

Salesforce

Exam Questions Salesforce-AI-Specialist

Salesforce Certified AI Specialist Exam



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

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 3

An administrator is responsible for ensuring the security and reliability of Universal Containers' (UC) CRM data. UC needs enhanced data protection and up-to-date AI capabilities. UC also needs to include relevant

information from a Salesforce record to be merged with the prompt. Which feature in the Einstein Trust Layer best supports UC's need?

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

Answer: B

Explanation:

Dynamic grounding with secure data retrieval is a key feature in Salesforce's Einstein Trust Layer, which provides enhanced data protection and ensures that AI-generated outputs are both accurate and securely sourced. This feature allows relevant Salesforce data to be merged into the AI-generated responses, ensuring that the AI outputs are contextually aware and aligned with real-time CRM data.

Dynamic grounding means that AI models are dynamically retrieving relevant information from Salesforce records (such as customer records, case data, or custom object data) in a secure manner. This ensures that any sensitive data is protected during AI processing and that the AI model's outputs are trustworthy and reliable for business use.

The other options are less aligned with the requirement:

? Data masking refers to obscuring sensitive data for privacy purposes and is not related to merging Salesforce records into prompts.

? Zero-data retention policy ensures that AI processes do not store any user data after processing, but this does not address the need to merge Salesforce record information into a prompt.

References:

? Salesforce Developer Documentation on Einstein Trust Layer

? Salesforce Security Documentation for AI and Data Privacy

NEW QUESTION 4

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 5

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 6

Universal Containers (UC) noticed an increase in customer contract cancellations in the last few months. UC is seeking ways to address this issue by implementing a proactive outreach program to customers before they cancel their contracts and is asking the Salesforce team to provide suggestions. Which use case functionality of Model Builder aligns with UC's request?

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

Answer: B

Explanation:

Customer churn prediction is the best use case for Model Builder in addressing Universal Containers' concerns about increasing customer contract cancellations. By implementing a model that predicts customer churn, UC can proactively identify customers who are at risk of canceling and take action to retain them before they decide to terminate their contracts. This functionality allows the business to forecast churn probability based on historical data and initiate timely outreach programs.
? Option B is correct because customer churn prediction aligns with UC's need to reduce cancellations through proactive measures.
? Option A (product recommendation prediction) is unrelated to contract cancellations.
? Option C (contract renewal date prediction) addresses timing but does not focus on predicting potential cancellations.
References:
? Salesforce Model Builder Use Case Overview: https://help.salesforce.com/s/articleView?id=sf.model_builder_use_cases.htm

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 to Salesforce documentation on grounding AI models and data quality best practices.

NEW QUESTION 9

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 10

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 10

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 11

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

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

Answer: A

Explanation:

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

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

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

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

NEW QUESTION 15

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

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

Answer: B

Explanation:

To enable Universal Containers service agents to query the current fulfillment status of an order using natural language and leverage an existing auto-launched flow that queries Oracle ERP, the best solution is to create a custom copilot action that calls the flow. This action will allow Einstein Copilot to interact with the flow and retrieve the required order fulfillment information seamlessly. Custom copilot actions can be tailored to call various backend systems or flows in response to user requests.

? Option B is correct because it enables integration between Einstein Copilot and the flow that connects to Oracle ERP.

? Option A (Flex prompt template) is more suited for static responses and not for invoking flows.

? Option C (Integration Flow Standard Action) is not directly related to creating a specific copilot action for this use case.

References:

? Salesforce Einstein Copilot

NEW QUESTION 17

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

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

Answer: A

Explanation:

To access externally-hosted models, such as a large language model (LLM) hosted on AWS, the Model Builder in Salesforce is the appropriate tool. Model Builder allows teams to integrate and deploy external AI models into the Salesforce platform, making it possible to leverage models hosted outside of Salesforce infrastructure while still benefiting from the platform's native AI capabilities.

? Option B, App Builder, is primarily used to build and configure applications in Salesforce, not to integrate AI models.

? Option C, Copilot Builder, focuses on building assistant-like tools rather than integrating external AI models.

Model Builder enables seamless integration with external systems and models, allowing Salesforce users to use external LLMs for generating AI-driven insights and automation. Salesforce AI Specialist References: For more details, check the Model Builder guide

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

NEW QUESTION 21

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

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

Answer: A

Explanation:

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

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

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

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

NEW QUESTION 25

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 timeconsuming for both UC's salesand 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:

ForUniversal Containers (UC), to reduce the time and ensure prompt consistency when using public generative AI models, the recommended feature isEinstein Prompt Builder and Prompt Templates. This feature allows teams to createreusable and consistent promptsfor generative AI tasks, ensuring that all users receive uniform responses without having to type in detailed prompts manually every time.

? Einstein Prompt Buildersimplifies the creation of prompts, andPrompt Templatesstandardize 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 28

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:

WhenUniversal Containers' AI data masking rulesdo not meet organizational privacy and security standards, the AI Specialist should configure thedata maskingrules within theEinstein Trust Layer. TheEinstein Trust Layerprovides 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 theEinstein 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 30

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 followingcapabilities 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 AIfeatures, particularly around preventing third-party LLMs from accessing or retaining their data. TheZero-Data Retention Policyin theEinstein Trust Layeris designed to address these concerns by ensuring that:

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

? No data is retainedoutside 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 DefenseandData Maskingare 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 32

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

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

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

Answer: A

Explanation:

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

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

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

NEW QUESTION 35

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 38

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

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

Answer: C

Explanation:

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

NEW QUESTION 39

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

Which step should UC take to meet this requirement?

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

Answer: B

Explanation:

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

? Service Replies generates responses automatically.

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

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

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

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

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

NEW QUESTION 42

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 47

Universal Containers wants to be able to detect with a high level confidence if content generated by a large language model (LLM) contains toxic language. Which action should an 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 51

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 55

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

How does Einstein Copilot handle user requests in Salesforce?

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

Answer: C

Explanation:

Einstein Copilot is designed to enhance user interaction within Salesforce by leveraging Large Language Models (LLMs) to process and respond to user inquiries. When a user submits a request, Einstein Copilot analyzes the input using natural language processing techniques. It then utilizes LLM technology to generate an appropriate and contextually relevant response, which is displayed directly to the user within the Salesforce interface.

Option C accurately describes this process. Einstein Copilot does not necessarily trigger a flow (Option A) or perform an HTTP callout to an LLM provider (Option B) for each user

request. Instead, it integrates LLM capabilities to provide immediate and intelligent responses, facilitating a broad range of user requests.

References:

? Salesforce AI Specialist Documentation - Einstein Copilot Overview:Details how Einstein Copilot employs LLMs to interpret user inputs and generate responses within the Salesforce ecosystem.
? Salesforce Help - How Einstein Copilot Works:Explains the underlying mechanisms of how Einstein Copilot processes user requests using AI technologies.

NEW QUESTION 57

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 inEinstein Trust Layer, the AI Specialist should request theEinstein Generative AI Audit Datafrom the Security sectionof 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 Bis correct because it allows for the retrieval of audit data that can be used to validate data masking.

? Option A(Flow Debugger) andOption C(Einstein Feedback) do not relate to validating field masking in the context of theEinstein Trust Layer.

References:

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

NEW QUESTION 60

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 improvingsales operations efficiency,Einstein Studiois ideal for creating AI-powered models that can predict outcomes based on data. One of the most valuable use cases is predictingcustomer 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 Bis 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, whileOption 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 62

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 Builderis designed to help organizations create and configure reusable prompts for large language models (LLMs). By integratinggenerative AI responses into workflows,Prompt Builderenables 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 forApex 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 64

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 theCall Summaryfeature.Call SummaryusesEinstein Generative Alto 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 65

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 67

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 70

An AI Specialist wants to include data from the response of an external service invocation (REST API callout) into the prompt template.

How should the AI Specialist meet this requirement?

A. Convert the JSON to an XML merge field.

B. Use External Service Record merge fields.

C. Use ??Add Prompt Instructions?? flow element.

Answer: B

Explanation:

An AI Specialist wants to include data from the response of an external service invocation (REST API callout) into a prompt template. The goal is to incorporate dynamic data retrieved from an external API into the AI-generated content.

Solution:

? Use External Service Record Merge Fields

? External Service Integration:

? External Service Record Merge Fields:

Implementation Steps:

? Register the External Service:

? Create a Named Credential:

? Use External Service in Flow:

? Configure the Prompt Template:

Why Other Options are Less Suitable:

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

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

References:

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

? Salesforce Help -Using Merge Fields with External Data:

? Salesforce Trailhead -External Services and Flow:

Conclusion:

By using External Service Record merge fields, the AI Specialist can effectively include data from external REST API responses into prompt templates, ensuring that the AI- generated content is enriched with up-to-date and relevant external data.

NEW QUESTION 73

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 Copilotis 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 anAI-powered assistantthat 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 Ais correct becauseEinstein Copilotbrings a conversational interface that caters to a wide range of users.

? Option BandOption Care more focused on developing and training AI models, which are not the primary functions ofEinstein Copilot.

References:

? Salesforce Einstein Copilot Overview:https://help.salesforce.com/s/articleView?id=einstein_copilot_overview.ht m

NEW QUESTION 77

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

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

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

Answer: B

Explanation:

WhenUniversal Containerscreates a new Sales Email prompt template using the"Save As"function, missing hyperparameters can result in different outputs. To ensure the new prompt produces comparable results to the standard Sales Email prompt, the AI Specialist shouldmanually add the necessary hyperparametersto the new template.

? Hyperparameters likeTemperature,Frequency Penalty, andPresence

Penaltydirectly affect how the AI generates responses. Ensuring that these are consistent with the standard template will result in similar outputs.

? Option A (Model Playground)is not necessary here, as it focuses on fine-tuning models, not adjusting templates directly.

? Option C (Reverting to the standard template)does not solve the issue of customizing the prompt template.

For more information, refer toPrompt Builder documentationon configuring hyperparameters in custom templates.

NEW QUESTION 81

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

Which feature provides insights about competitor mentions and coaching opportunities?

- A. Call Summaries
- B. Einstein Sales Insights
- C. Call Explorer

Answer: C

Explanation:

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

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

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

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

References:

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

NEW QUESTION 83

How does the Einstein Trust Layer ensure that sensitive data isprotected 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 88

A data scientist needs to view and manage models in Einstein Studio. The data scientist also needs to create prompt templates in Prompt Builder.

Which permission sets should an AI Specialist assign to the data scientist?

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

Answer: A

Explanation:

To allow a data scientist to view and manage models in Einstein Studio and create prompt templates in Prompt Builder, the AI Specialist should assign the Data Cloud Admin and Prompt Template Manager permission sets.

? Data Cloud Admin provides access to manage and oversee models within Einstein Studio.

? Prompt Template Manager gives the user the ability to create and manage prompt templates within Prompt Builder.

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

? Option B and Option C are incorrect as they do not provide the correct combination of permissions for managing models and building prompts.

References:

? Salesforce Permissions Documentation: https://help.salesforce.com/s/articleView?id=sf.perm_sets_overview.htm

NEW QUESTION 92

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

Salesforce-AI-Specialist Practice Exam Features:

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

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