

SERVICENOW CERTIFIED TECHNICAL ARCHITECT – WASHINGTON PRACTICE EXAM (SAMPLE)

STUDY GUIDE



SAMPLE STUDY GUIDE WITH LIMITED QUESTIONS

VISIT US ONLINE FOR
THE FULL VERSION STUDY GUIDE

<https://servicenowtecharchitectwa.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. PAD trigger**
 - A. It defines the end condition
 - B. It controls user permissions
 - C. It determines data sources
 - D. It specifies when to start running the process
- 2. Which scheme category is described as exact: dividing information into mutually exclusive sections?**
 - A. Subjective.
 - B. Exact.
 - C. Matrix.
 - D. Hierarchical.
- 3. Which pillar is responsible for ensuring program awareness and providing resources to begin the journey for citizen developers?**
 - A. Enable Pillar
 - B. Equip Pillar
 - C. Control Pillar
 - D. Governance Pillar
- 4. Which is a type of organizational structure that allows users to determine their own path?**
 - A. Hierarchical.
 - B. Open-ended.
 - C. Sequential.
 - D. Matrix.
- 5. Which statement about IDR's scheduling of near-real-time replication is true?**
 - A. It cannot be scheduled.
 - B. It can be scheduled nightly.
 - C. It runs every 10 minutes.
 - D. It runs only on demand.

- 6. How do you publish data from the producer instance in IDR?**
- A. Use only full data dumps.
 - B. Use only API calls to pull data.
 - C. Define a table, data filter, select attributes, and attachments to share.
 - D. Publish data automatically without filters.
- 7. Which statement describes IDR strategies?**
- A. They enable single-target, one-way replication only.
 - B. They operate with no message bus and manual configuration.
 - C. They scale out and split data across instances while providing a single source of truth across prod and sub-prod.
 - D. They require external CRM tools for every data set.
- 8. What is seeding in IDR replication?**
- A. Full live replication of all records.
 - B. Real-time streaming.
 - C. Seeding - less than 3 million records.
 - D. Incremental delta only.
- 9. In the change management policy, which question concerns the types of changes to use for change requests against the instance?**
- A. What change types should we use when creating change requests against the instance?
 - B. Does this instance fall under change management control?
 - C. What is the change management process for conducting a change to the instance?
 - D. Which approvals are required for changes to the instance?
- 10. When optimizing the solution, which steps should you consider?**
- A. Increase data redundancy and hardware capacity
 - B. Add more manual processes to ensure control
 - C. Leverage automation, integrate where possible, eliminate manual processes
 - D. Delay automation until security reviews complete

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. PAD trigger

- A. It defines the end condition
- B. It controls user permissions
- C. It determines data sources
- D. It specifies when to start running the process**

In Process Automation Designer, a trigger defines when the automation should begin. It watches for a specific event or a scheduled time, and when that condition occurs, the process starts executing the defined steps. This is why the trigger is about starting the process: it sets the starting point for the flow. End conditions are handled by the flow's internal logic, not by the trigger; permissions are controlled by access controls, and data sources are inputs used by actions, not the trigger itself. For example, a trigger can start a flow when a new record is created, on a daily schedule, or when a webhook arrives.

2. Which scheme category is described as exact: dividing information into mutually exclusive sections?

- A. Subjective.
- B. Exact.**
- C. Matrix.
- D. Hierarchical.

Exact categorization means you split information into precise, non-overlapping buckets. When a scheme is described as exact, it emphasizes mutual exclusivity: each item belongs to one and only one section, with clear boundaries between sections. The phrase "dividing information into mutually exclusive sections" matches that intent, because mutual exclusivity prevents overlap and keeps categorization unambiguous. Subjective relies on personal judgment or opinions, which can blur boundaries rather than keep them strict. Matrix uses a grid that crosses dimensions, allowing items to relate to multiple intersections, so it isn't about single, non-overlapping partitions. Hierarchical organizes data in levels and trees, focusing on order and containment rather than enforcing exclusive partitions across all sections. Therefore, exact best fits the described concept.

3. Which pillar is responsible for ensuring program awareness and providing resources to begin the journey for citizen developers?

- A. Enable Pillar**
- B. Equip Pillar
- C. Control Pillar
- D. Governance Pillar

Awareness and onboarding are handled by the Enable Pillar. This pillar focuses on opening the journey for citizen developers: making the program known, communicating its purpose and benefits, and providing the initial resources needed to get started—such as onboarding guides, starter templates, learning paths, and access to the community. Once people know about the program and have a clear starting path, they can move into the next phase (Equip) where tools, training, and environments are provided. Governance and Control manage policy, risk, and oversight, not the initial outreach. So, enabling participation by creating visibility and supplying starter resources makes the Enable Pillar the correct choice.

4. Which is a type of organizational structure that allows users to determine their own path?

- A. Hierarchical.
- B. Open-ended.
- C. Sequential.
- D. Matrix.**

In a matrix organization, people report to multiple managers and work across cross-functional teams, which creates a variety of roles and project opportunities. This setup gives individuals more control over their development by selecting assignments that fit their interests and goals, allowing them to shape their own path within the organization's framework. Hierarchical structures tend to lock people into a single chain of command with a more fixed progression, while sequential models imply a rigid, step-by-step path. Open-ended isn't a standard formal structure in most management frameworks.

5. Which statement about IDR's scheduling of near-real-time replication is true?

- A. It cannot be scheduled.**
- B. It can be scheduled nightly.
- C. It runs every 10 minutes.
- D. It runs only on demand.

Near-real-time replication with IDR is designed to push changes as soon as they happen, driven by data-change events rather than a timer. Because the process is continuous and event-driven, there isn't a time-based schedule to configure. That's why it cannot be scheduled—the whole point of near-real-time is immediate propagation, not periodic runs. The other options describe fixed or manual timings, which would introduce delays and don't align with the instant nature of near-real-time replication.

6. How do you publish data from the producer instance in IDR?

- A. Use only full data dumps.
- B. Use only API calls to pull data.
- C. Define a table, data filter, select attributes, and attachments to share.**
- D. Publish data automatically without filters.

In IDR, publishing data is a controlled, table-specific process that you configure on the producer side. You explicitly define which table to publish, apply a data filter to limit which records are included, choose which attributes (fields) to share, and decide whether to include attachments. This combination gives you precise control over what is replicated and sent to consumers, which is essential for governing data volume, privacy, and relevance. Full data dumps bypass these safeguards and can expose unnecessary or sensitive information; relying only on API calls shifts the data access model to on-demand pulls rather than a managed publish/subscribe flow; and publishing automatically without filters eliminates governance, likely delivering irrelevant or excess data to consumers. The chosen approach aligns with IDR's intent: selective, configurable data sharing from producer to consumer.

7. Which statement describes IDR strategies?

- A. They enable single-target, one-way replication only.
- B. They operate with no message bus and manual configuration.
- C. They scale out and split data across instances while providing a single source of truth across prod and sub-prod.**
- D. They require external CRM tools for every data set.

IDR strategies are about making data available across multiple ServiceNow instances in a coordinated way, so you can scale out by partitioning data across environments while keeping a single source of truth that spans production and sub-production environments. This means changes in one instance propagate to others in near real time, preserving consistency and enabling independent operation of each environment without losing a unified view of the data. That's why the statement describing IDR best is the one that emphasizes scaling out and splitting data across instances while maintaining a single source of truth across prod and sub-prod. The other descriptions don't fit because IDR relies on automated, message-driven data replication to synchronize across environments, not manual configuration or absence of a message bus; it isn't limited to a single-target, one-way replication; and it doesn't require external CRM tools for data replication.

8. What is seeding in IDR replication?

- A. Full live replication of all records.
- B. Real-time streaming.
- C. Seeding - less than 3 million records.**
- D. Incremental delta only.

Seeding in IDR replication is the initial bootstrap process that loads a subset of data into the target instance to start replication. It uses a smaller, representative dataset rather than the entire production load, typically a seed of under about 3 million records. This lets you validate mappings, transformations, and performance in a controlled way before switching to ongoing, full synchronization. After seeding, the replication usually proceeds with real-time or incremental updates to keep the target in sync. The other options describe different replication modes—full live replication pushes all data, real-time streaming is continuous data flow, and incremental delta only refers to ongoing changes after the seed, not the seed itself.

9. In the change management policy, which question concerns the types of changes to use for change requests against the instance?

- A. What change types should we use when creating change requests against the instance?**
- B. Does this instance fall under change management control?
- C. What is the change management process for conducting a change to the instance?
- D. Which approvals are required for changes to the instance?

The main idea is how change requests are categorized within change management. Change types establish the categories you apply to a request—such as standard, normal, or emergency—and they determine the workflow, risk assessment, and approval path that will be used. This question focuses on what change type to apply when creating change requests against the instance, which is exactly about selecting the appropriate classification for the request. The other options address whether the instance is under change management control (scope), the overall process for carrying out a change (the steps), or the approvals required (who must sign off). All of these are related aspects but do not specify how to classify the change request itself.

10. When optimizing the solution, which steps should you consider?

- A. Increase data redundancy and hardware capacity
- B. Add more manual processes to ensure control
- C. Leverage automation, integrate where possible, eliminate manual processes**
- D. Delay automation until security reviews complete

Optimizing a solution means removing manual work and ensuring tasks flow smoothly across systems through automation and integration. The best choice is to leverage automation to handle repeatable tasks, integrate where possible so data doesn't have to be reentered, and eliminate manual processes that create errors and slow delivery. This approach yields faster, more reliable, and scalable results with consistent controls and easier auditing. Security and governance should be embedded in the automation from the start, not treated as a separate later step. Adding manual processes wastes effort and creates bottlenecks, while simply increasing hardware or data redundancy doesn't improve how the work gets done. Delaying automation until security reviews are complete can stall progress. So, automate where possible, integrate across systems, and remove manual steps.

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://servicenowtecharchitectwa.examzify.com>

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

SAMPLE