

MuleSoft Integration Architect 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 9

Explanations 11

Next Steps 17

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. If you route a message with a flow reference to a sub flow, the flow message is processed how?**
 - A. The message is processed asynchronously**
 - B. The message is processed synchronously**
 - C. The message is discarded**
 - D. The system throws an error**

- 2. In Anypoint deployment, what is the relationship between a VPC and a business group?**
 - A. Each VPC is assigned to a business group.**
 - B. VPCs can belong to multiple business groups.**
 - C. Business groups determine the region of a VPC.**
 - D. A VPC has no relation to business groups.**

- 3. Which statement about Java and DataWeave in MuleSoft is true?**
 - A. Java is generally NOT recommended for data transformations**
 - B. DataWeave cannot call out to Java methods**
 - C. Java classes cannot be defined in Spring context**
 - D. MuleSoft discourages encapsulating Java transformations in classes**

- 4. Which statement about on Error Propagate is true?**
 - A. It rethrows the error up to the next level after executing the error scope**
 - B. The rest of the flow that threw the error is always executed**
 - C. At the end of the scope, the rest of the flow that threw the error is not executed and the error is re-thrown up to the next level**
 - D. It converts the error into an HTTP 200 response**

- 5. Why should you not use file-based persistence for a Mule application on CloudHub?**
 - A. Mule applications have unlimited and persistent file system access**
 - B. The file connector can access any folder such as /etc**
 - C. CloudHub provides built-in long-term file-based persistence**
 - D. CloudHub restricts file system access and persistence is ephemeral**

- 6. For regulatory compliance, which deployment options are appropriate?**
- A. CloudHub**
 - B. Runtime Fabric**
 - C. Private Cloud Edition or Pivotal Cloud Foundry**
 - D. CloudHub and Runtime Fabric**
- 7. What is true about the connector reconnection strategy in Mule?**
- A. It always retries immediately until success.**
 - B. There is no option to fail deployment if reconnection fails.**
 - C. A failsDeploy attribute can throw an exception if reconnection attempts fail and prevent startup.**
 - D. By default a failed connectivity test is logged and the application continues to run without trying to reconnect.**
- 8. Which statement best describes the difference between Edge policies and API policies?**
- A. Edge policies apply to many API instances; API policies apply to exactly one API instance**
 - B. Policies are applied in the Edge gateway or in the API Implementation / API proxy app**
 - C. Edge policies require tokenization**
 - D. There is no difference; both apply to the same scope**
- 9. What are the two steps described in the two-phase commit protocol?**
- A. First you check if the system is ready to commit; When it's ready to commit, you commit to the DB**
 - B. You always commit to the DB first, then check readiness**
 - C. Transactions are committed independently on each resource without coordination**
 - D. The protocol relies on a central coordinator to flood commit messages after preparing**

10. Which CIDR range is allowed for a VPC according to the material?

- A. Between /24 and /16**
- B. Only /23**
- C. Only /16**
- D. Between /28 and /12**

SAMPLE

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. If you route a message with a flow reference to a sub flow, the flow message is processed how?
 - A. The message is processed asynchronously
 - B. The message is processed synchronously**
 - C. The message is discarded
 - D. The system throws an error

Invoking a sub-flow with a flow reference processes the message synchronously. It runs on the same processing thread as the calling flow, and the parent flow waits for the sub-flow to finish before continuing. The payload returned after the flow-reference completes is whatever the sub-flow leaves as the message payload (subject to any further changes in the caller). If an error occurs in the sub-flow, it follows the defined error handling for the flow. If asynchronous behavior is needed, you'd structure the flows differently (for example, using an async scope or a separate flow invoked asynchronously).

2. In Anypoint deployment, what is the relationship between a VPC and a business group?
 - A. Each VPC is assigned to a business group.**
 - B. VPCs can belong to multiple business groups.
 - C. Business groups determine the region of a VPC.
 - D. A VPC has no relation to business groups.

A VPC is tied to a single business group, meaning it is created, owned, and managed within one specific business group. This setup provides a clear boundary for access, governance, and billing: only members of that group have deploy and configuration rights for the VPC, and resources inside the VPC stay isolated from other groups. This one-to-one relationship keeps responsibilities clean and prevents cross-group changes that could impact security or performance. It's not shared across multiple groups, and the VPC's region is not dictated by the business group—region choice is a deployment detail separate from group ownership.

3. Which statement about Java and DataWeave in MuleSoft is true?

- A. Java is generally NOT recommended for data transformations**
- B. DataWeave cannot call out to Java methods**
- C. Java classes cannot be defined in Spring context**
- D. MuleSoft discourages encapsulating Java transformations in classes**

Data transformation in MuleSoft is best expressed with DataWeave because it is purpose-built for mapping and shaping data across formats like JSON, XML, and CSV in a concise, readable way. DataWeave provides a declarative syntax, built-in functions, and operators that let you express complex transformations without writing procedural code. This makes maintenance and evolution of API data contracts much smoother. Java is typically kept for logic that isn't just about transforming data—things like custom business logic, integrations, or performance-critical code. Implementing data transformations in Java tends to be more verbose, harder to adjust when input shapes change, and harder to maintain across multiple APIs. While you can call Java methods from DataWeave or use Java components for non-transformation tasks, that isn't the standard approach for data transformation, which is why the statement in question is considered true. The other options don't fit because DataWeave can interact with Java when needed, Java classes can be defined and used within a Spring context, and while you could encapsulate transformation logic in Java classes, the recommended practice in MuleSoft is to rely on DataWeave for the transformation step.

4. Which statement about on Error Propagate is true?

- A. It rethrows the error up to the next level after executing the error scope**
- B. The rest of the flow that threw the error is always executed**
- C. At the end of the scope, the rest of the flow that threw the error is not executed and the error is re-thrown up to the next level**
- D. It converts the error into an HTTP 200 response**

On Error Propagate is about sending the fault up to the caller instead of letting the flow continue as if nothing happened. When an error occurs inside a scope and this error handler is invoked, the error path runs its actions, but the normal flow that started the error does not resume afterwards. Instead, the error is re-thrown to the enclosing flow or the next level up, so the original flow's remaining steps are not executed and the fault moves outward for further handling. This is why the statement describing that, at the end of the scope, the rest of the flow that threw the error is not executed and the error is re-thrown up to the next level is correct. It captures the upward propagation behavior rather than continuing processing or converting the error to a success response. For contrast, On Error Continue would allow the flow to resume processing after handling the error, and turning the error into an HTTP 200 would require separate response handling, not the default behavior of On Error Propagate.

5. Why should you not use file-based persistence for a Mule application on CloudHub?
- A. Mule applications have unlimited and persistent file system access
 - B. The file connector can access any folder such as /etc
 - C. CloudHub provides built-in long-term file-based persistence
 - D. CloudHub restricts file system access and persistence is ephemeral**

In CloudHub, the file system is not a durable storage option. The runtime runs in containers with ephemeral disks, and local files can be lost when a worker restarts, scales, or moves. Because each worker often has its own isolated filesystem and there's no guaranteed shared, long-term storage, relying on file-based persistence can lead to data loss and inconsistent state. To ensure durability, use external persistence such as databases or cloud storage (for example, a database, S3/Blob storage, etc.) and store state or large artifacts there instead of the local filesystem. CloudHub also restricts access to the underlying file system, so the idea of broad, unrestricted file access (like rooting into folders such as /etc) isn't allowed. That's why persistent file-based storage isn't suitable on CloudHub.

6. For regulatory compliance, which deployment options are appropriate?
- A. CloudHub
 - B. Runtime Fabric
 - C. Private Cloud Edition or Pivotal Cloud Foundry**
 - D. CloudHub and Runtime Fabric

Regulatory compliance hinges on having control over where data resides and how the environment is governed. For strict requirements, you want an on-premises or private cloud deployment where you set your own security controls, access policies, and audit trails. Private Cloud Edition provides an on-premises MuleSoft deployment, giving you full control over the data and the environment. Pivotal Cloud Foundry lets you run MuleSoft inside a private CF environment, offering the same level of data residency and governance within your private cloud. Public-cloud offerings like CloudHub place data in MuleSoft's public, multi-tenant cloud, which can complicate data residency and compliance requirements. Runtime Fabric is flexible and can be run on your own infrastructure, but when the goal is explicit regulatory compliance, the explicit private-cloud options are the most straightforward way to meet those controls and audits.

7. What is true about the connector reconnection strategy in Mule?

- A. It always retries immediately until success.**
- B. There is no option to fail deployment if reconnection fails.**
- C. A failsDeploy attribute can throw an exception if reconnection attempts fail and prevent startup.**
- D. By default a failed connectivity test is logged and the application continues to run without trying to reconnect.**

When Mule checks a connector's connectivity at startup, the behavior you'll observe is governed by the reconnection strategy. By default, if the connectivity test fails, Mule logs the issue and allows the application to keep running without attempting automatic reconnects. This makes the statement true: a failed connectivity test is logged and the application continues to run without trying to reconnect. There are ways to change this behavior (to fail startup or to configure retries), so the other options don't fit as the default behavior. The default is not endless immediate retries, and while you can enable startup failure under certain configurations, that isn't the default mode.

8. Which statement best describes the difference between Edge policies and API policies?

- A. Edge policies apply to many API instances; API policies apply to exactly one API instance**
- B. Policies are applied in the Edge gateway or in the API Implementation / API proxy app**
- C. Edge policies require tokenization**
- D. There is no difference; both apply to the same scope**

Edge policies operate at the gateway level and govern traffic for many APIs that pass through that gateway, so they're scoped to multiple API instances. API policies, by contrast, are applied to the API's implementation or its proxy app, so they're scoped to a single API instance. The statement reflects this difference in scope: edge policies cover multiple APIs, while API policies apply to exactly one API. The other options either don't address this scope distinction or are incomplete.

9. What are the two steps described in the two-phase commit protocol?

- A. First you check if the system is ready to commit; When it's ready to commit, you commit to the DB**
- B. You always commit to the DB first, then check readiness**
- C. Transactions are committed independently on each resource without coordination**
- D. The protocol relies on a central coordinator to flood commit messages after preparing**

Two-phase commit ensures a distributed transaction either commits everywhere or rolls back everywhere by coordinating across all participants. In the first phase, a central coordinator asks every resource if it can commit and each one responds after preparing to commit (a vote of yes or no). If all resources say yes, the second phase commits the transaction across all resources; if any resource cannot commit, the coordinator aborts across all of them. The option describing checking readiness first and then performing the commit matches this flow: you determine readiness (prepare) and only then issue the commit. The other choices either imply committing without coordination, or describe incorrect sequences, or oversimplify how the coordinator and messages work in 2PC.

10. Which CIDR range is allowed for a VPC according to the material?

- A. Between /24 and /16**
- B. Only /23**
- C. Only /16**
- D. Between /28 and /12**

A VPC CIDR block is sized to provide enough private IP space while keeping address management practical. The material states you can choose a VPC CIDR block that falls within the range from /16 to /24. That means you can use any prefix length from /16 (65,536 IPs) down to /24 (256 IPs). This range gives flexibility: a /16 supports large deployments with many subnets, while a /24 works for smaller networks, and you can still carve smaller subnets inside the VPC. Blocks outside this window—whether larger like /12 or smaller than /24—aren't covered as allowed in the material, so the move from /16 up to /24 is the intended scope.

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://mulesoftintegrationarchitect.examzify.com>

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

SAMPLE