

## Salesforce-AI-Specialist Dumps

### Salesforce Certified AI Specialist Exam

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



### NEW QUESTION 1

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

Which Einstein Copilot capability helps the agent accomplish this?

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

**Answer:** A

#### Explanation:

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

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

The other options are incorrect:

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

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

References:

? Salesforce Documentation on Einstein Copilot

? Trailhead Module on Einstein for Service

### NEW QUESTION 2

An AI Specialist needs to create a 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 3

An AI Specialist wants to ground a new prompt template with the User related list. What should the AI Specialist consider?

- A. The User related list should have View All access.
- B. The User related list needs to be included on the record page.
- C. The User related list is not supported in prompt templates.

**Answer:** C

#### Explanation:

An AI Specialist wants to ground a new prompt template with the User related list. Grounding in prompt templates involves using data from related lists to provide context or additional information to the Large Language Model (LLM) when generating responses.

Key Consideration:

? Unsupported Related Lists in Prompt Templates:

Why Options A and B are Incorrect:

? Option A (The User related list should have View All access):

? Option B (The User related list needs to be included on the record page):

References:

? Salesforce AI Specialist Documentation - Prompt Templates Limitations:

? Salesforce Help - Data Access in Prompt Templates:

? Salesforce Trailhead - Understanding Prompt Template Grounding:

Conclusion:

Since the User related list is not supported in prompt templates, the AI Specialist must

consider alternative approaches. They might need to redesign the prompt template to use supported objects or related lists, or explore other methods to incorporate necessary user information while adhering to Salesforce's data access policies.

### NEW QUESTION 4

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

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

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

**Answer:** B

**Explanation:**

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

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

The other options are less aligned with the requirement:

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

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

References:

? Salesforce Developer Documentation on Einstein Trust Layer

? Salesforce Security Documentation for AI and Data Privacy

**NEW QUESTION 5**

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 6**

Universal Containers (UC) noticed an increase in customer contract cancellations in the last few months. UC is seeking ways to address this issue by implementing a proactive outreach program to

customers before they cancel their contracts and is asking the Salesforce team to provide suggestions.

Which use case functionality of Model Builder aligns with UC's request?

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

**Answer:** B

**Explanation:**

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

? Option B is correct because customer churn prediction aligns with UC's need to

reduce cancellations through proactive measures.

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

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

References:

? Salesforce Model Builder Use Case Overview: [https://help.salesforce.com/s/articleView?id=sf.model\\_builder\\_use\\_case\\_s.htm](https://help.salesforce.com/s/articleView?id=sf.model_builder_use_case_s.htm)

**NEW QUESTION 7**

Northern Trail Outfitters (NTO) wants to configure Einstein Trust Layer in its production org but is unable to see the option on the Setup page.

After provisioning Data Cloud, which step must an AI Specialist take to make this option available to NTO?

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

**Answer:** B

**Explanation:**

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

? Option A (Turning on Einstein Copilot) is unrelated to the setup of the Einstein

Trust Layer, which focuses more on generative AI interactions and data handling.

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

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

steps: [https://help.salesforce.com/s/articleView?id=sf.einstein\\_trust\\_layer\\_overview.htm](https://help.salesforce.com/s/articleView?id=sf.einstein_trust_layer_overview.htm)

#### NEW QUESTION 8

Universal Containers (UC) recently rolled out Einstein Generative capabilities and has created a custom prompt to summarize case records. Users have reported that the case summaries generated are not returning the appropriate information. What is a possible explanation for the poor prompt performance?

- A. The data being used for grounding is incorrect or incomplete.
- B. The prompt template version is incompatible with the chosen LLM.
- C. The Einstein Trust Layer is incorrectly configured.

**Answer:** A

#### Explanation:

Poor prompt performance when generating case summaries is often due to the data used for grounding being incorrect or incomplete. Grounding involves feeding accurate, relevant data to the AI so it can generate appropriate outputs. If the data source is incomplete or contains errors, the generated summaries will reflect that by being inaccurate or insufficient.

? Option B (prompt template incompatibility with the LLM) is unlikely because such incompatibility usually results in more technical failures, not poor content quality.

? Option C (Einstein Trust Layer misconfiguration) is focused on data security and auditing, not the quality of prompt responses.

For more information, refer to Salesforce documentation on grounding AI models and data quality best practices.

#### NEW QUESTION 9

Universal Containers plans to enhance the customer support team's productivity using AI. Which specific use case necessitates the use of Prompt Builder?

- 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.

#### NEW QUESTION 10

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

Which API service should UC use to meet this requirement?

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

**Answer:** A

#### Explanation:

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

? Metadata API is used for retrieving and deploying metadata, not for data operations like generating summaries.

? SOAP API is an older API used for integration but is less flexible compared to REST for this specific use case.

For more details, refer to Salesforce REST API documentation regarding using REST for data integration and generating content.

#### NEW QUESTION 10

Universal Containers wants to utilize Einstein for Sales to help sales reps reach their sales quotas by providing AI-generated plans containing guidance and steps for closing deals.

Which feature should the AI Specialist recommend to the sales team?

- A. Find Similar Deals
- B. Create Account Plan
- C. Create Close Plan

**Answer:** C

#### Explanation:

The "Create Close Plan" feature is designed to help sales reps by providing AI-generated strategies and steps specifically focused on closing deals. This feature leverages AI to analyze the current state of opportunities and generate a plan that outlines the actions, timelines, and key steps required to move deals toward closure. It aligns directly with the sales team's need to meet quotas by offering actionable insights and structured plans.

? Find Similar Deals (Option A) helps sales reps discover opportunities similar to their current deals but doesn't offer a plan for closing.

? Create Account Plan (Option B) focuses on long-term strategies for managing accounts, which might include customer engagement and retention, but doesn't focus on deal closure.

Salesforce AI Specialist References: For more information on using AI for sales, visit: [https://help.salesforce.com/s/articleView?id=sf.einstein\\_for\\_sales\\_overview.htm](https://help.salesforce.com/s/articleView?id=sf.einstein_for_sales_overview.htm)

#### NEW QUESTION 14

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 15

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

#### NEW QUESTION 17

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 21

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 23

An AI Specialist turned on Einstein Generative AI in Setup. Now, the AI Specialist would like to create custom prompt templates in Prompt Builder. However, they cannot access Prompt Builder in the Setup menu. What is causing the problem?

- A. The Prompt Template User permission set was not assigned correctly.
- B. The Prompt Template Manager permission set was not assigned correctly.
- C. The large language model (LLM) was not configured correctly in Data Cloud.

**Answer: B**

#### Explanation:

In order to access and create custom prompt templates in Prompt Builder, the AI Specialist must have the Prompt Template Manager permission set assigned. Without this permission, they will not be able to access Prompt Builder in the Setup menu, even though Einstein Generative AI is enabled.

? Option B is correct because the Prompt Template Manager permission set is required to use Prompt Builder.

? Option A (Prompt Template User permission set) is incorrect because this permission allows users to use prompts, but not create or manage them.

? Option C (LLM configuration in Data Cloud) is unrelated to the ability to access Prompt Builder.

References:

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

### NEW QUESTION 27

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 28

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 31

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

### NEW QUESTION 33

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 36

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?

- 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.

### NEW QUESTION 37

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

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

**Answer:** C

#### Explanation:

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

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

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

References:

? Salesforce AI Specialist Documentation - Einstein Copilot Overview: Details how Einstein Copilot employs LLMs to interpret user inputs and generate responses within the Salesforce ecosystem.

? Salesforce Help - How Einstein Copilot Works: Explains the underlying mechanisms of how Einstein Copilot processes user requests using AI technologies.

### NEW QUESTION 38

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 40

Universal Containers (UC) has a mature Salesforce org with a lot of data in cases and Knowledge articles. UC is concerned that there are many legacy fields, with data that might not be applicable for Einstein AI to draft accurate email responses.

Which solution should UC use to ensure Einstein AI can draft responses from a defined data source?

- A. Service AI Grounding
- B. Work Summaries
- C. Service Replies

**Answer: A**

#### Explanation:

Service AI Grounding is the solution that Universal Containers should use to ensure Einstein AI drafts responses based on a well-defined data source. Service AI Grounding allows the AI model to be anchored in specific, relevant data sources, ensuring that any AI-generated responses (e.g., email replies) are accurate, relevant, and drawn from up-to-date information, such as Knowledge articles or cases.

Given that UC has legacy fields and outdated data, Service AI Grounding ensures that only the valid and applicable data is used by Einstein AI to craft responses. This helps improve the relevance of responses and avoids inaccuracies caused by outdated or irrelevant fields. Work Summaries and Service Replies are useful features but do not address the need for grounding AI outputs in specific, current data sources like Service AI Grounding does.

For more details, you can refer to Salesforce's Service AI Grounding documentation for managing AI-generated content based on accurate data sources.

#### NEW QUESTION 44

An AI Specialist at Universal Containers (UC) is tasked with creating a new custom prompt template to populate a field with generated output. UC enabled the Einstein Trust Layer to ensure AI Audit data is captured and monitored for adoption and possible enhancements.

Which prompt template type should the AI Specialist use and which consideration should they review?

- A. Flex, and that Dynamic Fields is enabled
- B. Field Generation, and that Dynamic Fields is enabled
- C. Field Generation, and that Dynamic Forms is enabled

**Answer: B**

#### Explanation:

When creating a custom prompt template to populate a field with generated output, the most appropriate template type is Field Generation. This template is specifically designed for generating field-specific outputs using generative AI.

Additionally, the AI Specialist must ensure that Dynamic Fields are enabled. Dynamic Fields allow the system to use real-time data inputs from related records or fields when generating content, ensuring that the AI output is contextually accurate and relevant. This is crucial when populating specific fields with AI-generated content, as it ensures the data source remains dynamic and up-to-date.

The Einstein Trust Layer will track and audit the interactions to ensure the organization can monitor AI adoption and make necessary enhancements based on AI usage patterns.

For further reading, refer to Salesforce's guidelines on Field Generation templates and the Einstein Trust Layer.

#### NEW QUESTION 46

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 51

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 52**

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

**NEW QUESTION 57**

What should an AI Specialist consider when using related list merge fields in a prompt template associated with an Account object in Prompt Builder?

- A. The Activities related list on the Account object is not supported because it is a polymorphic field.
- B. If person accounts have been enabled, merge fields will not be available for the Account object.
- C. Prompt generation will yield no response when there is no related list associated with an Account in runtime.

**Answer:** A

**Explanation:**

When using related list merge fields in a prompt template associated with the Account object in Prompt Builder, the Activities related list is not supported due to it being a polymorphic field. Polymorphic fields can reference multiple different types of objects, which makes them incompatible with some merge field operations in prompt generation.

? Option B is incorrect because person accounts do not limit the availability of merge fields for the Account object.

? Option C is irrelevant since even if no related lists are available at runtime, the prompt can still generate based on other available data fields.

For more information, refer to Salesforce documentation on supported fields and limitations in Prompt Builder.

**NEW QUESTION 62**

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

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

**Answer:** A

**Explanation:**

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

? Temperature: Controls the randomness of predictions. A higher temperature leads

to more diverse outputs, while a lower temperature makes the model's responses more focused and deterministic.

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

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

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

**NEW QUESTION 67**

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

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

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

**Answer:** A

**Explanation:**

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

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

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

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

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

References:

? Salesforce Permissions Documentation: [https://help.salesforce.com/s/articleView?id=sf.perm\\_sets\\_overview.htm](https://help.salesforce.com/s/articleView?id=sf.perm_sets_overview.htm)

**NEW QUESTION 70**

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 73**

.....

## 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>