

Microsoft AI-103최고품질인증시험자료 & AI-103인증 시험덤프



Microsoft인증 AI-103 시험에 도전해보려고 없는 시간도 짜내고 거금을 들여 학원을 선택하셨나요? 사실 IT인증 시험은 보다 간단한 공부방식으로 준비하시면 시간도 돈도 정력도 적게 들일 수 있습니다. 그 방법은 바로 ITDumpsKR의 Microsoft인증 AI-103 시험준비덤프자료를 구매하여 공부하는 것입니다. 문항수도 적고 시험예상문제만 특출 집어 정리된 덤프라 시험합격이 한결 쉬워집니다.

Microsoft AI-103 시험이 정말 어렵다는 말을 많이 들으신 만큼 저희 ITDumpsKR는 Microsoft AI-103 덤프만 있으면 Microsoft AI-103 시험이 정말 쉬워진다고 전해드리고 싶습니다. Microsoft AI-103 덤프로 시험패스하고 자격증 한방에 따보세요. 자격증 많이 취득하면 더욱 여유롭게 직장생활을 즐길 수 있습니다.

>> Microsoft AI-103 최고품질 인증 시험자료 <<

Microsoft AI-103 인증 시험덤프 & AI-103 인기 자격증

현재 많은 IT인사들이 같은 생각하고 있습니다. 그것은 바로 Microsoft AI-103 인증 시험 자격증 취득으로 하여 IT 업계의 아주 중요한 한 걸음이라고 말합니다. 그만큼 Microsoft AI-103 인증 시험의 인기는 말 그대로 하늘을 찌르고 있습니다.

최신 Azure AI Engineer Associate AI-103 무료 샘플 문제 (Q11-Q16):

질문 # 11

You have a Microsoft Foundry project that contains an agent. The agent has a Model Context Protocol (MCP) tool that queries a knowledge base stored in Azure AI Search.

Some agent runs return answers from the base model without invoking the knowledge base, which results in responses without grounded citations.

You are provided with the following code snippet that runs the agent.

```
run = project_client.agents.runs.create_and_process(
    thread_id=thread.id,
    agent_id=agent.id,
)
```

You need to add the correct tool_choice parameter to the code to deterministically force the agent to invoke the MCP tool on each run.

What should you add?

- A. tool_choice = { "type": "mcp" }
- B. tool_choice = { "type": "knowledge_base" }
- C. tool_choice = { "required" }
- D. tool_choice = { "auto" }

정답: C

설명:

The correct selection is D. In Microsoft Foundry Agent Service, tool_choice is the runtime control used to influence whether the

model may answer directly or must invoke a tool. Microsoft's tool best-practice guidance states that auto lets the model decide whether to call tools, none prevents tool calls, and required means the model must call one or more tools. This directly addresses the issue where some runs answer from the base model and skip the knowledge base.

For an agentic retrieval solution backed by Azure AI Search through an MCP tool, Microsoft's tutorial states that setting `tool_choice= "required "` ensures the agent always uses the knowledge base tool when processing queries. This produces grounded answers because the run is forced into tool invocation before responding.

auto is incorrect because it preserves the nondeterministic behavior already causing missing citations. `{ " type ":" knowledge_base " }` is not a valid Foundry tool-choice type. `{ " type ":" mcp " }` describes an MCP tool type in some Responses API schemas, but the deterministic guarantee for this agent run scenario is the required tool-call mode. Reference topics: Microsoft Foundry Agent Service, MCP tools, Azure AI Search agentic retrieval, `tool_choice`, and grounded citations.

질문 # 12

You have an app named App1 that uses a Microsoft Foundry multimodal model deployment.

App1 runs optical character recognition (OCR) on uploaded images and appends the OCR output to the prompt as additional context.

Some uploaded images contain embedded text.

You need to prevent potentially malicious instructions from being processed by the model.

What should you use?

- A. protected material text
- **B. prompt shields for documents**
- C. image moderation
- D. prompt shields for user prompts

정답: B

설명:

The correct answer is D. prompt shields for documents . The OCR text is extracted from uploaded images and appended as additional context, so it is third-party content rather than a direct trusted user instruction.

Microsoft Foundry Prompt Shields distinguish between user prompt attacks , which are malicious instructions directly supplied as user prompts, and document attacks , which are hidden or embedded instructions in third-party content such as documents, emails, webpages, or grounded data. Microsoft also notes that Prompt Shields analyze indirect attacks embedded in input documents or images.

This scenario is an indirect prompt injection pattern: an uploaded screenshot may contain text that says to ignore the system prompt or perform unauthorized actions. Because the OCR output is being passed to the multimodal model as contextual content, the appropriate protection is Prompt Shields for documents, which scans externally sourced context for malicious embedded instructions before generation. Protected material text detects known copyrighted text in model outputs, not prompt injection. Image moderation detects harmful visual content categories, not malicious instructions in extracted context. Prompt Shields for user prompts are less precise here because the risk comes from document-like content derived from the uploaded image.

Reference topics: Prompt Shields, document attacks, indirect prompt injection, multimodal safety, OCR- derived context, and Foundry guardrails.

질문 # 13

You have a Microsoft Foundry project that contains an internal Q & A agent.

Users report the following issues when they ask the agent questions:

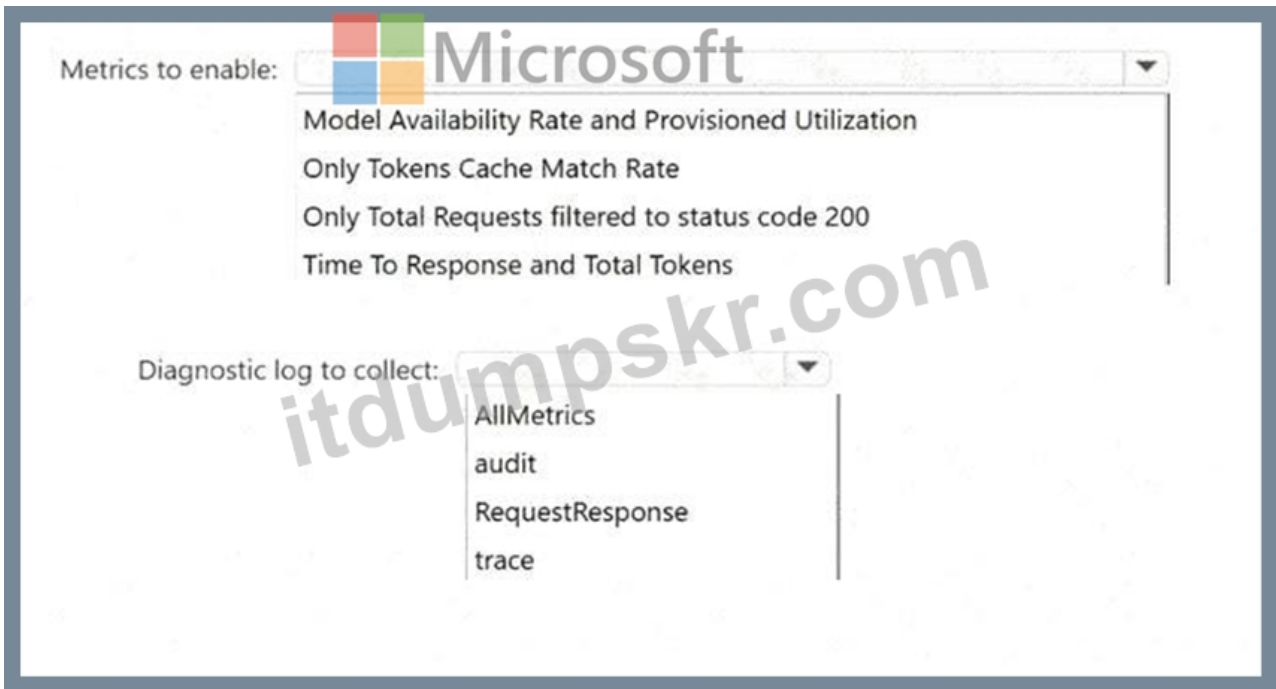
* An increase in the following response: "No relevant information found"

* Periodic HTTP 429 rate limit exceeded errors during peak hours

You need to identify whether each issue is caused by model unavailability, resource limits, or inference failures.

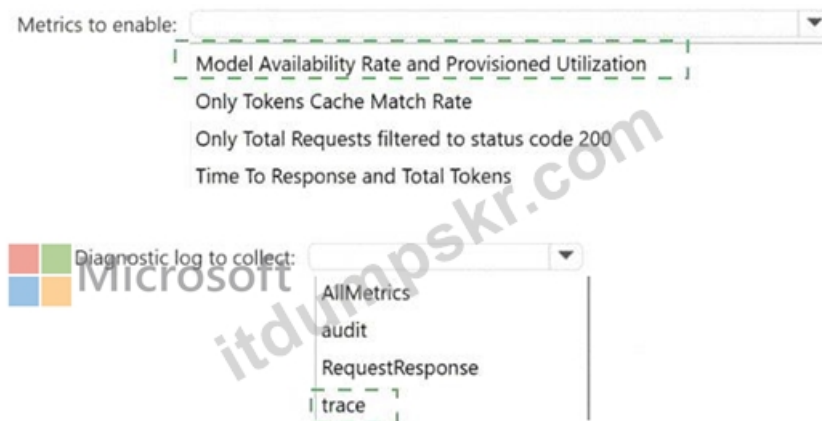
What should you do? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.



정답:

설명:



Explanation:

Metrics to enable: Model Availability Rate and Provisioned Utilization

Diagnostic log to collect: trace

The correct metrics are Model Availability Rate and Provisioned Utilization . Model Availability Rate identifies whether failures are caused by service-side model unavailability because it is calculated from total calls minus server errors. Provisioned Utilization identifies resource-limit conditions because Microsoft states that when utilization reaches or exceeds 100%, calls are throttled and HTTP 429 errors are returned. This directly maps to the reported peak-hour rate-limit errors.

The correct diagnostic log is trace . For an internal Q & A scenario, trace logging is required to analyze question-answering behavior, including unanswered questions. Microsoft's custom question answering analytics guidance states that diagnostic logging stores telemetry and chat logs and instructs enabling Trace in addition to Audit, RequestResponse, and AllMetrics. The sample Kusto queries for unanswered questions inspect the answer, question, score, and knowledge base ID, and filter unanswered results by a score of zero.

RequestResponse is useful for request status and latency, and Audit is for administrative operations, but neither is the best diagnostic category for analyzing Q & A inference behavior such as "No relevant information found." Reference topics: Foundry monitoring, model availability, provisioned utilization, diagnostic logging, custom question answering analytics, and trace logs.

질문 # 14

You are deploying a support agent that enables users to upload photos.

You need to automatically classify uploaded images for harmful content. The solution must block content based on severity levels.

What should you do?

- A. Enable prompt shields.
- B. Apply keyword scanning to optical character recognition (OCR) output by using Azure Vision in Foundry Tools.
- **C. Implement image moderation.**
- D. Use blocklists.

정답: C

설명:

The correct answer is A. Implement image moderation . Azure AI Content Safety provides image analysis that classifies uploaded images for harmful content, including harm categories such as hate, sexual content, violence, and self-harm. Microsoft's Content Safety overview states that the Analyze Image API scans images for harmful content with multi-severity levels, which directly matches the requirement to automatically classify uploaded photos and block content based on configured severity thresholds. Prompt Shields are intended to detect prompt injection and jailbreak-style attacks against generative models, not to classify image harm categories. Keyword scanning OCR output would only detect visible text extracted from the image and would miss visual harm in the image itself. Blocklists can help match known words or custom patterns, but they are not a complete image safety classifier and do not provide the built-in severity- based image harm classification required here. Image moderation is therefore the correct control for user- uploaded photos. Reference topics: Azure AI Content Safety, image moderation, harm categories, severity levels, Foundry guardrails, and responsible AI controls.

질문 # 15

You have a Microsoft Foundry project that contains a customer support agent. The agent calls an internal knowledge API tool before generating responses.

Users report the following issues:

- * Some requests take more than 15 seconds to complete.
- * Some responses are incorrect, even when the knowledge API returns the expected data.

You need to inspect individual agent runs to view the ordered sequence of large language model (LLM) calls, tool invocations, and timing information.

Which observability capability should you use?

- A. safety metrics
- **B. tracing**
- C. monitoring
- D. token usage

정답: B

설명:

The correct capability is tracing because the requirement is to inspect the execution path of an individual agent run. Microsoft Foundry tracing captures detailed telemetry for agent behavior, including LLM calls, tool invocations, agent decision flows, inputs, outputs, tool results, token consumption, duration, and latency.

This is the appropriate observability mechanism when you need to determine which step introduced a delay, whether the agent called the internal knowledge API, what data the tool returned, and how the model used that data before producing the final response. Microsoft's Foundry observability guidance describes distributed tracing as the mechanism that provides visibility into LLM calls, tool invocations, agent decisions, and inter-service dependencies.

Token usage is useful for cost analysis and prompt optimization, but it does not show ordered run steps or tool-call sequencing. Safety metrics evaluate risk-related output behavior, not latency or tool execution.

General monitoring provides aggregate health, latency, success-rate, and dashboard views, but the question asks for per-run sequence inspection and timing breakdowns. Foundry agent tracing specifically supports debugging unexpected behavior and monitoring latency across requests. Reference topics: Microsoft Foundry observability, agent tracing, OpenTelemetry-based traces, tool invocations, LLM call inspection, and latency diagnostics.

질문 # 16

.....

아직도 Microsoft인증 AI-103 시험준비를 어떻게 해야 할지 망설이고 계시나요? 고객님의 IT인증 시험준비길에는 언제나 ITDumpsKR가 곁을 지켜주고 있습니다. ITDumpsKR 시험공부자료를 선택하시면 자격증 취득의 소원이 이루어 집니다. Microsoft인증 AI-103 시험덤프는 ITDumpsKR가 최고의 선택입니다.

