

Exam : AI-103

**Title : Developing AI Apps and
Agents on Azure**

Version : DEMO

1.HOTSPOT

Case Study

This is a case study. Case studies are not timed separately from other exam sections. You can use as much exam time as you would like to complete each case study. However, there might be additional case studies or other exam sections. Manage your time to ensure that you can complete all the exam sections in the time provided. Pay attention to the Exam Progress at the top of the screen so you have sufficient time to complete any exam sections that follow this case study.

To answer the case study questions, you will need to reference information that is provided in the case. Case studies and associated questions might contain exhibits or other resources that provide more information about the scenario described in the case. Information provided in an individual question does not apply to the other questions in the case study.

A Review Screen will appear at the end of this case study. From the Review Screen, you can review and change your answers before you move to the next exam section. After you leave this case study, you will NOT be able to return to it.

To start the case study

To display the first question in this case study, select the “Next” button. To the left of the question, a menu provides links to information such as business requirements, the existing environment, and problem statements. Please read through all this information before answering any questions. When you are ready to answer a question, select the “Question” button to return to the question.

Overview

Company Information

Contoso, Ltd is a multinational retail company that builds, deploys, and manages generative AI and agent-based solutions by using Microsoft Foundry.

Existing Environment

Identity Environment

Contoso uses Microsoft Entra ID for identity management, authentication, and authorization capabilities that enable agents to access organizational resources and services.

Contoso recently formed a new AI engineering team named Agent1Dev Team to optimize and maintain existing AI solutions. The team collaborates with solution architects, DevOps engineers, and security engineers to design, implement, monitor, and secure AI applications.

Contoso also has a team named Agent1Test Team that is responsible for validating AI solutions before the solution deployments.

Generative Environment

Contoso has a Microsoft Foundry deployment that contains two projects named Project1 and Project2.

Project1

Project1 contains a customer support agent named Agent1 that assists customers with product inquiries

and troubleshooting requests.

Agent1 has the following configurations:

Agent1 uses a base model deployment.

A safety evaluation pipeline is NOT enabled.

Tool invocation approval workflows are NOT enabled.

Conversation memory constraints are NOT configured.

Agent1 interacts with customers by using digital support channels and answers general questions about Contoso products.

Project1 is deployed to an Azure region located in the European Union (EU).

Agent1Dev Team will use Project1 to optimize and maintain Agent1.

Project2

Project2 contains a deployed video generation model. The marketing department at Contoso has access to Project2 and plans to use the model to develop a video creation solution.

Development of the solution is incomplete.

Data Environment

Contoso stores product-related information in Azure resources that support AI applications.

The Azure environment contains an Azure Blob Storage account named storage1 that stores product detail sheets for all the Contoso products.

The product sheets include specifications, feature descriptions, and product support information that Agent1 can use to answer customer questions. The product sheets are stored in the PDF format.

Problem Statements

Contoso identifies the following issues:

Agent1 has only general knowledge of the Contoso products.

A recent chat interaction with Agent1 was analyzed for sentiment. The results of the analysis have NOT been processed yet. Agent1 does NOT use the detailed product information in the product sheets stored in storage1 when responding to customer questions.

The finance department at Contoso reports that vendor invoices must be reviewed manually to ensure that the invoices match the terms defined in the vendor contracts. The invoices contain tables, logos, and varied layouts that make the documents difficult to process consistently.

Requirements

Planned Changes

Contoso plans to implement the following changes:

Implement a solution for Project1 that analyzes the vendor invoices by evaluating both the visual layout and the textual content of the invoices, so that the invoice details can be verified against the vendor contract terms.

Update the base model deployment used by Agent1 and standardize the model version to ensure continuity and consistent responses.

Enable Agent1 to retrieve and use the detailed product information from the product sheets stored in storage1.

Implement an indexing solution for the product sheets that Agent1 can use to answer customer

questions.

Complete the development of the video creation solution.

Technical Requirements

Contoso identifies the following technical requirements:

The model deployment used by Agent1 must support scalable, high-throughput generative AI workloads and dynamically scale to handle variable customer support traffic, without requiring reserved throughput capacity.

The product sheets must be processed by using an indexing pipeline that enables semantic and vector search, so that Agent1 can retrieve the relevant product information.

Responses generated by using the product sheet information must be relevant, complete, and accurate. Agent1 must be able to use the product sheets to answer natural language questions about product details.

The model version used by Agent1 must remain consistent to ensure stable responses.

The data processed by the model must remain within the EU.

Security and Compliance Requirements

Contoso identifies the following security and compliance requirements:

API keys must NOT be used to access Foundry-deployed models.

Access to the Azure resources must follow the principle of least privilege.

The developers at Contoso must authenticate to Microsoft Foundry resources by using Microsoft Entra authentication. Access to Project1 must be assigned to the members of Agent1Dev Team by using a security group named SC_Agent1_Dev. Access to Project1 must be assigned to the members of Agent1Test Team by using a security group named SC_Agent1_Test. Agent1 must never reveal customer information, even if a document that contains customer data is added erroneously to the product sheet repository in storage1.

The product sheets might contain images that include embedded text. Agent1 must be protected from malicious instructions potentially hidden within the images.

Business Requirements


Contoso identifies the following business requirements:


Users that interact with Agent1 must have a personalized experience in future interactions, including the ability for Agent1 to retain conversation context and recall relevant information from previous interactions. Agent1 must answer questions only about the products sold by Contoso.

You need to configure the model deployment for Agent1 to meet the technical requirements.

What should you configure? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.


Answer Area


Deployment type: 

Version update policy: 

Answer:

Answer Area

Deployment type: 

Version update policy: 

2.HOTSPOT

Case Study

This is a case study. Case studies are not timed separately from other exam sections. You can use as much exam time as you would like to complete each case study. However, there might be additional case studies or other exam sections. Manage your time to ensure that you can complete all the exam sections in the time provided. Pay attention to the Exam Progress at the top of the screen so you have sufficient time to complete any exam sections that follow this case study.

To answer the case study questions, you will need to reference information that is provided in the case. Case studies and associated questions might contain exhibits or other resources that provide more information about the scenario described in the case. Information provided in an individual question does not apply to the other questions in the case study.

A Review Screen will appear at the end of this case study. From the Review Screen, you can review and change your answers before you move to the next exam section. After you leave this case study, you will

NOT be able to return to it.

To start the case study

To display the first question in this case study, select the “Next” button. To the left of the question, a menu provides links to information such as business requirements, the existing environment, and problem statements. Please read through all this information before answering any questions. When you are ready to answer a question, select the “Question” button to return to the question.

Overview

Company Information

Contoso, Ltd is a multinational retail company that builds, deploys, and manages generative AI and agent-based solutions by using Microsoft Foundry.

Existing Environment

Identity Environment

Contoso uses Microsoft Entra ID for identity management, authentication, and authorization capabilities that enable agents to access organizational resources and services.

Contoso recently formed a new AI engineering team named Agent1Dev Team to optimize and maintain existing AI solutions. The team collaborates with solution architects, DevOps engineers, and security engineers to design, implement, monitor, and secure AI applications.

Contoso also has a team named Agent1Test Team that is responsible for validating AI solutions before the solution deployments.

Generative Environment

Contoso has a Microsoft Foundry deployment that contains two projects named Project1 and Project2.

Project1

Project1 contains a customer support agent named Agent1 that assists customers with product inquiries and troubleshooting requests.

Agent1 has the following configurations:

Agent1 uses a base model deployment.

A safety evaluation pipeline is NOT enabled.

Tool invocation approval workflows are NOT enabled.

Conversation memory constraints are NOT configured.

Agent1 interacts with customers by using digital support channels and answers general questions about Contoso products.

Project1 is deployed to an Azure region located in the European Union (EU).

Agent1Dev Team will use Project1 to optimize and maintain Agent1.

Project2

Project2 contains a deployed video generation model. The marketing department at Contoso has access to Project2 and plans to use the model to develop a video creation solution.

Development of the solution is incomplete.

Data Environment

Contoso stores product-related information in Azure resources that support AI applications.

The Azure environment contains an Azure Blob Storage account named storage1 that stores product detail sheets for all the Contoso products.

The product sheets include specifications, feature descriptions, and product support information that Agent1 can use to answer customer questions. The product sheets are stored in the PDF format.

Problem Statements

Contoso identifies the following issues:

Agent1 has only general knowledge of the Contoso products.

A recent chat interaction with Agent1 was analyzed for sentiment. The results of the analysis have NOT been processed yet. Agent1 does NOT use the detailed product information in the product sheets stored in storage1 when responding to customer questions.

The finance department at Contoso reports that vendor invoices must be reviewed manually to ensure that the invoices match the terms defined in the vendor contracts. The invoices contain tables, logos, and varied layouts that make the documents difficult to process consistently.

Requirements

Planned Changes

Contoso plans to implement the following changes:

Implement a solution for Project1 that analyzes the vendor invoices by evaluating both the visual layout and the textual content of the invoices, so that the invoice details can be verified against the vendor contract terms.

Update the base model deployment used by Agent1 and standardize the model version to ensure continuity and consistent responses.

Enable Agent1 to retrieve and use the detailed product information from the product sheets stored in storage1.

Implement an indexing solution for the product sheets that Agent1 can use to answer customer questions.

Complete the development of the video creation solution.

Technical Requirements

Contoso identifies the following technical requirements:

The model deployment used by Agent1 must support scalable, high-throughput generative AI workloads and dynamically scale to handle variable customer support traffic, without requiring reserved throughput capacity.

The product sheets must be processed by using an indexing pipeline that enables semantic and vector search, so that Agent1 can retrieve the relevant product information.

Responses generated by using the product sheet information must be relevant, complete, and accurate. Agent1 must be able to use the product sheets to answer natural language questions about product details.

The model version used by Agent1 must remain consistent to ensure stable responses.

The data processed by the model must remain within the EU.

Security and Compliance Requirements

Contoso identifies the following security and compliance requirements:

API keys must NOT be used to access Foundry-deployed models.

Access to the Azure resources must follow the principle of least privilege.

The developers at Contoso must authenticate to Microsoft Foundry resources by using Microsoft Entra authentication. Access to Project1 must be assigned to the members of Agent1Dev Team by using a security group named SC_Agent1_Dev. Access to Project1 must be assigned to the members of Agent1Test Team by using a security group named SC_Agent1_Test. Agent1 must never reveal customer information, even if a document that contains customer data is added erroneously to the product sheet repository in storage1.

The product sheets might contain images that include embedded text. Agent1 must be protected from malicious instructions potentially hidden within the images.

Business Requirements

Contoso identifies the following business requirements:

Users that interact with Agent1 must have a personalized experience in future interactions, including the ability for Agent1 to retain conversation context and recall relevant information from previous interactions. Agent1 must answer questions only about the products sold by Contoso.

You need to configure Agent1 to meet the security and compliance requirements.

What should you use?

- A. self-harm content filtering
- B. prompt shields
- C. Personally identifiable information (PII) Detection
- D. violence content filtering

Answer: C

3.You are planning a Microsoft Foundry project named Project1 that will contain multiple agents. Each agent will access the same Azure AI Search resource.

You need to recommend a solution to centrally manage the Azure AI Search credentials within Project1. The solution must be implemented across all the agents.

What should you recommend?

- A. Enable role-based access control (RBAC) for the Azure AI Search resource.
- B. Disable key-based access control on the Azure AI Search resource.
- C. Add a connection to the Azure AI Search resource.
- D. Create a managed private endpoint that connects to the Azure AI Search resource.

Answer: C

4.HOTSPOT -

Your company is piloting a customer support agent in a Microsoft Foundry project name Project1.

Project1 is connected to an existing Application Insights resource, and the company's support team reviews runs in the Traces tab.

The Foundry Agent Service is configured to perform the following actions:

Retrieve the Application Insights connection string by calling

`project_client.telemetry.get_application_insights_connection_string()`.

Call `configure_azure_monitor(connection_string=...)` to enable telemetry.

A separate LangChain service is configured to use OpenTelemetry and has the following configurations:

Uses `AzureAIOpenTelemetryTracer(connection_string=..., enable_content_recording=False)`

Passes the tracer by using `config={"callbacks":[azure_tracer]}`

Company policy has the following requirements:

Telemetry from LangChain and OpenTelemetry must be distinguishable within the same Application Insights resource.

Secrets and credentials must NOT be stored in prompts, tool arguments, or span attributes.

For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE:

Each correct selection is worth one point.

Answer Area

Statements	Yes	No
The LangChain service will appear in Traces without configuring a tracer.	☐	☐
Setting different <code>OTEL_SERVICE_NAME</code> values separates the services in Application Insights.	☐	☐
When using <code>enable_content_recording=False</code> , prompts and tool data will be captured in the telemetry.	☐	☐

Answer:

Answer Area

Statements	Yes	No
The LangChain service will appear in Traces without configuring a tracer.	☐	☒
Setting different <code>OTEL_SERVICE_NAME</code> values separates the services in Application Insights.	☒	☐
When using <code>enable_content_recording=False</code> , prompts and tool data will be captured in the telemetry.	☐	☒

5.DRAG DROP

You have a Microsoft Foundry project that processes procurement documents submitted by suppliers.

You need to implement two pipelines by using Azure Content Understanding in Foundry Tools.

The solution must meet the following requirements:

Include a pipeline named Pipeline1 that supports cost-effective, high-volume processing of standalone PDF invoices.

Include a pipeline named Pipeline2 that supports cross-document validation by using multi-step

reasoning and reference data.

How should you configure each pipeline? To answer, drag the appropriate configurations to the correct pipelines. Each configuration may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content. NOTE: Each correct selection is worth one point.

Configurations

- ⋮ Multi-file task in pro mode
- ⋮ Multi-file task in standard mode
- ⋮ Single-file task in pro mode
- ⋮ Single-file task in standard mode

Answer Area

- Pipeline1: Configuration
- Pipeline2: Configuration

Answer:

Configurations

- ⋮ Multi-file task in pro mode
- ⋮ Multi-file task in standard mode
- ⋮ Single-file task in pro mode
- ⋮ Single-file task in standard mode

Answer Area

- Pipeline1: ⋮ Single-file task in pro mode
- Pipeline2: ⋮ Multi-file task in pro mode