# 100% Pass 2026 Amazon Perfect AIP-C01: Key AWS Certified Generative AI Developer - Professional Concepts


What's more, part of that Itbraindumps AIP-C01 dumps now are free: https://drive.google.com/open?id=1y-Ehp4MeQwaiZhoOJtqPL13i6eA6A_6X

Itbraindumps has designed AWS Certified Generative AI Developer - Professional which has actual exam Dumps questions, especially for the students who are willing to pass the Amazon AIP-C01 exam for the betterment of their future. The study material is available in three different formats. Amazon Practice Exam are also available so the students can test their preparation with unlimited tries and pass AWS Certified Generative AI Developer - Professional (AIP-C01) certification exam on the first try.

Today we use computers & internet every day, high-technology products bring our life convenient and benefits. Many positions have great demand. Itbraindumps releases valid AIP-C01 dumps torrent files to help workers go through exams and get certifications so that many dreaming young people can enter into this field and even get a good position. Amazon AIP-C01 Dumps Torrent files is the leading position in this field and can be your NO.1 choice.

**>> Key AIP-C01 Concepts <<**

## AIP-C01 Study Guide Pdf | AIP-C01 Braindumps Downloads

These days the Itbraindumps is providing you online Amazon AIP-C01 exam questions to crack the Amazon AIP-C01 certification exam which means you don't need to be physically present anywhere except the chair at your home. You need a laptop and an active internet connection to access the Itbraindumps Amazon AIP-C01 Exam Questions and practice exam.

## Amazon AIP-C01 Exam Syllabus Topics:

| Topic | Details |
|---|---|
| Topic 1 | • Testing, Validation, and Troubleshooting: This domain covers evaluating foundation model outputs, implementing quality assurance processes, and troubleshooting GenAI-specific issues including prompts, integrations, and retrieval systems. |
| Topic 2 | • Operational Efficiency and Optimization for GenAI Applications: This domain encompasses cost optimization strategies, performance tuning for latency and throughput, and implementing comprehensive monitoring systems for GenAI applications. |

| | |
|---|---|
| Topic 3 | • Implementation and Integration: This domain focuses on building agentic AI systems, deploying foundation models, integrating GenAI with enterprise systems, implementing FM APIs, and developing applications using AWS tools. |
| Topic 4 | • Foundation Model Integration, Data Management, and Compliance: This domain covers designing GenAI architectures, selecting and configuring foundation models, building data pipelines and vector stores, implementing retrieval mechanisms, and establishing prompt engineering governance. |
| Topic 5 | • AI Safety, Security, and Governance: This domain addresses input<br>• output safety controls, data security and privacy protections, compliance mechanisms, and responsible AI principles including transparency and fairness. |

## Amazon AWS Certified Generative AI Developer - Professional Sample Questions (Q49-Q54):

**NEW QUESTION # 49**
A company uses AWS Lambda functions to build an AI agent solution. A GenAI developer must set up a Model Context Protocol (MCP) server that accesses user information. The GenAI developer must also configure the AI agent to use the new MCP server. The GenAI developer must ensure that only authorized users can access the MCP server.
Which solution will meet these requirements?

- A. Use a Lambda function to host the MCP server. Grant the AI agent Lambda functions permission to invoke the Lambda function that hosts the MCP server. Configure the AI agent's MCP client to invoke the MCP server asynchronously.
- B. Use a Lambda function to host the MCP server. Grant the AI agent Lambda functions permission to invoke the Lambda function that hosts the MCP server. Configure the AI agent to use the STDIO transport with the MCP server.
- C. Use a Lambda layer to host the MCP server. Add the Lambda layer to the AI agent Lambda functions.Configure the agentic AI solution to use the STDIO transport to send requests to the MCP server. In the AI agent's MCP configuration, specify the Lambda layer ARN as the command. Specify the user credentials as environment variables.
- D. Use a Lambda function to host the MCP server. Create an Amazon API Gateway HTTP API that proxies requests to the Lambda function. Configure the AI agent solution to use the Streamable HTTP transport to make requests through the HTTP API. Use Amazon Cognito to enforce OAuth 2.1.

**Answer: D**

Explanation:
Option C is the correct solution because it provides a secure, scalable, and standards-compliant way to expose an MCP server to an AI agent while enforcing strong user authorization. The Model Context Protocol supports HTTP-based transports for remote MCP servers, making Streamable HTTP the appropriate choice when the server is hosted as a managed service rather than a local process.
Hosting the MCP server in AWS Lambda enables automatic scaling and cost-efficient execution. By placing Amazon API Gateway in front of the Lambda function, the company creates a secure, managed HTTP endpoint that the AI agent can invoke reliably. This architecture cleanly separates transport, authentication, and business logic, which aligns with AWS serverless best practices.
Using Amazon Cognito to enforce OAuth 2.1 ensures that only authenticated and authorized users can access the MCP server. This satisfies security and compliance requirements when the MCP server handles sensitive user information. Cognito integrates natively with API Gateway, removing the need for custom authentication logic and reducing operational overhead.
Option A lacks user-level authorization controls. Option B and Option D rely on STDIO transport, which is intended for local or tightly coupled processes and is not suitable for distributed, serverless architectures.
Option D also introduces security risks by handling credentials through environment variables.
Therefore, Option C best meets the requirements for secure access control, scalability, and correct MCP integration in an AWS-based AI agent architecture.

**NEW QUESTION # 50**
A company is designing a canary deployment strategy for a payment processing API. The system must support automated gradual traffic shifting between multiple Amazon Bedrock models based on real-time inference metrics, historical traffic patterns, and service health. The solution must be able to gradually increase traffic to new model versions. The system must increase traffic if metrics remain healthy and decrease traffic if the performance degrades below acceptable thresholds.
The company needs to comprehensively monitor inference latency and error rates during the deployment phase. The company must

also be able to halt deployments and revert to a previous model version without any manual intervention.
Which solution will meet these requirements?

- A. Use Amazon Bedrock with provisioned throughput to host model versions. Configure an Amazon EventBridge rule to invoke an AWS Step Functions workflow when a new model version is released.
  Configure the workflow to shift traffic in stages, wait for a specified time period, and invoke an AWS Lambda function to check Amazon CloudWatch performance metrics. Configure the workflow to increase traffic if metrics meet thresholds and to trigger a traffic rollback if performance metrics fall below thresholds.
- B. Use Amazon SageMaker AI endpoint variants to represent multiple Amazon Bedrock model versions.
  Use variant weights to shift traffic. Use Amazon CloudWatch and SageMaker Model Monitor to trigger rollbacks. Use EventBridge to roll back deployments if an anomaly is detected.
- C. Use Amazon OpenSearch Service to track inference logs. Configure OpenSearch Service to invoke an AWS Systems Manager Automation runbook to update Amazon Bedrock model endpoints to shift traffic based on inference logs.
- D. Use AWS Lambda functions to invoke various Amazon Bedrock model versions. Use an Amazon API Gateway HTTP API with stage variables and weighted routing to shift traffic gradually. Use Amazon CloudWatch to monitor performance. Use external logic to adjust traffic and roll back if performance falls below thresholds.

**Answer: A**

Explanation:
Option A is the most complete solution because it provides a fully automated canary strategy with staged traffic shifts, metric-based decisioning, and automatic rollback, all using managed AWS services. The requirement emphasizes automation, health-based traffic progression, and zero manual intervention to revert if performance degrades.
AWS Step Functions is well suited for orchestrating controlled deployment workflows with deterministic stages, waits, and conditional branches. By shifting traffic in stages and pausing for observation windows, the system can evaluate real-time inference latency and error rates before promoting more traffic to the new model version. Amazon CloudWatch provides the necessary real-time metrics and alarms for latency and error monitoring.
Invoking a Lambda function to evaluate CloudWatch metrics enables dynamic logic: increase traffic if thresholds remain healthy, reduce traffic or roll back if error rates rise or latency exceeds limits. Step Functions can halt the deployment by stopping progression or triggering rollback steps immediately, meeting the requirement for automated revert without human action.
Amazon EventBridge provides reliable automation triggers when a new model version is released, ensuring the deployment process is event-driven and repeatable.
Option B depends on "external logic," which introduces operational risk and does not guarantee automatic rollback without custom systems. Option C incorrectly uses SageMaker endpoint variants to represent Bedrock model versions, which is not the intended integration model. Option D is overly indirect and operationally complex, using log pipelines and automation runbooks instead of direct metric-based traffic control.
Therefore, Option A best meets the requirements for automated gradual traffic shifting, real-time monitoring, and automatic rollback for Amazon Bedrock model deployments in a canary strategy.

**NEW QUESTION # 51**
A company uses AWS Lake Formation to set up a data lake that contains databases and tables for multiple business units across multiple AWS Regions. The company wants to use a foundation model (FM) through Amazon Bedrock to perform fraud detection. The FM must ingest sensitive financial data from the data lake.
The data includes some customer personally identifiable information (PII).
The company must design an access control solution that prevents PII from appearing in a production environment. The FM must access only authorized data subsets that have PII redacted from specific data columns. The company must capture audit trails for all data access.
Which solution will meet these requirements?

- A. Configure the FM to request temporary credentials from AWS Security Token Service. Access the data by using presigned S3 URLs that are generated by an API that applies business unit and Regional filters. Use AWS CloudTrail to collect comprehensive audit trails of data access.
- B. Create a separate dataset in a separate Amazon S3 bucket for each business unit and Region combination. Configure S3 bucket policies to control access based on IAM roles that are assigned to FM training instances. Use S3 access logs to track data access.
- C. Configure the FM to authenticate by using AWS Identity and Access Management roles and Lake Formation permissions based on LF-Tag expressions. Define business units and Regions as LF-Tags that are assigned to databases and tables. Use AWS CloudTrail to collect comprehensive audit trails of data access.
- D. Use direct IAM principal grants on specific databases and tables in Lake Formation. Create a custom application layer that logs access requests and further filters sensitive columns before sending data to the FM.

**Answer: C**

Explanation:
Option B is the correct solution because it uses native AWS governance, access control, and auditing capabilities to protect PII while enabling controlled FM access to authorized data subsets. AWS Lake Formation is designed specifically to manage fine-grained permissions for data lakes, including column-level access control, which is critical when handling sensitive financial and PII data.
LF-Tags allow data administrators to define scalable, attribute-based access control policies. By tagging databases, tables, and columns with business unit and Region metadata, the company can enforce policies that ensure the foundation model only accesses approved datasets with PII-redacted columns. This eliminates the risk of sensitive data leaking into production inference workflows.
IAM role-based authentication ensures that the FM accesses data using least-privilege credentials. This integrates cleanly with Amazon Bedrock, which supports IAM-based authorization for service-to-service access. AWS CloudTrail provides immutable audit logs for all access attempts, satisfying compliance and regulatory requirements.
Option A introduces unnecessary data duplication and weak governance controls. Option C relies on custom application logic, increasing operational risk and complexity. Option D bypasses Lake Formation's fine- grained controls and relies on presigned URLs, which reduces governance visibility and control.
Therefore, Option B best meets the requirements for security, compliance, scalability, and auditability when integrating Amazon Bedrock with a Lake Formation-governed data lake.

**NEW QUESTION # 52**
A company is building a video analysis platform on AWS. The platform will analyze a large video archive by using Amazon Rekognition and Amazon Bedrock. The platform must comply with predefined privacy standards. The platform must also use secure model I/O, control foundation model (FM) access patterns, and provide an audit of who accessed what and when.
Which solution will meet these requirements?

- A. Configure AWS CloudTrail Insights to analyze API call patterns across accounts and detect anomalous activity in Amazon Bedrock, Amazon Rekognition, Amazon S3, and AWS KMS. Deploy Amazon Macie to scan and classify the video archive. Use server-side encryption with AWS KMS keys (SSE- KMS) to encrypt all stored data. Configure CloudTrail to capture KMS API usage events for audit purposes. Configure Amazon EventBridge rules to process CloudTrail Insights anomalies and Macie findings. Use CloudWatch alarms to trigger automated notifications and security responses when potential security issues are detected.
- B. Configure VPC endpoints for Amazon Bedrock model API calls. Implement Amazon Bedrock guardrails to filter harmful or unauthorized content in prompts and responses. Use Amazon Bedrock trace events to track all agent and model invocations for auditing purposes. Export the traces to Amazon CloudWatch Logs as an audit record of model usage. Store all prompts and outputs in Amazon S3 with server-side encryption with AWS KMS keys (SSE-KMS).
- C. Restrict access to services by using VPC endpoint policies. Use AWS Config to track resource changes and compliance with security rules. Use server-side encryption with AWS KMS keys (SSE-KMS) to encrypt data at rest. Store the model's I/O in separate Amazon S3 buckets. Enable S3 server access logging to track file-level interactions.
- D. Define access control by using IAM with attribute-based access control (ABAC) to map departments to specific permissions. Configure VPC endpoints for Amazon Bedrock model API calls. Use IAM condition keys to enforce specific GuardrailIdentifier and ModelId values. Configure AWS CloudTrail to capture management and data events for S3 objects and KMS key usage activities. Enable S3 server access logging to record detailed file-level interactions with the video archives. Send all CloudTrail logs to AWS CloudTrail Lake. Set up Amazon CloudWatch alarms to detect and alert on unexpected activity from Amazon Bedrock, Amazon Rekognition, and AWS KMS.

**Answer: D**

Explanation:
Option B is the correct solution because it delivers end-to-end governance, security, and auditability across Amazon Bedrock, Amazon Rekognition, and the underlying data layer while meeting strict privacy and compliance requirements.
Using IAM attribute-based access control (ABAC) allows the company to control access to foundation models and data based on department, role, or workload attributes rather than static permissions. This is critical for controlling FM access patterns at scale.
Enforcing specific ModelId and GuardrailIdentifier values with IAM condition keys ensures that only approved models and guardrails are used, which directly supports secure model I/O and governance requirements.
Configuring VPC endpoints for Amazon Bedrock ensures that all model invocations remain on private AWS network paths, reducing data exfiltration risk and supporting privacy standards. AWS CloudTrail captures both management and data events, providing a definitive audit trail of who accessed which resources and when. Sending logs to CloudTrail Lake enables centralized, long-term, queryable auditing across services.
Amazon S3 server access logging adds file-level visibility into video archive access, which is essential for compliance and forensic analysis. Amazon CloudWatch alarms provide near real-time detection of anomalous or unauthorized activity across Amazon

Bedrock, Amazon Rekognition, and AWS KMS.
Option A focuses primarily on model-level tracing but lacks comprehensive IAM governance and S3 access auditing. Option C provides partial controls but lacks identity-aware auditing and model governance. Option D focuses on anomaly detection and classification but does not explicitly control FM access patterns.
Therefore, Option B best satisfies all stated requirements in a unified, auditable, and security-first architecture.

**NEW QUESTION # 53**
A company is using AWS Lambda and REST APIs to build a reasoning agent to automate support workflows.
The system must preserve memory across interactions, share relevant agent state, and support event-driven invocation and synchronous invocation. The system must also enforce access control and session-based permissions.
Which combination of steps provides the MOST scalable solution? (Select TWO.)

- A. Use Amazon Bedrock Agents for reasoning and conversation management. Use AWS Step Functions and Amazon SQS for orchestration. Store agent state in Amazon DynamoDB.
- B. Build a custom RAG pipeline by using Amazon Kendra and Amazon Bedrock. Use AWS Lambda to orchestrate tool invocations. Store agent state in Amazon S3.
- C. Use Amazon Bedrock AgentCore to manage memory and session-aware reasoning. Deploy the agent with built-in identity support, event handling, and observability.
- D. Deploy the reasoning logic as a container on Amazon ECS behind API Gateway. Use Amazon Aurora to store memory and identity data.
- E. Register the Lambda functions and REST APIs as actions by using Amazon API Gateway and Amazon EventBridge. Enable Amazon Bedrock AgentCore to invoke the Lambda functions and REST APIs without custom orchestration code.

**Answer: C,E**

Explanation:
The combination of Options A and B provides the most scalable and AWS-native architecture for building reasoning agents with persistent memory, session awareness, secure access control, and flexible invocation models.
Amazon Bedrock AgentCore is purpose-built to manage agent memory, session context, and identity-aware reasoning across interactions. It eliminates the need for developers to manually store and retrieve agent state, manage session lifecycles, or implement custom memory layers. AgentCore natively supports both synchronous requests and event-driven execution, making it ideal for support workflow automation.
Option B complements AgentCore by enabling seamless tool invocation. By registering AWS Lambda functions and REST APIs as agent actions through API Gateway and EventBridge, the agent can invoke tools reactively or synchronously without custom orchestration code. EventBridge enables event-driven execution, while API Gateway supports synchronous request-response patterns.
This combination provides built-in security, observability, and scaling, while avoiding the operational burden of managing queues, databases, or custom workflow engines.
Option C introduces unnecessary orchestration complexity. Option D increases infrastructure management and cost. Option E stores agent state in S3, which is not suitable for low-latency, session-based reasoning.
Therefore, A and B together deliver the most scalable, secure, and low-overhead solution for production- grade reasoning agents on AWS.

**NEW QUESTION # 54**
.....

AIP-C01 □✔□ and download it for free immediately on { www.examcollectionpass.com } □AIP-C01 Latest Guide Files

- Reliable AIP-C01 Test Sample □ AIP-C01 New Cram Materials □ Test AIP-C01 Price □ Search for { AIP-C01 } and easily obtain a free download on ➥ www.pdfvce.com □ □New AIP-C01 Test Vce Free
- AIP-C01 Latest Guide Files □ AIP-C01 Latest Exam Discount □ AIP-C01 Latest Exam Discount □ Open website 《 www.dumpsquestion.com 》 and search for ▸ AIP-C01 ◂ for free download □AIP-C01 Exam Prep
- AIP-C01 Exams Dumps □ Updated AIP-C01 Demo □ Latest AIP-C01 Test Pdf □ Search for □ AIP-C01 □ and download it for free on ➠ www.pdfvce.com □ website □Original AIP-C01 Questions
- AIP-C01 New Cram Materials ✡ AIP-C01 Exam Test □ AIP-C01 Exams Dumps □ Simply search for ➤ AIP-C01 □ for free download on （ www.torrentvce.com ） □AIP-C01 New Cram Materials
- Free PDF 2026 Amazon Useful AIP-C01: Key AWS Certified Generative AI Developer - Professional Concepts ↪ Easily obtain ⇛ AIP-C01 ⇚ for free download through [ www.pdfvce.com ] □AIP-C01 New Cram Materials
- 100% Pass Amazon - AIP-C01 –Trustable Key Concepts □ Open " www.prepawayete.com " and search for ⇛ AIP-C01 ⇚ to download exam materials for free ◂Reliable AIP-C01 Test Sample
- Latest AIP-C01 Practice Materials: AWS Certified Generative AI Developer - Professional offer you the most accurate Exam Questions - Pdfvce □ Immediately open ➤ www.pdfvce.com □ and search for ➡ AIP-C01 □□□ to obtain a free download □Test AIP-C01 Price
- AIP-C01 Latest Exam Discount □ Updated AIP-C01 Demo □ AIP-C01 Exams Dumps □ Search for ➽ AIP-C01 □ □ and download it for free on ▹ www.examcollectionpass.com ◃ website □Test AIP-C01 Pattern
- www.zazzle.com, k12.instructure.com, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, trialzone.characterzstore.com, www.educateonlinengr.com, courses.code-maze.com, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, 888.8337.net, telegra.ph, 91xiaojie.com, Disposable vapes

BTW, DOWNLOAD part of Itbraindumps AIP-C01 dumps from Cloud Storage: https://drive.google.com/open?id=1y-Ehp4MeQwaiZhoOJtqPL13i6eA6A_6X