

# Salesforce

## Exam Questions Salesforce-AI-Specialist

Salesforce Certified AI Specialist Exam



### NEW QUESTION 1

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.

**Answer: B**

#### 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.

References:

? Salesforce Email Template and Permissions Documentation: [https://help.salesforce.com/s/articleView?id=sf.email\\_templates\\_field\\_permissions.htm](https://help.salesforce.com/s/articleView?id=sf.email_templates_field_permissions.htm)

### NEW QUESTION 2

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.

**Answer: A**

#### 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:

? Refers to invoking a flow to create a Knowledge article, which is unrelated to the task of retrieving existing Knowledge articles.

? Focuses on generating Knowledge articles, which is not the immediate need for this situation where the agent requires guidance on existing procedures.

References:

? Salesforce Documentation on Einstein Copilot

? Trailhead Module on Einstein for Service

### NEW QUESTION 3

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 4

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 5

Universal Containers is planning a marketing email about products that most closely match a customer's expressed interests.

What should an AI Specialist recommend to generate this email?

- A. Standard email marketing template using Apex or flows for matching interest in products
- B. Custom sales email template which is grounded with interest and product information
- C. Standard email draft with Einstein and choose standard email template

**Answer: B**

#### Explanation:

To generate an email about products that closely match a customer's expressed interests, an AI Specialist should recommend using a custom sales email template that is grounded with interest and product information. This ensures that the email content is personalized based on the customer's preferences, increasing the relevance of the marketing message.

Using grounding ensures that the generative AI pulls the correct data related to customer interests and product matches, making the email more effective.

For more information, refer to Salesforce documentation on grounding AI-generated content and email personalization strategies.

#### NEW QUESTION 6

Universal Containers (UC) noticed an increase in customer contract cancellations in the last few months. UC is seeking ways to address this issue by implementing a proactive outreach program to

customers before they cancel their contracts and is asking the Salesforce team to provide suggestions.

Which use case functionality of Model Builder aligns with UC's request?

- A. Product recommendation prediction
- B. Customer churn prediction
- C. Contract Renewal Date prediction

**Answer: B**

#### Explanation:

Customer churn prediction is the best use case for Model Builder in addressing Universal Containers' concerns about increasing customer contract cancellations. By implementing a model that predicts customer churn, UC can proactively identify customers who are at risk of canceling and take action to retain them before they decide to terminate their contracts. This functionality allows the business to forecast churn probability based on historical data and initiate timely outreach programs.

? Option B is correct because customer churn prediction aligns with UC's need to reduce cancellations through proactive measures.

? Option A (product recommendation prediction) is unrelated to contract cancellations.

? Option C (contract renewal date prediction) addresses timing but does not focus on predicting potential cancellations.

References:

? Salesforce Model Builder Use Case Overview: [https://help.salesforce.com/s/articleView?id=sf.model\\_builder\\_use\\_cases.htm](https://help.salesforce.com/s/articleView?id=sf.model_builder_use_cases.htm)

#### NEW QUESTION 7

Universal Containers plans to enhance the customer support team's productivity using AI. Which specific use case necessitates the use of Prompt Builder?

- A. Creating a draft of a support bulletin post for new product patches
- B. Creating an AI-generated customer support agent performance score
- C. Estimating support ticket volume based on historical data and seasonal trends

**Answer: A**

#### Explanation:

The use case that necessitates the use of Prompt Builder is creating a draft of a support bulletin post for new product patches. Prompt Builder allows the AI Specialist to create and refine prompts that generate specific, relevant outputs, such as drafting support communication based on product information and patch details.

? Option B (agent performance score) would likely involve predictive modeling, not prompt generation.

? Option C (estimating support ticket volume) would require data analysis and predictive tools, not prompt building.

For more details, refer to Salesforce's Prompt Builder documentation for generative AI content creation.

#### NEW QUESTION 8

When a customer chat is initiated, which functionality in Salesforce provides generative AI replies or draft emails based on recommended Knowledge articles?

- A. Einstein Reply Recommendations
- B. Einstein Service Replies
- C. Einstein Grounding

**Answer: B**

#### Explanation:

When a customer chat is initiated, Einstein Service Replies provides generative AI replies or draft emails based on recommended Knowledge articles. This feature

uses the information from the Salesforce Knowledge base to generate responses that are relevant to the customer's query, improving the efficiency and accuracy of customer support interactions.

? Option B is correct because Einstein Service Replies is responsible for generating AI-driven responses based on knowledge articles.

? Option A (Einstein Reply Recommendations) is focused on recommending replies but does not generate them.

? Option C (Einstein Grounding) refers to grounding responses in data but is not directly related to drafting replies.

References:

? Einstein Service Replies Overview: [https://help.salesforce.com/s/articleView?id=sf.einstein\\_service\\_replies.htm](https://help.salesforce.com/s/articleView?id=sf.einstein_service_replies.htm)

### NEW QUESTION 9

Universal Containers (UC) has recently received an increased number of support cases. As a result, UC has hired more customer support reps and has started to assign some of the ongoing cases to newer reps.

Which generative AI solution should the new support reps use to understand the details of a case without reading through each case comment?

- A. Einstein Copilot
- B. Einstein Sales Summaries
- C. Einstein Work Summaries

**Answer: C**

#### Explanation:

New customer support reps at Universal Containers can use Einstein Work Summaries to quickly understand the details of a case without reading through each case comment. Work Summaries leverage generative AI to provide a concise overview of ongoing cases, summarizing all relevant information in an easily digestible format.

? Einstein Copilot can assist with a variety of tasks but is not specifically designed for summarizing case details.

? Einstein Sales Summaries are focused on summarizing sales-related activities, which is not applicable for support cases.

For more details, refer to Salesforce documentation on Einstein Work Summaries.

### NEW QUESTION 10

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

**Answer: A**

#### 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 References: For more details, check the Model Builder guide

here: [https://help.salesforce.com/s/articleView?id=sf.model\\_builder\\_external\\_models.htm](https://help.salesforce.com/s/articleView?id=sf.model_builder_external_models.htm)

### NEW QUESTION 10

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

**Answer: B**

#### 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 References: For more information on Einstein Reply

Recommendations: [https://help.salesforce.com/s/articleView?id=sf.einstein\\_reply\\_recommendations\\_overview.htm](https://help.salesforce.com/s/articleView?id=sf.einstein_reply_recommendations_overview.htm)

### NEW QUESTION 13

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](https://help.salesforce.com/s/articleView?id=sf.prompt_builder_overview.htm)

**NEW QUESTION 18**

An AI Specialist at Universal Containers is working on a prompt template to generate personalized emails for product demonstration requests from customers. It is important for the AI-generated email to adhere strictly to the guidelines, using only associated opportunity information, and to encourage the recipient to take the desired action.

How should the AI Specialist include these instructions on a new line in the prompt template?

- A. Surround them with triple quotes (""").
- B. Make sure merged fields are defined.
- C. Use curly brackets {} to encapsulate instructions.

**Answer:** A

**Explanation:**

In Salesforce prompt templates, instructions that guide how the Large Language Model (LLM) should generate content (in this case, personalized emails) can be included by surrounding the instruction text with triple quotes ("""). This formatting ensures that the LLM adheres to the specific instructions while generating the email content.

The use of triple quotes allows the AI to understand that the enclosed text is a directive for how to approach the task, such as limiting the content to associated opportunity information or encouraging a specific action from the recipient.

Refer to Salesforce Prompt Builder documentation for detailed instructions on how to structure prompts for generative AI.

**NEW QUESTION 22**

An AI Specialist wants to use the related lists from an account in a custom prompt template.

What should the AI Specialist consider when configuring the prompt template?

- A. The text encoding (for example, UTF-8, ASCII) option
- B. The maximum number of related list merge fields
- C. The choice between XML and JSON rendering formats for the list

**Answer:** B

**Explanation:**

When configuring a custom prompt template to use related lists, the AI Specialist must be aware of the maximum number of related list merge fields that can be included. Salesforce enforces limits to ensure prompt templates perform efficiently and do not overload the system with too much data. As a best practice, it's important to monitor and optimize the number of merge fields used.

? Option B is correct because there is a limit on how many related list merge fields can be included in a prompt template.

? Option A (text encoding) and Option C (XML/JSON rendering) are not key considerations in this context.

References:

? Salesforce Prompt Builder Documentation: [https://help.salesforce.com/s/articleView?id=sf.prompt\\_builder.htm](https://help.salesforce.com/s/articleView?id=sf.prompt_builder.htm)

**NEW QUESTION 26**

Universal Containers is considering leveraging the Einstein Trust Layer in conjunction with Einstein Generative AI Audit Data.

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

**Answer:** C

**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.

References:

? 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.

**NEW QUESTION 30**

Universal Containers is very concerned about security compliance and wants to understand:

Which prompt text is sent to the large language model (LLM)

\* How it is masked

\* The masked response

What should the AI Specialist recommend?

- A. Ingest the Einstein Shield Event logs into CRM Analytics.
- B. Review the debug logs of the running user.
- C. Enable audit trail in the Einstein Trust Layer.

**Answer:** C

**Explanation:**

To address security compliance concerns and provide visibility into the prompt text sent to the LLM, how it is masked, and the masked response, the AI Specialist should recommend enabling the audit trail in the Einstein Trust Layer. This feature captures and logs the prompts sent to the large language model (LLM) along with the masking of sensitive information and the AI's response. This audit trail ensures full transparency and compliance with security requirements.

? Option A (Einstein Shield Event logs) is focused on system events rather than specific AI prompt data.

? Option B (debug logs) would not provide the necessary insight into AI prompt masking or responses.

For further details, refer to Salesforce's Einstein Trust Layer documentation about auditing and security measures.

**NEW QUESTION 32**

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 detailed reports on all Copilot conversations over any time period.

**Answer:** B

**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 References: You can refer to this documentation for further

insights: [https://help.salesforce.com/s/articleView?id=sf.einstein\\_copilot\\_event\\_logging.htm](https://help.salesforce.com/s/articleView?id=sf.einstein_copilot_event_logging.htm)

**NEW QUESTION 33**

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.

**Answer:** B

**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.

References:

? Salesforce Prompt Builder Documentation: [https://help.salesforce.com/s/articleView?id=sf.prompt\\_builder\\_overview.htm](https://help.salesforce.com/s/articleView?id=sf.prompt_builder_overview.htm)

**NEW QUESTION 37**

An AI Specialist is considering using a Field Generation prompt template type.

What should the AI Specialist check before creating the Field Generation prompt to ensure it is possible for the field to be enabled for generative AI?

- A. That the field chosen must be a rich text field with 255 characters or more.
- B. That the org is set to API version 59 or higher
- C. That the Lightning page layout where the field will reside has been upgraded to Dynamic Forms

**Answer:** B

**Explanation:**

Before creating a Field Generation prompt template, the AI Specialist must ensure that the Salesforce org is set to API version 59 or higher. This version of the API introduces support for advanced generative AI features, such as enabling fields for generative AI outputs. This is a critical technical requirement for the Field Generation prompt template to function correctly.

? Option A (rich text field requirement) is not necessary for generative AI functionality.

? Option C (Dynamic Forms) does not impact the ability of a field to be generative AI-enabled, although it might enhance the user interface.

For more information, refer to Salesforce documentation on API versioning and Field Generation templates.

**NEW QUESTION 42**

The AI Specialist of Northern Trail Outfitters reviewed the organization's data masking settings within the Configure Data Masking menu within Setup. Upon assessing all of the fields, a few additional fields were deemed sensitive and have been masked within Einstein's Trust Layer. Which steps should the AI Specialist take upon modifying the masked fields?

- A. Turn off the Einstein Trust Layer and turn it on again.
- B. Test and confirm that the responses generated from prompts that utilize the data and masked data do not adversely affect the quality of the generated response
- C. Turn on Einstein Feedback so that end users can report if there are any negative side effects on AI features.

**Answer: B**

**Explanation:**

After modifying masked fields in Einstein's Trust Layer, the next important step is to test and confirm that the responses generated by prompts utilizing the newly masked data still meet quality standards. This ensures that masking sensitive information does not negatively impact the usefulness or accuracy of the AI-generated content. Thorough testing helps identify any issues in prompt performance that could arise due to masking, and adjustments can be made if needed.

? Option B is correct because testing the effects of masking on AI responses is a critical step in ensuring AI continues to function as expected.

? Option A (turning off and on the Einstein Trust Layer) is unnecessary after changing the masked fields.

? Option C (turning on Einstein Feedback) allows for user feedback but is not a direct step following field masking modifications.

References:

? Salesforce Einstein Trust Layer Overview: [https://help.salesforce.com/s/articleView?id=sf.einstein\\_trust\\_layer.htm](https://help.salesforce.com/s/articleView?id=sf.einstein_trust_layer.htm)

**NEW QUESTION 45**

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.

**Answer: B**

**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.

References:

? Salesforce Einstein Trust Layer Documentation: [https://help.salesforce.com/s/articleView?id=sf.einstein\\_trust\\_layer\\_audit.htm](https://help.salesforce.com/s/articleView?id=sf.einstein_trust_layer_audit.htm)

**NEW QUESTION 48**

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 52**

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 57**

An AI Specialist is creating a custom action in Einstein Copilot. Which option is available for the AI Specialist to choose for the custom copilot action?

- A. Apex trigger
- B. SOQL
- C. Flows

**Answer: C**

**Explanation:**

When creating a custom action in Einstein Copilot, one of the available options is to use Flows. Flows are a powerful automation tool in Salesforce, allowing the AI Specialist to define custom logic and actions within the Copilot system. This makes it easy to extend Copilot's functionality without needing custom code. While Apex triggers and SOQL are important Salesforce tools, Flows are the recommended method for creating custom actions within Einstein Copilot because they are declarative and highly adaptable. For further guidance, refer to Salesforce Flow documentation and Einstein Copilot customization resources.

**NEW QUESTION 59**

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

**NEW QUESTION 63**

Which use case is best supported by Salesforce Einstein Copilot's capabilities?

- A. Bring together a conversational interface for interacting with AI for all Salesforce users, such as developers and ecommerce retailers.
- B. Enable Salesforce admin users to create and train custom large language models (LLMs) using CRM data.
- C. Enable data scientists to train predictive AI models with historical CRM data using built-in machine learning capabilities

**Answer: A**

**Explanation:**

Salesforce Einstein Copilot is designed to provide a conversational AI interface that can be utilized by different types of Salesforce users, such as developers, sales agents, and retailers. It acts as an AI-powered assistant that facilitates natural interactions with the system, enabling users to perform tasks and access data easily. This includes tasks like pulling reports, updating records, and generating personalized responses in real time.

- ? Option A is correct because Einstein Copilot brings a conversational interface that caters to a wide range of users.
- ? Option B and Option C are more focused on developing and training AI models, which are not the primary functions of Einstein Copilot.

References:

- ? Salesforce Einstein Copilot Overview: [https://help.salesforce.com/s/articleView?id=einstein\\_copilot\\_overview.htm](https://help.salesforce.com/s/articleView?id=einstein_copilot_overview.htm)

**NEW QUESTION 66**

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.

**Answer: B**

**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.

References:

? Salesforce Prompt Builder Overview: [https://help.salesforce.com/s/articleView?id=sf.prompt\\_builder\\_overview.htm](https://help.salesforce.com/s/articleView?id=sf.prompt_builder_overview.htm)

**NEW QUESTION 68**

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 73**

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](https://help.salesforce.com/s/articleView?id=einstein_conversation_insights.htm)

**NEW QUESTION 75**

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 77**

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

**Answer:** A

**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.

References:

? Salesforce Permissions Documentation: [https://help.salesforce.com/s/articleView?id=sf.perm\\_sets\\_overview.htm](https://help.salesforce.com/s/articleView?id=sf.perm_sets_overview.htm)

**NEW QUESTION 78**

An AI 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.

**Answer:** B

**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.

**NEW QUESTION 80**

.....

## **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](#)**