



Free Questions for [Salesforce-AI-Specialist](#) by [certsinside](#)

Shared by [Gillespie](#) on [23-09-2024](#)

For More Free Questions and Preparation Resources

[Check the Links on Last Page](#)

Question 1

Question Type: MultipleChoice

Which feature in the Einstein Trust Layer helps to minimize the risks of jailbreaking and prompt injection attacks?

Options:

- A- Secure Data Retrieval and Grounding
- B- Data Masking
- C- Prompt Defense

Answer:

C

Explanation:

Prompt Defense is a feature in the Einstein Trust Layer that helps minimize the risks of jailbreaking and prompt injection attacks. These attacks occur when malicious users try to manipulate the AI model by providing unintended inputs. Prompt Defense ensures that the prompts are processed securely, protecting the system from such vulnerabilities.

Option A (Secure Data Retrieval and Grounding) relates to ensuring that data used by AI is securely retrieved but does not address prompt security.

Option B (Data Masking) focuses on protecting sensitive information but does not prevent injection attacks.

For more information, refer to Salesforce's Einstein Trust Layer documentation on Prompt Defense and security features.

Question 2

Question Type: MultipleChoice

Universal Containers plans to enhance the customer support team's productivity using AI.

Which specific use case necessitates the use of Prompt Builder?

Options:

- A- Creating a draft of a support bulletin post for new product patches
- B- Creating an AI-generated customer support agent performance score
- C- Estimating support ticket volume based on historical data and seasonal trends

Answer:

A

Explanation:

The use case that necessitates the use of Prompt Builder is creating a draft of a support bulletin post for new product patches. Prompt Builder allows the AI Specialist to create and refine prompts that generate specific, relevant outputs, such as drafting support communication based on product information and patch details.

Option B (agent performance score) would likely involve predictive modeling, not prompt generation.

Option C (estimating support ticket volume) would require data analysis and predictive tools, not prompt building.

For more details, refer to Salesforce's Prompt Builder documentation for generative AI content creation.

Question 3

Question Type: MultipleChoice

An AI Specialist is considering using a Field Generation prompt template type.

What should the AI Specialist check before creating the Field Generation prompt to ensure it is possible for the field to be enabled for generative AI?

Options:

- A-** That the field chosen must be a rich text field with 255 characters or more.
- B-** That the org is set to API version 59 or higher
- C-** That the Lightning page layout where the field will reside has been upgraded to Dynamic Forms

Answer:

B

Explanation:

Before creating a Field Generation prompt template, the AI Specialist must ensure that the Salesforce org is set to API version 59 or higher. This version of the API introduces support for advanced generative AI features, such as enabling fields for generative AI outputs. This is a critical technical requirement for the Field Generation prompt template to function correctly.

Option A (rich text field requirement) is not necessary for generative AI functionality.

Option C (Dynamic Forms) does not impact the ability of a field to be generative AI-enabled, although it might enhance the user interface.

For more information, refer to Salesforce documentation on API versioning and Field Generation templates.

Question 4

Question Type: MultipleChoice

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?

Options:

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.

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

Question 5

Question Type: MultipleChoice

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

Which tool should the AI Specialist recommend?

Options:

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

Answer:

C

Explanation:

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

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

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

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

Question 6

Question Type: MultipleChoice

The AI Specialist of Northern Trail Outfitters reviewed the organization's data masking settings within the Configure Data Masking menu within Setup. Upon assessing all of the fields, a few additional fields

were deemed sensitive and have been masked within Einstein's Trust Layer.

Which steps should the AI Specialist take upon modifying the masked fields?

Options:

- A-** Turn off the Einstein Trust Layer and turn it on again.
- B-** Test and confirm that the responses generated from prompts that utilize the data and masked data do not adversely affect the quality of the generated response
- C-** Turn on Einstein Feedback so that end users can report if there are any negative side effects on AI features.

Answer:

B

Explanation:

After modifying masked fields in Einstein's Trust Layer, the next important step is to test and confirm that the responses generated by prompts utilizing the newly masked data still meet quality standards. This ensures that masking sensitive information does not negatively impact the usefulness or accuracy of the AI-generated content. Thorough testing helps identify any issues in prompt performance that could arise due to masking, and adjustments can be made if needed.

Option B is correct because testing the effects of masking on AI responses is a critical step in ensuring AI continues to function as expected.

Option A (turning off and on the Einstein Trust Layer) is unnecessary after changing the masked fields.

Option C (turning on Einstein Feedback) allows for user feedback but is not a direct step following field masking modifications.

[Salesforce Einstein Trust Layer Overview: https://help.salesforce.com/s/articleView?id=sf.einstein_trust_layer.htm](https://help.salesforce.com/s/articleView?id=sf.einstein_trust_layer.htm)

Question 7

Question Type: MultipleChoice

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?

Options:

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.

Question 8

Question Type: MultipleChoice

Universal Containers (UC) has recently received an increased number of support cases. As a result, UC has hired more customer support reps and has started to assign some of the ongoing cases to newer reps.

Which generative AI solution should the new support reps use to understand the details of a case without reading through each case comment?

Options:

- A- Einstein Copilot
- B- Einstein Sales Summaries
- C- Einstein Work Summaries

Answer:

C

Explanation:

New customer support reps at Universal Containers can use Einstein Work Summaries to quickly understand the details of a case without reading through each case comment. Work Summaries leverage generative AI to provide a concise overview of ongoing cases, summarizing all relevant information in an easily digestible format.

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

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

For more details, refer to Salesforce documentation on Einstein Work Summaries.

Question 9

Question Type: MultipleChoice

Universal Containers (UC) plans to send one of three different emails to its customers based on the customer's lifetime value score and their market segment.

Considering that UC are required to explain why an e-mail was selected, which AI model should UC use to achieve this?

Options:

A- Predictive model and generative model

B- Generative model

C- Predictive model

Answer:

C

Explanation:

Universal Containers should use a Predictive model to decide which of the three emails to send based on the customer's lifetime value score and market segment. Predictive models analyze data to forecast outcomes, and in this case, it would predict the most appropriate email to send based on customer attributes. Additionally, predictive models can provide explainability to show why a certain email was chosen, which is crucial for UC's requirement to explain the decision-making process.

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

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

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

Question 10

Question Type: MultipleChoice

Amid their busy schedules, sales reps at Universal Containers dedicate time to follow up with prospects and existing clients via email regarding renewals or new deals. They spend many hours throughout the

week reviewing past communications and details about their customers before performing their outreach.

Which standard Copilot action helps sales reps draft personalized emails to prospects by generating text based on previous successful communications?

Options:

A- Einstein Copilot Action: Find Similar Opportunities

B- Einstein Copilot Action: Draft or Revise Sales Email

C- Einstein Copilot Action: Summarize Record

Answer:

B

Explanation:

For sales reps who need to draft personalized emails based on previous communications, the AI Specialist should recommend the Einstein Copilot Action: Draft or Revise Sales Email. This action uses AI to generate or revise email content, leveraging past successful communications to create personalized and relevant outreach to prospects or clients.

Find Similar Opportunities is used for opportunity matching, not email drafting.

Summarize Record provides a summary of customer data but does not directly help with drafting emails.

For more information, refer to Salesforce's Einstein Copilot documentation on standard actions for sales teams.

Question 11

Question Type: MultipleChoice

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

Which features should the AI Specialist recommend?

Options:

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.

Question 12

Question Type: MultipleChoice

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?

Options:

- 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.

To Get Premium Files for Salesforce-AI-Specialist Visit

<https://www.p2pexams.com/products/salesforce-ai-specialist>

For More Free Questions Visit

<https://www.p2pexams.com/salesforce/pdf/salesforce-ai-specialist>

