 - B. It stores only metadata about data sources
 - C. It is a database containing analytically useful information
 - D. It stores user credentials
- 8. Which function converts a string to uppercase?**
- A. Upper
 - B. Lower
 - C. InitCap
 - D. Position
- 9. Which set operations are used to combine results of two SELECT statements that are union compatible?**
- A. UNION
 - B. INTERSECT
 - C. UNION, INTERSECT, and EXCEPT
 - D. ALL of the above
- 10. Which statement about the LIKE operator is true?**
- A. It requires an exact match
 - B. It compares numbers only
 - C. It sorts results by pattern
 - D. It searches for patterns in text and is case sensitive

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. Which term reverts the information to its original state?

- A. Recovery
- B. Encryption
- C. Authentication
- D. Decryption Key**

Reversing encrypted data back to its original form is decryption. The decryption key is what enables that reversal, letting you recover the plaintext from the ciphertext. In symmetric systems, the same key decrypts the data; in public-key setups, a private decryption key performs the reverse operation. The other terms describe different ideas: recovery is about restoring data after loss, encryption is the process of turning plaintext into ciphertext, and authentication is about verifying identity. Since the focus is on what enables returning the data to its original state, the decryption key is the best fit.

2. Which term identifies a point in the recovery process when updates are written to disk?

- A. recovery log
- B. checkpoint**
- C. commit
- D. rollback

Checkpointing marks a point in the recovery process where the system flushes all dirty pages to disk and records that state. By writing updated pages from memory to disk and logging this event, the database creates a known-good starting point for crash recovery. When the system restarts, recovery uses the most recent checkpoint plus the log to redo actions from that point and to undo any transactions that didn't commit before the crash. The recovery log itself is the record of changes used during recovery, but it is not the moment updates are written to disk. A commit signals that a transaction's results are durable, and rollback undoes uncommitted work; neither inherently defines the disk-write checkpoint. Thus, this term correctly identifies the moment in the recovery process when disk writes occur.

3. Which statement defines an Index in databases?

- A. A mechanism to ensure the uniqueness of records.
- B. A mechanism for increasing the speed of data search and data retrieval on relations with a large number of records.**
- C. A feature to enforce referential integrity across all transactions.
- D. A technique to normalize data.

An index is a data structure that helps the database locate rows quickly by providing fast lookups on one or more columns. It works like a separate map that stores the indexed column values together with pointers to the actual rows, so the system can jump directly to the relevant data instead of scanning every row in a large table. This accelerates data search and retrieval, especially for queries in WHERE clauses, joins, or sorts on big relations. Keep in mind that indexes add some overhead for writes and extra storage, and they don't by themselves enforce uniqueness or referential integrity (those are handled by constraints and keys), nor do they define how data is structured across tables, which is what normalization is about.

4. What does a commit operation do in a database?

- A. Marks updates as permanent but not necessarily written to disk
- B. Inserts a new transaction log entry without persisting changes
- C. Causes all the updates to be recorded on the disk**
- D. Reverts updates since the last commit

A commit finalizes a transaction and makes its changes durable. When you commit, the database records all the operations of that transaction in the persistent log and ensures those changes are written to durable storage. This guarantees that after a crash, the committed updates survive and are visible to other transactions. The idea is that a commit moves changes from being provisional to permanently stored on disk. The other options miss this durability aspect or describe rollback behavior, which is not what a commit does.

5. Which concept best describes the cost associated with a query's execution, used by the optimizer to compare plans?

- A. recovery log
- B. query optimization
- C. query cost**
- D. operational information

In query optimization, the optimizer assigns a numeric estimate of the work required to run a plan, called the query cost. This cost reflects expected resource use—primarily I/O and CPU, and sometimes memory—based on statistics and a cost model. Plans with lower estimated costs are preferred because they are expected to execute more efficiently, so the optimizer can compare options and pick the cheapest one. The recovery log is about durability and crash recovery, not about planning or comparing query plans. The broader process is query optimization itself, not the single metric used to evaluate plans. Operational information isn't the standard term for the plan-cost measure either.

6. Which function returns the current date or current time?

- A. current_date/time**
- B. NOW
- C. make_date/time
- D. Date_Trunc

The idea being tested is how to obtain the current date or the current time in SQL. A shorthand option like current_date/time directly expresses getting the present date or the present time, so it matches the request most closely. It encapsulates the idea of retrieving either component without specifying a concrete value. The other choices behave differently: NOW (or CURRENT_TIMESTAMP) yields the exact current moment as a full timestamp (date and time together); make_date or make_time are constructors that build a date or a time from given parts, not the current moment; Date_Trunc takes an existing timestamp and rounds it to a specified precision, which is a transformation rather than retrieving the current value.

7. Which statement best describes the data stored in a data warehouse?

- A. It contains analytically useful information
- B. It stores only metadata about data sources
- C. It is a database containing analytically useful information**
- D. It stores user credentials

A data warehouse is a database optimized for analytical querying, designed to store integrated, historical data from multiple sources to support decision making. The statement that best describes it is that it is a database containing analytically useful information. This captures both what it is (a database) and its purpose (holding information geared for analysis across time and sources). The other ideas don't fit as well: metadata-only stores aren't the warehouse itself, and storing user credentials is not the warehouse's function. In short, a data warehouse is a database built to hold analytically useful, integrated, historical data for analysis.

8. Which function converts a string to uppercase?

- A. Upper**
- B. Lower
- C. InitCap
- D. Position

The function being tested is the one that turns all letters in a string into uppercase. When you apply this function to a string, every alphabetic character becomes uppercase while other characters stay the same. For example, applying it to 'Database' yields 'DATABASE'. This is exactly what converting to uppercase means. The other options do different things: one converts all letters to lowercase, another capitalizes the first letter of each word (sometimes leaving the rest lowercase), and the last finds the position of a substring within a string rather than changing the case.

9. Which set operations are used to combine results of two SELECT statements that are union compatible?

- A. UNION
- B. INTERSECT
- C. UNION, INTERSECT, and EXCEPT**
- D. ALL of the above

When two SELECT results are union compatible—meaning they return the same number of columns with compatible data types—you can blend them using set operations. The standard ones you can use are UNION, INTERSECT, and EXCEPT. UNION combines all distinct rows from both results, giving a combined list without duplicates. INTERSECT gives you only the rows that appear in both result sets. EXCEPT yields rows that are in the first result but not in the second. These are the standard ways to merge two result sets, which is why they're the correct set of operations for this scenario. Note that while UNION ALL exists in many databases to preserve duplicates, it's not part of the standard trio described here, so the statement focusing on the standard set operations is the best fit.

10. Which statement about the LIKE operator is true?

- A. It requires an exact match
- B. It compares numbers only
- C. It sorts results by pattern
- D. It searches for patterns in text and is case sensitive**

The main idea here is that the LIKE operator is all about pattern matching in text. It lets you describe a pattern and find strings that fit it, using wildcards such as % for any sequence of characters and _ for a single character. Because of that, it enables flexible searches like finding names that start with a letter, contain a substring, or end with a suffix, without requiring an exact match. It isn't about comparing numbers, and it isn't a tool for sorting results by pattern. Whether the matching is case sensitive depends on the database and the text's collation. Some systems treat LIKE as case sensitive, others treat it as case insensitive, unless you use a variant like ILIKE or apply a case-conversion function. Given that, the statement that LIKE searches for patterns in text is true, and in many contexts it is described as case sensitive, which is why that option is considered the best fit among the choices.

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://databasesystems.examzify.com>

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

SAMPLE