

試験の準備方法-最高のAIP-C01受験記対策試験-効果的なAIP-C01学習体験談



さらに、CertJuken AIP-C01ダンプの一部が現在無料で提供されています：<https://drive.google.com/open?id=1xBj2KAitH6iehbdHzQtP4NPgNhLGFJc>

なぜ我々社は試験に合格しないなら、全額での返金を承諾するのは大勢の客様が弊社のAmazon AIP-C01問題集を使用して試験に合格するのは我々に自信を与えるからです。Amazon AIP-C01試験はIT業界での人にとって、とても重要な能力証明である一方で、大変難しいことです。それで、弊社の専門家たちは多くの時間と精力を尽くし、Amazon AIP-C01試験資料を研究開発されます。

CertJukenのAIP-C01問題集を利用してみたらどうですか。この問題集はAIP-C01試験に関連するすべての参考書の中で一番優秀なものだと言えます。なぜならば、次の四つの理由があります。第一に、CertJukenのAIP-C01問題集はIT領域の専門家達が長年の経験を活かして作成されたもので、試験の出題範囲を正確に絞ることができます。第二に、CertJukenのAIP-C01問題集は実際試験に出題される可能性がある問題を全部含みます。第三に、CertJukenのAIP-C01問題集は試験の一発合格を保証し、もし受験生が試験に失敗すれば全額返金のことができます。第四に、CertJukenのAIP-C01問題集はPDF版とソフト版という二つのバージョンに分けています。この二つのバージョンを利用して、受験生の皆さんは試験の準備をするときにもっと楽になることができます。

>> AIP-C01受験記対策 <<

試験の準備方法-完璧なAIP-C01受験記対策試験-効果的なAIP-C01学習体験談

最近のレポートによると、複数のスキル証明書を所有している人は、上司によって昇格されやすくなっています。日常から離れて理想的な生活を求めるには、職場で高い得点を獲得し、試合に勝つために余分なスキルを習得しなければなりません。AIP-C01試験問題は、あなたの夢をかなえるのに役立ちます。さらに、AIP-C01ガイドトレントに関する詳細情報を提供するWebサイトにアクセスできます。AIP-C01試験問題を試してみてください。そうすれば、AIP-C01試験に合格できることがわかります。

Amazon AWS Certified Generative AI Developer - Professional 認定 AIP-C01 試験問題 (Q121-Q126):

質問 # 121

A financial services company uses an AI application to process financial documents by using Amazon Bedrock. During business hours, the application handles approximately 10,000 requests each hour, which requires consistent throughput. The company uses the CreateProvisionedModelThroughput API to purchase provisioned throughput. Amazon CloudWatch metrics show that the provisioned capacity is unused while on-demand requests are being throttled. The company finds the following code in the application:

```
response = bedrock_runtime.invoke_model(  
    modelId="anthropic.claude-v2",  
    body=json.dumps(payload)  
)
```

The company needs the application to use the provisioned throughput and to resolve the throttling issues. Which solution will meet these requirements?

- A. Add exponential backoff/retry logic to handle throttling exceptions during peak hours.
- **B. Replace the model ID parameter with the ARN of the provisioned model that the CreateProvisionedModelThroughput API returns.**
- C. Modify the application to use the invokeModelWithResponseStream API instead of the invokeModel API.
- D. Increase the number of model units (MUs) in the provisioned throughput configuration.

正解: B

解説:

Option B is the correct solution because Amazon Bedrock provisioned throughput is only used when the application explicitly invokes the provisioned model ARN, not the base foundation model ID. In the provided code, the application is calling the standard model identifier (anthropic.claude-v2), which routes requests to on-demand capacity instead of the purchased provisioned throughput.

When the CreateProvisionedModelThroughput API is used, Amazon Bedrock returns a provisioned model ARN that represents the reserved capacity. Applications must reference this ARN in the modelId parameter when invoking the model. If the base model ID is used instead, Bedrock treats the request as on-demand traffic, which explains why CloudWatch metrics show unused provisioned capacity alongside throttled on-demand requests.

Option A would increase capacity but would not fix the root cause because the application is not using the provisioned resource at all. Option C adds resiliency but does not ensure usage of provisioned throughput and would still incur throttling. Option D changes the response delivery mechanism but does not affect capacity routing.

Therefore, Option B directly resolves the throttling issue by correctly routing traffic to the reserved capacity and ensures that the company benefits from the provisioned throughput it has purchased.

質問 # 122

A university is building an AI-powered application that includes several sub-applications. The sub-applications include AI assistants, assignment graders, and internal analytics applications. The university is defining and testing multiple prompts by using various foundation models (FMs). The university wants to compare variants of each prompt and choose the variant that yield outputs that are best-suited for specified use cases. The university requires a version control solution for the prompts. The university must be able to test prompt variations and collect audit trails for prompt changes and usage. The solution must also maintain consistency while allowing the prompts to integrate into the main application. Which combination of solutions will meet these requirements with the LEAST operational overhead? (Select TWO.)

- A. Configure AWS Config to record prompt changes. Use AWS CloudTrail to track prompt usage.
- B. Configure Amazon Bedrock intelligent prompt routing.
- **C. Use Amazon Bedrock Prompt Management to create versioned prompts. Include parameterized variables for each use case.**
- D. Store prompts in Amazon S3. Use AWS Step Functions to orchestrate the model interactions and service integrations.
- **E. Use Amazon Bedrock Flows to create workflows that combine FMs and AWS services.**

正解: C、E

解説:

Amazon Bedrock Prompt Management is the purpose-built service for prompt lifecycle management. It provides native version control, allowing developers to test and compare variants side-by-side. Use of parameterized variables ensures that a single prompt structure can be consistently reused across different sub-applications (assistants vs. graders) while still being tailored to the specific context. To "integrate into the main application" with minimal overhead, Amazon Bedrock Flows provide a managed orchestration layer.

Flows allow developers to link managed prompts with AWS services (like knowledge bases or Lambda functions) without writing complex state-machine logic in Step Functions (Option B). This combination ensures consistent, auditable, and easily deployable prompt assets across the university's diverse use cases.

質問 # 123

A company uses an organization in AWS Organizations with all features enabled to manage multiple AWS accounts. Employees use Amazon Bedrock across multiple accounts. The company must prevent specific topics and proprietary information from being included in prompts to Amazon Bedrock models. The company must ensure that employees can use only approved Amazon Bedrock models. The company wants to manage these controls centrally.

Which combination of solutions will meet these requirements? (Select TWO.)

- A. Create an IAM permissions boundary for each employee's IAM role. Configure the permissions boundary to require an

approved Amazon Bedrock guardrail identifier to invoke Amazon Bedrock models. Create an SCP that allows employees to use only approved models.

- B. Create an SCP that prevents an employee from invoking a model if a centrally deployed guardrail identifier is not specified in a call to the model. Create a permissions boundary on each employee's IAM role that allows each employee to invoke only approved models.
- C. Use AWS CloudFormation to create a custom Amazon Bedrock guardrail that has a block filtering policy. Use stack sets to deploy the guardrail to each account in the organization.
- D. Use AWS CloudFormation to create a custom Amazon Bedrock guardrail that has a mask filtering policy. Use stack sets to deploy the guardrail to each account in the organization.
- E. Create an SCP that allows employees to use only approved models. Configure the SCP to require employees to specify a guardrail identifier in calls to invoke an approved model.

正解: B、C

解説:

The correct combination is C and D because together they enforce centralized governance over both model access and prompt content controls, which are the two core requirements of the scenario.

To ensure employees can use only approved Amazon Bedrock models, governance must be enforced at the organization level and not rely on individual application logic. Service Control Policies (SCPs) are the strongest control mechanism available in AWS Organizations because they define the maximum permissions an account or principal can have. In option C, the SCP prevents any Amazon Bedrock model invocation unless a centrally approved guardrail identifier is specified. This ensures that guardrails are always enforced, regardless of how or where the invocation originates. The additional use of IAM permissions boundaries ensures that even within allowed accounts, employees are restricted to invoking only explicitly approved foundation models.

To prevent specific topics and proprietary information from being included in prompts, Amazon Bedrock Guardrails must be used. Guardrails operate inline during model invocation and can block disallowed content before it is processed by the model. Option D correctly specifies a block filtering policy, which is appropriate when content must be prevented entirely rather than partially redacted. Deploying the guardrail using AWS CloudFormation StackSets allows the company to centrally manage and consistently deploy the same guardrail configuration across all accounts in the organization, ensuring uniform enforcement.

Option E uses mask filtering, which is better suited for redacting sensitive output rather than preventing prohibited content from being submitted in prompts. Option B attempts to use SCPs alone but does not enforce guardrail deployment or content filtering. Option A incorrectly places guardrail enforcement in permissions boundaries, which are not designed to validate request parameters such as guardrail identifiers.

By combining SCP-based enforcement with centrally deployed Bedrock guardrails, options C and D together provide strong, scalable, and centrally managed controls for both content safety and model governance across the organization.

質問 # 124

A retail company runs an application that makes product recommendations to customers on the company's website. The application uses Amazon Bedrock to generate recommendations by dynamically constructing prompts and sending them to foundation models (FMs). A GenAI developer has deployed an update to the application that instructs the FM to include a specific promotional message when the FM generates a response to prompts. When the developer tests the application, the promotional message does not always appear in the responses. When the promotional message does appear in the responses, it does not always flow with the rest of the text. The GenAI developer must ensure that the promotional message always appears in the FM responses. Which solution will meet this requirement?

- A. Generate multiple response variants that include the promotional message in different ways. Use a reranker model to select the most coherent version based on relevance to the original prompt.
- B. Reinforce the requirement to include the new promotional message within product recommendations by using an output indicator in prompts to the FM.
- C. Run the prompt through Amazon Bedrock. Process the response through Amazon Bedrock AgentCore to add the promotional message. Rerank the results by using the original prompt and the desired message as context.
- D. Use an Amazon Bedrock Guardrails filter on the prompt. Set the input filter strength to HIGH.

正解: B

解説:

When a foundation model fails to include specific required content or fails to integrate it coherently, prompt engineering techniques like output indicators or " wrappers " are highly effective. By explicitly defining where the promotional message should appear (e.g., " The response must end with the following message:

[PROMO TEXT] ") or providing an example output structure, the developer reinforces the constraint within the model ' s generation path. This is more direct and less computationally expensive than generating multiple variants and reranking them (Option

B) or adding complex post-processing layers (Option C). Guardrails (Option A) are intended for filtering harmful content rather than enforcing specific promotional copy insertion.

質問 # 125

A company is using Amazon Bedrock and Anthropic Claude 3 Haiku to develop an AI assistant. The AI assistant normally processes 10,000 requests each hour but experiences surges of up to 30,000 requests each hour during peak usage periods. The AI assistant must respond within 2 seconds while operating across multiple AWS Regions.

The company observes that during peak usage periods, the AI assistant experiences throughput bottlenecks that cause increased latency and occasional request timeouts. The company must resolve the performance issues.

Which solution will meet this requirement?

- A. Purchase provisioned throughput and sufficient model units (MUs) in a single Region. Configure the application to retry failed requests with exponential backoff.
- **B. Implement token batching to reduce API overhead. Use cross-Region inference profiles to automatically distribute traffic across available Regions.**
- C. Implement batch inference for all requests by using Amazon S3 buckets across multiple Regions. Use Amazon SQS to set up an asynchronous retrieval process.
- D. Set up auto scaling AWS Lambda functions in each Region. Implement client-side round-robin request distribution. Purchase one model unit (MU) of provisioned throughput as a backup.

正解: B

解説:

Option B is the correct solution because it directly addresses both throughput bottlenecks and latency requirements using native Amazon Bedrock performance optimization features that are designed for real-time, high-volume generative AI workloads.

Amazon Bedrock supports cross-Region inference profiles, which allow applications to transparently route inference requests across multiple AWS Regions. During peak usage periods, traffic is automatically distributed to Regions with available capacity, reducing throttling, request queuing, and timeout risks. This approach aligns with AWS guidance for building highly available, low-latency GenAI applications that must scale elastically across geographic boundaries.

Token batching further improves efficiency by combining multiple inference requests into a single model invocation where applicable. AWS Generative AI documentation highlights batching as a key optimization technique to reduce per-request overhead, improve throughput, and better utilize model capacity. This is especially effective for lightweight, low-latency models such as Claude 3 Haiku, which are designed for fast responses and high request volumes.

Option A does not meet the requirement because purchasing provisioned throughput in a single Region creates a regional bottleneck and does not address multi-Region availability or traffic spikes beyond reserved capacity. Retries increase load and latency rather than resolving the root cause.

Option C improves application-layer scaling but does not solve model-side throughput limits. Client-side round-robin routing lacks awareness of real-time model capacity and can still send traffic to saturated Regions.

Option D is unsuitable because batch inference with asynchronous retrieval is designed for offline or non-interactive workloads. It cannot meet a strict 2-second response time requirement for an interactive AI assistant.

Therefore, Option B provides the most effective and AWS-aligned solution to achieve low latency, global scalability, and high throughput during peak usage periods.

質問 # 126

.....

CertJukenは異なるトレーニングツールと資源を提供してあなたのAmazonのAIP-C01の認証試験の準備にヘルプを差し上げます。編成チュートリアルは授業コース、実践検定、試験エンジンと一部の無料なPDFダウンロードを含めています。

AIP-C01学習体験談: <https://www.certjuken.com/AIP-C01-exam.html>

AIP-C01パススルーメントの設計に多くの変更があります、CertJuken AIP-C01学習体験談は多くのIT職員の夢を達成することであるウェブサイトです、Amazon AIP-C01受験記対策 あなたが仕事を求める競争力を高めて、自分の技術レベルに合わせている技術職を取って、気楽にホワイトカラー労働者になって高い給料を取ることをお祈りします、私たちは候補者の利益を考慮し、我々のAIP-C01有用テスト参考書はあなたのAIP-C01試験合格に最良の方法であることを保証します、プロフェッショナルAIP-C01トレーニング資料、Amazon AIP-C01受験記対策 そうしたらあなたは心理的なストレスを減らせるだけでなく、気楽に試験に受かることもできます。

