Number: AI Specialist
Passing Score: 800
Time Limit: 120
File Version: 4.0

Exam Code: AI Specialist

Exam Name: Salesforce Certified AI Specialist

Exam A

QUESTION 1
Universal Containers' data science team is hosting a generative large language model (LLM) on Amazon Web Services (AWS).
What should the team use to access externally-hosted models in the Salesforce Platform?

A. Model Builder
B. App Builder
C. Copilot Builder

Correct Answer: A
Section:
Explanation:
To access externally-hosted models, such as a large language model (LLM) hosted on AWS, the Model Builder in Salesforce is the appropriate tool. Model Builder allows teams to integrate and deploy external AI models into the Salesforce platform, making it possible to leverage models hosted outside of Salesforce infrastructure while still benefiting from the platform's native AI capabilities.
Option B, App Builder, is primarily used to build and configure applications in Salesforce, not to integrate AI models.
Option C, Copilot Builder, focuses on building assistant-like tools rather than integrating external AI models.
Model Builder enables seamless integration with external systems and models, allowing Salesforce users to use external LLMs for generating AI-driven insights and automation.
Salesforce AI Specialist
Reference: For more details, check the Model Builder guide here: https://help.salesforce.com/s/articleView?id=sf.model_builder_external_models.htm

QUESTION 2
How does the Einstein Trust Layer ensure that sensitive data is protected while generating useful and meaningful responses?

A. Masked data will be de-masked during response journey.
B. Masked data will be de-masked during request journey.
C. Responses that do not meet the relevance threshold will be automatically rejected.

Correct Answer: A
Section:
Explanation:
The Einstein Trust Layer ensures that sensitive data is protected while generating useful and meaningful responses by masking sensitive data before it is sent to the Large Language Model (LLM) and then de-masking it during the response journey.
How It Works:
Data Masking in the Request Journey:
Sensitive Data Identification: Before sending the prompt to the LLM, the Einstein Trust Layer scans the input for sensitive data, such as personally identifiable information (PII), confidential business information, or any other data deemed sensitive.
Masking Sensitive Data: Identified sensitive data is replaced with placeholders or masks. This ensures that the LLM does not receive any raw sensitive information, thereby protecting it from potential exposure.
Processing by the LLM:
Masked Input: The LLM processes the masked prompt and generates a response based on the masked data.
No Exposure of Sensitive Data: Since the LLM never receives the actual sensitive data, there is no risk of it inadvertently including that data in its output.
De-masking in the Response Journey:
Re-insertion of Sensitive Data: After the LLM generates a response, the Einstein Trust Layer replaces the placeholders in the response with the original sensitive data.
Providing Meaningful Responses: This de-masking process ensures that the final response is both meaningful and complete, including the necessary sensitive information where appropriate.
Maintaining Data Security: At no point is the sensitive data exposed to the LLM or any unintended recipients, maintaining data security and compliance.
Why Option A is Correct:

De-masking During Response Journey: The de-masking process occurs after the LLM has generated its response, ensuring that sensitive data is only reintroduced into the output at the final stage, securely and appropriately.
Balancing Security and Utility: This approach allows the system to generate useful and meaningful responses that include necessary sensitive information without compromising data security.
Why Options B and C are Incorrect:
Option B (Masked data will be de-masked during request journey):
Incorrect Process: De-masking during the request journey would expose sensitive data before it reaches the LLM, defeating the purpose of masking and compromising data security.
Option C (Responses that do not meet the relevance threshold will be automatically rejected):
Irrelevant to Data Protection: While the Einstein Trust Layer does enforce relevance thresholds to filter out inappropriate or irrelevant responses, this mechanism does not directly relate to the protection of sensitive data. It addresses response quality rather than data security.
Salesforce AI Specialist Documentation - Einstein Trust Layer Overview:
Explains how the Trust Layer masks sensitive data in prompts and re-inserts it after LLM processing to protect data privacy.
Salesforce Help - Data Masking and De-masking Process:
Details the masking of sensitive data before sending to the LLM and the de-masking process during the response journey.
Salesforce AI Specialist Exam Guide - Security and Compliance in AI:
Outlines the importance of data protection mechanisms like the Einstein Trust Layer in AI implementations.
Conclusion:
The Einstein Trust Layer ensures sensitive data is protected by masking it before sending any prompts to the LLM and then de-masking it during the response journey. This process allows Salesforce to generate useful and meaningful responses that include necessary sensitive information without exposing that data during the AI processing, thereby maintaining data security and compliance.

QUESTION 3
Universal Containers (UC) wants to enable its sales team to get insights into product and competitor names mentioned during calls.
How should UC meet this requirement?

A. Enable Einstein Conversation Insights, assign permission sets, define recording managers, and customize insights with up to 50 competitor names.
B. Enable Einstein Conversation Insights, connect a recording provider, assign permission sets, and customize insights with up to 25 products.
C. Enable Einstein Conversation Insights, enable sales recording, assign permission sets, and customize insights with up to 50 products.

Correct Answer: C
Section:
Explanation:
To provide the sales team with insights into product and competitor names mentioned during calls, Universal Containers should:
Enable Einstein Conversation Insights: Activates the feature that analyzes call recordings for valuable insights.
Enable Sales Recording: Allows calls to be recorded within Salesforce without needing an external recording provider.
Assign Permission Sets: Grants the necessary permissions to sales team members to access and utilize conversation insights.
Customize Insights: Configure the system to track mentions of up to 50 products and 50 competitors, providing tailored insights relevant to the organization's needs.
Option C accurately reflects these steps. Option A mentions defining recording managers but omits enabling sales recording within Salesforce. Option B suggests connecting a recording provider and limits customization to 25 products, which does not fully meet UC's requirements.
Salesforce AI Specialist Documentation - Setting Up Einstein Conversation Insights: Provides instructions on enabling conversation insights and sales recording.
Salesforce Help - Customizing Conversation Insights: Details how to customize insights with up to 50 products and competitors.
Salesforce AI Specialist Exam Guide: Outlines best practices for implementing AI features like Einstein Conversation Insights in a sales context.
==========

QUESTION 4
What is the role of the large language model (LLM) in executing an Einstein Copilot Action?

A. Find similar requests and provide actions that need to be executed
B. Identify the best matching actions and correct order of execution
C. Determine a user's access and sort actions by priority to be executed

Correct Answer: B

Section:
Explanation:
In Einstein Copilot, the role of the Large Language Model (LLM) is to analyze user inputs and identify the best matching actions that need to be executed. It uses natural language understanding to break down the user's request and determine the correct sequence of actions that should be performed.
By doing so, the LLM ensures that the tasks and actions executed are contextually relevant and are performed in the proper order. This process provides a seamless, AI-enhanced experience for users by matching their requests to predefined Salesforce actions or flows.
The other options are incorrect because:
A mentions finding similar requests, which is not the primary role of the LLM in this context.
C focuses on access and sorting by priority, which is handled more by security models and governance than by the LLM.
Salesforce Einstein Documentation on Einstein Copilot Actions
Salesforce AI Documentation on Large Language Models

QUESTION 5
A service agent is looking at a custom object that stores travel information. They recently received a weather alert and now need to cancel flights for the customers that are related with this itinerary. The service agent needs to review the Knowledge articles about canceling and rebooking the customer flights.
Which Einstein Copilot capability helps the agent accomplish this?

A. Execute tasks based on available actions, answering questions using information from accessible Knowledge articles.
B. Invoke a flow which makes a call to external data to create a Knowledge article.
C. Generate a Knowledge article based off the prompts that the agent enters to create steps to cancel flights.

Correct Answer: A
Section:
Explanation:
In this scenario, the Einstein Copilot capability that best helps the agent is its ability to execute tasks based on available actions and answer questions using data from Knowledge articles. Einstein Copilot can assist the service agent by providing relevant Knowledge articles on canceling and rebooking flights, ensuring that the agent has access to the correct steps and procedures directly within the workflow.
This feature leverages the agent's existing context (the travel itinerary) and provides actionable insights or next steps from the relevant Knowledge articles to help the agent quickly resolve the customer's needs.
The other options are incorrect:
B refers to invoking a flow to create a Knowledge article, which is unrelated to the task of retrieving existing Knowledge articles.
C focuses on generating Knowledge articles, which is not the immediate need for this situation where the agent requires guidance on existing procedures.
Salesforce Documentation on Einstein Copilot
Trailhead Module on Einstein for Service

QUESTION 6
An AI Specialist has created a copilot custom action using flow as the reference action type. However, it is not delivering the expected results to the conversation preview, and therefore needs troubleshooting.
What should the AI Specialist do to identify the root cause of the problem?

A. In Copilot Builder within the Dynamic Panel, turn on dynamic debugging to show the inputs and outputs.
B. Copilot Builder within the Dynamic Panel, confirm selected action and observe the values in Input and Output sections.
C. In Copilot Builder, verify the utterance entered by the user and review session event logs for debug information.

Correct Answer: A
Section:
Explanation:
When troubleshooting a copilot custom action using flow as the reference action type, enabling dynamic debugging within Copilot Builder's Dynamic Panel is the most effective way to identify the root cause. By turning on dynamic debugging, the AI Specialist can see detailed logs showing both the inputs and outputs of the flow, which helps identify where the action might be failing or not delivering the expected results.
Option B, confirming selected actions and observing the Input and Output sections, is useful for monitoring flow configuration but does not provide the deep diagnostic details available with dynamic debugging.
Option C, verifying the user utterance and reviewing session event logs, could provide helpful context, but dynamic debugging is the primary tool for identifying issues with inputs and outputs in real time.
Salesforce AI Specialist

Reference: To explore more about dynamic debugging in Copilot Builder, see: https://help.salesforce.com/s/articleView?id=sf.copilot_custom_action_debugging.htm

QUESTION 7
A support team handles a high volume of chat interactions and needs a solution to provide quick, relevant responses to customer inquiries.
Responses must be grounded in the organization's knowledge base to maintain consistency and accuracy.
Which feature in Einstein for Service should the support team use?

A. Einstein Service Replies
B. Einstein Reply Recommendations
C. Einstein Knowledge Recommendations

Correct Answer: B
Section:
Explanation:
The support team should use Einstein Reply Recommendations to provide quick, relevant responses to customer inquiries that are grounded in the organization's knowledge base. This feature leverages AI to recommend accurate and consistent replies based on historical interactions and the knowledge stored in the system, ensuring that responses are aligned with organizational standards.
Einstein Service Replies (Option A) is focused on generating replies but doesn't have the same emphasis on grounding responses in the knowledge base.
Einstein Knowledge Recommendations (Option C) suggests knowledge articles to agents, which is more about assisting the agent in finding relevant articles than providing automated or AI-generated responses to customers.
Salesforce AI Specialist
Reference: For more information on Einstein Reply Recommendations: https://help.salesforce.com/s/articleView?id=sf.einstein_reply_recommendations_overview.htm

QUESTION 8
Universal Containers implemented Einstein Copilot for its users.
One user complains that Einstein Copilot is not deleting activities from the past 7 days.
What is the reason for this issue?

A. Einstein Copilot Delete Record Action permission is not associated to the user.
B. Einstein Copilot does not have the permission to delete the user's records.
C. Einstein Copilot does not support the Delete Record action.

Correct Answer: C
Section:
Explanation:
Einstein Copilot currently supports various actions like creating and updating records but does not support the Delete Record action. Therefore, the user's request to delete activities from the past 7 days cannot be fulfilled using Einstein Copilot.
Unsupported Action: The inability to delete records is due to the current limitations of Einstein Copilot's supported actions. It is designed to assist with tasks like data retrieval, creation, and updates, but for security and data integrity reasons, it does not facilitate the deletion of records.
User Permissions: Even if the user has the necessary permissions to delete records within Salesforce, Einstein Copilot itself does not have the capability to execute delete operations.
Salesforce AI Specialist Documentation - Einstein Copilot Supported Actions:
Lists the actions that Einstein Copilot can perform, noting the absence of delete operations.
Salesforce Help - Limitations of Einstein Copilot:
Highlights current limitations, including unsupported actions like deleting records.

QUESTION 9
Universal Containers' service team wants to customize the standard case summary response from Einstein Copilot.
What should the AI Specialist do to achieve this?

A. Customize the standard Record Summary template for the Case object,
B. Summarize the Case with a standard copilot action.

C. Create a custom Record Summary prompt template for the Case object.

Correct Answer: C
Section:
Explanation:
To customize the case summary response from Einstein Copilot, the AI Specialist should create a custom Record Summary prompt template for the Case object. This allows Universal Containers to tailor the way case data is summarized, ensuring the output aligns with specific business requirements or user preferences.
Option A (customizing the standard Record Summary template) does not provide the flexibility required for deep customization.
Option B (standard Copilot action) won't allow customization; it will only use default settings.
Refer to Salesforce Prompt Builder documentation for guidance on creating custom templates for record summaries.

QUESTION 10
Universal Containers wants to be able to detect with a high level confidence if content generated by a large language model (LLM) contains toxic language.
Which action should an Al Specialist take in the Trust Layer to confirm toxicity is being appropriately managed?

A. Access the Toxicity Detection log in Setup and export all entries where isToxicityDetected is true.
B. Create a flow that sends an email to a specified address each time the toxicity score from the response exceeds a predefined threshold.
C. Create a Trust Layer audit report within Data Cloud that uses a toxicity detector type filter to display toxic responses and their respective scores.

Correct Answer: C
Section:
Explanation:
To ensure that content generated by a large language model (LLM) is appropriately screened for toxic language, the AI Specialist should create a Trust Layer audit report within Data Cloud. By using the toxicity detector type filter, the report can display toxic responses along with their respective toxicity scores, allowing Universal Containers to monitor and manage any toxic content generated with a high level of confidence.
Option C is correct because it enables visibility into toxic language detection within the Trust Layer and allows for auditing responses for toxicity.
Option A suggests checking a toxicity detection log, but Salesforce provides more comprehensive options via the audit report.
Option B involves creating a flow, which is unnecessary for toxicity detection monitoring.
Salesforce Trust Layer Documentation: https://help.salesforce.com/s/articleView?id=sf.einstein_trust_layer_audit.htm

QUESTION 11
Universal Containers (UC) is using Einstein Generative AI to generate an account summary. UC aims to ensure the content is safe and inclusive, utilizing the Einstein Trust Layer's toxicity scoring to assess the content's safety level.
What does a safety category score of 1 indicate in the Einstein Generative Toxicity Score?

A. Not safe
B. Safe
C. Moderately safe

Correct Answer: B
Section:
Explanation:
In the Einstein Trust Layer, the toxicity scoring system is used to evaluate the safety level of content generated by AI, particularly to ensure that it is non-toxic, inclusive, and appropriate for business contexts. A toxicity score of 1 indicates that the content is deemed safe.
The scoring system ranges from 0 (unsafe) to 1 (safe), with intermediate values indicating varying degrees of safety. In this case, a score of 1 means that the generated content is fully safe and meets the trust and compliance guidelines set by the Einstein Trust Layer.
For further reference, check Salesforce's official Einstein Trust Layer documentation regarding toxicity scoring for AI-generated content.

QUESTION 12
Universal Containers has an active standard email prompt template that does not fully deliver on the business requirements.

Which steps should an AI Specialist take to use the content of the standard prompt email template in question and customize it to fully meet the business requirements?

A. Save as New Template and edit as needed.
B. Clone the existing template and modify as needed.
C. Save as New Version and edit as needed.

Correct Answer: B
Section:
Explanation:
When an active standard email prompt template doesn't meet the business requirements, the best approach is to clone the existing template and modify it as needed. Cloning allows the AI Specialist to preserve the original template while making adjustments to fit specific business needs. This ensures that any customizations are applied without altering the original standard template.
Saving as a new version is typically used for versioning changes in the same template, while Save as New Template creates a brand-new template without linking to the existing one. Cloning provides a balance, allowing modifications while retaining the original structure for future reference.
For more details, refer to Salesforce Prompt Builder documentation for guidance on cloning and modifying templates.

QUESTION 13
The marketing team at Universal Containers is looking for a way personalize emails based on customer behavior, preferences, and purchase history.
Why should the team use Einstein Copilot as the solution?

A. To generate relevant content when engaging with each customer
B. To analyze past campaign performance
C. To send automated emails to all customers

Correct Answer: A
Section:
Explanation:
Einstein Copilot is designed to assist in generating personalized, AI-driven content based on customer data such as behavior, preferences, and purchase history. For the marketing team at Universal Containers, this is the perfect solution to create dynamic and relevant email content. By leveraging Einstein Copilot, they can ensure that each customer receives tailored communications, improving engagement and conversion rates.
Option A is correct as Einstein Copilot helps generate real-time, personalized content based on comprehensive data about the customer.
Option B refers more to Einstein Analytics or Marketing Cloud Intelligence, and Option C deals with automation, which isn't the primary focus of Einstein Copilot.
Salesforce Einstein Copilot Overview: https://help.salesforce.com/s/articleView?id=einstein_copilot_overview.htm

QUESTION 14
Leadership needs to populate a dynamic form field with a summary or description created by a large language model (LLM) to facilitate more productive conversations with customers. Leadership also wants to keep a human in the loop to be considered in their AI strategy.
Which prompt template type should the AI Specialist recommend?

A. Sales Email
B. Field Generation
C. Record Summary

Correct Answer: B
Section:
Explanation:
The correct answer is Field Generation because this template type is designed to dynamically populate form fields with content generated by a large language model (LLM). In this scenario, leadership wants a dynamic form field that contains a summary or description generated by AI to aid customer interactions. Additionally, they want to keep a human in the loop, meaning the generated content will likely be reviewed or edited by a person before it's finalized, which aligns with the Field Generation prompt template.
Field Generation: This prompt type allows you to generate content for specific fields in Salesforce, leveraging large language models to create dynamic and contextual information. It ensures that AI content is available within

the record where needed, but it allows human oversight or review, supporting the 'human-in-the-loop' strategy.
Sales Email: This prompt type is mainly used for generating email content for outreach or responses, which doesn't align directly with populating fields in a form.
Record Summary: While this option might seem close, it is typically used to summarize entire records for high-level insights rather than filling specific fields with dynamic content based on AI generation.
Salesforce AI Specialist
Reference:
You can explore more about these prompt templates and AI capabilities through Salesforce documentation and official resources on Prompt Builder:
https://help.salesforce.com/s/articleView?id=sf.prompt_builder_templates_overview.htm

QUESTION 15
Universal Containers is considering leveraging the Einstein Trust Layer in conjunction with Einstein Generative AI Audit Dat a.
Which audit data is available using the Einstein Trust Layer?

A. Response accuracy and offensiveness score
B. Hallucination score and bias score
C. Masked data and toxicity score

Correct Answer: C
Section:
Explanation:
Universal Containers is considering the use of the Einstein Trust Layer along with Einstein Generative AI Audit Data. The Einstein Trust Layer provides a secure and compliant way to use AI by offering features like data masking and toxicity assessment.
The audit data available through the Einstein Trust Layer includes information about masked data---which ensures sensitive information is not exposed---and the toxicity score, which evaluates the generated content for inappropriate or harmful language.
Salesforce AI Specialist Documentation - Einstein Trust Layer: Details the auditing capabilities, including logging of masked data and evaluation of generated responses for toxicity to maintain compliance and trust.

QUESTION 16
Universal Containers wants to make a sales proposal and directly use data from multiple unrelated objects (standard and custom) in a prompt template.
What should the AI Specialist recommend?

A. Create a Flex template to add resources with standard and custom objects as inputs.
B. Create a prompt template passing in a special custom object that connects the records temporarily,
C. Create a prompt template-triggered flow to access the data from standard and custom objects.

Correct Answer: A
Section:
Explanation:
Universal Containers needs to generate a sales proposal using data from multiple unrelated standard and custom objects within a prompt template. The most effective way to achieve this is by using a Flex template.
Flex templates in Salesforce allow AI specialists to create prompt templates that can accept inputs from multiple sources, including various standard and custom objects. This flexibility enables the direct use of data from unrelated objects without the need to create intermediary custom objects or complex flows.
Salesforce AI Specialist Documentation - Flex Templates: Explains how Flex templates can be utilized to incorporate data from multiple sources, providing a flexible solution for complex data requirements in prompt templates.

QUESTION 17
What is an AI Specialist able to do when the ''Enrich event logs with conversation data' setting in Einstein Copilot is enabled?

A. View the user click path that led to each copilot action.
B. View session data including user Input and copilot responses for sessions over the past 7 days.
C. Generate details reports on all Copilot conversations over any time period.

Correct Answer: B
Section:
Explanation:
When the 'Enrich event logs with conversation data' setting is enabled in Einstein Copilot, it allows an AI Specialist or admin to view session data, including both the user input and copilot responses from interactions over the past 7 days. This data is crucial for monitoring how the copilot is being used, analyzing its performance, and improving future interactions based on past inputs.
This setting enriches the event logs with detailed conversational data for better insights into the interaction history, helping AI specialists track AI behavior and user engagement.
Option A, viewing the user click path, focuses on navigation but is not part of the conversation data enrichment functionality.
Option C, generating detailed reports over any time period, is incorrect because this specific feature is limited to data for the past 7 days.
Salesforce AI Specialist
Reference: You can refer to this documentation for further insights: https://help.salesforce.com/s/articleView?id=sf.einstein_copilot_event_logging.htm

QUESTION 18
Universal Containers' current AI data masking rules do not align with organizational privacy and security policies and requirements.
What should an AI Specialist recommend to resolve the issue?

A. Enable data masking for sandbox refreshes.
B. Configure data masking in the Einstein Trust Layer setup.
C. Add new data masking rules in LLM setup.

Correct Answer: B
Section:
Explanation:
When Universal Containers' AI data masking rules do not meet organizational privacy and security standards, the AI Specialist should configure the data masking rules within the Einstein Trust Layer. The Einstein Trust Layer provides a secure and compliant environment where sensitive data can be masked or anonymized to adhere to privacy policies and regulations.
Option A, enabling data masking for sandbox refreshes, is related to sandbox environments, which are separate from how AI interacts with production data.
Option C, adding masking rules in the LLM setup, is not appropriate because data masking is managed through the Einstein Trust Layer, not the LLM configuration.
The Einstein Trust Layer allows for more granular control over what data is exposed to the AI model and ensures compliance with privacy regulations.
Salesforce AI Specialist
Reference: For more information, refer to: https://help.salesforce.com/s/articleView?id=sf.einstein_trust_layer_data_masking.htm

QUESTION 19
An administrator wants to check the response of the Flex prompt template they've built, but the preview button is greyed out.
What is the reason for this?

A. The records related to the prompt have not been selected.
B. The prompt has not been saved and activated,
C. A merge field has not been inserted in the prompt.

Correct Answer: A
Section:
Explanation:
When the preview button is greyed out in a Flex prompt template, it is often because the records related to the prompt have not been selected. Flex prompt templates pull data dynamically from Salesforce records, and if there are no records specified for the prompt, it can't be previewed since there is no content to generate based on the template.
Option B, not saving or activating the prompt, would not necessarily cause the preview button to be greyed out, but it could prevent proper functionality.
Option C, missing a merge field, would cause issues with the output but would not directly grey out the preview button.
Ensuring that the related records are correctly linked is crucial for testing and previewing how the prompt will function in real use cases.
Salesforce AI Specialist
Reference: Refer to the documentation on troubleshooting Flex templates here: https://help.salesforce.com/s/articleView?id=sf.flex_prompt_builder_troubleshoot.htm

QUESTION 20
What is the correct process to leverage Prompt Builder in a Salesforce org?

A. Select the appropriate prompt template type to use, select one of Salesforce's standard prompts, determine the object to associate the prompt, select a record to validate against, and associate the prompt to an action.
B. Select the appropriate prompt template type to use, develop the prompt within the prompt workspace, select resources to dynamically insert CRM-derived grounding data, pick the model to use, and test and validate the generated responses.
C. Enable the target object for generative prompting, develop the prompt within the prompt workspace, select records to fine-tune and ground the response, enable the Trust Layer, and associate the prompt to an action.

Correct Answer: B
Section:
Explanation:
When using Prompt Builder in a Salesforce org, the correct process involves several important steps:
Select the appropriate prompt template type based on the use case.
Develop the prompt within the prompt workspace, where the template is created and customized.
Select CRM-derived grounding data to be dynamically inserted into the prompt, ensuring that the AI-generated responses are based on accurate and relevant data.
Pick the model to use for generating responses, either using Salesforce's built-in models or custom ones.
Test and validate the generated responses to ensure accuracy and effectiveness.
Option B is correct as it follows the proper steps for using Prompt Builder.
Option A and Option C do not capture the full process correctly.
Salesforce Prompt Builder Documentation: https://help.salesforce.com/s/articleView?id=sf.prompt_builder_overview.htm

QUESTION 21
An AI Specialist wants to include data from the response of external service invocation (REST API callout) into the prompt template.
How should the AI Specialist meet this requirement?

A. Convert the JSON to an XML merge field.
B. Use External Service Record merge fields.
C. Use ''Add Prompt Instructions'' flow element.

Correct Answer: B
Section:
Explanation:
An AI Specialist wants to include data from the response of an external service invocation (REST API callout) into a prompt template. The goal is to incorporate dynamic data retrieved from an external API into the AI-generated content.
Solution:
Use External Service Record Merge Fields
External Service Integration:
Definition: External Services in Salesforce allow the integration of external REST APIs into Salesforce without custom code.
Registration: The external service must be registered in Salesforce, defining the API's schema and methods.
External Service Record Merge Fields:
Purpose: Enables the inclusion of data from external service responses directly into prompt templates using merge fields.
Functionality:
Dynamic Data Inclusion: Allows prompt templates to access and use data returned from REST API callouts.
Merge Fields Syntax: Use merge fields in the prompt template to reference specific data points from the API response.
Implementation Steps:
Register the External Service:
Use External Services to register the REST API in Salesforce.
Define the API's schema, including methods and data structures.
Create a Named Credential:

Configure authentication and endpoint details for the external API.
Use External Service in Flow:
Build a Flow that invokes the external service and captures the response.
Ensure the flow outputs the necessary data for use in the prompt template.
Configure the Prompt Template:
Use External Service Record merge fields in the prompt template to reference data from the flow's output.
Syntax Example: {{flowOutputVariable.fieldName}}
Why Other Options are Less Suitable:
Option A (Convert the JSON to an XML merge field):
Irrelevance: Converting JSON to XML merge fields is unnecessary and complicates the process.
Unsupported Method: Salesforce prompt templates do not support direct inclusion of XML merge fields from JSON conversion.
Option C (Use ''Add Prompt Instructions'' flow element):
Purpose of Add Prompt Instructions:
Allows adding instructions to the prompt within a flow but does not facilitate including external data.
Limitation: Does not directly help in incorporating external service responses into the prompt template.
Salesforce AI Specialist Documentation - Integrating External Services with Prompt Templates:
Explains how to use External Services and merge fields in prompt templates.
Salesforce Help - Using Merge Fields with External Data:
Provides guidance on referencing external data in templates using merge fields.
Salesforce Trailhead - External Services and Flow:
Offers a practical understanding of integrating external APIs using External Services and Flow.
Conclusion:
By using External Service Record merge fields, the AI Specialist can effectively include data from external REST API responses into prompt templates, ensuring that the AI-generated content is enriched with up-to-date and relevant external data.

QUESTION 22
Universal Containers (UC) has a legacy system that needs to integrate with Salesforce. UC wishes to create a digest of account action plans using the generative API feature.
Which API service should UC use to meet this requirement?

A. REST API
B. Metadata API
C. SOAP API

Correct Answer: A
Section:
Explanation:
To create a digest of account action plans using the generative API feature, Universal Containers should use the REST API. The REST API is ideal for integrating Salesforce with external systems and enabling interaction with Salesforce data, including generative capabilities like creating summaries or digests. It supports modern web standards and is suitable for flexible, lightweight interactions between Salesforce and legacy systems.
Metadata API is used for retrieving and deploying metadata, not for data operations like generating summaries.
SOAP API is an older API used for integration but is less flexible compared to REST for this specific use case.
For more details, refer to Salesforce REST API documentation regarding using REST for data integration and generating content.

QUESTION 23
The sales team at a hotel resort would like to generate a guest summary about the guests' interests and provide recommendations based on their activity preferences captured in each guest profile. They want the summary to be available only on the contact record page.
Which AI capability should the team use?

A. Einstein Copilot
B. Prompt Builder

C. Model Builder

Correct Answer: B
Section:
Explanation:
The sales team at a hotel resort wants to generate a guest summary about guests' interests and provide recommendations based on their activity preferences captured in each guest profile. They require the summary to be available only on the contact record page.
Solution:
Use Prompt Builder to create a prompt template that generates the desired summary and displays it on the contact record page.
Prompt Builder:
Purpose: Allows the creation of custom prompt templates that leverage AI to generate content based on Salesforce data.
Functionality:
Field Generation Templates: Can be used to populate fields on records with AI-generated summaries.
Customization: Enables the AI Specialist to design prompts that utilize data from the guest profiles to produce personalized summaries and recommendations.
Relevance to the Use Case:
The sales team wants the summary to be available on the contact record page, which aligns with the capabilities of Prompt Builder to generate and display content on specific record pages.
Implementation Steps:
Create a Field Generation Prompt Template:
Use Prompt Builder to create a new prompt template of type Field Generation.
Design the prompt to instruct the AI to generate a summary based on the guest's interests and activity preferences.
Include Relevant Data:
Use merge fields to include data from the guest profile in the prompt.
Ensure that the prompt accesses the necessary fields to generate accurate recommendations.
Configure the Contact Page Layout:
Add the field that will display the AI-generated summary to the contact record page layout.
Ensure that the field is only visible where appropriate, adhering to the requirement of availability only on the contact record page.
Why Not Einstein Copilot or Model Builder:
Option A (Einstein Copilot):
Purpose: Einstein Copilot is a conversational AI assistant designed to interact with users through natural language.
Mismatch with Requirements:
The team wants a static summary displayed on the contact record page, not an interactive conversational experience.
Option C (Model Builder):
Purpose: Model Builder is used to create custom AI models for predictions and classifications.
Inapplicability:
Building a custom model is unnecessary for generating text summaries based on existing data.
Model Builder does not directly provide functionality to generate and display summaries on record pages.
Salesforce AI Specialist Documentation - Prompt Builder Overview:
Provides an introduction to Prompt Builder and its capabilities.
Salesforce Help - Creating Field Generation Prompt Templates:
Guides on creating prompt templates that generate content for fields on records.
Salesforce Trailhead - Customize AI Content with Prompt Builder:
Offers hands-on experience in building and customizing prompt templates.
Conclusion:
By utilizing Prompt Builder, the sales team can create a customized prompt template that generates personalized guest summaries and recommendations based on activity preferences. This solution meets the requirement of displaying the summary only on the contact record page, enhancing the team's ability to engage with guests effectively.

QUESTION 24
An Al Specialist is tasked with configuring a generative model to create personalized sales emails using customer data stored in Salesforce. The AI Specialist has already fine-tuned a large language model (LLM) on the OpenAI platform. Security and data privacy are critical concerns for the client.
How should the AI Specialist integrate the custom LLM into Salesforce?

A. Create an application of the custom LLM and embed it in Sales Cloud via iFrame.
B. Add the fine-tuned LLM in Einstein Studio Model Builder.
C. Enable model endpoint on OpenAI and make callouts to the model to generate emails.

Correct Answer: B
Section:
Explanation:
Since security and data privacy are critical, the best option for the AI Specialist is to integrate the fine-tuned LLM (Large Language Model) into Salesforce by adding it to Einstein Studio Model Builder. Einstein Studio allows organizations to bring their own AI models (BYOM), ensuring the model is securely managed within Salesforce's environment, adhering to data privacy standards.
Option A (embedding via iFrame) is less secure and doesn't integrate deeply with Salesforce's data and security models.
Option C (making callouts to OpenAI) raises concerns about data privacy, as sensitive Salesforce data would be sent to an external system.
Einstein Studio provides the most secure and seamless way to integrate custom AI models while maintaining control over data privacy and compliance. More details can be found in Salesforce's Einstein Studio documentation on integrating external models.

QUESTION 25
What should an AI Specialist consider when using related list merge fields in a prompt template associated with an Account object in Prompt Builder?

A. The Activities related list on the Account object is not supported because it is a polymorphic field.
B. If person accounts have been enabled, merge fields will not be available for the Account object.
C. Prompt generation will yield no response when there is no related list associated with an Account in runtime.

Correct Answer: A
Section:
Explanation:
When using related list merge fields in a prompt template associated with the Account object in Prompt Builder, the Activities related list is not supported due to it being a polymorphic field. Polymorphic fields can reference multiple different types of objects, which makes them incompatible with some merge field operations in prompt generation.
Option B is incorrect because person accounts do not limit the availability of merge fields for the Account object.
Option C is irrelevant since even if no related lists are available at runtime, the prompt can still generate based on other available data fields.
For more information, refer to Salesforce documentation on supported fields and limitations in Prompt Builder.

QUESTION 26
Universal Containers (UC) wants to use the Draft with Einstein feature in Sales Cloud to create a personalized introduction email.
After creating a proposed draft email, which predefined adjustment should UC choose to revise the draft with a more casual tone?

A. Make Less Formal
B. Enhance Friendliness
C. Optimize for Clarity

Correct Answer: A
Section:
Explanation:
When Universal Containers uses the Draft with Einstein feature in Sales Cloud to create a personalized email, the predefined adjustment to Make Less Formal is the correct option to revise the draft with a more casual tone. This option adjusts the wording of the draft to sound less formal, making the communication more approachable while still maintaining professionalism.
Enhance Friendliness would make the tone more positive, but not necessarily more casual.
Optimize for Clarity focuses on making the draft clearer but doesn't adjust the tone.
For more details, see Salesforce documentation on Einstein-generated email drafts and tone adjustments.

QUESTION 27
Universal Containers recently launched a pilot program to integrate conversational AI into its CRM business operations with Einstein Copilot.
How should the AI Specialist monitor Copilot's usability and the assignment of actions?

A. Run a report on the Platform Debug Logs.
B. Query the Copilot log data using the metadata API.
C. Run Einstein Copilot Analytics.

Correct Answer: C
Section:
Explanation:
To monitor Einstein Copilot's usability and the assignment of actions, the AI Specialist should run Einstein Copilot Analytics. This feature provides insights into how often Copilot is used, the types of actions it is handling, and overall user engagement with the system. It's the most effective way to track Copilot's performance and usage patterns.
Platform Debug Logs are not relevant for tracking user behavior or the assignment of Copilot actions.
Querying the Copilot log data via the Metadata API would not provide the necessary insights in a structured manner.
For more details, refer to Salesforce's Copilot Analytics documentation for tracking AI-driven interactions.

QUESTION 28
Universal Containers (UC) is experimenting with using public Generative AI models and is familiar with the language required to get the information it needs. However, it can be time consuming for both UC's sales and service reps to type in the prompt to get the information they need, and ensure prompt consistency.
Which Salesforce feature should a Salesforce AI Specialist recommend to address these concerns?

A. Einstein Recommendation Builder
B. Einstein Copilot Action: Query Records
C. Einstein Prompt Builder and Prompt Templates

Correct Answer: C
Section:
Explanation:
For Universal Containers (UC), to reduce the time and ensure prompt consistency when using public generative AI models, the recommended feature is Einstein Prompt Builder and Prompt Templates. This feature allows teams to create reusable and consistent prompts for generative AI tasks, ensuring that all users receive uniform responses without having to type in detailed prompts manually every time.
Einstein Prompt Builder simplifies the creation of prompts, and Prompt Templates standardize the inputs, saving time for sales and service reps.
Option A (Einstein Recommendation Builder) is more focused on recommendations, not prompt standardization.
Option B (Einstein Copilot Action: Query Records) is for querying records, not generating AI-driven prompts.
Salesforce Prompt Builder Overview: https://help.salesforce.com/s/articleView?id=sf.prompt_builder_overview.htm

QUESTION 29
Universal Containers tests out a new Einstein Generative AI feature for its sales team to create personalized and contextualized emails for its customers. Sometimes, users find that the draft email contains placeholders for attributes that could have been derived from the recipient's contact record.
What is the most likely explanation for why the draft email shows these placeholders?

A. The user does not have Einstein Sales Emails permission assigned.
B. The user does not have permission to access the fields.
C. The user's locale language is not supported by Prompt Builder.

Correct Answer: B
Section:
Explanation:

When using Einstein Generative AI to create personalized emails, if placeholders appear in the draft email where data from a recipient's Contact record should be, the most likely reason is that the user lacks permission to access the necessary fields. Salesforce's field-level security may prevent users from viewing or utilizing certain data fields, resulting in placeholders being shown instead of the actual values.
Option B is correct because missing field permissions will cause placeholders in email drafts.
Option A (missing Einstein Sales Emails permission) is unlikely, as this would prevent email generation altogether, not just placeholders.
Option C (locale language issues) would more likely affect language-specific issues, not field placeholders.
Salesforce Email Template and Permissions Documentation: https://help.salesforce.com/s/articleView?id=sf.email_templates_field_permissions.htm

QUESTION 30
Universal Containers (UC) has implemented Generative AI within Salesforce to enable summarization of a custom object called Guest. Users have reported mismatches in the generated information.
In refining its prompt design strategy, which key practices should UC prioritize?

A. Enable prompt test mode, allocate different prompt variations to a subset of users for evaluation, and standardize the most effective model based on performance feedback.
B. Create concise, clear, and consistent prompt templates with effective grounding, contextual role-playing, clear instructions, and iterative feedback.
C. Submit a prompt review case to Salesforce and conduct thorough testing In the playground to refine outputs until they meet user expectations.

Correct Answer: B
Section:
Explanation:
For Universal Containers (UC) to refine its Generative AI prompt design strategy and improve the accuracy of the generated summaries for the custom object Guest, the best practice is to focus on crafting concise, clear, and consistent prompt templates. This includes:
Effective grounding: Ensuring the prompt pulls data from the correct sources.
Contextual role-playing: Providing the AI with a clear understanding of its role in generating the summary.
Clear instructions: Giving unambiguous directions on what to include in the response.
Iterative feedback: Regularly testing and adjusting prompts based on user feedback.
Option B is correct because it follows industry best practices for refining prompt design.
Option A (prompt test mode) is useful but less relevant for refining prompt design itself.
Option C (prompt review case with Salesforce) would be more appropriate for technical issues or complex prompt errors, not general design refinement.
Salesforce Prompt Design Best Practices: https://help.salesforce.com/s/articleView?id=sf.prompt_design_best_practices.htm

QUESTION 31
An AI Specialist needs to create a Sales Email with a custom prompt template. They need to ground on the following data.
Opportunity Products Events near the customer Tone and voice examples
How should the AI Specialist obtain related items?

A. Call prompt initiated flow to fetch and ground the required data.
B. Create a flex template that takes the records in question as inputs.
C. Utilize a standard email template and manually insert the required data fields.

Correct Answer: A
Section:
Explanation:
To ground a sales email on Opportunity Products, Events near the customer, and Tone and voice examples, the AI Specialist should use a prompt-initiated flow. This flow can dynamically fetch the necessary data from related records in Salesforce and ground the generative AI output with contextually accurate information.
Option B (flex template) does not provide the ability to fetch dynamic data from Salesforce records automatically.
Option C (manual insertion) would not allow for the dynamic and automated grounding of data required for custom prompts.
Refer to Salesforce documentation on flows and grounding for more details on integrating data into custom prompt templates.

QUESTION 32
Universal Containers (UC) wants to create a new Sales Email prompt template in Prompt Builder using the 'Save As' function. However, UC notices that the new template produces different results compared to the standard

Sales Email prompt due to missing hyperparameters.
What should UC do to ensure the new prompt template produces results comparable to the standard Sales Email prompts?

A. Use Model Playground to create a model configuration with the specified parameters.
B. Manually add the hyperparameters to the new template.
C. Revert to using the standard template without modifications.

Correct Answer: B
Section:
Explanation:
When Universal Containers creates a new Sales Email prompt template using the 'Save As' function, missing hyperparameters can result in different outputs. To ensure the new prompt produces comparable results to the standard Sales Email prompt, the AI Specialist should manually add the necessary hyperparameters to the new template.
Hyperparameters like Temperature, Frequency Penalty, and Presence Penalty directly affect how the AI generates responses. Ensuring that these are consistent with the standard template will result in similar outputs.
Option A (Model Playground) is not necessary here, as it focuses on fine-tuning models, not adjusting templates directly.
Option C (Reverting to the standard template) does not solve the issue of customizing the prompt template.
For more information, refer to Prompt Builder documentation on configuring hyperparameters in custom templates.

QUESTION 33
Universal Containers (UC) uses Salesforce Service Cloud to support its customers and agents handling cases. UC is considering implementing Einstein Copilot and extending Service Cloud to mobile users.
When would Einstein Copilot implementation be most advantageous?

A. When the goal is to streamline customer support processes and improve response times
B. When the main objective is to enhance data security and compliance measures
C. When the focus is on optimizing marketing campaigns and strategies

Correct Answer: A
Section:
Explanation:
Einstein Copilot implementation would be most advantageous in Salesforce Service Cloud when the goal is to streamline customer support processes and improve response times. Einstein Copilot can assist agents by providing real-time suggestions, automating repetitive tasks, and generating contextual responses, thus enhancing service efficiency.
Option B (data security) is not the primary focus of Einstein Copilot, which is more about improving operational efficiency.
Option C (marketing campaigns) falls outside the scope of Service Cloud and Einstein Copilot's primary benefits, which are aimed at improving customer service and case management.
For further reading, refer to Salesforce documentation on Einstein Copilot for Service Cloud and how it improves support processes.

QUESTION 34
An AI Specialist needs to create a prompt template to fill a custom field named Latest Opportunities Summary on the Account object with information from the three most recently opened opportunities.
How should the AI Specialist gather the necessary data for the prompt template?

A. Create a flow to retrieve the opportunity information.
B. Select the Account Opportunity object as a resource when creating the prompt template.
C. Select the latest Opportunities related list as a merge field.

Correct Answer: A
Section:
Explanation:
To gather the necessary data for populating the Latest Opportunities Summary custom field on the Account object with information from the three most recently opened opportunities, the AI Specialist should create a flow. A flow can be configured to query and retrieve the required opportunity records based on criteria such as their open date. Once the flow has gathered the necessary data, it can be used in a prompt template or other automation processes to populate the custom field on the Account record.

Option A is correct because creating a flow allows for dynamic data retrieval and control over the logic for selecting the most recent opportunities.
Option B and Option C do not provide sufficient control or data retrieval capabilities needed for this scenario.
Salesforce Flow Documentation: https://help.salesforce.com/s/articleView?id=sf.flow.htm

QUESTION 35
A data scientist needs to view and manage models in Einstein Studio. The data scientist also needs to create prompt templates in Prompt Builder.
Which permission sets should an AI Specialist assign to the data scientist?

A. Data Cloud Admin and Prompt Template Manager
B. Prompt Template Manager and Prompt Template User
C. Prompt Template User and Data Cloud Admin

Correct Answer: A
Section:
Explanation:
To allow a data scientist to view and manage models in Einstein Studio and create prompt templates in Prompt Builder, the AI Specialist should assign the Data Cloud Admin and Prompt Template Manager permission sets.
Data Cloud Admin provides access to manage and oversee models within Einstein Studio.
Prompt Template Manager gives the user the ability to create and manage prompt templates within Prompt Builder.
Option A is correct because it assigns the necessary permissions for both managing models and creating prompt templates.
Option B and Option C are incorrect as they do not provide the correct combination of permissions for managing models and building prompts.
Salesforce Permissions Documentation: https://help.salesforce.com/s/articleView?id=sf.perm_sets_overview.htm

QUESTION 36
An AI Specialist configured Data Masking within the Einstein Trust Layer.
How should the AI Specialist begin validating that the correct fields are being masked?

A. Use a Flow-based resource in Prompt Builder to debug the fields' merge values using Flow Debugger.
B. Request the Einstein Generative AI Audit Data from the Security section of the Setup menu.
C. Enable the collection and storage of Einstein Generative AI Audit Data on the Einstein Feedback setup page.

Correct Answer: B
Section:
Explanation:
To begin validating that the correct fields are being masked in Einstein Trust Layer, the AI Specialist should request the Einstein Generative AI Audit Data from the Security section of the Salesforce Setup menu. This audit data allows the AI Specialist to see how data is being processed, including which fields are being masked, providing transparency and validation that the configuration is working as expected.
Option B is correct because it allows for the retrieval of audit data that can be used to validate data masking.
Option A (Flow Debugger) and Option C (Einstein Feedback) do not relate to validating field masking in the context of the Einstein Trust Layer.
Salesforce Einstein Trust Layer Documentation: https://help.salesforce.com/s/articleView?id=sf.einstein_trust_layer_audit.htm

QUESTION 37
When configuring a prompt template, an AI Specialist previews the results of the prompt template they've written. They see two distinct text outputs: Resolution and Response.
Which information does the Resolution text provide?

A. It shows the full text that is sent to the Trust Layer.
B. It shows the response from the LLM based on the sample record.
C. It shows which sensitive data is masked before it is sent to the LLM.

Correct Answer: B

Section:
Explanation:
When previewing a prompt template in Salesforce, the Resolution text provides the response from the LLM (Large Language Model) based on the data from a sample record. This output shows what the AI model generated in response to the prompt, giving the AI Specialist a chance to review and adjust the response before finalizing the template.
Option B is correct because Resolution displays the actual response generated by the LLM.
Option A refers to sending the text to the Trust Layer, but that's not what Resolution represents.
Option C relates to data masking, which is shown elsewhere, not under Resolution.
Salesforce Prompt Builder Overview: https://help.salesforce.com/s/articleView?id=sf.prompt_builder_overview.htm

QUESTION 38
What is the primary function of the planner service in the Einstein Copilot system?

A. Generating record queries based on conversation history
B. Offering real-time language translation during conversations
C. Identifying copilot actions to respond to user utterances

Correct Answer: C
Section:
Explanation:
The primary function of the planner service in the Einstein Copilot system is to identify copilot actions that should be taken in response to user utterances. This service is responsible for analyzing the conversation and determining the appropriate actions (such as querying records, generating a response, or taking another action) that the Einstein Copilot should perform based on user input.

QUESTION 39
Universal Containers (UC) wants to enable its sales team with automatic post-call visibility into mention of competitors, products, and other custom phrases.
Which feature should the AI Specialist set up to enable UC's sales team?

A. Call Summaries
B. Call Explorer
C. Call Insights

Correct Answer: C
Section:
Explanation:
To enable Universal Containers' sales team with automatic post-call visibility into mentions of competitors, products, and custom phrases, the AI Specialist should set up Call Insights. Call Insights analyzes voice and video calls for key phrases, topics, and mentions, providing insights into critical aspects of the conversation. This feature automatically surfaces key details such as competitor mentions, product discussions, and custom phrases specified by the sales team.
Call Summaries provide a general overview of the call but do not specifically highlight keywords or topics.
Call Explorer is a tool for navigating through call data but does not focus on automatic insights.
For more information, refer to Salesforce's Call Insights documentation regarding the analysis of call content and extracting actionable information.

QUESTION 40
A sales rep at Universal Containers is extremely busy and sometimes will have very long sales calls on voice and video calls and might miss key details. They are just starting to adopt new generative AI features.
Which Einstein Generative AI feature should an AI Specialist recommend to help the rep get the details they might have missed during a conversation?

A. Call Summary
B. Call Explorer
C. Sales Summary

Correct Answer: A
Section:
Explanation:
For a sales rep who may miss key details during long sales calls, the AI Specialist should recommend the Call Summary feature. Call Summary uses Einstein Generative AI to automatically generate a concise summary of important points discussed during the call, helping the rep quickly review the key information they might have missed.
Call Explorer is designed for manually searching through call data but doesn't summarize.
Sales Summary is focused more on summarizing overall sales activity, not call-specific content.
For more details, refer to Salesforce's Call Summary documentation on how AI-generated summaries can improve sales rep productivity.

QUESTION 41
Universal Containers wants to allow its service agents to query the current fulfillment status of an order with natural language. There is an existing auto launched flow to query the information from Oracle ERP, which is the system of record for the order fulfillment process.
How should an AI Specialist apply the power of conversational AI to this use case?

A. Create a Flex prompt template in Prompt Builder.
B. Create a custom copilot action which calls a flow.
C. Configure the Integration Flow Standard Action in Einstein Copilot.

Correct Answer: B
Section:
Explanation:
To enable Universal Containers service agents to query the current fulfillment status of an order using natural language and leverage an existing auto-launched flow that queries Oracle ERP, the best solution is to create a custom copilot action that calls the flow. This action will allow Einstein Copilot to interact with the flow and retrieve the required order fulfillment information seamlessly. Custom copilot actions can be tailored to call various backend systems or flows in response to user requests.
Option B is correct because it enables integration between Einstein Copilot and the flow that connects to Oracle ERP.
Option A (Flex prompt template) is more suited for static responses and not for invoking flows.
Option C (Integration Flow Standard Action) is not directly related to creating a specific copilot action for this use case.
Salesforce Einstein Copilot Actions: https://help.salesforce.com/s/articleView?id=einstein_copilot_actions.htm

QUESTION 42
An AI Specialist built a Field Generation prompt template that worked for many records, but users are reporting random failures with token limit errors.
What is the cause of the random nature of this error?

A. The number of tokens generated by the dynamic nature of the prompt template will vary by record.
B. The template type needs to be switched to Flex to accommodate the variable amount of tokens generated by the prompt grounding.
C. The number of tokens that can be processed by the LLM varies with total user demand.

Correct Answer: A
Section:
Explanation:
The reason behind the token limit errors lies in the dynamic nature of the prompt template used in Field Generation. In Salesforce's AI generative models, each prompt and its corresponding output are subject to a token limit, which encompasses both the input and output of the large language model (LLM). Since the prompt template dynamically adjusts based on the specific data of each record, the number of tokens varies per record. Some records may generate longer outputs based on their data attributes, pushing the token count beyond the allowable limit for the LLM, resulting in token limit errors.
This behavior explains why users experience random failures---it is dependent on the specific data used in each case. For certain records, the combined input and output may fall within the token limit, while for others, it may exceed it. This variation is intrinsic to how dynamic templates interact with large language models.
Salesforce provides guidance in their documentation, stating that prompt template design should take into account token limits and suggests testing with varied records to avoid such random errors. It does not mention switching to Flex template type as a solution, nor does it suggest that token limits fluctuate with user demand. Token limits are a constant defined by the model itself, independent of external user load.
Salesforce Developer Documentation on Token Limits for Generative AI Models
Salesforce AI Best Practices on Prompt Design (Trailhead or Salesforce blog resources)

QUESTION 43
An administrator is responsible for ensuring the security and reliability of Universal Containers' (UC) CRM data. UC needs enhanced data protection and up-to-date AI capabilities. UC also needs to include relevant information from a Salesforce record to be merged with the prompt.
Which feature in the Einstein Trust Layer best supports UC's need?

A. Data masking
B. Dynamic grounding with secure data retrieval
C. Zero-data retention policy

Correct Answer: B
Section:
Explanation:
Dynamic grounding with secure data retrieval is a key feature in Salesforce's Einstein Trust Layer, which provides enhanced data protection and ensures that AI-generated outputs are both accurate and securely sourced. This feature allows relevant Salesforce data to be merged into the AI-generated responses, ensuring that the AI outputs are contextually aware and aligned with real-time CRM data.
Dynamic grounding means that AI models are dynamically retrieving relevant information from Salesforce records (such as customer records, case data, or custom object data) in a secure manner. This ensures that any sensitive data is protected during AI processing and that the AI model's outputs are trustworthy and reliable for business use.
The other options are less aligned with the requirement:
Data masking refers to obscuring sensitive data for privacy purposes and is not related to merging Salesforce records into prompts.
Zero-data retention policy ensures that AI processes do not store any user data after processing, but this does not address the need to merge Salesforce record information into a prompt.
Salesforce Developer Documentation on Einstein Trust Layer
Salesforce Security Documentation for AI and Data Privacy

QUESTION 44
A Salesforce Administrator is exploring the capabilities of Einstein Copilot to enhance user interaction within their organization. They are particularly interested in how Einstein Copilot processes user requests and the mechanism it employs to deliver responses. The administrator is evaluating whether Einstein Copilot directly interfaces with a large language model (LLM) to fetch and display responses to user inquiries, facilitating a broad range of requests from users.
How does Einstein Copilot handle user requests In Salesforce?

A. Einstein Copilot will trigger a flow that utilizes a prompt template to generate the message.
B. Einstein Copilot will perform an HTTP callout to an LLM provider.
C. Einstein Copilot analyzes the user's request and LLM technology is used to generate and display the appropriate response.

Correct Answer: C
Section:
Explanation:
Einstein Copilot is designed to enhance user interaction within Salesforce by leveraging Large Language Models (LLMs) to process and respond to user inquiries. When a user submits a request, Einstein Copilot analyzes the input using natural language processing techniques. It then utilizes LLM technology to generate an appropriate and contextually relevant response, which is displayed directly to the user within the Salesforce interface.
Option C accurately describes this process. Einstein Copilot does not necessarily trigger a flow (Option A) or perform an HTTP callout to an LLM provider (Option B) for each user request. Instead, it integrates LLM capabilities to provide immediate and intelligent responses, facilitating a broad range of user requests.
Salesforce AI Specialist Documentation - Einstein Copilot Overview: Details how Einstein Copilot employs LLMs to interpret user inputs and generate responses within the Salesforce ecosystem.
Salesforce Help - How Einstein Copilot Works: Explains the underlying mechanisms of how Einstein Copilot processes user requests using AI technologies.

QUESTION 45
Universal Containers wants to utilize Einstein for Sales to help sales reps reach their sales quotas by providing AI-generated plans containing guidance and steps for closing deals.
Which feature should the AI Specialist recommend to the sales team?

A. Find Similar Deals
B. Create Account Plan

C. Create Close Plan

Correct Answer: C
Section:
Explanation:
The 'Create Close Plan' feature is designed to help sales reps by providing AI-generated strategies and steps specifically focused on closing deals. This feature leverages AI to analyze the current state of opportunities and generate a plan that outlines the actions, timelines, and key steps required to move deals toward closure. It aligns directly with the sales team's need to meet quotas by offering actionable insights and structured plans.
Find Similar Deals (Option A) helps sales reps discover opportunities similar to their current deals but doesn't offer a plan for closing.
Create Account Plan (Option B) focuses on long-term strategies for managing accounts, which might include customer engagement and retention, but doesn't focus on deal closure.
Salesforce AI Specialist
Reference: For more information on using AI for sales, visit: https://help.salesforce.com/s/articleView?id=sf.einstein_for_sales_overview.htm