

## Salesforce-AI-Specialist Dumps

### Salesforce Certified AI Specialist Exam

<https://www.certleader.com/Salesforce-AI-Specialist-dumps.html>



#### NEW QUESTION 1

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 2

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

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

**Answer:** A

#### Explanation:

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

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

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

#### NEW QUESTION 3

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 4

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](https://help.salesforce.com/s/articleView?id=sf.einstein_trust_layer_overview.htm)

#### NEW QUESTION 5

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 6**

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](https://help.salesforce.com/s/articleView?id=sf.prompt_template_flow_integration.htm)

**NEW QUESTION 7**

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](https://help.salesforce.com/s/articleView?id=sf.model_builder_external_models.htm)

**NEW QUESTION 8**

Universal Containers wants to use an external large language model (LLM) in Prompt Builder. What should an AI Specialist recommend?

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

**Answer:** B

**Explanation:**

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

? Option B is the correct choice as Einstein Studio provides a built-in feature to work with external models.  
? Option A suggests using Apex, but BYO-LLM functionality offers a more streamlined solution.  
? Option C focuses on Flow and External Services, which is more about data integration and isn't ideal for working with LLMs.  
References:  
? Salesforce Einstein Studio BYO-LLM Documentation: [https://help.salesforce.com/s/articleView?id=sf.einstein\\_studio\\_llm](https://help.salesforce.com/s/articleView?id=sf.einstein_studio_llm).

#### NEW QUESTION 9

Universal Containers (UC) wants to use Flow to bring data from unified Data Cloud objects to prompt templates. Which type of flow should UC use?

- A. Data Cloud-triggered flow
- B. Template-triggered prompt flow
- C. Unified-object linking flow

**Answer:** A

#### Explanation:

In this scenario, Universal Containers wants to bring data from unified Data Cloud objects into prompt templates, and the best way to do that is through a Data Cloud-triggered flow. This type of flow is specifically designed to trigger actions based on data changes within Salesforce Data Cloud objects. Data Cloud-triggered flows can listen for changes in the unified data model and automatically bring relevant data into the system, making it available for prompt templates. This ensures that the data is both real-time and up-to-date when used in generative AI contexts. For more detailed guidance, refer to Salesforce documentation on Data Cloud-triggered flows and Data Cloud integrations with generative AI solutions.

#### NEW QUESTION 10

Universal Containers (UC) uses Salesforce Service Cloud to support its customers and agents handling cases. UC is considering implementing Einstein Copilot and extending Service Cloud to mobile users. When would Einstein Copilot implementation be most advantageous?

- A. When the goal is to streamline customer support processes and improve response times
- B. When the main objective is to enhance data security and compliance measures
- C. When the focus is on optimizing marketing campaigns and strategies

**Answer:** A

#### Explanation:

Einstein Copilot implementation would be most advantageous in Salesforce Service Cloud when the goal is to streamline customer support processes and improve response times. Einstein Copilot can assist agents by providing real-time suggestions, automating repetitive tasks, and generating contextual responses, thus enhancing service efficiency.  
? Option B (data security) is not the primary focus of Einstein Copilot, which is more about improving operational efficiency.  
? Option C (marketing campaigns) falls outside the scope of Service Cloud and Einstein Copilot's primary benefits, which are aimed at improving customer service and case management.  
For further reading, refer to Salesforce documentation on Einstein Copilot for Service Cloud and how it improves support processes.

#### NEW QUESTION 10

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](https://help.salesforce.com/s/articleView?id=sf.prompt_builder_overview.htm)

#### NEW QUESTION 11

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](https://help.salesforce.com/s/articleView?id=sf.einstein_trust_layer_data_masking.htm)

#### NEW QUESTION 14

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

Which tool should the AI Specialist recommend?

- 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.

References:

? Salesforce Copilot Builder Overview: [https://help.salesforce.com/s/articleView?id=sf.einstein\\_copilot\\_builder.htm](https://help.salesforce.com/s/articleView?id=sf.einstein_copilot_builder.htm)

#### NEW QUESTION 17

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](https://help.salesforce.com/s/articleView?id=sf.einstein_trust_layer.htm)

#### NEW QUESTION 18

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 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 24**

Universal Containers is considering leveraging the Einstein Trust Layer in conjunction with Einstein Generative AI Audit Data. Which audit data is available using the Einstein Trust Layer?

- A. Response accuracy and offensiveness score
- B. Hallucination score and bias score
- C. Masked data and toxicity score

**Answer:** C

**Explanation:**

Universal Containers is considering the use of the Einstein Trust Layer along with Einstein Generative AI Audit Data. The Einstein Trust Layer provides a secure and compliant way to use AI by offering features like data masking and toxicity assessment.

The audit data available through the Einstein Trust Layer includes information about

masked data—which ensures sensitive information is not exposed—and the toxicity score

, which evaluates the generated content for inappropriate or harmful language.

References:

? Salesforce AI Specialist Documentation - Einstein Trust Layer: Details the auditing capabilities, including logging of masked data and evaluation of generated responses for toxicity to maintain compliance and trust.

**NEW QUESTION 26**

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 28**

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

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

**Answer:** B

**Explanation:**

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

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

The other options are incorrect because:

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

References:

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

### NEW QUESTION 29

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

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

**Answer: B**

#### Explanation:

When the "Enrich event logs with conversation data" setting is enabled in Einstein Copilot, it allows an AI Specialist or admin to view session data, including both the user input and copilot responses from interactions over the past 7 days. This data is crucial for monitoring how the copilot is being used, analyzing its performance, and improving future interactions based on past inputs.

? This setting enriches the event logs with detailed conversational data for better insights into the interaction history, helping AI specialists track AI behavior and user engagement.

? Option A, viewing the user click path, focuses on navigation but is not part of the conversation data enrichment functionality.

? Option C, generating detailed reports over any time period, is incorrect because this specific feature is limited to data for the past 7 days.

Salesforce AI Specialist References: You can refer to this documentation for further insights: [https://help.salesforce.com/s/articleView?id=sf.einstein\\_copilot\\_event\\_logging.htm](https://help.salesforce.com/s/articleView?id=sf.einstein_copilot_event_logging.htm)

### NEW QUESTION 32

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?

- 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.

### NEW QUESTION 34

Universal Containers' service team wants to customize the standard case summary response from Einstein Copilot.

What should the AI Specialist do to achieve this?

- A. Customize the standard Record Summary template for the Case object,
- B. Summarize the Case with a standard copilot action.
- C. Create a custom Record Summary prompt template for the Case object.

**Answer: C**

#### Explanation:

To customize the case summary response from Einstein Copilot, the AI Specialist should create a custom Record Summary prompt template for the Case object. This allows Universal Containers to tailor the way case data is summarized, ensuring the output aligns with specific business requirements or user preferences.

? Option A (customizing the standard Record Summary template) does not provide the flexibility required for deep customization.

? Option B (standard Copilot action) won't allow customization; it will only use default settings.

Refer to Salesforce Prompt Builder documentation for guidance on creating custom templates for record summaries.

### NEW QUESTION 37

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 39

Universal Containers wants to be able to detect with a high level confidence if content generated by a large language model (LLM) contains toxic language. Which action should an AI Specialist take in the Trust Layer to confirm toxicity is being appropriately managed?

- A. Access the Toxicity Detection log in Setup and export all entries where isToxicityDetected is true.
- B. Create a flow that sends an email to a specified address each time the toxicity score from the response exceeds a predefined threshold.
- C. Create a Trust Layer audit report within Data Cloud that uses a toxicity detector type filter to display toxic responses and their respective scores.

**Answer: C**

#### Explanation:

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

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

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

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

References:

? Salesforce Trust Layer Documentation:[https://help.salesforce.com/s/articleView?id=sf.einstein\\_trust\\_layer\\_audit.htm](https://help.salesforce.com/s/articleView?id=sf.einstein_trust_layer_audit.htm)

### NEW QUESTION 41

The sales team at a hotel resort would like to generate a guest summary about the guests' interests and provide recommendations based on their activity preferences captured in each guest profile. They want the summary to be available only on the contact record page. Which AI capability should the team use?

- A. Einstein Copilot
- B. Prompt Builder
- C. Model Builder

**Answer: B**

#### Explanation:

The sales team at a hotel resort wants to generate a guest summary about guests' interests and provide recommendations based on their activity preferences captured in each guest profile. They require the summary to be available only on the contact record page.

Solution:

? Use Prompt Builder to create a prompt template that generates the desired summary and displays it on the contact record page.

? Prompt Builder:

? Implementation Steps:

? Why Not Einstein Copilot or Model Builder:

References:

? Salesforce AI Specialist Documentation - Prompt Builder Overview:

? Salesforce Help - Creating Field Generation Prompt Templates:

? Salesforce Trailhead - Customize AI Content with Prompt Builder:

Conclusion:

By utilizing Prompt Builder, the sales team can create a customized prompt template that generates personalized guest summaries and recommendations based on activity preferences. This solution meets the requirement of displaying the summary only on the contact record page, enhancing the team's ability to engage with guests effectively.

### NEW QUESTION 45

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](https://help.salesforce.com/s/articleView?id=sf.einstein_studio_overview)

### NEW QUESTION 48

What is the main purpose of Prompt Builder?

- A. A tool for developers to use in Visual Studio Code that creates prompts for Apex programming, assisting developers in writing code more efficiently.
- B. A tool that enables companies to create reusable prompts for large language models (LLMs), bringing generative AI responses to their flow of work.
- C. A tool within Salesforce offering real-time AI-powered suggestions and guidance to users, improving productivity and decision-making.

**Answer:** B

**Explanation:**

Prompt Builder is designed to help organizations create and configure reusable prompts for large language models (LLMs). By integrating generative AI responses into workflows, Prompt Builder enables customization of AI prompts that interact with Salesforce data and automate complex processes. This tool is especially useful for creating tailored and consistent AI-generated content in various business contexts, including customer service and sales.

? It is not a tool for Apex programming (as in option A).

? It is also not limited to real-time suggestions as mentioned in option C. Instead, it provides a flexible way for companies to manage and customize how AI-driven responses are generated and used in their workflows.

References:

? Salesforce Prompt Builder

Overview: [https://help.salesforce.com/s/articleView?id=sf.prompt\\_builder.htm](https://help.salesforce.com/s/articleView?id=sf.prompt_builder.htm)

**NEW QUESTION 51**

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 52**

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 57**

An AI Specialist built a Field Generation prompt template that worked for many records, but users are reporting random failures with token limit errors.

What is the cause of the random nature of this error?

- A. The number of tokens generated by the dynamic nature of the prompt template will vary by record.
- B. The template type needs to be switched to Flex to accommodate the variable amount of tokens generated by the prompt grounding.
- C. The number of tokens that can be processed by the LLM varies with total user demand.

**Answer:** A

**Explanation:**

The reason behind the token limit errors lies in the dynamic nature of the prompt template used in Field Generation. In Salesforce's AI generative models, each prompt and its corresponding output are subject to a token limit, which encompasses both the input and output of the large language model (LLM). Since the prompt template dynamically adjusts based on the specific data of each record, the number of tokens varies per record. Some records may generate longer outputs based on their data attributes, pushing the token count beyond the allowable limit for the LLM, resulting in token limit errors.

This behavior explains why users experience random failures—it is dependent on the specific data used in each case. For certain records, the combined input and output may fall within the token limit, while for others, it may exceed it. This variation is intrinsic to how dynamic templates interact with large language models. Salesforce provides guidance in their documentation, stating that prompt template design should take into account token limits and suggests testing with varied records to avoid such random errors. It does not mention switching to Flex template type as a solution, nor does it suggest that token limits fluctuate with user demand. Token limits are a constant defined by the model itself, independent of external user load.

References:

? Salesforce Developer Documentation on Token Limits for Generative AI Models

? Salesforce AI Best Practices on Prompt Design (Trailhead or Salesforce blog resources)

**NEW QUESTION 60**

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

- 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.

**NEW QUESTION 64**

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](https://help.salesforce.com/s/articleView?id=einstein_copilot_overview.htm)

**NEW QUESTION 68**

When configuring a prompt template, an AI Specialist previews the results of the prompt template they've written. They see two distinct text outputs: Resolution and Response.

Which information does the Resolution text provide?

A. It shows the full text that is sent to the Trust Layer.

B. It shows the response from the LLM based on the sample record.

C. It shows which sensitive data is masked before it is sent to the LLM.

**Answer: B**

**Explanation:**

When previewing a prompt template in Salesforce, the Resolution text provides the response from the LLM (Large Language Model) based on the data from a sample record. This output shows what the AI model generated in response to the prompt, giving the AI Specialist a chance to review and adjust the response before finalizing the template.

? Option B is correct because Resolution displays the actual response generated by the LLM.

? Option A refers to sending the text to the Trust Layer, but that's not what Resolution represents.

? Option C relates to data masking, which is shown elsewhere, not under Resolution.

References:

? Salesforce Prompt Builder Overview: [https://help.salesforce.com/s/articleView?id=sf.prompt\\_builder\\_overview.htm](https://help.salesforce.com/s/articleView?id=sf.prompt_builder_overview.htm)

**NEW QUESTION 72**

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 74**

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 77**

Universal Containers plans to implement prompt templates that utilize the standard foundation models. What should the AI Specialist consider when building prompt templates in Prompt Builder?

- A. Include multiple-choice questions within the prompt to test the LLM's understanding of the context.
- B. Ask it to role-play as a character in the prompt template to provide more context to the LLM.
- C. Train LLM with data using different writing styles including word choice, intensifiers, emojis, and punctuation.

**Answer: C**

**Explanation:**

When building prompt templates in Prompt Builder, it is essential to consider how the Large Language Model (LLM) processes and generates outputs. Training the LLM with various writing styles, such as different word choices, intensifiers, emojis, and punctuation, helps the model better understand diverse writing patterns and produce more contextually appropriate responses.

This approach enhances the flexibility and accuracy of the LLM when generating outputs for different use cases, as it is trained to recognize various writing conventions and styles. The prompt template should focus on providing rich context, and this stylistic variety helps improve the model's adaptability.

Options A and B are less relevant because adding multiple-choice questions or role-playing scenarios doesn't contribute significantly to improving the AI's output generation quality within standard business contexts.

For more details, refer to Salesforce's Prompt Builder documentation and LLM tuning strategies.

**NEW QUESTION 82**

An AI Specialist has created a copilot custom action using flow as the reference action type. However, it is not delivering the expected results to the conversation preview, and therefore needs troubleshooting. What should the AI Specialist do to identify the root cause of the problem?

- A. In Copilot Builder within the Dynamic Panel, turn on dynamic debugging to show the inputs and outputs.
- B. Copilot Builder within the Dynamic Panel, confirm selected action and observe the values in Input and Output sections.
- C. In Copilot Builder, verify the utterance entered by the user and review session event logs for debug information.

**Answer: A**

**Explanation:**

When troubleshooting a copilot custom action using flow as the reference action type, enabling dynamic debugging within Copilot Builder's Dynamic Panel is the most effective way to identify the root cause. By turning on dynamic debugging, the AI Specialist can see detailed logs showing both the inputs and outputs of the flow, which helps identify where the action might be failing or not delivering the expected results.

? Option B, confirming selected actions and observing the Input and Output

sections, is useful for monitoring flow configuration but does not provide the deep diagnostic details available with dynamic debugging.

? Option C, verifying the user utterance and reviewing session event logs, could

provide helpful context, but dynamic debugging is the primary tool for identifying issues with inputs and outputs in real time.

Salesforce AI Specialist References: To explore more about dynamic debugging in Copilot Builder,

see: [https://help.salesforce.com/s/articleView?id=sf.copilot\\_custom\\_action\\_debugging.htm](https://help.salesforce.com/s/articleView?id=sf.copilot_custom_action_debugging.htm)

**NEW QUESTION 87**

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 88**

.....

## Thank You for Trying Our Product

\* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

\* One year free update

You can enjoy free update one year. 24x7 online support.

\* Trusted by Millions

We currently serve more than 30,000,000 customers.

\* Shop Securely

All transactions are protected by VeriSign!

**100% Pass Your Salesforce-AI-Specialist Exam with Our Prep Materials Via below:**

<https://www.certleader.com/Salesforce-AI-Specialist-dumps.html>