

Exam Questions Salesforce-AI-Specialist

Salesforce Certified AI Specialist Exam

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



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

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

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

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

Answer: A

Explanation:

To ground a sales email on Opportunity Products, Events near the customer, and Tone and voice examples, the AI Specialist should use a prompt-initiated flow. This flow can dynamically fetch the necessary data from related records in Salesforce and ground the generative AI output with contextually accurate information.

? Option B (flex template) does not provide the ability to fetch dynamic data from Salesforce records automatically.

? Option C (manual insertion) would not allow for the dynamic and automated grounding of data required for custom prompts. Refer to Salesforce documentation on flows and grounding for more details on integrating data into custom prompt templates.

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

An AI Specialist wants to ground a new prompt template with the User related list. What should the AI Specialist consider?

- A. The User related list should have View All access.
- B. The User related list needs to be included on the record page.
- C. The User related list is not supported in prompt templates.

Answer: C

Explanation:

An AI Specialist wants to ground a new prompt template with the User related list. Grounding in prompt templates involves using data from related lists to provide context or additional information to the Large Language Model (LLM) when generating responses.

Key Consideration:

? Unsupported Related Lists in Prompt Templates:

Why Options A and B are Incorrect:

? Option A (The User related list should have View All access):

? Option B (The User related list needs to be included on the record page):

References:

? Salesforce AI Specialist Documentation -Prompt Templates Limitations:

? Salesforce Help -Data Access in Prompt Templates:

? Salesforce Trailhead -Understanding Prompt Template Grounding:

Conclusion:

Since the User related list is not supported in prompt templates, the AI Specialist must consider alternative approaches. They might need to redesign the prompt template to use supported objects or related lists, or explore other methods to incorporate necessary user information while adhering to Salesforce's data access policies.

NEW QUESTION 6

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

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

Answer: A

Explanation:

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

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

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

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

NEW QUESTION 7

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

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

Answer: B

Explanation:

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

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

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

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

NEW QUESTION 8

Universal Containers (UC) recently rolled out Einstein Generative capabilities and has created a custom prompt to summarize case records. Users have reported that the case summaries generated are not returning the appropriate information. What is a possible explanation for the poor prompt performance?

- A. The data being used for grounding is incorrect or incomplete.
- B. The prompt template version is incompatible with the chosen LLM.
- C. The Einstein Trust Layer is incorrectly configured.

Answer: A

Explanation:

Poor prompt performance when generating case summaries is often due to the data used for grounding being incorrect or incomplete. Grounding involves feeding accurate, relevant data to the AI so it can generate appropriate outputs. If the data source is incomplete or contains errors, the generated summaries will reflect

that by being inaccurate or insufficient.

? Option B(prompt template incompatibility with the LLM) is unlikely because such incompatibility usually results in more technical failures, not poor content quality.

? Option C(Einstein Trust Layer misconfiguration) is focused on data security and auditing, not the quality of prompt responses.

For more information, refer toSalesforce documentation on grounding AI modelsand data quality best practices.

NEW QUESTION 9

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 ofPrompt Builderiscreating a draft of a support bulletin postfor new product patches.Prompt Builderallows 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 toSalesforce's Prompt Builder documentationfor generative AI content creation.

NEW QUESTION 10

Universal Containers is using Einstein Copilot for Sales to find similar opportunities to help close deals faster. The team wants to understand the criteria used by the copilot to match opportunities.

What is one criteria that Einstein Copilot for Sales uses to match similar opportunities?

- A. Matched opportunities are limited to the same account.
- B. Matched opportunities were created in the last 12 months.
- C. Matched opportunities have a status of Closed Won from last 12 months.

Answer: C

Explanation:

WhenEinstein Copilot for Salesmatches similar opportunities, one of the primary criteria used is whether the opportunities have astatus of Closed Wonwithin the last 12 months. This is a key factor in identifying successful patterns that could help close current deals. By focusing on opportunities that have been recently successful, Einstein Copilot can provide relevant insights and suggestions to sales reps to help them close similar deals faster.

For more information, reviewSalesforce Einstein Copilot documentationrelated to opportunity matchingand sales success patterns.

NEW QUESTION 10

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 atUniversal Containerscan useEinstein Work Summariessto quickly understand the details of a case without reading through each case comment.Work Summariesleverage generative AI to provide a concise overview of ongoing cases, summarizing all relevant information in an easily digestible format.

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

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

For more details, refer toSalesforce documentation on Einstein Work Summaries.

NEW QUESTION 13

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 aFlex 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 14

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

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

Answer: C

Explanation:

To provide the sales team with insights into product and competitor names mentioned during calls, Universal Containers should:

? Enable Einstein Conversation Insights: Activates the feature that analyzes call recordings for valuable insights.

? Enable Sales Recording: Allows calls to be recorded within Salesforce without needing an external recording provider.

? Assign Permission Sets: Grants the necessary permissions to sales team members to access and utilize conversation insights.

? Customize Insights: Configure the system to track mentions of up to 50 products and 50 competitors, providing tailored insights relevant to the organization's needs. Option C accurately reflects these steps. Option A mentions defining recording managers but omits enabling sales recording within Salesforce. Option B suggests connecting a recording provider and limits customization to 25 products, which does not fully meet UC's requirements.

References:

? Salesforce AI Specialist Documentation - Setting Up Einstein Conversation Insights: Provides instructions on enabling conversation insights and sales recording.

? Salesforce Help - Customizing Conversation Insights: Details how to customize insights with up to 50 products and competitors.

? Salesforce AI Specialist Exam Guide: Outlines best practices for implementing AI features like Einstein Conversation Insights in a sales context.

=====

NEW QUESTION 17

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 21

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 24

Universal Containers wants to use an external large language model (LLM) in Prompt Builder.

What should an AI Specialist recommend?

- A. Use Apex to connect to an external LLM and ground the prompt.
- B. Use BYO-LLM functionality in Einstein Studio,
- C. Use Flow and External Services to bring data from an external LLM.

Answer: B

Explanation:

Bring Your Own Large Language Model (BYO-LLM) functionality in Einstein Studio allows organizations to integrate and use external large language models (LLMs) within the Salesforce ecosystem. Universal Containers can leverage this feature to connect and ground prompts with external LLMs, allowing for custom AI model use cases and seamless integration with Salesforce data.

? Option B is the correct choice as Einstein Studio provides a built-in feature to work with external models.

? Option A suggests using Apex, but BYO-LLM functionality offers a more streamlined solution.

? Option C focuses on Flow and External Services, which is more about data integration and isn't ideal for working with LLMs.

References:

? Salesforce Einstein Studio BYO-LLM Documentation: https://help.salesforce.com/s/articleView?id=sf.einstein_studio_llm.

NEW QUESTION 25

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

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

Answer: B

Explanation:

When Universal Containers' AI data masking rules do not meet organizational privacy and security standards, the AI Specialist should configure the data masking rules within the Einstein Trust Layer. The Einstein Trust Layer provides a secure and compliant environment where sensitive data can be masked or anonymized to adhere to privacy policies and regulations.

? Option A, enabling data masking for sandbox refreshes, is related to sandbox environments, which are separate from how AI interacts with production data.

? Option C, adding masking rules in the LLM setup, is not appropriate because data masking is managed through the Einstein Trust Layer, not the LLM configuration.

The Einstein Trust Layer allows for more granular control over what data is exposed to the AI model and ensures compliance with privacy regulations.

Salesforce AI Specialist References: For more information, refer to: https://help.salesforce.com/s/articleView?id=sf.einstein_trust_layer_data_masking.htm

NEW QUESTION 27

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 31

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) andOption 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 35

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 37

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

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

Answer: B

Explanation:

In Einstein Copilot, the role of the Large Language Model (LLM) is to analyze user inputs and identify the best matching actions that need to be executed. It uses natural language understanding to break down the user's request and determine the correct sequence of actions that should be performed.

By doing so, the LLM ensures that the tasks and actions executed are contextually relevant and are performed in the proper order. This process provides a seamless, AI-enhanced experience for users by matching their requests to predefined Salesforce actions or flows.

The other options are incorrect because:

A mentions finding similar requests, which is not the primary role of the LLM in this context. C focuses on access and sorting by priority, which is handled more by security models and governance than by the LLM.

References:

Salesforce Einstein Documentation on Einstein Copilot Actions Salesforce AI Documentation on Large Language Models

NEW QUESTION 38

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

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

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 40

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 45

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

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

A. Save as New Template and edit as needed.

B. Clone the existing template and modify as needed.

C. Save as New Version and edit as needed.

Answer: B

Explanation:

When an active standard email prompt template doesn't meet the business requirements, the best approach is to clone the existing template and modify it as needed. Cloning allows the AI Specialist to preserve the original template while making adjustments to fit specific business needs. This ensures that any customizations are applied without altering the original standard template.

Saving as a new version is typically used for versioning changes in the same template, while Save as New Template creates a brand-new template without linking to the existing one. Cloning provides a balance, allowing modifications while retaining the original structure for future reference.

For more details, refer to Salesforce Prompt Builder documentation for guidance on cloning and modifying templates.

NEW QUESTION 46

Universal Containers wants to be able to detect with a high level of confidence if content generated by a large language model (LLM) contains toxic language.

Which action should an AI Specialist take in the Trust Layer to confirm toxicity is being appropriately managed?

A. Access the Toxicity Detection log in Setup and export all entries where isToxicityDetected is true.

B. Create a flow that sends an email to a specified address each time the toxicity score from the response exceeds a predefined threshold.

C. Create a Trust Layer audit report within Data Cloud that uses a toxicity detector type filter to display toxic responses and their respective scores.

Answer: C

Explanation:

To ensure that content generated by a large language model (LLM) is appropriately screened for toxic language, the AI Specialist should create a Trust Layer audit report within Data Cloud. By using the toxicity detector type filter, the report can display toxic responses along with their respective toxicity scores, allowing Universal Containers to monitor and manage any toxic content generated with a high level of confidence.

? Option C is correct because it enables visibility into toxic language detection within the Trust Layer and allows for auditing responses for toxicity.

? Option A suggests checking a toxicity detection log, but Salesforce provides more comprehensive options via the audit report.

? Option B involves creating a flow, which is unnecessary for toxicity detection monitoring.

References:

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

NEW QUESTION 48

Universal Containers wants to reduce overall agent handling time minimizing the time spent typing routine answers for common questions in chat, and reducing the post-chat analysis by suggesting values for case fields.

Which combination of Einstein for Service features enables this effort?

A. Einstein Service Replies and Work Summaries

B. Einstein Reply Recommendations and Case Summaries

C. Einstein Reply Recommendations and Case Classification

Answer: C

Explanation:

Universal Containers aims to reduce overall agent handling time by minimizing the time agents spend typing routine answers for common questions during chats and by reducing post-chat analysis through suggesting values for case fields.

To achieve these objectives, the combination of Einstein Reply Recommendations and Case Classification is the most appropriate solution.

* 1. Einstein Reply Recommendations:

? Purpose: Helps agents respond faster during live chats by suggesting the best responses based on historical chat data and common customer inquiries.

? Functionality:

? Benefit: Significantly reduces the time agents spend typing routine answers, thus improving efficiency and reducing handling time.

* 2. Case Classification:

? Purpose: Automatically suggests or populates values for case fields based on historical data and patterns identified by AI.

? Functionality:

? Benefit: Reduces the time agents spend on post-chat analysis and data entry by automating the classification and field population process.

Why Options A and B are Less Suitable:

? Option A (Einstein Service Replies and Work Summaries):

? Option B (Einstein Reply Recommendations and Case Summaries):

References:

? Salesforce AI Specialist Documentation -Einstein Reply Recommendations:

? Salesforce AI Specialist Documentation -Einstein Case Classification:

? Salesforce Trailhead -Optimize Service with AI:

NEW QUESTION 51

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 55

Universal Containers Is Interested In Improving the sales operation efficiency by analyzing their data using AI-powered predictions in Einstein Studio. Which use case works for this scenario?

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

Answer: B

Explanation:

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

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

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

References:

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

NEW QUESTION 59

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 61

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 66

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 68

Universal Containers (UC) has implemented Generative AI within Salesforce to enable summarization of a custom object called Guest. Users have reported mismatches in the generated information.

In refining its prompt design strategy, which key practices should UC prioritize?

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

Answer: B

Explanation:

For Universal Containers (UC) to refine its Generative AI prompt design strategy and improve the accuracy of the generated summaries for the custom object Guest, the best practice is to focus on crafting concise, clear, and consistent prompt templates. This includes:

? Effective grounding: Ensuring the prompt pulls data from the correct sources.

? Contextual role-playing: Providing the AI with a clear understanding of its role in generating the summary.

? Clear instructions: Giving unambiguous directions on what to include in the response.

? Iterative feedback: Regularly testing and adjusting prompts based on user feedback.

? Option B is correct because it follows industry best practices for refining prompt design.

? Option A (prompt test mode) is useful but less relevant for refining prompt design itself.

? Option C (prompt review case with Salesforce) would be more appropriate for technical issues or complex prompt errors, not general design refinement.

References:

? Salesforce Prompt Design Best Practices: https://help.salesforce.com/s/articleView?id=sf.prompt_design_best_practices.htm

NEW QUESTION 70

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 72

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 74

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 predictiveAI 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

NEW QUESTION 78

Universal Containers implemented Einstein Copilot for its users.

One user complains that Einstein Copilot is not deleting activities from the past 7 days. What is the reason for this issue?

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

Answer: C

Explanation:

Einstein Copilot currently supports various actions like creating and updating records but does not support the Delete Record action. Therefore, the user's request to delete activities from the past 7 days cannot be fulfilled using Einstein Copilot.

? Unsupported Action: The inability to delete records is due to the current limitations of Einstein Copilot's supported actions. It is designed to assist with tasks like data retrieval, creation, and updates, but for security and data integrity reasons, it does not facilitate the deletion of records.

? User Permissions: Even if the user has the necessary permissions to delete records within Salesforce, Einstein Copilot itself does not have the capability to execute delete operations.

References:

? Salesforce AI Specialist Documentation -Einstein Copilot Supported Actions:

? Salesforce Help -Limitations of Einstein Copilot:

NEW QUESTION 80

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 85

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 89

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 92

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 97

.....

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