

Exam Questions Salesforce-AI-Specialist

Salesforce Certified AI Specialist Exam

<https://www.2passeasy.com/dumps/Salesforce-AI-Specialist/>



NEW QUESTION 1

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 2

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 3

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

Answer: C

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.

NEW QUESTION 4

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_case_s.htm

NEW QUESTION 5

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 6

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

NEW QUESTION 7

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

Answer: A

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.

NEW QUESTION 8

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.

Answer: A

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.

References:

? 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.

NEW QUESTION 9

Universal Containers implements Custom Copilot Actions to enhance its customer service operations. The development team needs to understand the core components of a Custom Copilot Action to ensure proper configuration and functionality.

What should the development team review in the Custom Copilot Action configuration to identify one of the core components of a Custom Copilot Action?

- A. Instructions
- B. Output Types
- C. Action Triggers

Answer: B

Explanation:

Universal Containers is enhancing its customer service operations with Custom Copilot Actions. The development team needs to understand the core components of a Custom Copilot Action to ensure proper configuration and functionality. One of these core components is the Output Types.

? Core Components of a Custom Copilot Action:

? Focus on Output Types:

Why Output Types are a Core Component:

? Integration with Copilot:

? Data Consistency:

? User Experience:

Why Other Options are Less Suitable:

? Option A (Instructions):

? Option C (Action Triggers):

References:

? Salesforce AI Specialist Documentation - Custom Copilot Actions:

? Salesforce Help - Defining Output Types in Custom Actions:

? Salesforce Trailhead - Building Custom Copilot Actions:

NEW QUESTION 10

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.

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 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 14

Before activating a custom copilot action, an AI Specialist would like to understand multiple real-world user utterances to ensure the action being selected appropriately.

Which tool should the AI Specialist recommend?

- A. Model Playground
- B. Einstein Copilot
- C. Copilot Builder

Answer: C

Explanation:

To understand multiple real-world user utterances and ensure the correct action is selected before activating a custom copilot action, the recommended tool is Copilot Builder. This tool allows AI Specialists to design and test conversational actions in response to user inputs, helping ensure the copilot can accurately handle different user queries and phrases. Copilot Builder provides the ability to test, refine, and improve actions based on real-world utterances.

? Option C is correct as Copilot Builder is designed for configuring and testing conversational actions.

? Option A (Model Playground) is used for testing models, not user utterances.

? Option B (Einstein Copilot) refers to the conversational interface but isn't the right tool for designing and testing actions.

References:

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

NEW QUESTION 17

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 18

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

NEW QUESTION 23

An AI Specialist turned on Einstein Generative AI in Setup. Now, the AI Specialist would like to create custom prompt templates in Prompt Builder. However, they cannot access Prompt Builder in the Setup menu.

What is causing the problem?

- A. The Prompt Template User permission set was not assigned correctly.
- B. The Prompt Template Manager permission set was not assigned correctly.
- C. The large language model (LLM) was not configured correctly in Data Cloud.

Answer: B

Explanation:

In order to access and create custom prompt templates in Prompt Builder, the AI Specialist must have the Prompt Template Manager permission set assigned. Without this permission, they will not be able to access Prompt Builder in the Setup menu, even though Einstein Generative AI is enabled.

? Option B is correct because the Prompt Template Manager permission set is required to use Prompt Builder.

? Option A (Prompt Template User permission set) is incorrect because this permission allows users to use prompts, but not create or manage them.

? Option C (LLM configuration in Data Cloud) is unrelated to the ability to access Prompt Builder.

References:

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

NEW QUESTION 26

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 27

Universal Containers is evaluating Einstein Generative AI features to improve the productivity of the service center operation.

Which features should the AI Specialist recommend?

- A. Service Replies and Case Summaries
- B. Service Replies and Work Summaries
- C. Reply Recommendations and Sales Summaries

Answer: A

Explanation:

To improve the productivity of the service center, the AI Specialist should recommend the Service Replies and Case Summaries features.

? Service Replies helps agents by automatically generating suggested responses to customer inquiries, reducing response time and improving efficiency.

? Case Summaries provide a quick overview of case details, allowing agents to get up to speed faster on customer issues.

? Work Summaries are not as relevant for direct customer service operations, and Sales Summaries are focused on sales processes, not service center productivity.

For more information, see Salesforce's Einstein Service Cloud documentation on the use of generative AI to assist customer service teams.

NEW QUESTION 28

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

NEW QUESTION 30

Universal Containers (UC) is looking to enhance its operational efficiency. UC has recently adopted Salesforce and is considering implementing Einstein Copilot to improve its processes.

What is a key reason for implementing Einstein Copilot?

- A. Improving data entry and data cleansing
- B. Allowing AI to perform tasks without user interaction
- C. Streamlining workflows and automating repetitive tasks

Answer: C

Explanation:

The key reason for implementing Einstein Copilot is its ability to streamline workflows and automate repetitive tasks. By leveraging AI, Einstein Copilot can assist users in handling mundane, repetitive processes, such as automatically generating insights, completing actions, and guiding users through complex processes, all of which significantly improve operational efficiency.

? Option A (Improving data entry and cleansing) is not the primary purpose of Einstein Copilot, as its focus is on guiding and assisting users through workflows.

? Option B (Allowing AI to perform tasks without user interaction) does not accurately describe the role of Einstein Copilot, which operates interactively to assist users in real time.

Salesforce AI Specialist References: More details can be found in the Salesforce documentation: https://help.salesforce.com/s/articleView?id=sf.einstein_copilot_overview.htm

NEW QUESTION 33

Universal Containers (UC) is implementing Einstein Generative AI to improve customer insights and interactions. UC needs audit and feedback data to be accessible for reporting purposes. What is a consideration for this requirement?

- A. Storing this data requires Data Cloud to be provisioned.
- B. Storing this data requires a custom object for data to be configured.
- C. Storing this data requires Salesforce big objects.

Answer: A

Explanation:

When implementing Einstein Generative AI for improved customer insights and interactions, the Data Cloud is a key consideration for storing and managing large-scale audit and feedback data. The Salesforce Data Cloud (formerly known as Customer 360 Audiences) is designed to handle and unify massive datasets from various sources, making it ideal for storing data required for AI-powered insights and reporting. By provisioning Data Cloud, organizations like Universal Containers (UC) can gain real-time access to customer data, making it a central repository for unified reporting across various systems.

? Audit and feedback data generated by Einstein Generative AI needs to be stored in

a scalable and accessible environment, and the Data Cloud provides this capability, ensuring that data can be easily accessed for reporting, analytics, and further model improvement.

? Custom objects or Salesforce Big Objects are not designed for the scale or the

specific type of real-time, unified data processing required in such AI-driven interactions. Big Objects are more suited for archival data, whereas Data Cloud ensures more robust processing, segmentation, and analysis capabilities.

References:

? Salesforce Data Cloud Documentation: <https://www.salesforce.com/products/data-cloud/overview/>

? Salesforce Einstein AI Overview: <https://www.salesforce.com/products/einstein/overview/>

NEW QUESTION 38

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

NEW QUESTION 40

Universal Containers (UC) plans to send one of three different emails to its customers based on the customer's lifetime value score and their market segment. Considering that UC are required to explain why an e-mail was selected, which AI model should UC use to achieve this?

- A. Predictive model and generative model
- B. Generative model
- C. Predictive model

Answer: C

Explanation:

Universal Containers should use a Predictive model to decide which of the three emails to send based on the customer's lifetime value score and market segment. Predictive models analyze data to forecast outcomes, and in this case, it would predict the most appropriate email to send based on customer attributes.

Additionally, predictive models can provide explainability to show why a certain email was chosen, which is crucial for UC's requirement to explain the decision-

making process.

? Generative models are typically used for content creation, not decision-making, and thus wouldn't be suitable for this requirement.

? Predictive models offer the ability to explain why a particular decision was made, which aligns with UC's needs.

Refer to Salesforce's Predictive AI model documentation for more insights on how predictive models are used for segmentation and decision making.

NEW QUESTION 41

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

NEW QUESTION 43

What is the main purpose of Prompt Builder?

A. A tool for developers to use in Visual Studio Code that creates prompts for Apex programming, assisting developers in writing code more efficiently.

B. A tool that enables companies to create reusable prompts for large language models (LLMs), bringing generative AI responses to their flow of work

C. A tool within Salesforce offering real-time AI-powered suggestions and guidance to users, improving productivity and decision-making.

Answer: B

Explanation:

Prompt Builder is designed to help organizations create and configure reusable prompts for large language models (LLMs). By integrating generative AI responses into workflows, Prompt Builder enables customization of AI prompts that interact with Salesforce data and automate complex processes. This tool is especially useful for creating tailored and consistent AI-generated content in various business contexts, including customer service and sales.

? It is not a tool for Apex programming (as in option A).

? It is also not limited to real-time suggestions as mentioned in option C. Instead, it provides a flexible way for companies to manage and customize how AI-driven responses are generated and used in their workflows.

References:

? Salesforce Prompt Builder

Overview: https://help.salesforce.com/s/articleView?id=sf.prompt_builder.htm

NEW QUESTION 45

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

Answer: A

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.

NEW QUESTION 47

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 51

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 52

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 56

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 61

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.

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 62

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

NEW QUESTION 64

Universal Containers has seen a high adoption rate of a new feature that uses generative AI to populate a summary field of a custom object, Competitor Analysis. All sales users have the same profile but one user cannot see the generative AI-enabled field icon next to the summary field.

What is the most likely cause of the issue?

- A. The user does not have the Prompt Template User permission set assigned.
- B. The prompt template associated with summary field is not activated for that user.
- C. The user does not have the field Generative AI User permission set assigned.

Answer: C

Explanation:

In Salesforce, Generative AI capabilities are controlled by specific permission sets. To use features such as generating summaries with AI, users need to have the correct permission sets that allow access to these functionalities.

? Generative AI User Permission Set: This is a key permission set required to enable the generative AI capabilities for a user. In this case, the missing Generative AI User permission set prevents the user from seeing the generative AI-enabled field icon. Without this permission, the generative AI feature in the Competitor Analysis custom object won't be accessible.

? Why not A? The Prompt Template User permission set relates specifically to users who need access to prompt templates for interacting with Einstein GPT, but it's not directly related to the visibility of AI-enabled field icons.

? Why not B? While a prompt template might need to be activated, this is not the primary issue here. The question states that other users with the same profile can see the icon, so the problem is more likely to be permissions-based for this particular user.

For more detailed information, you can review Salesforce documentation on permission sets related to AI capabilities at Salesforce AI Documentation and Einstein GPT permissioning guidelines.

NEW QUESTION 69

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.

Answer: A

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:

? Processing by the LLM:

? De-masking in the Response Journey:

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):

? Option C (Responses that do not meet the relevance threshold will be automatically rejected):

References:

? Salesforce AI Specialist Documentation -Einstein Trust Layer Overview:

? Salesforce Help -Data Masking and De-masking Process:

? Salesforce AI Specialist Exam Guide -Security and Compliance in AI:

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.

NEW QUESTION 72

In Model Playground, which hyperparameters of an existing Salesforce-enabled foundational model can an AI Specialist change?

A. Temperature, Frequency Penalty, Presence Penalty

B. Temperature, Top-k sampling, Presence Penalty

C. Temperature, Frequency Penalty, Output Tokens

Answer: A

Explanation:

In Model Playground, an AI specialist working with a Salesforce-enabled foundational model has control over specific hyperparameters that can directly affect the behavior of the generative model:

? Temperature: Controls the randomness of predictions. A higher temperature leads

to more diverse outputs, while a lower temperature makes the model's responses more focused and deterministic.

? Frequency Penalty: Reduces the likelihood of the model repeating the same phrases or outputs frequently.

? Presence Penalty: Encourages the model to introduce new topics in its responses, rather than sticking with familiar, previously mentioned content.

These hyperparameters are adjustable to fine-tune the model's responses, ensuring that it meets the desired behavior and use case requirements. Salesforce documentation confirms that these three are the key tunable hyperparameters in the Model Playground. For more details, refer to Salesforce AI Model Playground guidance from Salesforce's official documentation on foundational model adjustments.

NEW QUESTION 73

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 75

Where should the AI Specialist go to add/update actions assigned to a copilot?

A. Copilot Actions page, the record page for the copilot action, or the Copilot Action Library tab

B. Copilot Actions page or Global Actions

C. Copilot Detail page, Global Actions, or the record page for the copilot action

Answer: A

Explanation:

To add or update actions assigned to a copilot, an AI Specialist can manage this through several areas:

? Copilot Actions Page: This is the central location where copilot actions are managed and configured.

? Record Page for the Copilot Action: From the record page, individual copilot actions can be updated or modified.

? Copilot Action Library Tab: This tab serves as a repository where predefined or custom actions for Copilot can be accessed and modified.

These areas provide flexibility in managing and updating the actions assigned to Copilot, ensuring that the AI assistant remains aligned with business requirements and processes. The other options are incorrect:

- ? Bmisses the Copilot Action Library, which is crucial for managing actions.
- ? Cincludes the Copilot Detail page, which isn't the primary place for action management.

References:

- ? Salesforce Documentation onManaging Copilot Actions
- ? Salesforce AI Specialist Guide onCopilot Action Management

NEW QUESTION 76

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 inEinstein Studioand create prompt templates inPrompt Builder, the AI Specialist should assign theData Cloud AdminandPrompt Template Managerpermission sets.

? Data Cloud Adminprovides access to manage and oversee models withinEinstein Studio.

? Prompt Template Managergives the user the ability to create and manage prompt templates withinPrompt Builder.

? Option Ais correct because it assigns the necessary permissions for both managing models and creating prompt templates.

? Option BandOption Care 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

NEW QUESTION 77

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-tunedLLM (Large Language Model)into Salesforce by adding it toEinstein Studio Model Builder.Einstein Studioallows 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 Studioprovides the most secure and seamless way to integrate custom AI models while maintaining control over data privacy and compliance. More details can be found inSalesforce's Einstein Studio documentationon integrating external models.

NEW QUESTION 78

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual Salesforce-AI-Specialist Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the Salesforce-AI-Specialist Product From:

<https://www.2passeasy.com/dumps/Salesforce-AI-Specialist/>

Money Back Guarantee

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