ISACA ADVANCED IN AI SECURITY MANAGEMENT (AAISM) PRACTICE TEST (SAMPLE)

STUDY GUIDE



SAMPLE STUDY GUIDE WITH LIMITED QUESTIONS

VISIT US ONLINE FOR THE FULL VERSION STUDY GUIDE

https://isacaaaism.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 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.

SAMPLE

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!

SAMPLE

Questions

SAMPLE

1. What are the key components of the recovery process after an AI incident?

A. Ensure the threat is contained and eradicated.

B. Rebuild the entire system from scratch without validation.

C. Increase exposure to data.

D. Ignore residual threats.

2. What does Adversarial Inference aim to achieve?

A. To make the AI model more resilient using regularization techniques and defensive distillation.

B. To bypass security by authenticating attackers.

C. To increase model complexity without defense.

D. To ignore adversarial inputs.

3. Which sequence correctly outlines the four steps to creating a data inventory for a Privacy Impact Assessment (PIA)?

A. Plan: Define scope, stakeholders, and resources. Decide: Determine what information to collect. Populate: Collect information using various techniques. Publish: Collate and publish the metadata.

B. Decide: Plan, Publish, Populate

C. Plan, Publish, Decide, Populate

D. Populate, Plan, Decide, Publish

4. Lack of Explainability/Transparency in AI implies?

A. Inability To Share Details Of The AI Model, Making It Hard To Identify Errors, Biases, And Demonstrate Regulatory Compliance

B. Excessive Documentation Of Model Internals

C. Perfectly Understandable Models

D. Overly Transparent Data Pipelines

5. Which factor most directly contributes to training data leakage?

A. Weak access controls.

B. Model architecture complexity.

C. Overfitting during training.

D. Excessive data augmentation.

SAMPLE

6. A common benefit of AI in business is?

A. Decreased data accuracy

B. Increased productivity through automation

C. Slower decision making

D. Higher process cycle times

7. Which statement best describes the Assess phase in incident response?

A. The focus is to collect facts to establish the timeline, scope, and impact of the incident without damaging the investigation.

B. The focus is to deploy containment measures immediately.

C. The focus is to recalculate penalties for vendors.

D. The focus is to archive data without review.

8. In AI governance, what is the purpose of risk management?

A. To integrate AI risks into the enterprise risk register and manage ethical, security, and reputational risks.

B. To replace all regulatory oversight.

C. To outsource risk management to vendors.

D. To ignore ethical concerns.

9. What is required for AI models to remain effective against evolving threats?

A. Updated with the latest threat intelligence

B. Fixed once deployed

C. Implemented without monitoring

D. Ignored new data

10. What does the AI Acceptable Use Policy dictate?

A. Permitted and not permitted uses of AI within the enterprise

B. Data Retention Periods

C. Incident Response Procedures

D. Vendor Selection Criteria

SAMPLE

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. What are the key components of the recovery process after an AI incident?

A. Ensure the threat is contained and eradicated.

B. Rebuild the entire system from scratch without validation.

C. Increase exposure to data.

D. Ignore residual threats.

During the recovery phase after an AI incident, the priority is to contain and eradicate the threat so that operations can be restored safely. This involves stopping the attacker's access, removing malicious code or compromised models, and replacing or repairing affected components with clean, validated equivalents. It also includes validating the integrity of data and model outputs, applying patches or configuration fixes, and enhancing monitoring and controls to detect any signs of lingering risk. Once the environment is confirmed clean and stable, recovery proceeds with reestablishing services, validating performance, and implementing improvements to prevent recurrence, along with post-incident review and updates to incident response playbooks. Rebuilding the entire system from scratch without validation would ignore essential checks that ensure safety and integrity. Increasing exposure to data would raise risk and could violate governance and privacy controls. Ignoring residual threats leaves an open door for renewed compromise, undermining any recovery effort.

2. What does Adversarial Inference aim to achieve?

A. To make the AI model more resilient using regularization techniques and defensive distillation.

B. To bypass security by authenticating attackers.

C. To increase model complexity without defense.

D. To ignore adversarial inputs.

Adversarial inference aims to make AI models more robust against adversarial inputs during inference by applying defensive techniques such as regularization and defensive distillation. Regularization helps prevent the model from relying on fragile patterns, encouraging smoother decision boundaries and better generalization to small input perturbations. Defensive distillation trains a model using softened output probabilities from a teacher model, which tends to dampen gradients and reduce sensitivity to subtle input changes that attackers exploit. These approaches focus on resilience and reliable performance in the face of adversarial perturbations, rather than enabling attacks, increasing unhelpful complexity, or ignoring adversarial inputs.

SAMPLE

3. Which sequence correctly outlines the four steps to creating a data inventory for a Privacy Impact Assessment (PIA)?

A. Plan: Define scope, stakeholders, and resources. Decide: Determine what information to collect. Populate: Collect information using various techniques. Publish: Collate and publish the metadata.

B. Decide: Plan, Publish, Populate

C. Plan, Publish, Decide, Populate

D. Populate, Plan, Decide, Publish

Starting with planning sets the foundation for a robust data inventory. In a Privacy Impact Assessment, you need to define the scope, identify the right stakeholders, and allocate resources. This ensures everyone agrees on what's in scope, what data sources exist, and what governance is in place before you start collecting details. Next, deciding what information to collect is essential. This step translates the planning into concrete data elements: data categories, processing purposes, data flows, retention, data subject rights, legal basis, and any vendors or systems involved. By choosing the exact metadata to capture, you avoid gathering irrelevant information and ensure the inventory will support the PIA's analysis. Then you populate the inventory. With the scope and data elements defined, you gather the actual information through interviews, policy and system reviews, data mapping, data lineage analyses, questionnaires, and automated scans. This is where the metadata starts to take shape and the inventory becomes a true record of what data exists and how it's processed. Finally, publishing the metadata makes the inventory available to stakeholders and supports ongoing governance. Publishing includes compiling, validating, and distributing the data, along with versioning, access controls, and plans for regular updates and re-assessment. If the steps are out of order, you'd risk collecting the wrong data, skipping critical sources, or failing to provide a governance mechanism for the resulting inventory.

4. Lack of Explainability/Transparency in AI implies?

A. Inability To Share Details Of The AI Model, Making It Hard To Identify Errors, Biases, And Demonstrate Regulatory Compliance

B. Excessive Documentation Of Model Internals

C. Perfectly Understandable Models

D. Overly Transparent Data Pipelines

Explainability and transparency in AI are about making the decision process understandable and auditable. When explainability is lacking, you can't share details of how the model works or why it made a given decision. This obscurity makes it difficult to spot errors in the data or the model, identify biased outcomes, and demonstrate to regulators or stakeholders that the system complies with rules and standards. In practice, there is often a trade-off between complexity and interpretability; some powerful models are less transparent, so organizations rely on governance, testing, and, where possible, interpretable approximations to show accountability. Without this clarity, risk increases because wrong or biased decisions cannot be traced or corrected, and trust in the system declines.

SAMPLE

5. Which factor most directly contributes to training data leakage?

A. Weak access controls.

B. Model architecture complexity.

C. Overfitting during training.

D. Excessive data augmentation.

Training data leakage happens when training data is exposed to someone who should not have access. The most direct contributor is weak access controls, because if authentication, authorization, and data protection are insufficient, the data can be read, copied, or exfiltrated by unauthorized users or compromised systems, causing leakage regardless of how the model is trained. Strengthening access controls, enforcing least privilege, and auditing data access are key to preventing this risk. While other factors relate to how a model learns or handles data, they do not create an exposure path as directly. Model architecture complexity affects learning capacity but not the presence of an accessible data channel. Overfitting can increase the risk of the model memorizing and potentially revealing training data through outputs, but the leakage route hinges more on who can access the data in the first place. Excessive data augmentation expands the dataset without inherently creating a leakage channel.

6. A common benefit of AI in business is?

A. Decreased data accuracy

B. Increased productivity through automation

C. Slower decision making

D. Higher process cycle times

The main advantage AI brings to business is increased productivity through automation. By taking over repetitive, rules-based tasks—like data entry, report generation, monitoring, and routine customer interactions—AI can work faster and with fewer errors, around the clock. This reduces manual effort and frees people to focus on higher-value work, speeding up processes and enabling the organization to scale. While AI can also improve accuracy and provide real-time insights for better decisions, outcomes describing decreased accuracy, slower decision-making, or longer cycle times don't reflect the typical productivity gains AI delivers.

7. Which statement best describes the Assess phase in incident response?

A. The focus is to collect facts to establish the timeline, scope, and impact of the incident without damaging the investigation.

B. The focus is to deploy containment measures immediately.

C. The focus is to recalculate penalties for vendors.

D. The focus is to archive data without review.

In incident response, the Assess phase is about understanding what happened and its effects while preserving evidence. The focus is on collecting facts to reconstruct the timeline, determine the scope of affected systems and data, and gauge the potential impact, all without altering or destroying evidence. This careful gathering helps establish how the incident unfolded and what needs to be protected or restored, and it sets the direction for containment and remediation. Immediate containment actions or penalties calculation fall outside this phase, and archiving data without review would skip the crucial analysis that informs the response.

SAMPLE

8. In AI governance, what is the purpose of risk management?

A. To integrate AI risks into the enterprise risk register and manage ethical, security, and reputational risks.

B. To replace all regulatory oversight.

C. To outsource risk management to vendors.

D. To ignore ethical concerns.

In AI governance, the aim is to bring AI-related risks into the organization's overall risk management program, so they are identified, assessed, and mitigated as part of enterprise risk governance. This approach treats AI risk like other business risks, ensuring there is visibility, ownership, and ongoing monitoring. It also explicitly includes ethical considerations (such as bias and fairness), security concerns (like data protection and system integrity), and reputational risks (trust and public perception). Embedding AI risks in the risk register and governance processes helps leaders set risk appetite, implement controls, and track effectiveness throughout the AI lifecycle, aligning risk response with strategic objectives and regulatory expectations. Replacing regulatory oversight, outsourcing risk management, or ignoring ethics would undermine accountability and the holistic protection that proper AI governance requires.

9. What is required for AI models to remain effective against evolving threats?

A. Updated with the latest threat intelligence

B. Fixed once deployed

C. Implemented without monitoring

D. Ignored new data

Continual adaptation through up-to-date threat intelligence keeps AI models effective against evolving threats. Threat landscapes change quickly as attackers develop new techniques and indicators of compromise. To stay effective, AI security models must regularly incorporate fresh data, which means retraining, updating features, and adjusting decision policies based on current information. This ongoing process helps the model recognize newly observed attack patterns, reduces blind spots, and maintains resilience against adversarial attempts. Why this approach works is that staying current with threat intel lets the model learn from recent incidents and observations—new IOCs, behavior changes, and novel attack vectors—so it can detect and respond to threats it hasn't seen before. It also enables testing against fresh scenarios and confirms that the model continues to perform well as conditions shift. If a model is fixed after deployment, it can't learn from new threats, so its effectiveness erodes as attackers evolve. Without monitoring, you miss drift and new threat signals, delaying necessary updates. Ignoring new data prevents the model from improving, leaving gaps that adversaries can exploit.

SAMPLE

10. What does the AI Acceptable Use Policy dictate?

A. Permitted and not permitted uses of AI within the enterprise

B. Data Retention Periods

C. Incident Response Procedures

D. Vendor Selection Criteria

The main concept being tested is how an AI Acceptable Use Policy guides how AI tools may be used within an organization. It defines the permitted and prohibited uses of AI, setting boundaries to protect data privacy, intellectual property, and security, and it governs how employees interact with and deploy AI tools in real work. This makes the option describing allowed and disallowed AI uses the best fit, because the policy's core purpose is to outline what you can and cannot do with AI in the enterprise. Policies about data retention periods deal with how long data is kept and purged, which is part of data governance rather than how AI is allowed to be used. Incident response procedures cover steps for detecting and responding to security incidents, not the rules for AI usage. Vendor selection criteria relate to how vendors are chosen and managed, which is procurement policy, not AI usage guidelines.

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://isacaaaism.examzify.com

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

SAMPLE