

AAISM Domain 1: AI Governance, Program Management 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	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. Which professionals are involved in data preprocessing and training and interpret anomalous outputs?**
 - A. AI ethicists**
 - B. Data stewards/owners**
 - C. Data engineers and scientists**
 - D. Privacy experts**

- 2. In AI governance structures, which role typically funds or champions the project?**
 - A. Domain experts**
 - B. Project sponsors**
 - C. Compliance officers**
 - D. Technical experts**

- 3. Which organization published AI Principles that include transparency, inclusiveness, sustainability, and accountability?**
 - A. IEEE**
 - B. OECD**
 - C. ISO**
 - D. NIST**

- 4. What are common elements included in AI risk assessments that may not be in traditional IT risk assessments?**
 - A. Only hardware failures.**
 - B. Marketing risk and brand reputation.**
 - C. Model bias, adversarial threats, explainability issues, and ethical considerations.**
 - D. Financial risk and budgeting.**

- 5. Which term stores files as vectors with associated metadata?**
 - A. Vector Database**
 - B. Data Exploration and Training Platform**
 - C. Timeliness**
 - D. Accuracy**

- 6. In data governance, what is the role of a business glossary?**
- A. Store user passwords**
 - B. Standardize terms and definitions across data assets**
 - C. Encrypt data at rest**
 - D. Execute data backups automatically**
- 7. Which practice involves simulating real-world adversaries to identify vulnerabilities in AI applications?**
- A. Explainability in AI**
 - B. Red teaming in AI systems**
 - C. TEVV in AI system development**
 - D. Chairperson**
- 8. Identifies the data and users involved in each step.**
- A. Data Usage Diagrams**
 - B. Dataflow Diagrams**
 - C. Swim Lane Model**
 - D. Data Source**
- 9. NIST AI RMF helps organizations with what objective?**
- A. Data privacy**
 - B. Identifying, assessing, and mitigating risk in AI systems**
 - C. Speed of deployment**
 - D. Regulatory compliance**
- 10. Which term describes a branch of artificial intelligence that converts human language into data usable by computer systems?**
- A. Robotic automation**
 - B. Computer vision**
 - C. Machine learning**
 - D. Natural language processing (NLP)**

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. Which professionals are involved in data preprocessing and training and interpret anomalous outputs?

- A. AI ethicists
- B. Data stewards/owners
- C. Data engineers and scientists**
- D. Privacy experts

Handling the data preparation, model training, and interpretation of unusual outputs is a practical ML workflow led by data engineers and data scientists. Data engineers build and maintain the data pipelines, clean and transform data, handle missing values, and engineer features so models can learn reliably. Data scientists choose suitable algorithms, train the models on the prepared data, tune parameters, and assess performance, while also examining and interpreting any anomalous or unexpected outputs to diagnose issues. This combination directly covers both the data-side prep and the modeling-and-diagnostics side, making it the best fit for these tasks. Roles like data stewards/owners, privacy experts, or AI ethicists contribute important governance, privacy, and ethical considerations, but they aren't typically responsible for the hands-on preprocessing, training, and interpretation work.

2. In AI governance structures, which role typically funds or champions the project?

- A. Domain experts
- B. Project sponsors**
- C. Compliance officers
- D. Technical experts

The funding and championing of an AI governance project typically come from the project sponsor. This role is usually a senior leader or business owner who is accountable for the project's value and alignment with strategic goals. The sponsor secures the budget, allocates resources, communicates the project's importance to stakeholders, and helps navigate organizational obstacles. This combination of financial backing and active advocacy makes the sponsor the central driver who keeps the initiative moving and ensures it receives the necessary authority to progress. Domain experts provide specialized knowledge about the subject matter, but they don't own funding or political backing. Compliance officers focus on ensuring policies, risk, and regulatory requirements are met, not championing the project's resources. Technical experts implement the solution, translating requirements into a working system, but they don't typically secure the project's funding or act as its primary advocate at the executive level.

3. Which organization published AI Principles that include transparency, inclusiveness, sustainability, and accountability?

- A. IEEE
- B. OECD**
- C. ISO
- D. NIST

This question tests your knowledge of which organization publicly framed AI principles around transparency, inclusiveness, sustainability, and accountability. The organization is OECD. In 2019 the Organisation for Economic Co-operation and Development published AI Principles as part of a Council Recommendation, framed to promote inclusive growth, sustainable development, and well-being, while ensuring AI systems are transparent and accountable and governed in a human-centered way. The inclusiveness aspect aligns with the idea of broad benefits and participation (inclusive growth and broad societal well-being), and sustainability maps to the emphasis on long-term social and environmental impacts. Other organizations, like IEEE or NIST/ISO, issue important standards and ethical guidelines, but the concise set that explicitly pairs transparency and accountability with inclusiveness and sustainability is the OECD AI Principles.

4. What are common elements included in AI risk assessments that may not be in traditional IT risk assessments?

- A. Only hardware failures.
- B. Marketing risk and brand reputation.
- C. Model bias, adversarial threats, explainability issues, and ethical considerations.**
- D. Financial risk and budgeting.

AI risk assessments focus on how the system operates and affects people, data, and society, not just how well the tech runs. The key elements here are model bias, adversarial threats, explainability issues, and ethical considerations. Model bias arises because models learn from data that may reflect existing prejudices or imbalances. If the training data or the way it's labeled advantages or disadvantages certain groups, the model's outputs can be unfair or harmful in real use. This isn't just a technical flaw; it has real societal and regulatory implications. Adversarial threats involve attempts to trick or manipulate the model with carefully crafted inputs, causing incorrect or unsafe outputs. This risk is particular to AI systems that interpret or classify data, and it can be exploited in ways that don't resemble ordinary IT failures. Explainability issues concern how much of the model's reasoning can be understood and traced. When decisions affect users or customers, stakeholders need transparency to trust, audit, and comply with rules. Black-box models can create accountability gaps and hinder governance. Ethical considerations cover privacy, consent, fairness, and alignment with values and laws. AI decisions can impact rights and well-being, so governance must address what is ethical to implement and monitor. By contrast, factors like hardware failures, marketing risk, or broad financial budgeting fall outside the unique scope of AI risk—the first set captures risks that arise specifically because of learning from data, model behavior, and the social impact of automated decisions.

5. Which term stores files as vectors with associated metadata?

A. Vector Database

B. Data Exploration and Training Platform

C. Timeliness

D. Accuracy

Storing files as vectors with associated metadata is the function of a vector database. A vector database stores high-dimensional representations (embeddings) of content alongside metadata like source, date, or tags. This setup lets you perform semantic searches by comparing query vectors to stored vectors and quickly retrieving the most similar items along with their metadata. The other options describe qualities or generic platforms rather than a storage and retrieval system for vector embeddings: a Data Exploration and Training Platform is about analyzing data and building models, while Timeliness and Accuracy refer to speed and correctness, not where or how embeddings with metadata are stored. So the vector database is the component designed for storing files as vectors with metadata.

6. In data governance, what is the role of a business glossary?

A. Store user passwords

B. Standardize terms and definitions across data assets

C. Encrypt data at rest

D. Execute data backups automatically

A business glossary in data governance provides a single, authoritative vocabulary for terms used across data assets. It is a centralized, governed repository that captures business terms, their definitions, synonyms, owners, and accepted values. This standardization ensures everyone—analysts, data stewards, and business users—interprets key terms the same way, reducing miscommunication and making analytics more reliable. By tying terms to specific data assets and metadata, the glossary supports data literacy, consistent reporting, and easier impact analysis when definitions evolve. The other options relate to security or operations, not governance vocabulary. Storing user passwords is about authentication security. Encrypting data at rest is about protecting data confidentiality. Automatically backing up data is about resilience and data operations. These are important, but they serve different purposes than standardizing terms across data assets.

7. Which practice involves simulating real-world adversaries to identify vulnerabilities in AI applications?

- A. Explainability in AI
- B. Red teaming in AI systems**
- C. TEVV in AI system development
- D. Chairperson

Red teaming involves simulating real-world adversaries to identify vulnerabilities in AI applications. In this practice, a dedicated team acts like attackers, probing the system with realistic scenarios to test data pipelines, prompts, access controls, deployment environments, and safety guardrails. The aim is to uncover weaknesses that might not show up in normal testing, such as susceptibility to prompt injection, data poisoning, model misuse, or privacy leaks, and to evaluate the effectiveness of governance, monitoring, and incident response. The insights from red teaming guide concrete mitigations—strong input validation, robust authentication, adversarial training, enhanced monitoring, and governance updates—so the AI system remains secure and trustworthy. This approach is different from explainability, which focuses on understanding how the model makes decisions.

8. Identifies the data and users involved in each step.

- A. Data Usage Diagrams
- B. Dataflow Diagrams
- C. Swim Lane Model**
- D. Data Source

The ability to map a process by who performs each step and what data they use at that step is best shown by the Swim Lane Model. In a swim lane diagram, each lane represents an actor, role, or system, and the process steps are placed within the lanes that are responsible for them. This layout makes it immediate to see, for every step, which user or unit is involved and which data items flow through or are produced at that moment. It directly highlights responsibilities and data interactions across the end-to-end process, which is exactly what identifying the data and users involved in each step requires. Dataflow diagrams focus on how data moves between processes, stores, and sinks, rather than who is performing each step or owning the data at that step. Data sources point to where data originates, not to the per-step involvement of users or systems. Data usage diagrams aren't a standard, well-established way to map responsibilities per step. So the Swim Lane Model best fits the need to identify both the data and the users involved at every step of the process.

9. NIST AI RMF helps organizations with what objective?

- A. Data privacy
- B. Identifying, assessing, and mitigating risk in AI systems**
- C. Speed of deployment
- D. Regulatory compliance

The main idea being tested is that the NIST AI RMF is designed to help organizations manage risk in AI systems throughout their life cycle. It gives a structured way to identify where risks come from, assess how likely they are and how severe their impact could be, and then put in place controls to reduce those risks to acceptable levels. This focus on identifying, assessing, and mitigating risk is what makes it the best answer: the framework is built to ensure AI is trusted, safe, reliable, fair, and secure, rather than simply speeding up deployment or chasing regulatory checklists. Context helps: the RMF encourages ongoing governance and monitoring, so organizations continuously evaluate AI systems as they evolve and adapt mitigations as new risks emerge. While privacy, security, or regulatory considerations can be part of the risk landscape, they are components of the broader risk-management objective rather than the sole aim.

10. Which term describes a branch of artificial intelligence that converts human language into data usable by computer systems?

- A. Robotic automation
- B. Computer vision
- C. Machine learning
- D. Natural language processing (NLP)**

Natural language processing focuses on enabling computers to understand human language and convert it into data they can use. It covers tasks that turn spoken or written language into machine-readable representations—like breaking down sentences, grasping meaning, identifying entities, and enabling actions such as chat responses, voice assistants, or translation. This is exactly what happens when language is transformed into structured data that a computer system can analyze and act upon. Other terms describe different areas: robotic automation focuses on automating physical or business processes with robots; computer vision deals with interpreting visual data from images or video; machine learning is the broad set of algorithms that learn from data, not language specifically. NLP is the field that handles turning language into data a computer can process.

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://aaismdomain1.examzify.com>

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

SAMPLE