

Salesforce.Salesforce-AI-Specialist.v2025-07-25.q40

Exam Code:	Salesforce-AI-Specialist
Exam Name:	Salesforce Certified AI Specialist Exam
Certification Provider:	Salesforce
Free Question Number:	40
Version:	v2025-07-25
# of views:	108
# of Questions views:	400
https://www.freepdfdumps.com/Salesforce.Salesforce-AI-Specialist.v2025-07-25.q40.html	

NEW QUESTION: 1

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

Answer: A (LEAVE A REPLY)

NEW QUESTION: 2

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 (LEAVE A REPLY)

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: 3

Universal Containers is using Einstein Copilot for Sales to find similar opportunities to help close deals faster.

The team wants to understand the criteria used by the copilot to match opportunities.

What is one criteria that Einstein Copilot for Sales uses to match similar opportunities?

- A.** Matched opportunities are limited to the same account.
- B.** Matched opportunities were created in the last 12 months.
- C.** Matched opportunities have a status of Closed Won from last 12 months.

Answer: [\(SHOW ANSWER\)](#)

When Einstein Copilot for Sales matches similar opportunities, one of the primary criteria used is whether the opportunities have a status of Closed Won within the last 12 months. This is a key factor in identifying successful patterns that could help close current deals. By focusing on opportunities that have been recently successful, Einstein Copilot can provide relevant insights and suggestions to sales reps to help them close similar deals faster. For more information, review Salesforce Einstein Copilot documentation related to opportunity matching and sales success patterns.

NEW QUESTION: 4

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 \(LEAVE A REPLY\)](#)

To access externally-hosted models, such as a large language model (LLM) hosted on AWS, the Model Builder in Salesforce is the appropriate tool. Model Builder allows teams to integrate and deploy external AI models into the Salesforce platform, making it possible to leverage models hosted outside of Salesforce infrastructure while still benefiting from the platform's native AI capabilities.

* Option B, App Builder, is primarily used to build and configure applications in Salesforce, not to integrate AI models.

* Option C, Copilot Builder, focuses on building assistant-like tools rather than integrating external AI models.

Model Builder enables seamless integration with external systems and models, allowing Salesforce users to use external LLMs for generating AI-driven insights and automation.

Salesforce AI Specialist References: For more details, check the Model Builder guide here: https://help.salesforce.com/s/articleView?id=sf.model_builder_external_models.htm

NEW QUESTION: 5

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 (LEAVE A REPLY)

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

NEW QUESTION: 6

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: (SHOW ANSWER)

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: 7

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 (LEAVE A REPLY)

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

[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: 8

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: (SHOW ANSWER)

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: 9

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 (LEAVE A REPLY)

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: 10

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: (SHOW ANSWER)

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: 11

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 (LEAVE A REPLY)

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: 12

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 (LEAVE A REPLY)

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

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

NEW QUESTION: 13

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 (LEAVE A REPLY)

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.htm

NEW QUESTION: 14

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 (LEAVE A REPLY)

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>.

NEW QUESTION: 15

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: (SHOW ANSWER)

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:**
- * **Limitation:** The User related list is not supported in prompt templates for grounding purposes.
- * **Reason:** Salesforce restricts certain objects and related lists from being used in prompt templates to maintain data security and integrity. The User object often contains sensitive information and is subject to strict access controls.
- * **Impact:** Attempting to use the User related list in a prompt template will not work as expected because the system does not support it.

Why Options A and B are Incorrect:

- * **Option A (The User related list should have View All access):**
- * **Incorrect:** Even with View All access, the User related list is still not supported in prompt templates.
- * **Security Concerns:** Granting View All access to the User object is a significant security risk and not a recommended practice.
- * **Option B (The User related list needs to be included on the record page):**
- * **Incorrect:** Including the User related list on the record page does not affect its availability in prompt templates.
- * **Irrelevance:** The placement of the related list on the record page does not change the system's ability to access it in a prompt template.

References:

- * **Salesforce AI Specialist Documentation - Prompt Templates Limitations:**
- * **Details the objects and related lists that are not supported in prompt templates.**
- * **Salesforce Help - Data Access in Prompt Templates:**
- * **Explains how data access and security considerations affect the use of objects in prompt templates.**
- * **Salesforce Trailhead - Understanding Prompt Template Grounding:**
- * **Provides insights into grounding prompt templates and the limitations involved.**

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: 16

Universal Containers (UC) wants to offer personalized service experiences and reduce agent handling time with AI-generated email responses, grounded in Knowledge base. Which AI capability should UC use?

- A.** Einstein Email Replies
- B.** Einstein Service Replies for Email
- C.** Einstein Generative Service Replies for Email

Answer: ([SHOW ANSWER](#))

For Universal Containers (UC) to offer personalized service experiences and reduce agent handling time using AI-generated responses grounded in the Knowledge base, the best solution is Einstein Service Replies for Email. This capability leverages AI to automatically generate responses to service-related emails based on historical data and the Knowledge base, ensuring accuracy and relevance while saving time for service agents.

* Einstein Email Replies (option A) is more suited for sales use cases.

* Einstein Generative Service Replies for Email (option C) could be a future offering, but as of now, Einstein Service Replies for Email is the correct choice for grounded, knowledge-based responses.

References:

* Einstein Service Replies Overview: https://help.salesforce.com/s/articleView?id=sf.einstein_service_replies.htm

Valid Salesforce-AI-Specialist Dumps shared by Actual4test.com for Helping Passing Salesforce-AI-Specialist Exam! Actual4test.com now offer the **newest Salesforce-AI-Specialist exam dumps**, the Actual4test.com Salesforce-AI-Specialist exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com Salesforce-AI-Specialist dumps with Test Engine here: https://www.actual4test.com/Salesforce-AI-Specialist_examcollection.html (157 Q&As Dumps, **30%OFF Special Discount: Freepdfdumps**)

NEW QUESTION: 17

Amid their busy schedules, sales reps at Universal Containers dedicate time to follow up with prospects and existing clients via email regarding renewals or new deals. They spend many hours throughout the week reviewing past communications and details about their customers before performing their outreach.

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

- A. Einstein Copilot Action: Find Similar Opportunities
- B. Einstein Copilot Action: Draft or Revise Sales Email
- C. Einstein Copilot Action: Summarize Record

Answer: B (LEAVE A REPLY)

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

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

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

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

NEW QUESTION: 18

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

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

Answer: C (LEAVE A REPLY)

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

NEW QUESTION: 19

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 (LEAVE A REPLY)

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_cases.htm

NEW QUESTION: 20

Universal Containers (UC) wants to enable its sales team to use AI to suggest recommended products from its catalog.

Which type of prompt template should UC use?

- A. Record summary prompt template
- B. Email generation prompt template
- C. Flex prompt template

Answer: C (LEAVE A REPLY)

Universal Containers (UC) wants to enable its sales team to leverage AI to recommend products from its catalog. The best option for this use case is a Flex prompt template.

A Flex prompt template is designed to provide flexible, customizable AI-driven recommendations or responses based on specific data points, such as product information, customer needs, or sales history. This template type allows the AI to consider various inputs and parameters, making it ideal for generating product recommendations dynamically.

In contrast:

* A Record summary prompt template (Option A) is used to summarize data related to a specific record, such as generating a quick summary of a sales opportunity or account, but not for recommending products.

* An Email generation prompt template (Option B) is tailored for crafting email content and is not suitable for suggesting products based on a catalog.

Given the need for dynamic recommendations that pull from a product catalog and potentially other sales data, the Flex prompt template is the correct approach.

Salesforce References:

* Salesforce Prompt Templates Overview: <https://help.salesforce.com/s/articleView?id=000391407&type=1>

* Flex Prompt Template Usage: https://developer.salesforce.com/docs/atlas.en-us.salesforce_ai.meta/salesforce_ai/prompt_flex_template

NEW QUESTION: 21

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 (LEAVE A REPLY)

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: 22

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: (SHOW ANSWER)

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: 23

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 (LEAVE A REPLY)

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: 24

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 (LEAVE A REPLY)

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: 25

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: (SHOW ANSWER)

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

NEW QUESTION: 26

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: (SHOW ANSWER)

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

NEW QUESTION: 27

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 (LEAVE A REPLY)

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: 28

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 (LEAVE A REPLY)

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: 29

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 (LEAVE A REPLY)

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

NEW QUESTION: 30

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 (LEAVE A REPLY)

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: 31

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 (LEAVE A REPLY)

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:

- * B refers 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

Valid Salesforce-AI-Specialist Dumps shared by Actual4test.com for Helping Passing Salesforce-AI-Specialist Exam! Actual4test.com now offer the **newest Salesforce-AI-Specialist exam dumps**, the Actual4test.com Salesforce-AI-Specialist exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com Salesforce-AI-Specialist dumps with Test Engine here:
https://www.actual4test.com/Salesforce-AI-Specialist_examcollection.html (157 Q&As Dumps, **30%OFF Special Discount: Freepdfdumps**)

NEW QUESTION: 32

An AI Specialist implements Einstein Sales Emails for a sales team. The team wants to send personalized follow-up emails to leads based on their interactions and data stored in Salesforce. The AI Specialist needs to configure the system to use the most accurate and up-to-date information for email generation.

Which grounding technique should the AI Specialist use?

- A. Ground with Apex Merge Fields
- B. Ground with Record Merge Fields
- C. Automatic grounding using Draft with Einstein feature

Answer: B (LEAVE A REPLY)

For Einstein Sales Emails to generate personalized follow-up emails, it is crucial to ground the email content with the most up-to-date and accurate information. Grounding refers to connecting the AI model with real-time data. The most appropriate technique in this case is Ground with Record Merge Fields. This method ensures that the content in the emails pulls dynamic and accurate data directly from Salesforce records, such as lead or contact information, ensuring the follow-up is relevant and customized based on the specific record.

* Record Merge Fields ensure the generated emails are highly personalized using data like lead name, company, or other Salesforce fields directly from the records.

* Apex Merge Fields are typically more suited for advanced, custom logic-driven scenarios but are not the most straightforward for this use case.

* Automatic grounding using Draft with Einstein is a different feature where Einstein automatically drafts the email, but it does not specifically ground the content with record-specific data like Record Merge Fields.

References:

* Salesforce Einstein Sales Emails Documentation:

https://help.salesforce.com/s/articleView?id=release-notes.rn_einstein_sales_emails.htm

NEW QUESTION: 33

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 (LEAVE A REPLY)

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

einstein_studio_overview

NEW QUESTION: 34

Universal Containers implemented Einstein Copilot for its users.

One user complains that Einstein Copilot is not deleting activities from the past 7 days.

What is the reason for this issue?

A. Einstein Copilot Delete Record Action permission is not associated to the user.

B. Einstein Copilot does not have the permission to delete the user's records.

C. Einstein Copilot does not support the Delete Record action.

Answer: C (LEAVE A REPLY)

Einstein Copilot currently supports various actions like creating and updating records but does not support the Delete Record action. Therefore, the user's request to delete activities from the past 7 days cannot be fulfilled using Einstein Copilot.

* **Unsupported Action:** The inability to delete records is due to the current limitations of Einstein Copilot's supported actions. It is designed to assist with tasks like data retrieval, creation, and updates, but for security and data integrity reasons, it does not facilitate the deletion of records.

* **User Permissions:** Even if the user has the necessary permissions to delete records within Salesforce, Einstein Copilot itself does not have the capability to execute delete operations.

References:

* **Salesforce AI Specialist Documentation - Einstein Copilot Supported Actions:**

* Lists the actions that Einstein Copilot can perform, noting the absence of delete operations.

* **Salesforce Help - Limitations of Einstein Copilot:**

* Highlights current limitations, including unsupported actions like deleting records.

NEW QUESTION: 35

Universal Containers (UC) wants to enable its sales team with automatic post-call visibility into mention of competitors, products, and other custom phrases.

Which feature should the AI Specialist set up to enable UC's sales team?

A. Call Summaries

B. Call Explorer

C. Call Insights

Answer: (SHOW ANSWER)

To enable Universal Containers' sales team with automatic post-call visibility into mentions of competitors, products, and custom phrases, the AI Specialist should set up Call Insights. Call Insights analyzes voice and video calls for key phrases, topics, and mentions, providing insights into critical aspects of the conversation. This feature automatically

surfaces key details such as competitor mentions, product discussions, and custom phrases specified by the sales team.

* Call Summaries provide a general overview of the call but do not specifically highlight keywords or topics.

* Call Explorer is a tool for navigating through call data but does not focus on automatic insights.

For more information, refer to Salesforce's Call Insights documentation regarding the analysis of call content and extracting actionable information.

NEW QUESTION: 36

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: (SHOW ANSWER)

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:

* Instructions:

* Defines what the action should accomplish.

* Provides guidance to the AI assistant on how to execute the action.

* Input Parameters:

* The data inputs required for the action to perform its task.

* Specifies the parameters that users or systems need to provide.

* Output Types:

* Definition: Specifies the type of data the action will output after execution.

* Importance: Ensures that the Copilot and other components understand the format and structure of the action's results.

* Examples: Text, JSON, records, or other data structures.

* Action Reference:

* Points to the underlying implementation, such as an Apex class or Flow.

* Action Triggers:

* Conditions or events that initiate the action.

* Focus on Output Types:

* Relevance in Configuration:

- * The development team must define the Output Types correctly to ensure that the action's results are properly processed and displayed by Copilot.

- * Incorrect or undefined Output Types can lead to misinterpretation of data and failures in functionality.

Why Output Types are a Core Component:

- * Integration with Copilot:

- * Copilot relies on the Output Types to understand how to handle the data returned by the action.

- * Data Consistency:

- * Defines the structure and format of the output, ensuring consistent communication between the action and Copilot.

- * User Experience:

- * Proper Output Types ensure that users receive the expected results in an understandable format.

Why Other Options are Less Suitable:

- * Option A (Instructions):

- * While Instructions are a core component, the question asks for what should be reviewed in the configuration to identify one of the core components.

- * In this context, reviewing Output Types is more critical to ensuring proper configuration and functionality.

- * Option C (Action Triggers):

- * Action Triggers are important but are not always considered a core component within the configuration of a Custom Copilot Action.

- * Triggers often relate to when an action is initiated rather than the configuration of the action itself.

References:

- * Salesforce AI Specialist Documentation - Custom Copilot Actions:

- * Details the components and configuration of Custom Copilot Actions.

- * Salesforce Help - Defining Output Types in Custom Actions:

- * Explains the importance of Output Types and how to configure them.

- * Salesforce Trailhead - Building Custom Copilot Actions:

- * Provides a hands-on approach to creating and configuring Custom Copilot Actions, highlighting key components.

NEW QUESTION: 37

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 (LEAVE A REPLY)

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: 38

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 (LEAVE A REPLY)

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: 39

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 (LEAVE A REPLY)

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: 40

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: (SHOW ANSWER)

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.

Valid Salesforce-AI-Specialist Dumps shared by Actual4test.com for Helping Passing Salesforce-AI-Specialist Exam! Actual4test.com now offer the **newest Salesforce-AI-Specialist exam dumps**, the Actual4test.com Salesforce-AI-Specialist exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com Salesforce-AI-Specialist dumps with Test Engine here:
https://www.actual4test.com/Salesforce-AI-Specialist_examcollection.html (157 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)