

# AI-103덤프데모문제 - AI-103시험덤프



경쟁율이 치열한 IT업계에서 아무런 목표없이 아무런 희망없이 무미건조한 생활을 하고 계시나요? 다른 사람들이 모두 취득하고 있는 자격증에 관심도 없는 분은 치열한 경쟁속에서 살아남기 어렵습니다. Microsoft인증 AI-103시험패스가 힘들다한들ITDumpsKR덤프만 있으면 어려운 시험도 쉬워질수 밖에 없습니다. Microsoft인증 AI-103덤프에 있는 문제만 잘 이해하고 습득하신다면Microsoft인증 AI-103시험을 패스하여 자격증을 취득해 자신의 경쟁력을 업그레이드하여 경쟁시대에서 안전감을 보유할수 있습니다.

ITDumpsKR의 Microsoft 인증 AI-103시험덤프공부자료는 pdf버전과 소프트웨어버전 두가지 버전으로 제공되는데 Microsoft 인증 AI-103실제시험예상문제가 포함되어있습니다.덤프의 예상문제는 Microsoft 인증 AI-103실제시험의 대부분 문제를 적중하여 높은 통과율과 점유율을 자랑하고 있습니다. ITDumpsKR의 Microsoft 인증 AI-103덤프를 선택하시면 IT자격증 취득에 더할것 없는 힘이 될것입니다.

>> AI-103덤프데모문제 <<

## Microsoft AI-103시험덤프, AI-103최신 업데이트 인증덤프

Microsoft인증 AI-103시험을 패스해서 자격증을 취득하려고 하는데 시험비며 학원비며 공부자료비며 비용이 만만치 않아요? 제일 저렴한 가격으로 제일 효과좋은ITDumpsKR의 Microsoft인증 AI-103덤프를 알고 계시는지요? ITDumpsKR의 Microsoft인증 AI-103덤프는 최신 시험문제에 근거하여 만들어진 시험준비공부 가이드로서 학원공부 필요없이 덤프공부만으로도 시험을 한방에 패스할수 있습니다. 덤프를 구매하신분은 철저한 구매후 서비스도 받을수 있습니다.

## 최신 Azure AI Engineer Associate AI-103 무료 샘플문제 (Q38-Q43):

### 질문 # 38

You are building a speech processing solution in Microsoft Foundry for a customer support platform.

The platform will transcribe live phone calls, so that supervisors at your company can view call transcripts and detect issues while the calls are in progress. The call audio will arrive as a continuous stream from the telephony system.

You need to ensure that the call transcripts appear within only a few seconds of the audio stream.

What should you do?

- A. Run a batch transcription job on recorded audio files.
- **B. Use real-time speech to text to process streaming audio input.**
- C. Use text to speech by using a custom neural voice.
- D. Use speech translation to generate the transcripts into multiple languages.

정답: B

설명:

The correct answer is B. Use real-time speech to text to process streaming audio input . The scenario requires live transcription from a continuous telephony stream, with transcript text appearing within a few seconds while the call is still in progress. Azure Speech in Foundry Tools real-time speech recognition is specifically intended for immediate transcription scenarios such as call center assistance, dictation, and live meeting captioning. Microsoft's Speech guidance describes real-time speech to text as processing

audio input and returning transcriptions in real time, which matches the supervisor monitoring requirement. Batch transcription is inappropriate because it processes stored audio files after recording, not an active live stream. Speech translation is used when the primary goal is translating speech into another language, not simply producing live same-language call transcripts. Text to speech performs the reverse operation by generating spoken audio from text and does not transcribe inbound calls. Real-time speech to text provides the low-latency streaming recognition path required for live operational monitoring. Reference topics: Azure Speech in Foundry Tools, real-time speech recognition, streaming audio input, call center transcription, and live captions.

### 질문 # 39

You have a Microsoft Foundry project that ingests scanned PDF invoices stored in Azure Blob Storage. Each invoice contains printed line items and has a table-based layout.

Extracted results are stored as structured JSON and used as grounding data for an agent in a Retrieval Augmented Generation (RAG) solution.

You need to create a single analyzer that meets the following requirements:

\* Extracts the invoice number, invoice date, vendor name, and total amount across varying templates

\* Returns confidence scores so that results with confidence below 0.80 can be routed for supervisor review What should you use?

- A. a Foundry agent that has groundedness guardrails enabled to extract invoice fields and confidence scores
- B. the Azure Content Understanding in Foundry Tools prebuilt-documentSearch analyzer and search.score from the Azure AI Search results for routing
- C. a custom Azure Content Understanding in Foundry Tools analyzer that defines the required fields as the extracted fields and the returned confidence scores for routing
- D. the Azure Content Understanding in Foundry Tools prebuilt-layout analyzer

정답: C

#### 설명:

The correct answer is C because the requirement is structured field extraction from invoices across varying templates, not only OCR or layout preservation. Azure Content Understanding analyzers are reusable configurations that combine content extraction, AI-powered analysis, and structured data output, and Microsoft states that custom analyzers can be created for specific extraction needs. In this case, the analyzer schema should define fields such as invoice number, invoice date, vendor name, and total amount so the output can be returned as structured JSON for downstream RAG grounding.

The confidence-routing requirement also points to Content Understanding field confidence scores. Microsoft documentation states that every field can include a confidence score from 0 to 1, and that confidence scores can be used to automate high-confidence results while routing low-confidence results for human review. A threshold such as 0.80 is therefore an application routing rule based on the returned field confidence. The prebuilt-layout analyzer preserves layout but does not define invoice-specific business fields. Groundedness guardrails evaluate generated answers, not invoice field extraction. Azure AI Search search.score measures retrieval relevance, not extraction confidence. Reference topics: Content Understanding custom analyzers, document field extraction, structured JSON output, confidence scoring, and RAG grounding.

### 질문 # 40

You have a Python application named App1 that integrates with a Microsoft Foundry project named Project1.

You need to ensure that App1 meets the following requirements:

\* Authenticates by using a Microsoft Entra managed identity

\* Sends prompts to a deployed model by using the Azure OpenAI Responses API How should you complete the Python code? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

```
from azure.identity import DefaultAzureCredential
from azure.ai.projects import AIProjectClient

credential =  ()

project_client = AIProjectClient(
    endpoint="https://contosoai.services.ai.azure.com/api/projects/project1",
    credential=credential,
)

with project_client.get_openai_client() as openai_client:
    response = openai_client.responses. (
        model="trail-guide-chat",
        input="Create a 3-day hiking itinerary near Seattle.",
    )

print(response.output_text)
```

정답:

설명:

```

from azure.identity import DefaultAzureCredential
from azure.ai.projects import AIProjectClient

credential =  ()

project_client = AIProjectClient(
    endpoint="https://contosoai.services.ai.azure.com/api/projects/project1",
    credential=credential,
)

with project_client.get_openai_client() as openai_client:
    response = openai_client.responses. (
        model="trail-guide-chat",
        input="Create a 3-day hiking itinerary near Seattle.",
    )

print(response.output_text)

```

Explanation:

```
credential = DefaultAzureCredential
response = openai_client.responses.create
```

The correct authentication class is `DefaultAzureCredential` because the requirement is Microsoft Entra-based authentication, and this credential supports managed identity in hosted Azure environments. The Microsoft Foundry quickstart for Python shows the project client being created with `AIProjectClient (endpoint=PROJECT_ENDPOINT, credential=DefaultAzureCredential())`, which enables the Foundry SDK to authenticate without using API keys or embedded secrets. The same guidance shows creating an OpenAI-compatible client from the project by calling `project.get_openai_client()`.

The correct Responses API method is `create` because the application must send a new prompt to the deployed model and receive generated output. Microsoft's Foundry quickstart demonstrates the exact pattern: `response = openai_client.responses.create(model="gpt-5-mini", input="...")`, followed by reading `response.output_text`.

The `retrieve` operation is used to fetch an existing response, not submit a new inference request, and `compact` is not the correct method for generating a model response. `AzureKeyCredential` would violate the Microsoft Entra managed identity requirement, while `ClientSecretCredential` uses an application secret rather than managed identity. Reference topics: Microsoft Foundry SDK, `AIProjectClient`, Microsoft Entra authentication, `DefaultAzureCredential`, and Azure OpenAI Responses API.

#### 질문 # 41

You have a Microsoft Foundry project that uses Azure AI Search to ground an agent in internal documentation. After a recent content update, users report that the agent's answers have become less accurate. You need to identify whether the retrieved content is negatively influencing the model's generated responses. Which observability signal should you review?

- A. latency breakdown traces
- B. prediction drift metrics
- C. groundedness evaluation metrics
- D. indexer status and failure history

정답: C

설명:

The correct observability signal is B. groundedness evaluation metrics . In a RAG solution, the key diagnostic question is whether the generated answer is supported by the retrieved context. Microsoft Foundry' s built-in evaluator reference defines Groundedness as the metric that measures how grounded the response is in the retrieved context, with scoring that indicates whether the model's claims are supported by the provided source material.

This matches the issue after a content update. If retrieved chunks are stale, misleading, incomplete, or poorly aligned with the user query, groundedness results can show that generated responses are not reliably supported by the retrieved documentation. The RAG evaluator guidance explains that groundedness focuses on whether the response avoids content outside the grounding context, while other process metrics such as retrieval evaluate how relevant the retrieved chunks are. Latency traces are useful for performance troubleshooting, not response accuracy. Indexer status can reveal ingestion failures, but it does not show whether retrieved content is influencing generated answers negatively. Prediction drift is a model monitoring concept and is not the primary signal for RAG grounding quality. Reference topics: Microsoft Foundry observability, RAG evaluators, groundedness, retrieved context, and response quality evaluation.

질문 # 42

You have an Azure Speech in Foundry Tools resource that hosts a custom speech to text model deployed to a custom endpoint. An agent uses the endpoint to perform real-time speech recognition.

You are approaching the expiration date of the custom speech to text model.

What is the expected behavior when the model expires?

- A. Speech recognition requests will continue to use the expired custom model until the model is removed manually.
- B. Speech recognition requests will return a 4xx error until a new custom model is deployed.
- C. The custom model will be deleted automatically when the model expires.
- D. Speech recognition requests will fall back to the most recent base model for the same locale.

정답: D

설명:

The correct answer is A. Speech recognition requests will fall back to the most recent base model for the same locale . Microsoft's custom speech model lifecycle guidance states that when a model expires, it is no longer available for transcription. For the custom endpoint route, speech recognition requests fall back to the most recent base model for the same locale. The documentation also warns that recognition results might still be returned, but the transcription may no longer reflect the domain-specific adaptation of the custom model.

This distinguishes real-time custom endpoint behavior from batch transcription behavior. Batch transcription requests that specify an expired model fail with a 4xx error, but that is not the route described in this question.

The agent is using a custom endpoint for real-time recognition, so fallback to the latest base model is the expected behavior. The model is not automatically deleted merely because it expires, and it does not continue to use the expired custom model indefinitely.

The operational recommendation is to update the endpoint's model before expiration by redeploying the endpoint with a newer custom model. Reference topics: Azure Speech custom speech model lifecycle, custom endpoints, model expiration, real-time speech recognition, and endpoint redeployment.

질문 # 43

.....

많은 분들이 고난의도인 Microsoft관련인증 시험을 응시하고 싶어 하는데 이런 시험은 많은 전문적인 관련지식이 필요합니다. 시험은 당연히 완전히 전문적인 AI-103관련지식을 터득하자만이 패스할 가능성이 높습니다. 하지만 지금은 많은 방법들로 여러분의 부족한 면을 보충해드릴 수 있으며 또 힘든 Microsoft시험도 패스하실 수 있습니다. 혹은 여러분은 전문적인 Developing AI Apps and Agents on Azure관련지식을 터득하자들보다 더 간단히 더 빨리 시험을 패스하실 수 있습니다.

AI-103시험덤프 : <https://www.itdumpskr.com/AI-103-exam.html>

만약Microsoft AI-103인증시험 자격증이 있다면 일에서도 많은 변화가 있을 것입니다, 연봉상승은 물론, 자기자신만의 공간도 넓어집니다, Microsoft AI-103덤프데모문제 여러분의 성공을 빕니다, Microsoft AI-103덤프데모문제 문제가 많으면 고객들의 시간을 허비하게 됩니다, AI-103시험은 영어로 출제되는 만큼 시험난이도가 많이 높습니다. 하지만 AI-103인증덤프만 있다면 아무리 어려운 시험도 쉬워집니다, 저희가 제공하는 AI-103인증시험 덤프는 여러분이 AI-103시험을 안전하게 통과는 물론 관련 전문지식 장악에도 많은 도움이 될것입니다, AI-103인증시험을

패스하여 자격증을 취득하는데 가장 쉬운 방법은 ITDumpsKR에서 제공해드리는 AI-103덤프를 공부하는 것입니다.

예쁘긴, 내가 처음 공에 왔을 적에 사내인지 계집애인지 아리송해서 마마님들이 이름을 아리'라고 했다는데, 어떻게든 흠을 잡으려고 할 텐데, 만약Microsoft AI-103인증시험 자격증이 있다면 일에서도 많은 변화가 있을 것입니다, 연봉상승은 물론, 자기자신만의 공간도 넓어집니다.

## 시험패스 가능한 AI-103덤프데모문제 최신버전 덤프데모 문제 다운

여러분의 성공을 빕니다, 문제가 많으면 고객들의 시간을 허비하게 됩니다, AI-103시험은 영어로 출제되는 만큼 시험난이도가 많이 높습니다.하지만 AI-103인증덤프만 있다면 아무리 어려운 시험도 쉬워집니다.

저희가 제공하는 AI-103인증시험 덤프는 여러분이 AI-103시험을 안전하게 통과는 물론 관련 전문지식 장악에도 많은 도움이 될것입니다.

- AI-103최고품질 인증시험 기출문제 \ AI-103최신 인증시험 공부자료 □ AI-103최고품질 인증시험 기출문제 □ ☀ www.exampassdump.com □☀□을(를) 열고[ AI-103 ]를 검색하여 시험 자료를 무료로 다운로드하십시오 AI-103최신 업데이트 시험대비자료
- AI-103테스트자료 □ AI-103인증시험 덤프자료 □ AI-103최신버전 시험대비자료 □ 무료로 쉽게 다운로드하려면 □ www.itdumpskr.com □에서▶ AI-103 ◀를 검색하세요 AI-103예상문제
- AI-103적중을 높은 시험대비덤프 □ AI-103합격보장 가능 공부자료 □ AI-103퍼펙트 최신버전 자료 □ ➡ www.itdumpskr.com □□□은 ( AI-103 ) 무료 다운로드를 받을 수 있는 최고의 사이트입니다 AI-103합격보장 가능 공부자료
- AI-103퍼펙트 덤프샘플 다운로드 ♥ AI-103최신 덤프문제 □ AI-103인기자격증 시험대비자료 □ ☀ www.itdumpskr.com □☀□을 통해 쉽게 □ AI-103 □무료 다운로드 받기 AI-103최신버전 시험대비자료
- AI-103최신덤프문제 □ AI-103최신 업데이트 시험대비자료 □ AI-103인기자격증 시험대비자료 ➔ ⇒ www.exampassdump.com ◀에서 □ AI-103 □를 검색하고 무료 다운로드 받기 AI-103 Dumps
- AI-103덤프데모문제 기출문제 □ ⇒ www.itdumpskr.com ◀에서[ AI-103 ]를 검색하고 무료로 다운로드하세요 AI-103예상문제
- 시험준비에 가장 좋은 AI-103덤프데모문제 덤프 샘플문제 다운받기 □ 지금▶ www.koreadumps.com ◀을(를) 열고 무료 다운로드를 위해▶ AI-103 ◀를 검색하십시오 AI-103인증시험 덤프자료
- 최신버전 AI-103덤프데모문제 덤프공부문제 □ 【 www.itdumpskr.com 】에서 검색만 하면▶ AI-103 ◀를 무료로 다운로드할 수 있습니다 AI-103최고품질 인증시험 기출문제
- AI-103인기자격증 시험대비자료 □ AI-103예상문제 □ AI-103 Dumps □ [ AI-103 ]를 무료로 다운로드하려면 「 www.itdumpskr.com 」 웹사이트를 입력하세요 AI-103인기자격증 시험대비자료
- AI-103덤프데모문제 기출문제 □ ➡ www.itdumpskr.com □웹사이트에서“ AI-103 ”를 열고 검색하여 무료 다운로드 AI-103테스트자료
- AI-103최신 인증시험 공부자료 ♥ AI-103인증시험 덤프공부 □ AI-103최신버전 시험대비자료 □ 오픈 웹사이트 【 kr.fast2test.com 】 검색 《 AI-103 》 무료 다운로드 AI-103최신 업데이트 시험대비자료
- allenvasz401077.blogdal.com, kallumpvht267690.blogofchange.com, marcawwk398139.westexwiki.com, tasneemktr278950.buscawiki.com, exactlybookmarks.com, knowyourmeme.com, brontelfnd327346.bimmwiki.com, extrabookmarking.com, soocareer.com, bookmarkworm.com, Disposable vapes