

Salesforce

Exam Questions Salesforce-AI-Specialist

Salesforce Certified AI Specialist Exam



NEW QUESTION 1

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 2

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 3

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

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

Answer: B

Explanation:

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

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

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

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

References:

? Salesforce Model Builder Use Case Overview: https://help.salesforce.com/s/articleView?id=sf.model_builder_use_case_s.htm

NEW QUESTION 4

Leadership needs to populate a dynamic form field with a summary or description created by a large language model (LLM) to facilitate more productive conversations with customers. Leadership also wants to keep a human in the loop to be considered in their AI strategy. Which prompt template type should the AI Specialist recommend?

- A. Sales Email
- B. Field Generation
- C. Record Summary

Answer: B

Explanation:

The correct answer is Field Generation because this template type is designed to dynamically populate form fields with content generated by a large language model (LLM). In this scenario, leadership wants a dynamic form field that contains a summary or description generated by AI to aid customer interactions. Additionally, they want to keep a human in the loop, meaning the generated content will likely be reviewed or edited by a person before it's finalized, which aligns with the Field Generation prompt template.

? Field Generation: This prompt type allows you to generate content for specific fields in Salesforce, leveraging large language models to create dynamic and contextual information. It ensures that AI content is available within the record where needed, but it allows human oversight or review, supporting the "human-in-the-loop" strategy.

? Sales Email: This prompt type is mainly used for generating email content for outreach or responses, which doesn't align directly with populating fields in a form.

? Record Summary: While this option might seem close, it is typically used to summarize entire records for high-level insights rather than filling specific fields with dynamic content based on AI generation.

Salesforce AI Specialist References:

? You can explore more about these prompt templates and AI capabilities through Salesforce documentation and official resources on Prompt Builder: https://help.salesforce.com/s/articleView?id=sf.prompt_builder_templates_overview.htm

NEW QUESTION 5

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

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

Answer: B

Explanation:

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

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

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

Salesforce AI Specialist References: For more details on the Einstein Trust Layer and setup

steps: https://help.salesforce.com/s/articleView?id=sf.einstein_trust_layer_overview.htm

NEW QUESTION 6

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

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

Answer: B

Explanation:

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

Supported Objects for Service AI Grounding:

? Case

? Knowledge

? Case Object:

? Knowledge Object:

? Exclusion of Other Objects:

Why Options A and C are Incorrect:

? Option A (Case, Knowledge, and Case Notes):

? Option C (Case, Case Emails, and Knowledge):

References:

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

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

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

NEW QUESTION 7

Universal Containers (UC) has a legacy system that needs to integrate with Salesforce. UC wishes to create a digest of account action plans using the generative API feature.

Which API service should UC use to meet this requirement?

- A. REST API
- B. Metadata API
- C. SOAP API

Answer: A

Explanation:

To create a digest of account action plans using the generative API feature, Universal Containers should use the REST API. The REST API is ideal for integrating Salesforce with external systems and enabling interaction with Salesforce data, including generative capabilities like creating summaries or digests. It supports

modern web standards and is suitable for flexible, lightweight interactions between Salesforce and legacy systems.
? Metadata APIs are used for retrieving and deploying metadata, not for data operations like generating summaries.
? SOAP APIs are an older API used for integration but is less flexible compared to REST for this specific use case.
For more details, refer to Salesforce REST API documentation regarding using REST for data integration and generating content.

NEW QUESTION 8

Universal Containers 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:

When Einstein Copilot for Sales matches similar opportunities, one of the primary criteria used is whether the opportunities have a status of Closed Won within 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, review Salesforce Einstein Copilot documentation related to opportunity matching and sales success patterns.

NEW QUESTION 9

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 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 10

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

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

Answer: A

Explanation:

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

? Option B, App Builder, is primarily used to build and configure applications in Salesforce, not to integrate AI models.
? Option C, Copilot Builder, focuses on building assistant-like tools rather than integrating external AI models.
Model Builder enables seamless integration with external systems and models, allowing Salesforce users to use external LLMs for generating AI-driven insights and automation. Salesforce AI Specialist References:For more details, check the Model Builder guide here:https://help.salesforce.com/s/articleView?id=sf.model_builder_external_models.htm

NEW QUESTION 14

A support team handles a high volume of chat interactions and needs a solution to provide quick, relevant responses to customer inquiries. Responses must be grounded in the organization's knowledge base to maintain consistency and accuracy. Which feature in Einstein for Service should the support team use?

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

Answer: B

Explanation:

The support team should use Einstein Reply Recommendationsto provide quick, relevant responses to customer inquiries that are grounded in the organization's knowledge base. This feature leverages AI to recommend accurate and consistent replies based on historical interactions and the knowledge stored in the system, ensuring that responses are aligned with organizational standards.

? Einstein Service Replies (Option A) is focused on generating replies but doesn't have the same emphasis on grounding responses in the knowledge base.

? Einstein Knowledge Recommendations (Option C) suggests knowledge articles to agents, which is more about assisting the agent in finding relevant articles than providing automated or AI-generated responses to customers.

Salesforce AI Specialist References:For more information on Einstein Reply

Recommendations:https://help.salesforce.com/s/articleView?id=sf.einstein_reply_recommendations_overview.htm

NEW QUESTION 19

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

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

The data sent cannot be accessed by the LLM provider.

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

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

Answer: B

Explanation:

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

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

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

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

? The LLM provider cannot access any customer data.

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

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

References:

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

NEW QUESTION 20

How should an organization use the Einstein Trust layer to audit, track, and view masked data?

- A. Utilize the audit trail that captures and stores all LLM submitted prompts in Data Cloud.
- B. In Setup, use Prompt Builder to send a prompt to the LLM requesting for the masked data.
- C. Access the audit trail in Setup and export all user-generated prompts.

Answer: A

Explanation:

The Einstein Trust Layer is designed to ensure transparency, compliance, and security for organizations leveraging Salesforce's AI and generative AI capabilities. Specifically, for auditing, tracking, and viewing masked data, organizations can utilize:

? Audit Trail in Data Cloud: The audit trail captures and stores all prompts submitted

to large language models (LLMs), ensuring that sensitive or masked data interactions are logged. This allows organizations to monitor and audit all AI-generated outputs, ensuring that data handling complies with internal and regulatory guidelines. The Data Cloud provides the infrastructure for managing and accessing this audit data.

? Why not B? Using Prompt Builder in Setup to send prompts to the LLM is for

creating and managing prompts, not for auditing or tracking data. It does not interact directly with the audit trail functionality.

? Why not C? Although the audit trail can be accessed in Setup, the user-generated

prompts are primarily tracked in the Data Cloud for broader control, auditing, and analysis. Setup is not the primary tool for exporting or managing these audit logs. More information on auditing AI interactions can be found in the Salesforce AI Trust Layer documentation, which outlines how organizations can manage and track generative AI interactions securely.

NEW QUESTION 23

Universal Containers recently launched a pilot program to integrate conversational AI into its CRM business operations with Einstein Copilot.

How should the AI Specialist monitor Copilot's usability and the assignment of actions?

- A. Run a report on the Platform Debug Logs.
- B. Query the Copilot log data using the metadata API.
- C. Run Einstein Copilot Analytics.

Answer: C

Explanation:

To monitor Einstein Copilot's usability and the assignment of actions, the AI Specialist should run Einstein Copilot Analytics. This feature provides insights into how often Copilot is used, the types of actions it is handling, and overall user engagement with the system. It's the most effective way to track Copilot's performance and usage patterns.

? Platform Debug Logs are not relevant for tracking user behavior or the assignment of Copilot actions.

? Querying the Copilot log data via the Metadata API would not provide the necessary insights in a structured manner.

For more details, refer to Salesforce's Copilot Analytics documentation for tracking AI-driven interactions.

NEW QUESTION 26

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

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

* How it is masked

* The masked response

What should the AI Specialist recommend?

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

Answer: C

Explanation:

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

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

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

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

NEW QUESTION 30

An AI Specialist implements Einstein Sales Emails for a sales team. The team wants to send personalized follow-up emails to leads based on their interactions and data stored in Salesforce. The AI Specialist needs to configure the system to use the most accurate and up-to-date information for email generation.

Which grounding technique should the AI Specialist use?

- A. Ground with Apex Merge Fields
- B. Ground with Record Merge Fields
- C. Automatic grounding using Draft with Einstein feature

Answer: B

Explanation:

For Einstein Sales Emails to generate personalized follow-up emails, it is crucial to ground the email content with the most up-to-date and accurate information. Grounding refers to connecting the AI model with real-time data. The most appropriate technique in this case is Ground with Record Merge Fields. This method ensures that the content in the emails pulls dynamic and accurate data directly from Salesforce records, such as lead or contact information, ensuring the follow-up is relevant and customized based on the specific record.

? Record Merge Fields ensure the generated emails are highly personalized using data like lead name, company, or other Salesforce fields directly from the records.

? Apex Merge Fields are typically more suited for advanced, custom logic-driven scenarios but are not the most straightforward for this use case.

? Automatic grounding using Draft with Einstein is a different feature where Einstein automatically drafts the email, but it does not specifically ground the content with record-specific data like Record Merge Fields.

References:

? Salesforce Einstein Sales Emails Documentation: https://help.salesforce.com/s/articleView?id=release-notes.rn_einstein_sales_emails.htm

NEW QUESTION 34

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 39

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 44

An AI Specialist configured Data Masking within the Einstein Trust Layer.

How should the AI Specialist begin validating that the correct fields are being masked?

- A. Use a Flow-based resource in Prompt Builder to debug the fields?? merge values using Flow Debugger.
- B. Request the Einstein Generative AI Audit Data from the Security section of the Setup menu.
- C. Enable the collection and storage of Einstein Generative AI Audit Data on the Einstein Feedback setup page.

Answer: B

Explanation:

To begin validating that the correct fields are being masked in Einstein Trust Layer, the AI Specialist should request the Einstein Generative AI Audit Data from the Security section of the Salesforce Setup menu. This audit data allows the AI Specialist to see how data is being processed, including which fields are being masked, providing transparency and validation that the configuration is working as expected.

? Option B is correct because it allows for the retrieval of audit data that can be used to validate data masking.

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

References:

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

NEW QUESTION 46

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

Which use case works for this scenario?

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

Answer: B

Explanation:

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

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

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

References:

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

NEW QUESTION 49

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 50

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

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

Answer: C

Explanation:

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

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

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

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

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

NEW QUESTION 54

An administrator wants to check the response of the Flex prompt template they've built, but the preview button is greyed out. What is the reason for this?

- A. The records related to the prompt have not been selected.
- B. The prompt has not been saved and activated,
- C. A merge field has not been inserted in the prompt.

Answer: A

Explanation:

When the preview button is greyed out in a Flex prompt template, it is often because the records related to the prompt have not been selected. Flex prompt templates pull data dynamically from Salesforce records, and if there are no records specified for the prompt, it can't be previewed since there is no content to generate based on the template.

? Option B, not saving or activating the prompt, would not necessarily cause the preview button to be greyed out, but it could prevent proper functionality.

? Option C, missing a merge field, would cause issues with the output but would not directly grey out the preview button.

Ensuring that the related records are correctly linked is crucial for testing and previewing how the prompt will function in real use cases.

Salesforce AI Specialist References: Refer to the documentation on troubleshooting Flex templates

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

NEW QUESTION 59

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

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

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

Answer: B

Explanation:

When Universal Containers creates a new Sales Email prompt template using the "Save As" function, missing hyperparameters can result in different outputs. To ensure the new prompt produces comparable results to the standard Sales Email prompt, the AI Specialist should manually add the necessary hyperparameters to the new template.

? Hyperparameters like Temperature, Frequency Penalty, and Presence

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

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

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

For more information, refer to Prompt Builder documentation on configuring hyperparameters in custom templates.

NEW QUESTION 62

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

- A. Temperature, Frequency Penalty, Presence Penalty
- B. Temperature, Top-k sampling, Presence Penalty
- C. Temperature, Frequency Penalty, Output Tokens

Answer: A

Explanation:

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

? Temperature: Controls the randomness of predictions. A higher temperature leads to more diverse outputs, while a lower temperature makes the model's responses more focused and deterministic.

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

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

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

NEW QUESTION 65

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

- A. Copilot Actions page, the record page for the copilot action, or the Copilot Action Library tab
- B. Copilot Actions page or Global Actions
- C. Copilot Detail page, Global Actions, or the record page for the copilot action

Answer: A

Explanation:

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

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

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

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

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

? Bmisses the Copilot Action Library, which is crucial for managing actions.

? Cincludes the Copilot Detail page, which isn't the primary place for action management.

References:

? Salesforce Documentation on Managing Copilot Actions

? Salesforce AI Specialist Guide on Copilot Action Management

NEW QUESTION 66

.....

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 FirstTry
- * Salesforce-AI-Specialist Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

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