

Salesforce

Exam Questions Salesforce-AI-Specialist

Salesforce Certified AI Specialist Exam



NEW QUESTION 1

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

Answer: A

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.

NEW QUESTION 2

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.

Answer: A

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.

References:

? Salesforce Flow Documentation: <https://help.salesforce.com/s/articleView?id=sf.flow.htm>

NEW QUESTION 3

Based on the user utterance, "Show me all the customers in New York", which standard Einstein Copilot action will the planner service use?

- A. Query Records
- B. Select Records
- C. Fetch Records

Answer: A

Explanation:

The standard Einstein Copilot action that would be used in response to the user utterance, "Show me all the customers in New York", is Query Records. This action is responsible for retrieving a set of records from Salesforce based on a specified condition — in this case, filtering customers by location (New York).

? Query Records is the action that fetches relevant data based on the criteria provided in the user's input.

? Select Records is more about picking specific records from an already presented list.

? Fetch Records is not a standard term used in this context for the action. Refer to Einstein Copilot documentation on how Copilot actions work with natural language queries and data retrieval.

NEW QUESTION 4

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

Answer: B

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 References:

? 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

NEW QUESTION 5

Northern Trail Outfitters (NTO) wants to configure Einstein Trust Layer in its production org but is unable to see the option on the Setup page. After provisioning Data Cloud, which step must an AI Specialist take to make this option available to NTO?

- A. Turn on Einstein Copilot.
- B. Turn on Einstein Generative AI.
- C. Turn on Prompt Builder.

Answer: B

Explanation:

For Northern Trail Outfitters (NTO) to configure the Einstein Trust Layer, the Einstein Generative AI feature must be enabled. The Einstein Trust Layer is closely tied to generative AI capabilities, ensuring that AI-generated content complies with data privacy, security, and trust standards.

? Option A (Turning on Einstein Copilot) is unrelated to the setup of the Einstein Trust Layer, which focuses more on generative AI interactions and data handling.

? Option C (Turning on Prompt Builder) is used for configuring and building AI-driven prompts, but it does not enable the Einstein Trust Layer.

Salesforce AI Specialist References: For more details on the Einstein Trust Layer and setup steps:
https://help.salesforce.com/s/articleView?id=sf.einstein_trust_layer_overview.htm

NEW QUESTION 6

Universal Containers (UC) is Implementing Service AI Grounding to enhance its customer service operations. UC wants to ensure that its AI-generated responses are grounded in the most relevant data sources. The team needs to configure the system to include all supported objects for grounding. Which objects should UC select to configure Service AI Grounding?

- A. Case, Knowledge, and Case Notes
- B. Case and Knowledge
- C. Case, Case Emails, and Knowledge

Answer: B

Explanation:

Universal Containers (UC) is implementing Service AI Grounding to enhance its customer service operations. They aim to ensure that AI-generated responses are grounded in the most relevant data sources and need to configure the system to include all supported objects for grounding.

Supported Objects for Service AI Grounding:

- ? Case
- ? Knowledge
- ? Case Object:
- ? Knowledge Object:
- ? Exclusion of Other Objects:

Why Options A and C are Incorrect:

- ? Option A (Case, Knowledge, and Case Notes):
- ? Option C (Case, Case Emails, and Knowledge):

References:

? Salesforce AI Specialist Documentation -Service AI Grounding Configuration: Details the objects supported for grounding AI responses in Service Cloud.

? Salesforce Help -Implementing Service AI Grounding: Provides guidance on setting up grounding with Case and Knowledge objects.

? Salesforce Trailhead -Enhance Service with AI Grounding: Offers an interactive learning path on using AI grounding in service scenarios.

NEW QUESTION 7

Universal Containers wants to implement a solution in Salesforce with a custom UX that allows users to enter a sales order number. Subsequently, the system will invoke a custom prompt template to create and display a summary of the sales order header and sales order details. Which solution should an AI Specialist implement to meet this requirement?

- A. Create a screen flow to collect sales order number and invoke the prompt template using the standard "Prompt Template" flow action.
- B. Create a template-triggered prompt flow and invoke the prompt template using the standard ??Prompt Template?? flow action.
- C. Create an autolaunched flow and invoke the prompt template using the standard ??Prompt Template" flow action.

Answer: A

Explanation:

To implement a solution where users enter a sales order number and the system generates a summary, the AI Specialist should create a screen flow to collect the sales order number and invoke the prompt template. The standard "Prompt Template" flow action can then be used to trigger the custom prompt, providing a summary of the sales order header and details.

? Option B, creating a template-triggered prompt flow, is not necessary for this scenario because the requirement is to directly collect input through a screen flow.

? Option C, using an autolaunched flow, would be inappropriate here because the solution requires user interaction (entering a sales order number), which is best suited to a screen flow.

Salesforce AI Specialist References: For further guidance on creating prompt templates with flows:
https://help.salesforce.com/s/articleView?id=sf.prompt_template_flow_integration.htm

NEW QUESTION 8

Universal Containers wants to allow its service agents to query the current fulfillment status of an order with natural language. There is an existing autolaunched

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.

Answer: B

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.

References:

? Salesforce Einstein Copilot

NEW QUESTION 9

What is best practice when refining Einstein Copilot custom action instructions?

- A. Provide examples of user messages that are expected to trigger the action.
- B. Use consistent introductory phrases and verbs across multiple action instructions.
- C. Specify the persona who will request the action.

Answer: A

Explanation:

When refining Einstein Copilot custom action instructions, it is considered best practice to provide examples of user messages that are expected to trigger the action. This helps ensure that the custom action understands a variety of user inputs and can effectively respond to the intent behind the messages.

? Option B (consistent phrases) can improve clarity but does not directly refine the triggering logic.

? Option C (specifying a persona) is not as crucial as giving examples that illustrate how users will interact with the custom action.

For more details, refer to Salesforce's Einstein Copilot documentation on building and refining custom actions.

NEW QUESTION 10

Universal Containers (UC) wants to use Flow to bring data from unified Data Cloud objects to prompt templates. Which type of flow should UC use?

- A. Data Cloud-triggered flow
- B. Template-triggered prompt flow
- C. Unified-object linking flow

Answer: A

Explanation:

In this scenario, Universal Containers wants to bring data from unified Data Cloud objects into prompt templates, and the best way to do that is through a Data Cloud-triggered flow. This type of flow is specifically designed to trigger actions based on data changes within Salesforce Data Cloud objects.

Data Cloud-triggered flows can listen for changes in the unified data model and automatically bring relevant data into the system, making it available for prompt templates. This ensures that the data is both real-time and up-to-date when used in generative AI contexts.

For more detailed guidance, refer to Salesforce documentation on Data Cloud-triggered flows and Data Cloud integrations with generative AI solutions.

NEW QUESTION 10

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

Answer: C

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.

References:

? Salesforce Prompt Builder Overview: https://help.salesforce.com/s/articleView?id=sf.prompt_builder_overview.htm

NEW QUESTION 15

Universal Containers (UC) wants to assess Salesforce's generative features but has concerns over its company data being exposed to third- party large language models (LLMs). Specifically, UC wants the following capabilities to be part of Einstein's generative AI service.

No data is used for LLM training or product improvements by third- party LLMs. No data is retained outside of UC's Salesforce org.

The data sent cannot be accessed by the LLM provider.

Which property of the Einstein Trust Layer should the AI Specialist highlight to UC that addresses these requirements?

- A. Prompt Defense
- B. Zero-Data Retention Policy
- C. Data Masking

Answer: B

Explanation:

Universal Containers (UC) has concerns about data privacy when using

Salesforce's generative AI features, particularly around preventing third-party LLMs from accessing or retaining their data. The Zero-Data Retention Policy in the Einstein Trust Layer is designed to address these concerns by ensuring that:

? No data is used for training or product improvements by third-party LLMs.

? No data is retained outside of the customer's Salesforce organization.

? The LLM provider cannot access any customer data.

This policy aligns perfectly with UC's requirements for keeping their data safe while leveraging generative AI capabilities.

? Prompt Defense and Data Masking are also security features, but they do not directly address the concerns related to third-party data access and retention.

References:

? Salesforce Einstein Trust Layer Documentation: https://help.salesforce.com/s/articleView?id=sf.einstein_trust_layer.htm

NEW QUESTION 20

Universal Containers (UC) wants to enable its sales reps to explore opportunities that are similar to previously won opportunities by entering the utterance, "Show me other opportunities like this one."

How should UC achieve this in Einstein Copilot?

- A. Use the standard Copilot action.
- B. Create a custom Copilot action calling a flow.
- C. Create a custom Copilot action calling an Apex class.

Answer: A

Explanation:

Universal Containers can achieve the request to explore similar opportunities by using the standard Copilot action. Einstein Copilot has built-in actions to handle natural language queries, such as "? Show me other opportunities like this one.?" The standard action will process the query and return results based on predefined matching criteria like opportunity details and past Closed Won deals.

This approach avoids the need to create custom flows or Apex classes, leveraging out-of-the-box functionality.

For further details, refer to Einstein Copilot for Sales documentation regarding standard actions and natural language processing.

NEW QUESTION 24

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

Answer: B

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.

References:

Salesforce Einstein Documentation on Einstein Copilot Actions [Salesforce AI Documentation on Large Language Models](#)

NEW QUESTION 29

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.

Answer: C

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.

NEW QUESTION 34

Universal Containers (UC) wants to improve the efficiency of addressing customer questions and reduce agent handling time with AI-generated responses. The agents should be able to leverage their existing knowledge base and identify whether the responses are coming from the large language model (LLM) or from Salesforce Knowledge.

Which step should UC take to meet this requirement?

- A. Turn on Service AI Grounding, Grounding with Case, and Service Replies.
- B. Turn on Service Replies, Service AI Grounding, and Grounding with Knowledge.
- C. Turn on Service AI Grounding and Grounding with Knowledge.

Answer: B

Explanation:

To meet Universal Containers' goal of improving efficiency and reducing agent handling time with AI-generated responses, the best approach is to enable Service Replies, Service AI Grounding, and Grounding with Knowledge.

? Service Replies generates responses automatically.

? Service AI Grounding ensures that the AI is using relevant case data.

? Grounding with Knowledge ensures that responses are backed by Salesforce Knowledge articles, allowing agents to identify whether a response is coming from the LLM or Salesforce Knowledge.

? Option C does not include Service Replies, which is necessary for generating AI responses.

? Option A lacks the Grounding with Knowledge, which is essential for identifying response sources.

For more details, refer to Salesforce Service AI documentation on grounding and service replies.

NEW QUESTION 36

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.

Answer: C

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.

References:

? 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.

NEW QUESTION 41

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

Answer: B

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:

? Implementation Steps:

? Why Not Einstein Copilot or Model Builder:

References:

? Salesforce AI Specialist Documentation - Prompt Builder Overview:

? Salesforce Help - Creating Field Generation Prompt Templates:

? Salesforce Trailhead - Customize AI Content with Prompt Builder:

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.

NEW QUESTION 46

Universal Containers is interested in improving the sales operation efficiency by analyzing their data using AI-powered predictions in Einstein Studio.

Which use case works for this scenario?

- A. Predict customer sentiment toward a promotion message.
- B. Predict customer lifetime value of an account.
- C. Predict most popular products from new product catalog.

Answer: B

Explanation:

For improving sales operations efficiency, Einstein Studio is ideal for creating AI-powered models that can predict outcomes based on data. One of the most valuable use cases is predicting customer lifetime value, which helps sales teams focus on high-value accounts and make more informed decisions. Customer lifetime value (CLV) predictions can optimize strategies around customer retention, cross-selling, and long-term engagement.

? Option B is the correct choice as predicting customer lifetime value is a well-established use case for AI in sales.

? Option A (customer sentiment) is typically handled through NLP models, while Option C (product popularity) is more of a marketing analysis use case.

References:

? Salesforce Einstein Studio Use Case Overview: https://help.salesforce.com/s/articleView?id=sf.einstein_studio_overview

NEW QUESTION 48

Universal Containers (UC) wants to offer personalized service experiences and reduce agent handling time with AI-generated email responses, grounded in Knowledge base.

Which AI capability should UC use?

- A. Einstein Email Replies
- B. Einstein Service Replies for Email
- C. Einstein Generative Service Replies for Email

Answer: B

Explanation:

For Universal Containers (UC) to offer personalized service experiences and reduce agent handling time using AI-generated responses grounded in the Knowledge base, the best solution is Einstein Service Replies for Email. This capability leverages AI to automatically generate responses to service-related emails based on historical data and the Knowledge base, ensuring accuracy and relevance while saving time for service agents.

? Einstein Email Replies (option A) is more suited for sales use cases.

? Einstein Generative Service Replies for Email (option C) could be a future offering, but as of now, Einstein Service Replies for Email is the correct choice for grounded, knowledge-based responses.

References:

? Einstein Service Replies Overview: https://help.salesforce.com/s/articleView?id=sf.einstein_service_replies.htm

NEW QUESTION 50

Universal Containers (UC) has a mature Salesforce org with a lot of data in cases and Knowledge articles. UC is concerned that there are many legacy fields, with data that might not be applicable for Einstein AI to draft accurate email responses.

Which solution should UC use to ensure Einstein AI can draft responses from a defined data source?

- A. Service AI Grounding
- B. Work Summaries
- C. Service Replies

Answer: A

Explanation:

Service AI Grounding is the solution that Universal Containers should use to ensure Einstein AI drafts responses based on a well-defined data source. Service AI Grounding allows the AI model to be anchored in specific, relevant data sources, ensuring that any AI-generated responses (e.g., email replies) are accurate, relevant, and drawn from up-to-date information, such as Knowledge articles or cases.

Given that UC has legacy fields and outdated data, Service AI Grounding ensures that only the valid and applicable data is used by Einstein AI to craft responses. This helps improve the relevance of responses and avoids inaccuracies caused by outdated or irrelevant fields. Work Summaries and Service Replies are useful features but do not address the need for grounding AI outputs in specific, current data sources like Service AI Grounding does.

For more details, you can refer to Salesforce's Service AI Grounding documentation for managing AI-generated content based on accurate data sources.

NEW QUESTION 54

An AI Specialist at Universal Containers (UC) is tasked with creating a new custom prompt template to populate a field with generated output. UC enabled the Einstein Trust Layer to ensure AI Audit data is captured and monitored for adoption and possible enhancements.

Which prompt template type should the AI Specialist use and which consideration should they review?

- A. Flex, and that Dynamic Fields is enabled
- B. Field Generation, and that Dynamic Fields is enabled
- C. Field Generation, and that Dynamic Forms is enabled

Answer: B

Explanation:

When creating a custom prompt template to populate a field with generated output, the most appropriate template type is Field Generation. This template is specifically designed for generating field-specific outputs using generative AI.

Additionally, the AI Specialist must ensure that Dynamic Fields are enabled. Dynamic Fields allow the system to use real-time data inputs from related records or fields when generating content, ensuring that the AI output is contextually accurate and relevant. This is crucial when populating specific fields with AI-generated content, as it ensures the data source remains dynamic and up-to-date.

The Einstein Trust Layer will track and audit the interactions to ensure the organization can monitor AI adoption and make necessary enhancements based on AI usage patterns.

For further reading, refer to Salesforce's guidelines on Field Generation templates and the Einstein Trust Layer.

NEW QUESTION 55

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.

Answer: A

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.

References:

- ? Salesforce Developer Documentation onToken Limits for Generative AI Models
- ? Salesforce AI Best Practices on Prompt Design (Trailhead or Salesforce blog resources)

NEW QUESTION 59

An AI Specialist wants to include data from the response of an 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.

Answer: B

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:

? External Service Record Merge Fields:

Implementation Steps:

? Register the External Service:

? Create a Named Credential:

? Use External Service in Flow:

? Configure the Prompt Template:

Why Other Options are Less Suitable:

? Option A (Convert the JSON to an XML merge field):

? Option C (Use ??Add Prompt Instructions?? flow element):

References:

? Salesforce AI Specialist Documentation -Integrating External Services with Prompt Templates:

? Salesforce Help -Using Merge Fields with External Data:

? Salesforce Trailhead -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.

NEW QUESTION 63

Which feature in the Einstein Trust Layer helps to minimize the risks of jailbreaking and prompt injection attacks?

- A. Secure Data Retrieval and Grounding
- B. Data Masking
- C. Prompt Defense

Answer: C

Explanation:

Prompt Defense is a feature in the Einstein Trust Layer that helps minimize the risks of jailbreaking and prompt injection attacks. These attacks occur when malicious users try to manipulate the AI model by providing unintended inputs. Prompt Defense ensures that the prompts are processed securely, protecting the system from such vulnerabilities.

? Option A (Secure Data Retrieval and Grounding) relates to ensuring that data used

by AI is securely retrieved but does not address prompt security.

? Option B (Data Masking) focuses on protecting sensitive information but does not prevent injection attacks.

For more information, refer to Salesforce's Einstein Trust Layer documentation on Prompt Defense and security features.

NEW QUESTION 66

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.

Answer: A

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 References: Refer to the documentation on troubleshooting Flex templates

here: https://help.salesforce.com/s/articleView?id=sf.flex_prompt_builder_troubleshoot.htm

NEW QUESTION 70

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.

Answer: B

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.

NEW QUESTION 71

Universal Containers plans to implement prompt templates that utilize the standard foundation models.

What should the AI Specialist consider when building prompt templates in Prompt Builder?

A. Include multiple-choice questions within the prompt to test the LLM's understanding of the context.

B. Ask it to role-play as a character in the prompt template to provide more context to the LLM.

C. Train LLM with data using different writing styles including word choice, intensifiers, emojis, and punctuation.

Answer: C

Explanation:

When building prompt templates in Prompt Builder, it is essential to consider how the Large Language Model (LLM) processes and generates outputs. Training the LLM with various writing styles, such as different word choices, intensifiers, emojis, and punctuation, helps the model better understand diverse writing patterns and produce more contextually appropriate responses.

This approach enhances the flexibility and accuracy of the LLM when generating outputs for different use cases, as it is trained to recognize various writing conventions and styles. The prompt template should focus on providing rich context, and this stylistic variety helps improve the model's adaptability.

Options A and B are less relevant because adding multiple-choice questions or role-playing scenarios doesn't contribute significantly to improving the AI's output generation quality within standard business contexts.

For more details, refer to Salesforce's Prompt Builder documentation and LLM tuning strategies.

NEW QUESTION 75

Universal Containers needs a tool that can analyze voice and video call records to provide insights on competitor mentions, coaching opportunities, and other key information. The goal is to enhance the team's performance by identifying areas for improvement and competitive intelligence.

Which feature provides insights about competitor mentions and coaching opportunities?

A. Call Summaries

B. Einstein Sales Insights

C. Call Explorer

Answer: C

Explanation:

For analyzing voice and video call records to gain insights into competitor mentions, coaching opportunities, and other key information, Call Explorer is the most suitable feature. Call Explorer, a part of Einstein Conversation Insights, enables sales teams to analyze calls, detect patterns, and identify areas where improvements can be

made. It uses natural language processing (NLP) to extract insights, including competitor mentions and moments for coaching. These insights are vital for improving sales performance by providing a clear understanding of the interactions during calls.

? Call Summaries offer a quick overview of a call but do not delve deep into competitor mentions or coaching insights.

? Einstein Sales Insights focuses more on pipeline and forecasting insights rather than call-based analysis.

References:

? Salesforce Einstein Conversation Insights Documentation: https://help.salesforce.com/s/articleView?id=einstein_conversation_insights.htm

NEW QUESTION 80

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.

Answer: A

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.

NEW QUESTION 82

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.

Answer: A

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 References: To explore more about dynamic debugging in Copilot Builder,

see: https://help.salesforce.com/s/articleView?id=sf.copilot_custom_action_debugging.htm

NEW QUESTION 84

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

Salesforce-AI-Specialist Practice Exam Features:

- * Salesforce-AI-Specialist Questions and Answers Updated Frequently
- * Salesforce-AI-Specialist Practice Questions Verified by Expert Senior Certified Staff
- * Salesforce-AI-Specialist Most Realistic Questions that Guarantee you a Pass on Your First Try
- * Salesforce-AI-Specialist Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The Salesforce-AI-Specialist Practice Test Here](#)