

# 最新Amazon AIF-C01試験の練習問題と解答



P.S. CertShikenがGoogle Driveで共有している無料かつ新しいAIF-C01ダンプ: <https://drive.google.com/open?id=1iDIPvPUMi3O0Zd2ml5i2XpvIHNzgjSTk>

弊社のCertShikenはIT認定試験のソフトの一番信頼たるバンドになるという目標を達成するために、弊社はあなたに最新版のAmazonのAIF-C01試験問題集を提供いたします。弊社のソフトを使用して、ほとんどのお客様は難しいと思われているAmazonのAIF-C01試験に順調に剛角しました。これも弊社が自信的にあなたに商品を薦める原因です。もし弊社のソフトを使ってあなたは残念で試験に失敗したら、弊社は全額で返金することを保証いたします。すべてのことの目的はあなたに安心して試験に準備さされるということです。

## Amazon AIF-C01 認定試験の出題範囲:

トピック	出題範囲
トピック 1	<ul style="list-style-type: none"><li>• Fundamentals of AI and ML: This domain covers the fundamental concepts of artificial intelligence (AI) and machine learning (ML), including core algorithms and principles. It is aimed at individuals new to AI and ML, such as entry-level data scientists and IT professionals.</li></ul>
トピック 2	<ul style="list-style-type: none"><li>• Fundamentals of Generative AI: This domain explores the basics of generative AI, focusing on techniques for creating new content from learned patterns, including text and image generation. It targets professionals interested in understanding generative models, such as developers and researchers in AI.</li></ul>
トピック 3	<ul style="list-style-type: none"><li>• Security, Compliance, and Governance for AI Solutions: This domain covers the security measures, compliance requirements, and governance practices essential for managing AI solutions. It targets security professionals, compliance officers, and IT managers responsible for safeguarding AI systems, ensuring regulatory compliance, and implementing effective governance frameworks.</li></ul>
トピック 4	<ul style="list-style-type: none"><li>• Guidelines for Responsible AI: This domain highlights the ethical considerations and best practices for deploying AI solutions responsibly, including ensuring fairness and transparency. It is aimed at AI practitioners, including data scientists and compliance officers, who are involved in the development and deployment of AI systems and need to adhere to ethical standards.</li></ul>
トピック 5	<ul style="list-style-type: none"><li>• Applications of Foundation Models: This domain examines how foundation models, like large language models, are used in practical applications. It is designed for those who need to understand the real-world implementation of these models, including solution architects and data engineers who work with AI technologies to solve complex problems.</li></ul>

## AIF-C01勉強資料 & AIF-C01全真問題集

AIF-C01の有用なテストガイド資料は、最も重要な情報を最も簡単な方法でクライアントに提示するため、AIF-C01の有用なテストガイドを学習するための時間とエネルギーはほとんど必要ありません。クライアントは、テストの学習と準備に20~30時間しかかかりません。仕事や学習などで忙しい人にとっては、これは良いニュースです。なぜなら、テストの準備に十分な時間がないことを心配する必要がなく、主なことをゆっくりとできるからです。AIF-C01学習実践ガイドをご覧ください。ですから、AIF-C01試験の教材の大きな利点であり、クライアントにとって非常に便利です。

### Amazon AWS Certified AI Practitioner 認定 AIF-C01 試験問題 (Q187-Q192):

#### 質問 # 187

What does an F1 score measure in the context of foundation model (FM) performance?

- A. Financial cost of operating the model.
- B. Energy efficiency of the model's computations.
- **C. Model precision and recall.**
- D. Model speed in generating responses.

正解: C

解説:

The F1 score is the harmonic mean of precision and recall, making it a balanced metric for evaluating model performance when there is an imbalance between false positives and false negatives. Speed, cost, and energy efficiency are unrelated to the F1 score.

References: AWS Foundation Models Guide.

#### 質問 # 188

A company is using an Amazon Bedrock base model to summarize documents for an internal use case. The company trained a custom model to improve the summarization quality.

Which action must the company take to use the custom model through Amazon Bedrock?

- A. Register the model with the Amazon SageMaker Model Registry.
- B. Grant access to the custom model in Amazon Bedrock.
- **C. Deploy the custom model in an Amazon SageMaker endpoint for real-time inference.**
- D. Purchase Provisioned Throughput for the custom model.

正解: C

#### 質問 # 189

An online learning company with large volumes of education materials wants to use enterprise search.

- **A. Amazon Kendra**
- B. Amazon Comprehend
- C. Amazon Textract
- D. Amazon Personalize

正解: A

解説:

The correct answer is C - Amazon Kendra, AWS's enterprise search service designed for organizations with large, diverse document repositories. Kendra uses machine learning and natural language understanding (NLU) to provide semantic search, meaning it retrieves results based on meaning rather than keyword matching. According to AWS documentation, Kendra is ideal for educational, enterprise, and knowledge-management scenarios where users need fast, accurate retrieval across PDFs, HTML, Office documents, FAQs, and multimedia transcripts. Kendra connectors can index content from S3, SharePoint, LMS platforms,

and internal databases, making it perfect for large volumes of training materials. Amazon Comprehend (A) is for NLP tasks like entity extraction, not enterprise search. Amazon Textract (B) extracts text from PDFs and scanned materials but does not provide search capabilities. Amazon Personalize (D) is for personalized recommendations, not document retrieval. Kendra is purpose-built for enterprise search and aligns directly with the company's needs.

Referenced AWS Documentation:

- \* Amazon Kendra Developer Guide - Enterprise Search
- \* AWS ML Specialty Guide - Intelligent Search Systems

### 質問 # 190

How can companies use large language models (LLMs) securely on Amazon Bedrock?

- A. Use Amazon CloudWatch Logs to make models explainable and to monitor for bias.
- **B. Design clear and specific prompts. Configure AWS Identity and Access Management (IAM) roles and policies by using least privilege access.**
- C. Enable Amazon Bedrock automatic model evaluation jobs.
- D. Enable AWS Audit Manager for automatic model evaluation jobs.

正解: B

解説:

To securely use large language models (LLMs) on Amazon Bedrock, companies should design clear and specific prompts to avoid unintended outputs and ensure proper configuration of AWS Identity and Access Management (IAM) roles and policies with the principle of least privilege. This approach limits access to sensitive resources and minimizes the potential impact of security incidents.

\* Option A (Correct): "Design clear and specific prompts. Configure AWS Identity and Access Management (IAM) roles and policies by using least privilege access": This is the correct answer as it directly addresses both security practices in prompt design and access management.

\* Option B: "Enable AWS Audit Manager for automatic model evaluation jobs" is incorrect because Audit Manager is for compliance and auditing, not directly related to secure LLM usage.

\* Option C: "Enable Amazon Bedrock automatic model evaluation jobs" is incorrect because Bedrock does not provide automatic model evaluation jobs specifically for security purposes.

\* Option D: "Use Amazon CloudWatch Logs to make models explainable and to monitor for bias" is incorrect because CloudWatch Logs are used for monitoring and not directly for making models explainable or secure.

AWS AI Practitioner References:

\* Secure AI Practices on AWS: AWS recommends configuring IAM roles and using least privilege access to ensure secure usage of AI models.

### 質問 # 191

A social media company wants to use a large language model (LLM) to summarize messages. The company has chosen a few LLMs that are available on Amazon SageMaker JumpStart. The company wants to compare the generated output toxicity of these models.

Which strategy gives the company the ability to evaluate the LLMs with the LEAST operational overhead?

- A. Crowd-sourced evaluation
- B. Reinforcement learning from human feedback (RLHF)
- **C. Automatic model evaluation**
- D. Model evaluation with human workers

正解: C

解説:

The least operational overhead comes from automated tools that can scan and evaluate LLM outputs for toxicity. AWS and SageMaker JumpStart support integrations with automatic evaluation tools and APIs (such as Amazon Comprehend or third-party toxicity classifiers).

\* B is correct: Automated evaluation provides quick, scalable, and repeatable analysis, requiring minimal human intervention.

\* A and C require manual effort, increasing operational overhead.

\* D (RLHF) is resource-intensive and not designed for rapid, automated model comparison.

"Automated evaluation can quickly assess generated text for specific attributes like toxicity, sentiment, or compliance using pre-trained classifiers, reducing human involvement and operational complexity." (Reference: AWS SageMaker JumpStart Evaluation, AWS AI Practitioner Guide)

