



## **Salesforce**

### **Exam Questions Salesforce-AI-Specialist**

Salesforce Certified AI Specialist Exam

## About ExamBible

### *Your Partner of IT Exam*

## Found in 1998

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

## Our Advances

### \* 99.9% Uptime

All examinations will be up to date.

### \* 24/7 Quality Support

We will provide service round the clock.

### \* 100% Pass Rate

Our guarantee that you will pass the exam.

### \* Unique Gurantee

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

### NEW QUESTION 1

Universal Containers (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

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 3

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 4

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 5

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

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

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

**Answer:** C

**Explanation:**

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

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

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

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

**NEW QUESTION 6**

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 7**

Universal Containers implements Custom Copilot Actions to enhance its customer service operations. The development team needs to understand the core components of a Custom Copilot Action to ensure proper configuration and functionality.

What should the development team review in the Custom Copilot Action configuration to identify one of the core components of a Custom Copilot Action?

- A. Instructions
- B. Output Types
- C. Action Triggers

**Answer:** B

**Explanation:**

Universal Containers is enhancing its customer service operations with Custom Copilot Actions. The development team needs to understand the core components of a Custom Copilot Action to ensure proper configuration and functionality. One of these core components is the Output Types.

? Core Components of a Custom Copilot Action:

? Focus on Output Types:

Why Output Types are a Core Component:

? Integration with Copilot:

? Data Consistency:

? User Experience:

Why Other Options are Less Suitable:

? Option A (Instructions):

? Option C (Action Triggers):

References:

? Salesforce AI Specialist Documentation - Custom Copilot Actions:

? Salesforce Help - Defining Output Types in Custom Actions:

? Salesforce Trailhead - Building Custom Copilot Actions:

**NEW QUESTION 8**

What is best practice when refining Einstein Copilot custom action instructions?

- A. Provide examples of user messages that are expected to trigger the action.
- B. Use consistent introductory phrases and verbs across multiple action instructions.
- C. Specify the persona who will request the action.

**Answer:** A

**Explanation:**

When refining Einstein Copilot custom action instructions, it is considered best practice to provide examples of user messages that are expected to trigger the action. This helps ensure that the custom action understands a variety of user inputs and can effectively respond to the intent behind the messages.

? Option B (consistent phrases) can improve clarity but does not directly refine the triggering logic.

? Option C (specifying a persona) is not as crucial as giving examples that illustrate how users will interact with the custom action.

For more details, refer to Salesforce's Einstein Copilot documentation on building and refining custom actions.

#### NEW QUESTION 9

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 10

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 10

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 11

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 12**

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 17**

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

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

**Answer: B**

**Explanation:**

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

? Select the appropriate prompt template type based on the use case.

? Develop the prompt within the prompt workspace, where the template is created and customized.

? Select CRM-derived grounding data to be dynamically inserted into the prompt, ensuring that the AI-generated responses are based on accurate and relevant data.

? Pick the model to use for generating responses, either using Salesforce's built-in models or custom ones.

? Test and validate the generated responses to ensure accuracy and effectiveness.

? Option B is correct as it follows the proper steps for using Prompt Builder.

? Option A and Option C do not capture the full process correctly.

References:

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

**NEW QUESTION 20**

Universal Containers (UC) wants to improve the efficiency of addressing customer questions and reduce agent handling time with AI-generated responses. The agents should be able to leverage their existing knowledge base and identify whether the responses are coming from the large language model (LLM) or from Salesforce Knowledge.

Which step should UC take to meet this requirement?

- A. Turn on Service AI Grounding, Grounding with Case, and Service Replies.
- B. Turn on Service Replies, Service AI Grounding, and Grounding with Knowledge.
- C. Turn on Service AI Grounding and Grounding with Knowledge.

**Answer: B**

**Explanation:**

To meet Universal Containers' goal of improving efficiency and reducing agent handling time with AI-generated responses, the best approach is to enable Service Replies, Service AI Grounding, and Grounding with Knowledge.

? Service Replies generates responses automatically.

? Service AI Grounding ensures that the AI is using relevant case data.

? Grounding with Knowledge ensures that responses are backed by Salesforce Knowledge articles, allowing agents to identify whether a response is coming from the LLM or Salesforce Knowledge.

? Option C does not include Service Replies, which is necessary for generating AI responses.

? Option A lacks the Grounding with Knowledge, which is essential for identifying response sources.

For more details, refer to Salesforce Service AI documentation on grounding and service replies.

#### NEW QUESTION 24

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 27

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

were deemed sensitive and have been masked within Einstein's Trust Layer. Which steps should the AI Specialist take upon modifying the masked fields?

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

**Answer: B**

#### Explanation:

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

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

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

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

References:

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

#### NEW QUESTION 28

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 29

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 34**

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 35**

How does the Einstein Trust Layer ensure that sensitive data is protected while generating useful and meaningful responses?

- A. Masked data will be de-masked during response journey.
- B. Masked data will be de-masked during request journey.
- C. Responses that do not meet the relevance threshold will be automatically rejected.

**Answer:** A

**Explanation:**

The Einstein Trust Layer ensures that sensitive data is protected while generating useful and meaningful responses by masking sensitive data before it is sent to the Large Language Model (LLM) and then de-masking it during the response journey.

How It Works:

? Data Masking in the Request Journey:

? Processing by the LLM:

? De-masking in the Response Journey:

Why Option A is Correct:

? De-masking During Response Journey: The de-masking process occurs after the LLM has generated its response, ensuring that sensitive data is only reintroduced into the output at the final stage, securely and appropriately.

? Balancing Security and Utility: This approach allows the system to generate useful and meaningful responses that include necessary sensitive information without compromising data security.

Why Options B and C are Incorrect:

? Option B (Masked data will be de-masked during request journey):

? Option C (Responses that do not meet the relevance threshold will be automatically rejected):

References:

? Salesforce AI Specialist Documentation - Einstein Trust Layer Overview:

? Salesforce Help - Data Masking and De-masking Process:

? Salesforce AI Specialist Exam Guide - Security and Compliance in AI:

Conclusion:

The Einstein Trust Layer ensures sensitive data is protected by masking it before sending any prompts to the LLM and then de-masking it during the response journey. This process allows Salesforce to generate useful and meaningful responses that include necessary sensitive information without exposing that data during the AI processing, thereby maintaining data security and compliance.

**NEW QUESTION 38**

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 41

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

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

**Answer:** A

#### Explanation:

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

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

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

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

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

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

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

References:

? Salesforce Documentation on Managing Copilot Actions

? Salesforce AI Specialist Guide on Copilot Action Management

#### NEW QUESTION 42

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 43

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 44

.....

## Relate Links

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

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

## Contact us

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

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