



Salesforce

Exam Questions Salesforce-AI-Specialist

Salesforce Certified AI Specialist Exam

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

Universal Containers tests out a new Einstein Generative AI feature for its sales team to create personalized and contextualized emails for its customers. Sometimes, users find that the draft email contains placeholders for attributes that could have been derived from the recipient's contact record. What is the most likely explanation for why the draft email shows these placeholders?

- A. The user does not have Einstein Sales Emails permission assigned.
- B. The user does not have permission to access the fields.
- C. The user's locale language is not supported by Prompt Builder.

Answer: B

Explanation:

When using Einstein Generative AI to create personalized emails, if placeholders appear in the draft email where data from a recipient's Contact record should be, the most likely reason is that the user lacks permission to access the necessary fields. Salesforce's field-level security may prevent users from viewing or utilizing certain data fields, resulting in placeholders being shown instead of the actual values.

? Option B is correct because missing field permissions will cause placeholders in email drafts.

? Option A (missing Einstein Sales Emails permission) is unlikely, as this would prevent email generation altogether, not just placeholders.

? Option C (locale language issues) would more likely affect language-specific issues, not field placeholders.

References:

? Salesforce Email Template and Permissions Documentation: https://help.salesforce.com/s/articleView?id=sf.email_templates_field_permissions.htm

NEW QUESTION 2

A service agent is looking at a custom object that stores travel information. They recently received a weather alert and now need to cancel flights for the customers that are related with this itinerary. The service agent needs to review the Knowledge articles about canceling and rebooking the customer flights.

Which Einstein Copilot capability helps the agent accomplish this?

- A. Execute tasks based on available actions, answering questions using information from accessible Knowledge articles.
- B. Invoke a flow which makes a call to external data to create a Knowledge article.
- C. Generate a Knowledge article based off the prompts that the agent enters to create steps to cancel flights.

Answer: A

Explanation:

In this scenario, the Einstein Copilot capability that best helps the agent is its ability to execute tasks based on available actions and answer questions using data from Knowledge articles. Einstein Copilot can assist the service agent by providing relevant Knowledge articles on canceling and rebooking flights, ensuring that the agent has access to the correct steps and procedures directly within the workflow.

This feature leverages the agent's existing context (the travel itinerary) and provides actionable insights or next steps from the relevant Knowledge articles to help the agent quickly resolve the customer's needs.

The other options are incorrect:

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

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

References:

? Salesforce Documentation on Einstein Copilot

? Trailhead Module on Einstein for Service

NEW QUESTION 3

Universal Containers (UC) wants to use the Draft with Einstein feature in Sales Cloud to create a personalized introduction email. After creating a proposed draft email, which predefined adjustment should UC choose to revise the draft with a more casual tone?

- A. Make Less Formal
- B. Enhance Friendliness
- C. Optimize for Clarity

Answer: A

Explanation:

When Universal Containers uses the Draft with Einstein feature in Sales Cloud to create a personalized email, the predefined adjustment to Make Less Formal is the correct option to revise the draft with a more casual tone. This option adjusts the wording of the draft to sound less formal, making the communication more approachable while still maintaining professionalism.

? Enhance Friendliness would make the tone more positive, but not necessarily more casual.

? Optimize for Clarity focuses on making the draft clearer but doesn't adjust the tone. For more details, see Salesforce documentation on Einstein-generated email drafts and tone adjustments.

NEW QUESTION 4

An AI Specialist needs to create a 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 5

An AI Specialist needs to create a prompt template to fill a custom field named Latest Opportunities Summary on the Account object with information from the three most recently opened opportunities.
How should the AI Specialist gather the necessary data for the prompt template?

- A. Create a flow to retrieve the opportunity information.
- B. Select the Account Opportunity object as a resource when creating the prompt template.
- C. Select the latest Opportunities related list as a merge field.

Answer: A

Explanation:

To gather the necessary data for populating the Latest Opportunities Summary custom field on the Account object with information from the three most recently opened opportunities, the AI Specialist should create a flow. A flow can be configured to query and retrieve the required opportunity records based on criteria such as their open date. Once the flow has gathered the necessary data, it can be used in a prompt template or other automation processes to populate the custom field on the Account record.

? Option A is correct because creating a flow allows for dynamic data retrieval and control over the logic for selecting the most recent opportunities.

? Option B and Option C do not provide sufficient control or data retrieval capabilities needed for this scenario.

References:

? Salesforce Flow Documentation: <https://help.salesforce.com/s/articleView?id=sf.flow.htm>

NEW QUESTION 6

Universal Containers is planning a marketing email about products that most closely match a customer's expressed interests.
What should an AI Specialist recommend to generate this email?

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

Answer: B

Explanation:

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

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

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

NEW QUESTION 7

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 8

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 9

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 10

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 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 15

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 18

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

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

Answer: B

Explanation:

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

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

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

Salesforce AI Specialist References: For more information on Einstein Reply

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

NEW QUESTION 20

Universal Containers (UC) is experimenting with using public Generative AI models and is familiar with the language required to get the information it needs. However, it can be time-consuming for both UC's sales and service reps to type in the prompt to get the information they need, and ensure prompt consistency. Which Salesforce feature should a Salesforce AI Specialist recommend to address these concerns?

- A. Einstein Recommendation Builder
- B. Einstein Copilot Action: Query Records
- C. Einstein Prompt Builder and Prompt Templates

Answer: C

Explanation:

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

? Einstein Prompt Builder simplifies the creation of prompts, and Prompt Templates standardize the inputs, saving time for sales and service reps.

? Option A (Einstein Recommendation Builder) is more focused on recommendations, not prompt standardization.

? Option B (Einstein Copilot Action: Query Records) is for querying records, not generating AI-driven prompts.

References:

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

NEW QUESTION 24

Universal Containers (UC) is using Einstein Generative AI to generate an account summary. UC aims to ensure the content is safe and inclusive, utilizing the Einstein Trust Layer's toxicity scoring to assess the content's safety level.

What does a safety category score of 1 indicate in the Einstein Generative Toxicity Score?

- A. Not safe
- B. Safe
- C. Moderately safe

Answer: B

Explanation:

In the Einstein Trust Layer, the toxicity scoring system is used to evaluate the safety level of content generated by AI, particularly to ensure that it is non-toxic, inclusive, and appropriate for business contexts. A toxicity score of 1 indicates that the content is deemed safe. The scoring system ranges from 0 (unsafe) to 1 (safe), with intermediate values indicating varying degrees of safety. In this case, a score of 1 means that the generated content is fully safe and meets the trust and compliance guidelines set by the Einstein Trust Layer. For further reference, check Salesforce's official Einstein Trust Layer documentation regarding toxicity scoring for AI-generated content.

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:

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

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

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

The Einstein Trust Layer allows for more granular control over what data is exposed to the AI model and ensures compliance with privacy regulations. Salesforce AI Specialist References: For more information, refer to: https://help.salesforce.com/s/articleView?id=sf.einstein_trust_layer_data_masking.htm

NEW QUESTION 33

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

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

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

Answer: A

Explanation:

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

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

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

NEW QUESTION 37

Universal Containers is evaluating Einstein Generative AI features to improve the productivity of the service center operation. Which features should the AI Specialist recommend?

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

Answer: A

Explanation:

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

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

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

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

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

NEW QUESTION 42

What is the correct process to leverage Prompt Builder in a Salesforce org?

- A. Select the appropriate prompt template type to use, select one of Salesforce's standard prompts, determine the object to associate the prompt, select a record to validate against, and associate the prompt to an action.
- B. Select the appropriate prompt template type to use, develop the prompt within the prompt workspace, select resources to dynamically insert CRM-derived grounding data, pick the model to use, and test and validate the generated responses.
- C. Enable the target object for generative prompting, develop the prompt within the prompt workspace, select records to fine-tune and ground the response, enable the Trust Layer, and associate the prompt to an action.

Answer: B

Explanation:

When using Prompt Builder in a Salesforce org, the correct process involves several important steps:

- ? Select the appropriate prompt template type based on the use case.
- ? Develop the prompt within the prompt workspace, where the template is created and customized.
- ? Select CRM-derived grounding data to be dynamically inserted into the prompt, ensuring that the AI-generated responses are based on accurate and relevant data.
- ? Pick the model to use for generating responses, either using Salesforce's built-in models or custom ones.
- ? Test and validate the generated responses to ensure accuracy and effectiveness.
- ? Option B is correct as it follows the proper steps for using Prompt Builder.
- ? Option A and Option C do not capture the full process correctly.

References:

- ? Salesforce Prompt Builder Documentation: https://help.salesforce.com/s/articleView?id=sf.prompt_builder_overview.htm

NEW QUESTION 45

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 46

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 47

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

What is a key reason for implementing Einstein Copilot?

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

Answer: C

Explanation:

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

- ? Option A (Improving data entry and cleansing) is not the primary purpose of Einstein Copilot, as its focus is on guiding and assisting users through workflows.
- ? Option B (Allowing AI to perform tasks without user interaction) does not accurately describe the role of Einstein Copilot, which operates interactively to assist users in real time.

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

NEW QUESTION 51

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

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

Answer: A

Explanation:

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

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

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

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

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

References:

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

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

NEW QUESTION 53

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 54

Universal Containers wants to reduce overall agent handling time by 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 57

A sales rep at Universal Containers is extremely busy and sometimes will have very long sales calls on voice and video calls and might miss key details. They are just starting to adopt new generative AI features.

Which Einstein Generative AI feature should an AI Specialist recommend to help the rep get the details they might have missed during a conversation?

- A. Call Summary

- B. Call Explorer
- C. Sales Summary

Answer: A

Explanation:

For a sales rep who may miss key details during long sales calls, the AI Specialist should recommend the Call Summary feature. Call Summary uses Einstein Generative AI to automatically generate a concise summary of important points discussed during the call, helping the rep quickly review the key information they might have missed.

? Call Explorer is designed for manually searching through call data but doesn't summarize.

? Sales Summary is focused more on summarizing overall sales activity, not call-specific content.

For more details, refer to Salesforce's Call Summary documentation on how AI-generated summaries can improve sales rep productivity.

NEW QUESTION 59

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 60

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 65

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

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

Answer: A

Explanation:

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

? Option A is correct because Einstein Copilot brings a conversational interface that caters to a wide range of users.

? Option B and Option C are more focused on developing and training AI models, which are not the primary functions of Einstein Copilot.

References:

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

NEW QUESTION 68

Universal Containers needs a tool that can analyze voice and video call records to provide insights on competitor mentions, coaching opportunities, and other key information. The goal is to enhance the team's performance by identifying areas for improvement and competitive intelligence. Which feature provides insights about competitor mentions and coaching opportunities?

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

Answer: C

Explanation:

For analyzing voice and video call records to gain insights into competitor mentions, coaching opportunities, and other key information, Call Explorer is the most suitable feature. Call Explorer, a part of Einstein Conversation Insights, enables sales teams to analyze calls, detect patterns, and identify areas where improvements can be made. It uses natural language processing (NLP) to extract insights, including competitor mentions and moments for coaching. These insights are vital for improving sales performance by providing a clear understanding of the interactions during calls.

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

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

References:

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

NEW QUESTION 73

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 78

An AI Specialist is tasked with configuring a generative model to create personalized sales emails using customer data stored in Salesforce. The AI Specialist has already fine-tuned a large language model (LLM) on the OpenAI platform. Security and data privacy are critical concerns for the client. How should the AI Specialist integrate the custom LLM into Salesforce?

- A. Create an application of the custom LLM and embed it in Sales Cloud via iFrame.
- B. Add the fine-tuned LLM in Einstein Studio Model Builder.
- C. Enable model endpoint on OpenAI and make callouts to the model to generate emails.

Answer: B

Explanation:

Since security and data privacy are critical, the best option for the AI Specialist is to integrate the fine-tuned LLM (Large Language Model) into Salesforce by adding it to Einstein Studio Model Builder. Einstein Studio allows organizations to bring their own AI models (BYOM), ensuring the model is securely managed within Salesforce's environment, adhering to data privacy standards.

? Option A (embedding via iFrame) is less secure and doesn't integrate deeply with Salesforce's data and security models.

? Option C (making callouts to OpenAI) raises concerns about data privacy, as sensitive Salesforce data would be sent to an external system.

Einstein Studio provides the most secure and seamless way to integrate custom AI models while maintaining control over data privacy and compliance. More details can be found in Salesforce's Einstein Studio documentation on integrating external models.

NEW QUESTION 81

.....

Relate Links

100% Pass Your Salesforce-AI-Specialist Exam with Exambible Prep Materials

<https://www.exambible.com/Salesforce-AI-Specialist-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>