


AIP-C01 PDF, AIP-C01 Musterprüfungsfragen



Sicherlich kennen Sie ZertPrüfung, weil es die Webseite mit höchster Bestehensrate für die Amazon AIP-C01 Zertifizierungsprüfung auf dem derzeitigen Markt ist. Sie können durch die Webseite ZertPrüfung ein paar kostenlosen Zertifizierungsantworten herunterladen und proben. Dann können Sie herausfinden, dass die Genauigkeit unserer Schulungsunterlagen zur Amazon AIP-C01 Zertifizierungsprüfung extrem hoch ist. Außerdem können Sie einjährige Aktualisierung genießen, nachdem Sie unsere Examsfragen gekauft haben.

Amazon AIP-C01 Prüfungsplan:

Thema	Einzelheiten
Thema 1	• 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.
Thema 2	• 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.
Thema 3	• 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.
Thema 4	• 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.

Thema 5	• AI Safety, Security, and Governance: This domain addresses input • output safety controls, data security and privacy protections, compliance mechanisms, and responsible AI principles including transparency and fairness.
---------	--

>> AIP-C01 PDF <<

AIP-C01 Mit Hilfe von uns können Sie bedeutendes Zertifikat der AIP-C01 einfach erhalten!

In den letzten Jahren ist die Konkurrenz in der IT-Branche immer heftiger geworden. IT-Zertifizierung ist ganz notwendig in der IT-Branche. Wenn Sie im Beruf eine gute Beförderungsmöglichkeit bekommen wollen, können Sie die Schulungsunterlagen zur Amazon AIP-C01 Zertifizierungsprüfung von ZertPruefung wählen, um die AIP-C01 Zertifizierungsprüfung zu bestehen. Unsere Schulungsunterlagen enthalten alle Fragen, die die Amazon AIP-C01 Zertifizierungsprüfung erfordert. So können Sie die AIP-C01 Prüfung 100% bestehen.

Amazon AWS Certified Generative AI Developer - Professional AIP-C01 Prüfungsfragen mit Lösungen (Q74-Q79):

74. Frage

A financial services company is deploying a generative AI (GenAI) application that uses Amazon Bedrock to assist customer service representatives to provide personalized investment advice to customers. The company must implement a comprehensive governance solution that follows responsible AI practices and meets regulatory requirements.

The solution must detect and prevent hallucinations in recommendations. The solution must have safety controls for customer interactions. The solution must also monitor model behavior drift in real time and maintain audit trails of all prompt-response pairs for regulatory review. The company must deploy the solution within 60 days. The solution must integrate with the company's existing compliance dashboard and respond to customers within 200 ms.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Deploy Amazon Bedrock and use AWS PrivateLink to access the application securely. Use AWS Lambda functions to implement custom prompt validation. Store prompt-response pairs in an Amazon S3 bucket and configure S3 Lifecycle policies. Create custom Amazon CloudWatch dashboards to monitor model performance metrics.
- B. Configure Amazon Bedrock guardrails to apply custom content filters and toxicity detection. Use Amazon Bedrock Model Evaluation to detect hallucinations. Store prompt-response pairs in Amazon DynamoDB to capture audit trails and set a TTL. Integrate Amazon CloudWatch custom metrics with the existing compliance dashboard.
- C. Use Amazon SageMaker Model Monitor to detect model behavior drift. Use AWS WAF to filter content. Store customer interactions in an encrypted Amazon RDS database. Use Amazon API Gateway to create custom HTTP APIs to integrate with the compliance dashboard.
- D. Use Amazon Bedrock Agents and Amazon Bedrock Knowledge Bases to ground responses. Use Amazon Bedrock Guardrails to enforce content safety. Use Amazon OpenSearch Service to store and index prompt-response pairs. Integrate OpenSearch Service with Amazon QuickSight to create compliance reports and to detect model behavior drift.

Antwort: B

Begründung:

Option A is the correct solution because it uses native Amazon Bedrock governance and evaluation capabilities to meet regulatory, performance, and deployment timeline requirements with the least operational overhead.

Amazon Bedrock guardrails provide built-in safety controls that enforce responsible AI policies directly during inference. Custom content filters and toxicity detection protect customer interactions and prevent disallowed investment guidance patterns without requiring custom application logic. Guardrails operate inline and are optimized for low latency, which helps meet the strict 200 ms response-time requirement.

Hallucination detection is addressed through Amazon Bedrock Model Evaluation, which supports automated evaluation at scale using LLM-as-a-judge techniques. This enables the company to detect factual inaccuracies and policy violations systematically, without building custom evaluation pipelines or requiring extensive human review. Evaluation outputs can be surfaced as metrics. Storing all prompt-response pairs in Amazon DynamoDB provides a low-latency, highly scalable audit store that aligns with financial regulatory requirements. Using TTL enforces data retention policies automatically, reducing compliance risk and storage overhead. Amazon CloudWatch custom metrics integrate seamlessly with existing compliance dashboards, allowing near-real-time monitoring of safety interventions, hallucination rates, and drift indicators. CloudWatch anomaly detection can be applied to these metrics to

surface behavior changes quickly.

Option B relies on custom Lambda logic and S3-based auditing, increasing latency and operational complexity. Option C introduces additional services that increase setup time and may exceed the 60-day deployment window. Option D uses non-Bedrock-native monitoring and adds unnecessary infrastructure layers.

Therefore, Option A provides the most complete, compliant, and low-overhead governance solution for a regulated GenAI financial services application.

75. Frage

A company is using Amazon Bedrock to build a customer-facing AI assistant that handles sensitive customer inquiries. The company must use defense-in-depth safety controls to block sophisticated prompt injection attacks. The company must keep audit logs of all safety interventions. The AI assistant must have cross-Region failover capabilities.

Which solution will meet these requirements?

- A. Configure Amazon Bedrock guardrails with content filters set to high. Use AWS WAF to block suspicious inputs. Use AWS CloudTrail to log API calls.
- **B. Configure Amazon Bedrock guardrails with content filters set to high to protect against prompt injection attacks. Use a guardrail profile to implement cross-Region guardrail inference. Use Amazon CloudWatch Logs with custom metrics to capture detailed guardrail intervention events.**
- C. Deploy Amazon Comprehend custom classifiers to detect prompt injection attacks. Use Amazon API Gateway request validation. Use CloudWatch Logs to capture intervention events.
- D. Configure Amazon Bedrock guardrails with custom content filters and word filters set to high. Configure cross-Region guardrail replication for failover. Store logs in AWS CloudTrail for compliance auditing.

Antwort: B

Begründung:

Option A provides the most complete, AWS-native defense-in-depth solution for protecting against prompt injection attacks while meeting audit and resiliency requirements. Amazon Bedrock guardrails are designed specifically to enforce safety policies on both user inputs and model outputs, including protections against prompt injection and jailbreak attempts.

Setting content filters to high increases sensitivity to malicious or manipulative inputs. Guardrail profiles allow the same guardrail configuration to be applied consistently across multiple Regions, enabling cross-Region inference and failover without configuration drift. This directly satisfies the requirement for regional resilience.

Amazon CloudWatch Logs captures detailed guardrail intervention events, including when content is blocked, modified, or flagged. Custom metrics derived from these logs enable fine-grained auditing, alerting, and reporting on safety enforcement actions. This provides a more detailed audit trail of safety interventions than API-level logs alone.

Option B adds WAF protection but lacks detailed guardrail intervention logging. Option C introduces additional services and custom logic that increase complexity and may miss model-specific injection patterns.

Option D references replication concepts that are not aligned with Bedrock guardrail operational models and relies on word filters, which are insufficient against sophisticated prompt injection techniques.

Therefore, Option A best meets the requirements for layered protection, auditability, and cross-Region resilience using managed Amazon Bedrock safety controls.

76. Frage

A company is building a generative AI (GenAI) application that processes financial reports and provides summaries for analysts. The application must run two compute environments. In one environment, AWS Lambda functions must use the Python SDK to analyze reports on demand. In the second environment, Amazon EKS containers must use the JavaScript SDK to batch process multiple reports on a schedule. The application must maintain conversational context throughout multi-turn interactions, use the same foundation model (FM) across environments, and ensure consistent authentication.

Which solution will meet these requirements?

- A. Create a centralized Amazon API Gateway REST API endpoint that handles all model interactions by using the InvokeModel API. Store interaction history in application process memory in each Lambda function or EKS container. Use environment variables to configure model parameters.
- B. Use the Amazon Bedrock InvokeModel API with a separate authentication method for each environment. Store conversation states in Amazon DynamoDB. Use custom I/O formatting logic for each programming language.
- **C. Use the Amazon Bedrock Converse API and IAM roles for authentication. Pass previous messages in the request messages array to maintain conversational context. Use programming language-specific SDKs to establish consistent API interfaces.**

- D. Use the Amazon Bedrock Converse API directly in both environments with a common authentication mechanism that uses IAM roles. Store conversation states in Amazon ElastiCache. Create programming language-specific wrappers for model parameters.

Antwort: C

Begründung:

Option D is the correct solution because the Amazon Bedrock Converse API is purpose-built for multi-turn conversational interactions and is designed to work consistently across SDKs and compute environments. The Converse API standardizes how messages, roles, and context are represented, which ensures consistent behavior whether the application is running in AWS Lambda with Python or in Amazon EKS with JavaScript.

By passing previous messages in the messages array, the application explicitly maintains conversational context across turns without relying on external state stores. This approach is recommended by AWS for conversational GenAI workflows because it avoids state synchronization complexity and ensures deterministic model behavior across environments.

Using IAM roles for authentication provides a single, consistent security model for both Lambda and EKS.

IAM roles integrate natively with AWS SDKs, eliminating the need for custom authentication logic or environment-specific credentials. This aligns with AWS best practices for least privilege and simplifies governance.

Option A introduces inconsistent authentication and custom formatting logic, increasing complexity. Option B unnecessarily introduces ElastiCache for state management, which is not required when using the Converse API correctly. Option C stores state in process memory, which is unsafe and unreliable for serverless and containerized workloads.

Therefore, Option D best satisfies the requirements for conversational consistency, multi-environment support, shared model usage, and consistent authentication with minimal operational overhead.

77. Frage

A company wants to select a new FM for its AI assistant. A GenAI developer needs to generate evaluation reports to help a data scientist assess the quality and safety of various foundation models FMs. The data scientist provides the GenAI developer with sample prompts for evaluation. The GenAI developer wants to use Amazon Bedrock to automate report generation and evaluation. Which solution will meet this requirement?

- A. Combine the sample prompts into a single JSON document. Create an Amazon Bedrock knowledge base from the document. Create an Amazon Bedrock evaluation job that uses the retrieval and response generation evaluation type. Specify an Amazon S3 bucket as the output. Run an evaluation job for each FM.
- B. Combine the sample prompts into a single JSONL document. Store the document in an Amazon S3 bucket. Create an Amazon Bedrock evaluation job that uses a judge model. Specify the S3 location as input and a different S3 location as output. Run an evaluation job for each FM and select the FM as the generator.
- C. Combine the sample prompts into a single JSONL document. Store the document in an Amazon S3 bucket. Create an Amazon Bedrock evaluation job that uses a judge model. Specify the S3 location as input and Amazon QuickSight as output. Run an evaluation job for each FM and select the FM as the evaluator.
- D. Combine the sample prompts into a single JSON document. Create an Amazon Bedrock knowledge base with the document. Write a prompt that asks the FM to generate a response to each sample prompt. Use the RetrieveAndGenerate API to generate a report for each model.

Antwort: B

Begründung:

Option B is correct because it uses the managed evaluation capability in Amazon Bedrock that is intended specifically for comparing foundation models using a consistent prompt set and producing structured results with minimal custom tooling. In a Bedrock evaluation workflow, you provide an input dataset of prompts, typically in JSON Lines format so each line represents one evaluation record. Storing the JSONL file in Amazon S3 allows Bedrock to read the dataset at scale and write standardized evaluation outputs back to S3 for downstream analysis, sharing, and retention.

The key requirement is to assess both quality and safety across multiple models. A Bedrock evaluation job can use a judge model to score the generated outputs against defined criteria. This approach supports repeatable, apples-to-apples comparisons because the same judge model and scoring rubric can be applied to every candidate foundation model. The candidate models are configured as generators, meaning each evaluation job run uses one selected FM to produce answers for the same prompt set, and the judge model evaluates those answers. That matches the requirement to generate evaluation reports that help a data scientist select the best FM.

Option A does not use Bedrock evaluation jobs, and a knowledge base plus RetrieveAndGenerate is a RAG pattern, not an evaluation framework. It would produce responses but not standardized scoring and reporting suitable for model selection. Option C is incorrect because Bedrock evaluation outputs are delivered to S3, not directly to a BI destination, and selecting the candidate FM as the evaluator conflicts with the intended pattern of using a stable judge model. Option D misuses knowledge bases and retrieval

evaluation types when the requirement is prompt-based model assessment rather than evaluating retrieval quality.

78. Frage

A company is using Amazon Bedrock to develop an AI-powered application that uses a foundation model that supports cross-Region inference and provisioned throughput. The application must serve users in Europe and North America with consistently low latency. The application must comply with data residency regulations that require European user data to remain within Europe-based AWS Regions.

During testing, the application experiences service degradation when Regional traffic spikes reach service quotas. The company needs a solution that maintains application resilience and minimizes operational complexity.

Which solution will meet these requirements?

- A. Use Amazon Bedrock cross-Region inference profiles by specifying geographical codes in profile IDs when the application calls the InvokeModel API. Configure separate Amazon API Gateway HTTP APIs to direct European and North American users to the appropriate Regional endpoints.
- B. Configure provisioned throughput for Amazon Bedrock in multiple Regions. Implement failover logic in the application code to switch between Regions when throttling occurs. Use AWS Global Accelerator to route traffic to the appropriate endpoints based on user location.
- C. Deploy separate Amazon Bedrock instances in North American and European Regions. Use a custom routing layer that directs traffic based on user location. Configure Amazon CloudWatch alarms to monitor Regional service usage. Use Amazon SNS to send email alerts to the company when usage approaches specified thresholds.
- D. Deploy a multi-Region Amazon API Gateway HTTP API and AWS Lambda functions that implement retry logic to handle throttling. Configure the Lambda functions to call the foundation model in the nearest secondary Region when the application reaches service quotas in the primary Region. Use intelligent routing to ensure compliance with data residency requirements.

Antwort: A

Begründung:

Option B best meets the latency, resilience, and data residency requirements while keeping operational complexity low by using built-in Amazon Bedrock cross-Region inference behavior through inference profiles. Cross-Region inference profiles are designed to provide higher availability and better traffic absorption when a single Region experiences throttling, transient capacity constraints, or quota-related degradation. By selecting the appropriate geography-scoped inference profile (for example, a Europe-scoped profile for European users and a North America-scoped profile for North American users), the application can keep inference traffic within the required geographic boundary. This directly supports EU data residency needs because European requests can be served only by Europe-based Regions while still benefiting from multi-Region resilience inside Europe.

The question also highlights degradation when Regional traffic spikes hit quotas. Cross-Region inference profiles help mitigate these conditions by allowing Bedrock to serve requests from another Region within the same geography, improving continuity during spikes without requiring the company to implement custom retry-and-failover logic across Regions. This reduces development and operational burden compared to building and maintaining a bespoke routing and fallback system.

Using separate Amazon API Gateway HTTP APIs to direct European and North American users to the correct endpoints simplifies request routing and provides a clean boundary for compliance controls, logging, and monitoring. It also allows each geography to scale independently and maintain consistently low latency by keeping users close to the entry point and the Bedrock geography they must use.

Option A requires custom routing and manual operational monitoring and does not inherently solve quota-driven degradation.

Option C adds significant complexity by embedding throttling retries and cross-Region selection logic in Lambda while still needing careful controls to prevent cross-border routing mistakes. Option D introduces the highest operational complexity and can inadvertently violate residency if failover crosses geographies unless additional safeguards are implemented.

79. Frage

.....

Wenn Sie die Schulungsunterlagen zur Amazon AIP-C01 Zertifizierungsprüfung haben, dann werden Sie sicherlich erfolgreich sein. Nachdem Sie unsere Lehrbücher gekauft haben, werden Sie einjährige Aktualisierung kostenlos genießen. Die Bestehensrate von Amazon AIP-C01 ist 100%. Wenn Sie die Zertifizierungsprüfung nicht bestehen oder die Schulungsunterlagen zur Amazon AIP-C01 Zertifizierungsprüfung irgend ein Problem haben, geben wir Ihnen eine bedingungslose volle Rückerstattung.

AIP-C01 Musterprüfungsfragen: https://www.zertpruefung.ch/AIP-C01_exam.html

- AIP-C01 Schulungsunterlagen ☐ AIP-C01 Zertifizierungsfragen ☐ AIP-C01 Prüfungsfrage ☐ Sie müssen nur zu {

[illegible]