

Salesforce

Exam Questions Agentforce-Specialist

Salesforce Certified Agentforce Specialist



NEW QUESTION 1

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:

The Einstein Trust Layer is designed to ensure responsible and compliant AI usage. Data Masking (B) is the mechanism that directly addresses compliance with data protection regulations like GDPR by obscuring or anonymizing sensitive personal data (e.g., names, emails, phone numbers) before it is processed by AI models. This prevents unauthorized exposure of personally identifiable information (PII) and ensures adherence to privacy laws.

Salesforce documentation explicitly states that Data Masking is a core component of the Einstein Trust Layer, enabling organizations to meet GDPR requirements by automatically redacting sensitive fields during AI interactions. For example, masked data ensures that PII is not stored or used in AI model training or inference without explicit consent.

In contrast:

? Toxicity Scoring (A) identifies harmful or inappropriate content in outputs but does not address data privacy.

? Prompt Defense (C) guards against malicious prompts or injection attacks but focuses on security rather than data protection compliance.

Reference:

Salesforce Help Article: Einstein Trust Layer ("Data Masking" section).

Einstein Trust Layer Overview: "Data Protection and Compliance Features" (GDPR alignment via Data Masking).

NEW QUESTION 2

An administrator is responsible for ensuring the security and reliability of Universal Containers' (UC) CRM data. UC needs enhanced data protection and up-to-date AI capabilities. UC also needs to include relevant information from a Salesforce record to be merged with the prompt. Which feature in the Einstein Trust Layer best supports UC's need?

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

Answer: B

Explanation:

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

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

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

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

References:

? Salesforce Developer Documentation on Einstein Trust Layer

? Salesforce Security Documentation for AI and Data Privacy

NEW QUESTION 3

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: C

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 4

Universal Containers Is Interested In Improving the sales operation efficiency by analyzing their data using AI-powered predictions in Einstein Studio. Which use case works for this scenario?

- A. Predict customer sentiment toward a promotion message.
- B. Predict customer lifetime value of an account.
- C. Predict most popular products from new product catalog.

Answer: B

Explanation:

For improving sales operations efficiency, Einstein Studio is ideal for creating AI-powered models that can predict outcomes based on data. One of the most valuable use cases is predicting customer lifetime value, which helps sales teams focus on high-value accounts and make more informed decisions. Customer lifetime value (CLV) predictions can optimize strategies around customer retention, cross-selling, and long-term engagement.

? Option B is the correct choice as predicting customer lifetime value is a well-established use case for AI in sales.

? Option A (customer sentiment) is typically handled through NLP models, while Option C (product popularity) is more of a marketing analysis use case.

References:

Salesforce Einstein Studio Use Case Overview: https://help.salesforce.com/s/articleView?id=sf.einstein_studio_overview

NEW QUESTION 5

An administrator wants to check the response of the Flex prompt template they've built, but the preview button is greyed out. What is the reason for this?

- A. The records related to the prompt have not been selected.
- B. The prompt has not been saved and activated,
- C. A merge field has not been inserted in the prompt.

Answer: A

Explanation:

When the preview button is greyed out in a Flex prompt template, it is often because the records related to the prompt have not been selected. Flex prompt templates pull data dynamically from Salesforce records, and if there are no records specified for the prompt, it can't be previewed since there is no content to generate based on the template.

? Option B, not saving or activating the prompt, would not necessarily cause the preview button to be greyed out, but it could prevent proper functionality.

? Option C, missing a merge field, would cause issues with the output but would not directly grey out the preview button.

Ensuring that the related records are correctly linked is crucial for testing and previewing how the prompt will function in real use cases.

Salesforce Agentforce Specialist References: Refer to the documentation on troubleshooting Flex templates here:

https://help.salesforce.com/s/articleView?id=sf.flex_prompt_builder_troubleshoot.htm

NEW QUESTION 6

Universal Container's internal auditing team asks An Agentforce to verify that address information is properly masked in the prompt being generated. How should the Agentforce Specialist verify the privacy of the masked data in the Einstein Trust Layer?

- A. Enable data encryption on the address field
- B. Review the platform event logs
- C. Inspect the AI audit trail

Answer: C

Explanation:

The AI audit trail in Salesforce provides a detailed log of AI activities, including the data used, its handling, and masking procedures applied in the Einstein Trust Layer. It allows the Agentforce Specialist to inspect and verify that sensitive data, such as addresses, is appropriately masked before being used in prompts or outputs.

? Enable data encryption on the address field: While encryption ensures data security at rest or in transit, it does not verify masking in AI operations.

? Review the platform event logs: Platform event logs capture system events but do not specifically focus on the handling or masking of sensitive data in AI processes.

? Inspect the AI audit trail: This is the most relevant option, as it provides visibility into how data is processed and masked in AI activities.

Reference:

"How Salesforce Ensures Trust in AI with Einstein Trust Layer | Salesforce" .

NEW QUESTION 7

The marketing team at Universal Containers is looking for a way to personalize emails based on customer behavior, preferences, and purchase history. Why should the team use Agent as the solution?

- A. To generate relevant content when engaging with each customer
- B. To analyze past campaign performance
- C. To send automated emails to all customers

Answer: A

Explanation:

Agent is designed to assist in generating personalized, AI-driven content based on customer data such as behavior, preferences, and purchase history. For the marketing team at Universal Containers, this is the perfect solution to create dynamic and relevant email content. By leveraging Agent, they can ensure that each customer receives tailored communications, improving engagement and conversion rates.

? Option A is correct as Agent helps generate real-time, personalized content based on comprehensive data about the customer.

? Option B refers more to Einstein Analytics or

? Marketing Cloud Intelligence, and Option C deals with automation, which isn't the primary focus of Agent.

References:

? Salesforce Agent Overview: https://help.salesforce.com/s/articleView?id=einstein_copilot_overview.htm

NEW QUESTION 8

Universal Containers implemented Agentforce for its users. One user complains that an Agent is not deleting activities from the past 7 days. What is the reason for this issue?

- A. Agentforce does not have the permission to delete the user's records.
- B. Agentforce Delete Record Action permission is not associated to the user.
- C. Agentforce does not have a standard Delete Record action.

Answer: C

Explanation:

- ? Context of the Question Universal Containers (UC) uses Agentforce, a specialized AI-driven assistant for Salesforce. A user reports that an Agent is unable to delete recent activities.
- ? Why Agentforce Cannot Delete Records
- ? Why Other Options Are Incorrect
- ? Conclusion The core reason for the issue is that Agentforce does not support a standard Delete Record action (Choice C).
- Salesforce Agentforce Specialist References & Documents
- ? Salesforce Official Documentation – Agentforce (Note: Agentforce may be a pilot or specialized feature; check pilot release notes or official docs for standard actions.)
- ? Salesforce Agentforce Specialist Study Guide Covers the limitations of certain AI-enabled features regarding record operations.

NEW QUESTION 9

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 is Toxicity Detected is true.
- B. Create a flow that sends an email to a specified address each time the toxicity score from the response exceeds a predefined threshold.
- C. Create a Trust Layer audit report within Data Cloud that uses a toxicity detector type filter to display toxic responses and their respective scores.

Answer: C

Explanation:

- To ensure that content generated by a large language model (LLM) is appropriately screened for toxic language, the Agentforce Specialist should create a Trust Layer audit report within Data Cloud. By using the toxicity detector type filter, the report can display toxic responses along with their respective toxicity scores, allowing Universal Containers to monitor and manage any toxic content generated with a high level of confidence.
- ? Option C is correct because it enables visibility into toxic language detection within the Trust Layer and allows for auditing responses for toxicity.
- ? Option A suggests checking a toxicity detection log, but Salesforce provides more comprehensive options via the audit report.
- ? Option B involves creating a flow, which is unnecessary for toxicity detection monitoring.
- References:
- ? Salesforce Trust Layer Documentation: https://help.salesforce.com/s/articleView?id=sf.einstein_trust_layer_audit.htm

NEW QUESTION 10

An Agentforce is creating a custom action in Agent. Which option is available for the Agentforce Specialist to choose for the custom copilot action?

- A. Apex trigger
- B. SOQL
- C. Flows

Answer: C

Explanation:

- When creating a custom action in Agent, one of the available options is to use Flows. Flows are a powerful automation tool in Salesforce, allowing the Agentforce Specialist to define custom logic and actions within the Copilot system. This makes it easy to extend Copilot's functionality without needing custom code. While Apex triggers and SOQL are important Salesforce tools, Flows are the recommended method for creating custom actions within Agent because they are declarative and highly adaptable.
- For further guidance, refer to Salesforce Flow documentation and Agent customization resources.

NEW QUESTION 10

Universal Containers would like to route a service agent conversation to a human agent queue. Which tool connects the service agent to the human agent queue for escalation?

- A. Outbound Omni-Channel Flow
- B. Screen Flow
- C. Prompt Flow

Answer: A

Explanation:

- Why is Outbound Omni-Channel Flow the Correct Answer?
- In Agentforce, when a service agent's conversation needs to be escalated to a human agent queue, Outbound Omni-Channel Flow is the appropriate tool to facilitate this process.
- Key Features of Outbound Omni-Channel Flow in Agentforce:
- ? Automates Escalation to a Human Agent Queue
- ? Seamless Transition from AI to Human Agents
- ? Ensures Proper Prioritization & Load Balancing
- ? Integration with Agentforce and Service Cloud
- Why Not the Other Options?
- * B. Screen Flow
- ? Screen Flow is used for interactive guided processes where users manually enter data in predefined steps.
- ? It does not support automated case routing to human agents in real time.
- * C. Prompt Flow
- ? Prompt Flow is designed to enhance AI-generated responses and workflows rather than routing service agent interactions to human agents.

? It lacks Omni-Channel integration, which is necessary for real-time queue management.

Agentforce Specialist References

The importance of using Omni-Channel Flow for routing AI-generated interactions to human agents is supported in the Agentforce Specialist exam objectives and documentation:

? Salesforce AI Specialist Material: Covers the importance of Omni-Channel routing for managing AI and human agent interactions.

? Salesforce Instructions for the Certification: Mentions routing AI-driven cases to human agents using automated flows.

? Agentforce Tools Documentation: Highlights Omni-Channel capabilities in Service AI.

NEW QUESTION 12

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

Explanation:

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 16

Universal Containers (UC) is rolling out an AI-powered support assistant to help customer service agents quickly retrieve relevant troubleshooting steps and policy guidelines. The assistant relies on a search index in Data Cloud that contains product manuals, policy documents, and past case resolutions. During testing, UC notices that agents are receiving too many irrelevant results from older product versions that no longer apply. How should UC address this issue?

- A. Modify the search index to only store documents from the last year and remove older records.
- B. Create a custom retriever in Einstein Studio, and apply filters for publication date and product line.
- C. Use the default retriever, as it already searches the entire search index and provides broad coverage.

Answer: C

Explanation:

Comprehensive and Detailed In-Depth Explanation: UC's support assistant uses a Data Cloud search index for grounding, but irrelevant results from outdated product versions are an issue. Let's evaluate the options.

? Option A: Modify the search index to only store documents from the last year and remove older records. While limiting the index to recent documents could reduce irrelevant results, this requires ongoing maintenance (e.g., purging older data) and risks losing valuable historical context from past resolutions. It's a blunt approach that doesn't leverage Data Cloud's filtering capabilities, making it less optimal and incorrect.

? Option B: Create a custom retriever in Einstein Studio, and apply filters for publication date and product line. There's no "Einstein Studio" in Salesforce—possibly a typo for Agentforce Studio or Data Cloud. Custom retrievers can be created in Data Cloud, but this requires advanced configuration (e.g., custom code or Data Cloud APIs) beyond standard Agentforce setup. This is overcomplicated compared to native options, making it incorrect.

? Option C: Use the default retriever, as it already searches the entire search index and provides broad coverage. This option seems misaligned at first glance, as the default retriever's broad coverage is causing the issue. However, the intent (based on typical Salesforce question patterns) likely implies using the default retriever with additional configuration. In Data Cloud, the default retriever searches the index, but you can apply filters (e.g., publication date, relevance) via the Data

Library or prompt grounding settings to prioritize current documents. Since the question lacks an explicit filtering option, this is interpreted as the closest correct choice with refinement assumed, making it the answer by elimination and context.

Why Option C is Correct (with Caveat): The default retriever, when paired with filters (assumed intent), allows UC to refine results without custom development.

Salesforce documentation emphasizes refining retriever scope over rebuilding indexes, though the question's phrasing is suboptimal. Option C is selected as the least incorrect, assuming filter application.

References:

? Salesforce Data Cloud Documentation: Search Indexes > Retrievers – Notes filter options for relevance.

? Trailhead: Data Cloud for Agentforce – Covers refining search results.

? Salesforce Help: Grounding with Data Cloud – Suggests default retriever with customization.

NEW QUESTION 20

Universal Containers' data science team is hosting a generative large language model (LLM) on Amazon Web Services (AWS). What should the team use to access externally-hosted models in the Salesforce Platform?

- A. Model Builder
- B. App Builder
- C. Copilot Builder

Answer: A

Explanation:

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

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

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

Model Builder enables seamless integration with external systems and models, allowing Salesforce users to use external LLMs for generating AI-driven insights and automation. Salesforce Agentforce 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 21

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 Agentforce 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 24

Universal Containers (UC) is experimenting with using public Generative AI models and is familiar with the language required to get the information it needs. However, it can be time-consuming for both UC's sales and service reps to type in the prompt to get the information they need, and ensure prompt consistency. Which Salesforce feature should the company use to address these concerns?

- A. Agent Builder and Action: Query Records.
- B. Einstein Prompt Builder and Prompt Templates.
- C. Einstein Recommendation Builder.

Answer: B

Explanation:

Comprehensive and Detailed In-Depth Explanation:UC wants to streamline the use of Generative AI by reducing the time reps spend typing prompts and ensuring consistency, leveraging their existing prompt knowledge. Let's evaluate the options.

? Option A: Agent Builder and Action: Query Records.Agent Builder in Agentforce Studio creates autonomous AI agents with actions like "Query Records" to fetch data. While this could retrieve information, it's designed for agent-driven workflows, not for simplifying manual prompt entry or ensuring consistency across user inputs. This doesn't directly address UC's concerns and is incorrect.

? Option B: Einstein Prompt Builder and Prompt Templates.Einstein Prompt Builder, part of Agentforce Studio, allows users to create reusable prompt templates that encapsulate specific instructions and grounding for Generative AI (e.g., using public models via the Atlas Reasoning Engine). UC can predefine prompts based on their known language, saving time for reps by eliminating repetitive typing and ensuring consistency across sales and service teams. Templates can be embedded in flows, Lightning pages, or agent interactions, perfectly addressing UC's needs. This is the correct answer.

? Option C: Einstein Recommendation Builder.Einstein Recommendation Builder generates personalized recommendations (e.g., products, next best actions) using predictive AI, not Generative AI for freeform prompts. It doesn't support custom prompt creation or address time/consistency issues for reps, making it incorrect.

Why Option B is Correct:Einstein Prompt Builder's prompt templates directly tackle UC's challenges by standardizing prompts and reducing manual effort, leveraging their familiarity with Generative AI language. This is a core feature for such use cases, as per Salesforce documentation.

References:

? Salesforce Agentforce Documentation: Einstein Prompt Builder – Details prompt templates for consistency and efficiency.

? Trailhead: Build Prompt Templates in Agentforce – Explains time-saving benefits of templates.

? Salesforce Help: Generative AI with Prompt Builder – Confirms use for streamlining rep interactions.

NEW QUESTION 27

Universal Containers (UC) plans to send one of three different emails to its customers based on the customer's lifetime value score and their market segment. Considering that UC are required to explain why an e-mail was selected, which AI model should UC use to achieve this?

- A. Predictive model and generative model
- B. Generative model
- C. Predictive model

Answer: C

Explanation:

Universal Containers should use a Predictive model to decide which of the three emails to send based on the customer's lifetime value score and market segment. Predictive models analyze data to forecast outcomes, and in this case, it would predict the most appropriate email to send based on customer attributes.

Additionally, predictive models can provide explainability to show why a certain email was chosen, which is crucial for UC's requirement to explain the decision-making process.

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

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

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

NEW QUESTION 31

An Agentforce is considering using a Field Generation prompt template type.

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

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

Answer: B

Explanation:

Before creating a Field Generation prompt template, the Agentforce 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 32

Universal Containers wants to incorporate the current order fulfillment status into a prompt for a large language model (LLM). The order status is stored in the external enterprise resource planning (ERP) system.

Which data grounding technique should the Agentforce Specialist recommend?

- A. External Object Record Merge Fields
- B. External Services Merge Fields
- C. Apex Merge Fields

Answer: A

Explanation:

? Context of the Requirement: Universal Containers wants to pull in real-time order status data from an external ERP system into an LLM prompt.

? Data Grounding in LLM Prompts: Data grounding ensures the Large Language

Model has access to the most current and relevant information. In Salesforce, one recommended approach is to use External Objects (via Salesforce Connect) when data resides outside of Salesforce.

? Why External Object Record Merge Fields:

? Why Not External Services Merge Fields or Apex Merge Fields:

? References and Study Resources:

NEW QUESTION 37

Universal Containers (UC) is implementing Einstein Generative AI to improve customer insights and interactions. UC needs audit and feedback data to be accessible for reporting purposes. What is a consideration for this requirement?

- A. Storing this data requires Data Cloud to be provisioned.
- B. Storing this data requires a custom object for data to be configured.
- C. Storing this data requires Salesforce big objects.

Answer: A

Explanation:

When implementing Einstein Generative AI for improved customer insights and interactions, the Data Cloud is a key consideration for storing and managing large-scale audit and feedback data. The Salesforce Data Cloud (formerly known as Customer 360 Audiences) is designed to handle and unify massive datasets from various sources, making it ideal for storing data required for AI-powered insights and reporting. By provisioning Data Cloud, organizations like Universal Containers (UC) can gain real-time access to customer data, making it a central repository for unified reporting across various systems.

? Audit and feedback data generated by Einstein Generative AI needs to be stored

in a scalable and accessible environment, and the Data Cloud provides this capability, ensuring that data can be easily accessed for reporting, analytics, and further model improvement.

? Custom objects or Salesforce Big Objects are not designed for the scale or the

specific type of real-time, unified data processing required in such AI-driven interactions. Big Objects are more suited for archival data, whereas Data Cloud ensures more robust processing, segmentation, and analysis capabilities.

References:

? Salesforce Data Cloud Documentation: <https://www.salesforce.com/products/data-cloud/overview/>

? Salesforce Einstein AI Overview:

<https://www.salesforce.com/products/einstein/overview/>

NEW QUESTION 42

Universal Containers needs a tool that can analyze voice and video call records to provide insights on competitor mentions, coaching opportunities, and other key information. The goal is to enhance the team's performance by identifying areas for improvement and competitive intelligence.

Which feature provides insights about competitor mentions and coaching opportunities?

- A. Call Summaries
- B. Einstein Sales Insights
- C. Call Explorer

Answer: C

Explanation:

For analyzing voice and video call records to gain insights into competitor mentions, coaching opportunities, and other key information, Call Explorer is the most suitable feature. Call Explorer, a part of Einstein Conversation Insights, enables sales teams to analyze calls, detect patterns, and identify areas where improvements can be made. It uses natural language processing (NLP) to extract insights, including competitor mentions and moments for coaching. These

insights are vital for improving sales performance by providing a clear understanding of the interactions during calls.

? Call Summaries offer a quick overview of a call but do not delve deep into competitor mentions or coaching insights.

? Einstein Sales Insights focuses more on pipeline and forecasting insights rather than call-based analysis.

References:

? Salesforce Einstein Conversation Insights Documentation: https://help.salesforce.com/s/articleView?id=einstein_conversation_insights.htm

NEW QUESTION 44

How does an Agent respond when it can't understand the request or find any requested information?

- A. With a preconfigured message, based on the action type.
- B. With a general message asking the user to rephrase the request.
- C. With a generated error message.

Answer: B

Explanation:

Comprehensive and Detailed In-Depth Explanation: Agentforce Agents are designed to handle situations where they cannot interpret a request or retrieve requested data gracefully. Let's assess the options based on Agentforce behavior.

? Option A: With a preconfigured message, based on the action type. While Agentforce allows customization of responses, there's no specific mechanism tying preconfigured messages to action types for unhandled requests. Fallback responses are more general, not action-specific, making this incorrect.

? Option B: With a general message asking the user to rephrase the request. When an Agentforce Agent fails to understand a request or find information, it defaults to a general fallback response, typically asking the user to rephrase or clarify their input (e.g., "I didn't quite get that—could you try asking again?"). This is configurable in Agent Builder but defaults to a user-friendly prompt to encourage retry, aligning with Salesforce's focus on conversational UX. This is the correct answer per documentation.

? Option C: With a generated error message. Agentforce Agents prioritize user experience over technical error messages. While errors might log internally (e.g., in Event Logs), the user-facing response avoids jargon and focuses on retry prompts, making this incorrect.

Why Option B is Correct: The default behavior of asking users to rephrase aligns with Agentforce's conversational design principles, ensuring a helpful response when comprehension fails, as noted in official resources.

References:

? Salesforce Agentforce Documentation: Agent Builder > Fallback Responses – Describes general retry messages.

? Trailhead: Build Agents with Agentforce – Covers handling ununderstood requests.

? Salesforce Help: Agentforce Interaction Design – Confirms user-friendly fallback behavior.

NEW QUESTION 49

Leadership needs to populate a dynamic form field with a summary or description created by a large language model (LLM) to facilitate more productive conversations with customers. Leadership also wants to keep a human in the loop to be considered in their AI strategy. Which prompt template type should the Agentforce Specialist recommend?

- A. Field Generation
- B. Sales Email
- C. Record Summary

Answer: A

Explanation:

Why is "Field Generation" the correct answer?

In Agentforce, the Field Generation prompt template type is designed to populate dynamic form fields with AI-generated content, such as summaries or descriptions created by a large language model (LLM).

Key Considerations for Using Field Generation in Dynamic Forms:

? AI-Powered Summarization in Form Fields

? Human-in-the-Loop AI Strategy

? Works with Salesforce Dynamic Forms

Why Not the Other Options?

* B. Sales Email

? Incorrect because Sales Email templates are designed for AI-generated email content, not for populating form fields.

* C. Record Summary

? Incorrect because Record Summary templates generate high-level summaries of entire records, but do not populate individual form fields dynamically.

Agentforce Specialist References

? Salesforce AI Specialist Material confirms that Field Generation templates are used for AI-powered dynamic form population.

NEW QUESTION 53

A sales manager needs to contact leads at scale with hyper-relevant solutions and customized communications in the most efficient manner possible. Which Salesforce solution best suits this need?

- A. Einstein Sales Assistant
- B. Prompt Builder
- C. Einstein Lead follow-up

Answer: B

Explanation:

Step 1: Define the Requirements

The question specifies a sales manager's need to:

? Contact leads at scale: Handle a large volume of leads simultaneously.

? Hyper-relevant solutions: Deliver tailored solutions based on lead-specific data (e.g., CRM data, behavior).

? Customized communications: Personalize outreach (e.g., emails, messages) for each lead.

? Most efficient manner possible: Minimize manual effort and maximize automation. This suggests a solution that leverages AI for personalization and automation for scale, ideally within the Salesforce ecosystem.

Step 2: Evaluate the Provided Options

*** A. Einstein Sales Assistant**

? Description: Einstein Sales Assistant is not a distinct, standalone product in Salesforce documentation as of March 2025 but is often associated with features in Sales Cloud Einstein or Einstein Copilot for Sales. It typically acts as an AI-powered assistant embedded in the sales workflow, offering suggestions (e.g., next best actions), drafting emails, or summarizing calls.

? Analysis Against Requirements:

? Conclusion: Einstein Sales Assistant is a productivity tool for reps, not a solution for autonomous, large-scale lead contact. It's not the best fit.

*** B. Prompt Builder**

? Description: Prompt Builder is a low-code tool within the Einstein 1 Platform that allows users to create reusable AI prompts for generating personalized content (e.g., emails, summaries) based on Salesforce CRM data. It integrates with generative AI models and can be embedded in workflows (e.g., via Flow) to automate content creation.

? Analysis Against Requirements:

: Salesforce documentation states, "Prompt Builder lets you create prompt templates that generate AI content grounded in your CRM data" (Salesforce Help: "Creating Prompt Templates").

Conclusion: Prompt Builder is a strong candidate for generating hyper-relevant, customized content efficiently. However, it requires additional tools for scale, making it a partial but viable solution.

*** C. Einstein Lead Follow-Up**

Description: There is no explicit product named "Einstein Lead Follow-Up" in Salesforce's official documentation as of March 08, 2025. This could be a misnomer or a hypothetical reference to features like Einstein Lead Scoring (prioritizing leads) or Agentforce SDR (autonomous lead nurturing). For fairness, let's assume it implies an AI-driven follow-up mechanism for leads.

Analysis Against Requirements:

Scale: If interpreted as part of Agentforce (e.g., SDR Agent), it could autonomously contact leads at scale, handling thousands of interactions 24/7.

Hyper-relevance: It could use CRM and external data to tailor follow-ups, aligning with the need for relevant solutions.

Customization: It might generate personalized messages or actions (e.g., booking meetings), depending on implementation.

Efficiency: An autonomous agent would maximize efficiency by offloading outreach tasks from reps.

Issue: Without a verified product called "Einstein Lead Follow-Up," we can't confirm its capabilities. Einstein Lead Scoring, for example, prioritizes leads but doesn't contact them. Agentforce SDR fits better but isn't listed.

Conclusion: If this were Agentforce SDR, it'd be ideal. Given the option's ambiguity, it's unreliable as a verified answer.

Step 3: Identify the Best Fit Among Options

Einstein Sales Assistant: Enhances rep productivity but lacks scale and autonomy.

Prompt Builder: Generates hyper-relevant, customized content efficiently and can scale when paired with automation tools like Flow or Agentforce. It's a verifiable, existing tool that partially meets the need.

Einstein Lead Follow-Up: Potentially ideal if it implies autonomous follow-up (e.g., Agentforce), but it's not a recognized product, making it speculative.

Among the given options, Prompt Builder stands out because:

It directly addresses hyper-relevance and customization via AI-generated content tied to CRM data.

It can be scaled with Salesforce automation (e.g., Flow to send emails to thousands of leads), though this requires additional setup.

It's efficient for content creation, a key bottleneck in lead outreach.

Step 4: Consider the Ideal Solution (Agentforce Context)

The question aligns closely with Agentforce Sales Agents (e.g., SDR), which autonomously contacts leads at scale, delivers hyper-relevant solutions, and customizes communications using Data Cloud and the Atlas Reasoning Engine. Salesforce documentation notes, "Agentforce SDR autonomously nurtures inbound leads" crafting personalized responses on preferred channels" (Salesforce.com: "Agentforce for Sales"). However, Agentforce isn't an option here, so we must choose from A, B, or C.

Step 5: Final Verification

Prompt Builder Reference: "Use Prompt Builder to generate personalized sales emails or summaries in bulk, integrated with Flow for automation" (Trailhead: "Customize AI Content with Prompt Builder"). This confirms its capability for relevance and customization, with scale achievable via integration.

No other option fully meets all criteria standalone. Einstein Sales Assistant lacks scale, and Einstein Lead Follow-Up lacks definition.

Thus, Prompt Builder (B) is the best choice among the provided options, assuming it's paired with automation for execution. Without that assumption, none fully suffice, but Prompt Builder is the most verifiable and closest fit.

NEW QUESTION 54

Universal Containers plans to enhance its sales team's productivity using AI. Which specific requirement necessitates the use of Prompt Builder?

- A. Creating a draft newsletter for an upcoming tradeshow.
- B. Predicting the likelihood of customers churning or discontinuing their relationship with the company.
- C. Creating an estimated Customer Lifetime Value (CLV) with historical purchase data.

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation: UC seeks an AI solution for sales productivity. Let's determine which requirement aligns with Prompt Builder.

? Option A: Creating a draft newsletter for an upcoming tradeshow. Prompt Builder excels at generating text outputs (e.g., newsletters) using Generative AI. UC can create a prompt template to draft personalized, context-rich newsletters based on sales data, boosting productivity. This matches Prompt Builder's capabilities, making it the correct answer.

? Option B: Predicting the likelihood of customers churning or discontinuing their relationship with the company. Churn prediction is a predictive AI task, suited for Einstein Prediction Builder or Data Cloud models, not Prompt Builder, which focuses on generative tasks. This is incorrect.

? Option C: Creating an estimated Customer Lifetime Value (CLV) with historical purchase data. CLV estimation involves predictive analytics, not text generation, and is better handled by Einstein Analytics or custom models, not Prompt Builder. This is incorrect.

Why Option A is Correct: Drafting newsletters is a generative task uniquely suited to Prompt Builder, enhancing sales productivity as per Salesforce documentation.

References:

? Salesforce Agentforce Documentation: Prompt Builder > Use Cases – Lists text generation like newsletters.

? Trailhead: Build Prompt Templates in Agentforce – Covers productivity-enhancing text outputs.

? Salesforce Help: Generative AI with Prompt Builder – Confirms drafting capabilities.

NEW QUESTION 55

Universal Containers (UC) configured a new PDF file ingestion in Data Cloud with all the required fields, and also created the mapping and the search Index. UC is now setting up the retriever and notices a required field is missing.

How should UC resolve this?

- A. Create a new custom Data Cloud object that includes the desired field.
- B. Update the search index to include the desired field.
- C. Modify the retriever's configuration to include the desired field..

Answer: B

Explanation:

Why is "Update the search index to include the desired field" the correct answer? When configuring a retriever in Data Cloud for PDF file ingestion, all necessary fields must be included in the search index. If a required field is missing, the correct action is to update the search index to ensure it is available for retrieval.

Key Considerations for Fixing Missing Fields in Data Cloud Retrievers:

? Search Index Controls Which Fields Are Searchable

? Ensures Complete and Accurate Data Retrieval

? Supports AI-Generated Responses

Why Not the Other Options?

* A. Create a new custom Data Cloud object that includes the desired field.

? Incorrect because the issue is with indexing, not with Data Cloud object structure.

? The field already exists in Data Cloud; it just needs to be indexed.

* C. Modify the retriever's configuration to include the desired field.

? Incorrect because retriever configurations only define query rules; they do not modify the index itself.

? Updating the search index is the required step to ensure the field is retrievable.

Agentforce Specialist References

? Salesforce AI Specialist Material confirms that search indexing is required for retrievers to access specific fields in Data Cloud.

NEW QUESTION 57

When a customer chat is initiated, which functionality in Salesforce provides generative AI replies or draft emails based on recommended Knowledge articles?

A. Einstein Reply Recommendations

B. Einstein Service Replies

C. Einstein Grounding

Answer: B

Explanation:

When a customer chat is initiated, Einstein Service Replies provides generative AI replies or draft emails based on recommended Knowledge articles. This feature uses the information from the Salesforce Knowledge base to generate responses that are relevant to the customer's query, improving the efficiency and accuracy of customer support interactions.

? Option B is correct because Einstein Service Replies is responsible for generating

AI-driven responses based on knowledge articles.

? Option A (Einstein Reply Recommendations) is focused on recommending replies but does not generate them.

? Option C (Einstein Grounding) refers to grounding responses in data but is not directly related to drafting replies.

References:

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

NEW QUESTION 61

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

Which tool should the Agentforce Specialist recommend?

A. Model Playground

B. Agent

C. Copilot Builder

Answer: C

Explanation:

To understand multiple real-world user utterances and ensure the correct action is selected before activating a custom copilot action, the recommended tool is Copilot Builder. This tool allows Agentforce 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 (Agent) 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 63

Once a data source is chosen for an Agentforce Data Library, what is true about changing that data source later?

A. The data source can be changed through the Data Cloud settings.

B. The Data Retriever can be reconfigured to use a different data source.

C. The data source cannot be changed after it is selected.

Answer: C

Explanation:

Why is "The data source cannot be changed after it is selected" the correct answer? When configuring an Agentforce Data Library, the data source selection is permanent. Once a data source is set, it cannot be modified or replaced. This design ensures data consistency, security, and reliability within Salesforce's AI-driven environment.

Key Considerations in Agentforce Data Library

? Data Source Lock-In

? Why Can't the Data Source Be Changed?

? Workarounds for Changing Data Sources

Why Not the Other Options?

- * A. The data source can be changed through the Data Cloud settings.
 - ? Incorrect because once the data source is linked to an Agentforce Data Library, it cannot be altered, even via Data Cloud settings.
 - * B. The Data Retriever can be reconfigured to use a different data source.
 - ? Incorrect as the Data Retriever works within the constraints of the selected data source and does not provide an option to swap data sources post-selection.
- Agentforce Specialist References
The Salesforce AI Specialist Material and Salesforce Instructions for the Certification confirm that once a data source is set for an Agentforce Data Library, it cannot be changed.

NEW QUESTION 64

In a Knowledge-based data library configuration, what is the primary difference between the identifying fields and the content fields?

- A. Identifying fields help locate the correct Knowledge article, while content fields enrich AI responses with detailed information.
- B. Identifying fields categorize articles for indexing purposes, while content fields provide a brief summary for display.
- C. Identifying fields highlight key terms for relevance scoring, while content fields store the full text of the article for retrieval.

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation: In Agentforce, a Knowledge-based data library (e.g., via Salesforce Knowledge or Data Cloud grounding) uses identifying fields and content fields to support AI responses. Let's analyze their roles.

? Option A: Identifying fields help locate the correct Knowledge article, while content fields enrich AI responses with detailed information. In a Knowledge-based data library, identifying fields (e.g., Title, Article Number, or custom metadata) are used to search and pinpoint the relevant Knowledge article based on user input or context. Content fields (e.g., Article Body, Details) provide the substantive data that the AI uses to generate detailed, enriched responses. This distinction is critical for grounding Agentforce prompts and aligns with Salesforce's documentation on Knowledge integration, making it the correct answer.

? Option B: Identifying fields categorize articles for indexing purposes, while content fields provide a brief summary for display. Identifying fields do more than categorize—they actively locate articles, not just index them. Content fields aren't limited to summaries; they include full article content for response generation, not just display. This option underrepresents their roles and is incorrect.

? Option C: Identifying fields highlight key terms for relevance scoring, while content fields store the full text of the article for retrieval. While identifying fields contribute to relevance (e.g., via search terms), their primary role is locating articles, not just scoring. Content fields do store full text, but their purpose is to enrich responses, not merely enable retrieval. This option shifts focus inaccurately, making it incorrect.

Why Option A is Correct: The primary difference—identifying fields for locating articles and content fields for enriching responses—reflects their roles in Knowledge-based grounding, as per official Agentforce documentation.

References:

? Salesforce Agentforce Documentation: Grounding with Knowledge > Data Library Setup – Defines identifying vs. content fields.

? Trailhead: Ground Your Agentforce Prompts – Explains field roles in Knowledge integration.

? Salesforce Help: Knowledge in Agentforce – Confirms locating and enriching functions.

NEW QUESTION 69

What is an appropriate use case for leveraging Agentforce Sales Agent in a sales context?

- A. Enable a sales team to use natural language to invoke defined sales tasks grounded in relevant data and be able to ensure company policies are applied conversationally and in the now or work.
- B. Enable a sales team by providing them with an interactive step-by-step guide based on business rules to ensure accurate data entry into Salesforce and help close deals faster.
- C. Instantly review and read incoming messages or emails that are then logged to the correct opportunity, contact, and account records to provide a full view of customer interactions and communications.

Answer: A

Explanation:

Agentforce Sales Agent is designed to let sales teams perform tasks via natural language commands, leveraging Salesforce data while adhering to policies. For example, agents can ask the AI to "update the opportunity stage to Closed Won" or "generate a quote," with the system enforcing validations and data security. This use case aligns with Salesforce's vision of conversational AI streamlining workflows without compromising compliance.

? Step-by-step guides (B) are typically handled by tools like Dynamic Forms or Guided Selling, not Agentforce.

? Logging messages/emails (C) is managed by Email-to-Case or Service Cloud, not a sales-specific AI agent.

Reference:

Salesforce Help Article: Agentforce for Sales ("Use Cases and Capabilities" section).

Einstein Agentforce Specialist Trailhead: "Sales Automation with Agentforce" (Natural Language Task Execution).

NEW QUESTION 72

Universal Containers implemented Agent for its users.

One user complains that Agent is not deleting activities from the past 7 days. What is the reason for this issue?

- A. Agent Delete Record Action permission is not associated to the user.
- B. Agent does not have the permission to delete the user's records.
- C. Agent does not support the Delete Record action.

Answer: C

Explanation:

Agent 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 Agent.

? Unsupported Action: The inability to delete records is due to the current limitations of Agent'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, Agent itself does not have the capability to execute delete operations.

References:

? Salesforce Agentforce Specialist Documentation - Agent Supported Actions:

? Salesforce Help - Limitations of Agent:

NEW QUESTION 77

In addition to Recipient and Sender, which object should An Agentforce utilize for inserting merge fields into a Sales email template prompt?

- A. Recipient Opportunities
- B. Recipient Account
- C. User Organization

Answer: B

Explanation:

? Sales Email Template Use Case:When creating a Sales email template (especially for outreach or follow-up), you often need to reference relevant details about the Account linked to the recipient.

? Standard Merge Fields in Salesforce Email Templates:

? Why Recipient Account?

? References and Study Resources:

NEW QUESTION 78

Universal Containers deployed the new Agentforce Sales Development Representative (SDR) Into production, but sales reps are saying they can't find it. What is causing this issue?

- A. Sales rep users profiles are missing the Allow SDR Agent permission.
- B. Sales rep users do not have access to the SDR Agent object.
- C. Sales rep users are missing the Use SDR Agent permission set.

Answer: C

Explanation:

Why is "Sales rep users are missing the Use SDR Agent permission set" the correct answer?

If sales reps are unable to find the Agentforce Sales Development Representative (SDR) Agent, the most likely cause is missing permissions. The "Use SDR Agent" permission set is required for users to access and interact with the SDR Agent in Agentforce.

Key Considerations for This Issue:

? Permission Set Restriction

? Agentforce Role-Based Access Control

? Fixing the Issue

Why Not the Other Options?

* A. Sales rep users' profiles are missing the Allow SDR Agent permission.

? Incorrect because "Allow SDR Agent" is not a standard permission setting in Agentforce.

? Permission is granted via permission sets, not profile-level settings.

* B. Sales rep users do not have access to the SDR Agent object.

? Incorrect because there is no separate "SDR Agent object" in Salesforce.

? SDR Agents are AI-driven features, not standard CRM objects that require object- level access.

Agentforce Specialist References

? Salesforce AI Specialist Material confirms that users require specific permission sets to access Agentforce SDR Agents.

? Salesforce Instructions for Certification highlight the role of permission sets in controlling Agentforce access.

NEW QUESTION 83

An Agentforce needs to enable the use of Sales Email prompt templates for the sales team. The Agentforce Specialist has already created the templates in Prompt Builder.

According to best practices, which steps should the Agentforce Specialist take to ensure the sales team can use these templates?

- A. Assign the Prompt Template User permission set and enable Sales Emails in Setup.
- B. Assign the Prompt Template Manager permission set and enable Sales Emails in setup.
- C. Assign the Data Cloud Admin permission set and enable Sales Emails in Setup.

Answer: A

Explanation:

To enable Sales Email prompt templates:

? Permission Set: Assign the Prompt Template User permission set to the sales team to grant access to use pre-built templates.

? Feature Activation: Enable Sales Emails in Salesforce Setup to activate the integration between prompt templates and email workflows.

? Option B (Manager permission set): Required for creating/modifying templates, not for usage.

? Option C (Data Cloud Admin): Unrelated to prompt template access.

References:

? Salesforce Help: Prompt Template Permissions

? Specifies that "Prompt Template User" is required to leverage templates in workflows.

? Sales Email Setup outlines enabling the feature in Setup.

NEW QUESTION 88

Universal Containers (UC) wants to ensure the effectiveness, reliability, and trust of its agents prior to deploying them in production. UC would like to efficiently test a large and repeatable number of utterances. What should the Agentforce Specialist recommend?

- A. Leverage the Agent Large Language Model (LLM) UI and test UC's agents with different utterances prior to activating the agent.
- B. Deploy the agent in a QA sandbox environment and review the Utterance Analysis reports to review effectiveness.
- C. Create a CSV file with UC's test cases in Agentforce Testing Center using the testing template.

Answer: C

Explanation:

Comprehensive and Detailed In-Depth Explanation:The goal of Universal Containers (UC) is to test its Agentforce agents for effectiveness, reliability, and trust

before production deployment, with a focus on efficiently handling a large and repeatable number of utterances. Let's evaluate each option against this requirement and Salesforce's official Agentforce tools and best practices.

? Option A: Leverage the Agent Large Language Model (LLM) UI and test UC's agents with different utterances prior to activating the agent. While Agentforce leverages advanced reasoning capabilities (powered by the Atlas Reasoning Engine), there's no specific "Agent Large Language Model (LLM) UI" referenced in Salesforce documentation for testing agents. Testing utterances directly within an LLM interface might imply manual experimentation, but this approach lacks scalability and repeatability for a large number of utterances. It's better suited for ad-hoc testing of individual responses rather than systematic evaluation, making it inefficient for UC's needs.

? Option B: Deploy the agent in a QA sandbox environment and review the Utterance Analysis reports to review effectiveness. Deploying an agent in a QA sandbox is a valid step in the development lifecycle, as sandboxes allow testing in a production-like environment without affecting live data. However, "Utterance Analysis reports" is not a standard term in Agentforce documentation. Salesforce provides tools like Agent Analytics or User Utterances dashboards for post-deployment analysis, but these are more about monitoring live performance than pre-deployment testing. This option doesn't explicitly address how to efficiently test a large and repeatable number of utterances before deployment, making it less precise for UC's requirement.

? Option C: Create a CSV file with UC's test cases in Agentforce Testing Center using the testing template. The Agentforce Testing Center is a dedicated tool within Agentforce Studio designed specifically for testing autonomous AI agents. According to Salesforce documentation, Testing Center allows users to upload a CSV file containing test cases (e.g., utterances and expected outcomes) using a provided template. This enables the generation and execution of hundreds of synthetic interactions in parallel, simulating real-world scenarios. The tool evaluates how the agent interprets utterances, selects topics, and executes actions, providing detailed results for iteration. This aligns perfectly with UC's need for efficiency (bulk testing via CSV), repeatability (standardized test cases), and reliability (systematic validation), ensuring the agent is production-ready. This is the recommended approach per official guidelines.

Why Option C is Correct: The Agentforce Testing Center is explicitly built for pre-deployment validation of agents. It supports bulk testing by allowing users to upload a CSV with utterances, which is then processed by the Atlas Reasoning Engine to assess accuracy and reliability. This method ensures UC can systematically test a large dataset, refine agent instructions or topics based on results, and build trust in the agent's performance—all before production deployment. This aligns with Salesforce's emphasis on testing non-deterministic AI systems efficiently, as noted in Agentforce setup documentation and Trailhead modules.

References:

? Salesforce Trailhead: Get Started with Salesforce Agentforce Specialist Certification Prep – Details the use of Agentforce Testing Center for testing agents with synthetic interactions.

? Salesforce Agentforce Documentation: Agentforce Studio > Testing Center – Explains how to upload CSV files with test cases for parallel testing.

? Salesforce Help: Agentforce Setup > Testing Autonomous AI Agents – Recommends Testing Center for pre-deployment validation of agent effectiveness and reliability.

NEW QUESTION 89

Universal Containers wants to incorporate CRM data as well-formatted JSON in a prompt to a large language model (LLM). What is an important consideration for this requirement?

- A. "CRM data to JSON" checkbox must be selected when creating a prompt template.
- B. Apex code can be used to return a JSON formatted merge field.
- C. JSON format should be enabled in Prompt Builder Settings.

Answer: B

Explanation:

? Context of the Question

? Why Apex Code for JSON Formatting?

? Conclusion The practical solution to pass CRM data in JSON format to an LLM is to use Apex code (or a specialized Flow approach) to produce a JSON string, which the prompt can then merge and pass along. Hence, Option B is correct.

Salesforce Agentforce Specialist References & Documents

? Salesforce Documentation: Working with JSON in Apex Describes how to serialize and deserialize data using Apex for integration or AI prompts.

? Salesforce Agentforce Specialist Study Guide Emphasizes the need for custom logic (often in Apex) when complex data transformations (like JSON formatting) are required.

NEW QUESTION 90

What is the role of the large language model (LLM) in executing an Agent 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 Agent, 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 Agent Actions

Salesforce AI Documentation on Large Language Models

NEW QUESTION 94

Universal Containers' Agent Action includes several Apex classes for the new Agentforce Agent. What is an important consideration when deploying Apex that is invoked by an Agent Action?

- A. The Apex classes must have at least 75% code coverage from unit tests, and all dependencies must be in the deployment package.
- B. Apex classes invoked by an Agent Action may be deployed with less than 75% test coverage as long as the agent is not activated in production.
- C. The Apex classes may bypass the 75% code coverage requirement as long as they are only used by the agent.

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation: Universal Containers (UC) is using Apex classes within an Agent Action for their Agentforce Agent. Deploying Apex in Salesforce has specific requirements, especially when tied to Agentforce functionality. Let's evaluate the options.

? Option A: The Apex classes must have at least 75% code coverage from unit tests, and all dependencies must be in the deployment package. Salesforce enforces a strict requirement that all Apex classes must achieve at least 75% code coverage from unit tests for deployment to production, regardless of their use case (e.g., Agentforce, triggers, or web services). Additionally, when Apex is invoked by an Agent Action (e.g., via a Flow or direct invocation), all dependencies (e.g., referenced classes, objects) must be included in the deployment package to ensure functionality. This is a standard deployment consideration in Salesforce and applies to Agentforce, making this the correct answer.

? Option B: Apex classes invoked by an Agent Action may be deployed with less than 75% test coverage as long as the agent is not activated in production. Salesforce's 75% code coverage requirement is mandatory for production deployment, regardless of whether the agent is activated. There's no exemption based on activation status—coverage is enforced at the deployment stage. This option is incorrect and contradicts Salesforce's Apex deployment rules.

? Option C: The Apex classes may bypass the 75% code coverage requirement as long as they are only used by the agent. No such bypass exists in Salesforce. The 75% code coverage rule applies universally to all Apex in production, including classes used by Agentforce. Agent-specific usage doesn't waive this requirement, making this incorrect.

Why Option A is Correct: The 75% code coverage requirement and inclusion of dependencies are fundamental Salesforce deployment rules, applicable to Apex in Agent Actions. This ensures reliability and functionality in production, as per official documentation.

References:

? Salesforce Agentforce Documentation: Agent Builder > Custom Actions > Apex – Notes standard Apex deployment rules apply.

? Salesforce Developer Guide: Apex Testing – Confirms 75% coverage requirement.

? Trailhead: Deploy Apex Code – Emphasizes coverage and dependencies for production.

NEW QUESTION 99

Universal Containers (UC) recently rolled out Einstein Generative AI 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 prompt template version is incompatible with the chosen LLM.
- B. The data being used for grounding is incorrect or incomplete.
- C. The Einstein Trust Layer is incorrectly configured.

Answer: B

Explanation:

Comprehensive and Detailed In-Depth Explanation: UC's custom prompt for summarizing case records is underperforming, and we need to identify a likely cause. Let's evaluate the options based on Agentforce and Einstein Generative AI mechanics.

? Option A: The prompt template version is incompatible with the chosen LLM. Prompt templates in Agentforce are designed to work with the Atlas Reasoning Engine, which abstracts the underlying large language model (LLM). Salesforce manages compatibility between prompt templates and LLMs, and there's no user-facing versioning that directly ties to LLM compatibility. This option is unlikely and not a common issue per documentation.

? Option B: The data being used for grounding is incorrect or incomplete. Grounding is the process of providing context (e.g., case record data) to the AI via prompt templates. If the grounding data—sourced from Record Snapshots, Data Cloud, or other integrations—is incorrect (e.g., wrong fields mapped) or incomplete (e.g., missing key case details), the summaries will be inaccurate. For example, if the prompt relies on Case.Subject but the field is empty or not included, the output will miss critical information. This is a frequent cause of poor performance in generative AI and aligns with Salesforce troubleshooting guidance, making it the correct answer.

? Option C: The Einstein Trust Layer is incorrectly configured. The Einstein Trust Layer enforces guardrails (e.g., toxicity filtering, data masking) to ensure safe and compliant AI outputs. Misconfiguration might block content or alter tone, but it's unlikely to cause summaries to lack appropriate information unless specific fields are masked unnecessarily. This is less probable than grounding issues and not a primary explanation here.

Why Option B is Correct: Incorrect or incomplete grounding data is a well-documented reason for subpar AI outputs in Agentforce. It directly affects the quality of case summaries, and specialists are advised to verify grounding sources (e.g., field mappings, Data Cloud queries) when troubleshooting, as per official guidelines.

References:

? Salesforce Agentforce Documentation: Prompt Templates > Grounding – Links poor outputs to grounding issues.

? Trailhead: Troubleshoot Agentforce Prompts – Lists incomplete data as a common problem.

? Salesforce Help: Einstein Generative AI > Debugging Prompts – Recommends checking grounding data first.

NEW QUESTION 100

Universal Containers wants to allow its service agents to query the current fulfillment status of an order with natural language. There is an existing auto launched flow to query the information from Oracle ERP, which is the system of record for the order fulfillment process.

How should An Agentforce apply the power of conversational AI to this use case?

- A. Create a Flex prompt template in Prompt Builder.
- B. Create a custom copilot action which calls a flow.
- C. Configure the Integration Flow Standard Action in Agent.

Answer: B

Explanation:

To enable Universal Containers service agents to query the current fulfillment status of an order using natural language and leverage an existing auto-launched flow that queries Oracle ERP, the best solution is to create a custom copilot action that calls the flow. This action will allow Agent to interact with the flow and retrieve the required order fulfillment information seamlessly. Custom copilot actions can be tailored to call various backend systems or flows in response to user requests.

? Option B is correct because it enables integration between Agent and the flow that connects to Oracle ERP.

? Option A (Flex prompt template) is more suited for static responses and not for invoking flows.

? Option C (Integration Flow Standard Action) is not directly related to creating a specific copilot action for this use case.

References:

? Salesforce Agent Actions: https://help.salesforce.com/s/articleView?id=einstein_copilot_actions.htm

NEW QUESTION 102

Universal Containers is rolling out a new generative AI initiative.

Which Prompt Builder limitations should the Agentforce Specialist be aware of?

- A. Rich text area fields are only supported in Flex template types.
- B. Creations or updates to the prompt templates are not recorded in the Setup Audit Trail.
- C. Custom objects are supported only for Flex template types.

Answer: C

Explanation:

The Prompt Builder in Salesforce has some specific limitations, one of which is that custom objects are supported only for Flex template types. This means that users must rely on Flex templates to integrate custom objects into their prompts.

? Option A: While rich text area fields have certain restrictions, this does not pertain to the core limitation of integrating custom objects.

? Option B: Updates and creations for prompt templates are indeed recorded in the Setup Audit Trail, so this statement is incorrect.

? Option C: This is the correct answer as it reflects a documented limitation of the Prompt Builder.

Reference:

"Prompt Builder Limitations | Salesforce Documentation" .

NEW QUESTION 106

After a successful implementation of Agentforce Sates Agent with sales users. Universal Containers now aims to deploy it to the service team. Which key consideration should the Agentforce Specialist keep in mind for this deployment?

- A. Assign the Agentforce for Service permission to the Service Cloud users.
- B. Assign the standard service actions to Agentforce Service Agent.
- C. Review and test standard and custom Agent topics and actions for Service Center usecases.

Answer: C

Explanation:

When deploying Einstein Agent (formerly Agentforce) from Sales to Service Cloud:

? Agent Topics and Actions are context-specific. Service Cloud use cases (e.g., case resolution, knowledge retrieval) require validation of existing topics/actions to ensure alignment with service workflows.

? Option A: Permissions like "Agentforce for Service" are necessary but secondary to functional compatibility.

? Option B: Standard service actions must be mapped to Agentforce, but testing ensures they function as intended.

References:

? Salesforce Help: Einstein Agent Setup

? Emphasizes reviewing "topics and actions for different user groups (Sales vs. Service)."

NEW QUESTION 107

Universal Containers (UC) wants to implement an AI-powered customer service agent that can:

? Retrieve proprietary policy documents that are stored as PDFs.

? Ensure responses are grounded in approved company data, not generic LLM knowledge.

What should UC do first?

- A. Set up an Agentforce Data Library for AI retrieval of policy documents.
- B. Expand the AI agent's scope to search all Salesforce records.
- C. Add the files to the content, and then select the data library option.

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation: To implement an AI- powered customer service agent that retrieves proprietary policy documents (stored as PDFs) and ensures responses are grounded in approved company data, UC must first establish a foundation for the AI to access and use this data. The Agentforce Data Library (Option A) is the correct starting point. A Data Library allows UC to upload PDFs containing policy documents, index them into Salesforce Data Cloud's vector database, and make them available for AI retrieval. This setup ensures the agent can perform Retrieval- Augmented Generation (RAG), grounding its responses in the specific, approved content from the PDFs rather than relying on generic LLM knowledge, directly meeting UC's requirements.

? Option B: Expanding the AI agent's scope to search all Salesforce records is too broad and unnecessary at this stage. The requirement focuses on PDFs with policy documents, not all Salesforce data (e.g., cases, accounts), making this premature and irrelevant as a first step.

? Option C: "Add the files to the content, and then select the data library option" is vague and not a precise process in Agentforce. While uploading files is part of setting up a Data Library, the phrasing suggests adding files to Salesforce Content (e.g., ContentDocument) without indexing, which doesn't enable AI retrieval. Setting up the Data Library (A) encompasses the full process correctly.

? Option A: This is the foundational step—creating a Data Library ensures the PDFs are uploaded, indexed, and retrievable by the agent, fulfilling both retrieval and grounding needs.

Option A is the correct first step for UC to achieve its goals.

References:

? Salesforce Agentforce Documentation: "Set Up a Data Library" (Salesforce Help: https://help.salesforce.com/s/articleView?id=sf.agentforce_data_library.htm&type=5)

? Salesforce Data Cloud Documentation: "Ground AI Responses with Data Cloud" (https://help.salesforce.com/s/articleView?id=sf.data_cloud_agentforce.htm&type=5)

NEW QUESTION 110

Universal Containers (UC) implements a custom retriever to improve the accuracy of AI-

generated responses. UC notices that the retriever is returning too many irrelevant results, making the responses less useful. What should UC do to ensure only relevant data is retrieved?

- A. Define filters to narrow the search results based on specific conditions.
- B. Change the search index to a different data model object (DMO).
- C. Increase the maximum number of results returned to capture a broader dataset.

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation: In Salesforce Agentforce, a custom retriever is used to fetch relevant data (e.g., from Data Cloud's vector database or Salesforce records) to ground AI responses. UC's issue is that their retriever returns too many irrelevant results, reducing response accuracy. The best solution is to define filters (Option A) to refine the retriever's search criteria. Filters allow UC to specify conditions (e.g., "only retrieve documents from the Policy category or records created after a certain date") that narrow the dataset, ensuring the retriever returns only relevant results. This directly improves the precision of AI-generated responses by excluding extraneous data, addressing UC's problem effectively.

Option B: Changing the search index to a different data model object (DMO) might be relevant if the retriever is querying the wrong object entirely (e.g., Accounts instead of Policies). However, the question implies the retriever is functional but unrefined, so adjusting the existing setup with filters is more appropriate than switching DMOs.

Option C: Increasing the maximum number of results would worsen the issue by returning even more data, including more irrelevant entries, contrary to UC's goal of improving relevance.

Option A: Filters are a standard feature in custom retrievers, allowing precise control over retrieved data, making this the correct action.

Option A is the most effective step to ensure relevance in retrieved data.

References:

Salesforce Agentforce Documentation: "Create Custom Retrievers" (Salesforce Help:

https://help.salesforce.com/s/articleView?id=sf.agentforce_custom_retrievers.htm&type=5)

Salesforce Data Cloud Documentation: "Filter Data for AI Retrieval" (https://help.salesforce.com/s/articleView?id=sf.data_cloud_retrieval_filters.htm&type=5)

NEW QUESTION 115

After creating a foundation model in Einstein Studio, which hyperparameter should An Agentforce use to adjust the balance between consistency and randomness of a response?

- A. Presence Penally
- B. Variability
- C. Temperature

Answer: C

Explanation:

The Temperature hyperparameter controls the randomness of model outputs:

Low Temperature (e.g., 0.2): More deterministic, consistent responses.

High Temperature (e.g., 1.0): More creative, varied responses.

Presence Penalty (Option A): Discourages repetition of tokens, unrelated to randomness.

Variability (Option B): Not a standard hyperparameter in Einstein Studio.

References:

Einstein Studio Documentation: Model Hyperparameters

Explicitly states "Temperature adjusts the balance between predictable and random outputs."

NEW QUESTION 117

Universal Containers wants to allow its service agents to query the current fulfillment status of an order with natural language. There is an existing autolaunched flow to query the Information from Oracle ERP, which is the system of record for the order fulfillment process.

How should an Agentforce Specialist apply the power of conversational AI to this use case?

- A. Create a custom Agent action which calls a flow.
- B. Configure the Integration Flow Standard Action in Agent Builder.
- C. Create a Flex prompt template in Prompt Builder.

Answer: A

Explanation:

Why is "Create a custom Agent action which calls a flow" the correct answer?

In Agentforce, the best way to allow service agents to query order fulfillment status from an external system (Oracle ERP) using natural language is to create a custom Agent action that invokes an existing autolaunched flow.

Key Considerations for This Approach:

Custom Agent Action Triggers the Flow

Enables AI-Powered Query Execution

No Need for Manual Data Entry

Why Not the Other Options?

* B. Configure the Integration Flow Standard Action in Agent Builder

Incorrect because Integration Flow Standard Actions are for predefined use cases, not custom ERP integrations.

They do not provide the flexibility needed to connect with Oracle ERP dynamically.

* C. Create a Flex Prompt Template in Prompt Builder

Incorrect because Flex prompts are used for structuring AI-generated responses, not executing queries on external systems.

This approach does not enable the AI to retrieve live fulfillment status from Oracle ERP.

Agentforce Specialist References

Salesforce AI Specialist Material confirms that custom Agent actions allow integration with external systems through Salesforce flows.

Salesforce Instructions for Certification mention that Agentforce supports custom Agent actions for external data retrieval.

NEW QUESTION 119

Universal Containers (UC) is using standard Service AI Grounding. UC created a custom rich text field to be used with Service AI Grounding.

What should UC consider when using standard Service AI Grounding?

- A. Service AI Grounding only works with Case and Knowledge objects.
- B. Service AI Grounding only supports String and Text Area type fields.
- C. Service AI Grounding visibility works in system mode.

Answer: B

Explanation:

Service AI Grounding retrieves data from Salesforce objects to ground AI-generated responses. Key considerations:

- ? Field Types: Standard Service AI Grounding supports String and Text Area fields. Custom rich text fields (e.g., RichTextArea) are not supported, making Option B correct.
 - ? Objects: While Service AI Grounding primarily uses Case and Knowledge objects (Option A), the limitation here is the field type, not the object.
 - ? Visibility: Service AI Grounding respects user permissions and sharing settings unless overridden (Option C is incorrect).
- References:
- ? Salesforce Help: Service AI Grounding Requirements
 - ? Explicitly states support for "Text Area and String fields" only.

NEW QUESTION 124

Universal Containers (UC) wants its AI agent to return responses quickly. UC needs to optimize the retriever's configuration to ensure minimal latency when grounding AI responses.

Which configuration aspect should UC prioritize?

- A. Configure the retriever to operate in dynamic mode so that it modifies the search Index structure at runtime.
- B. Ensure the retriever's filters are defined to limit the scope of each search efficiently.
- C. Increase the recency bias setting for the retriever limiting scope to more recent data.

Answer: B

Explanation:

Why is "Ensure the retriever's filters are defined to limit the scope of each search efficiently" the correct answer?

In Agentforce, when optimizing a retriever's configuration to ensure minimal latency in AI-generated responses, the most effective approach is narrowing the scope of searches

by applying specific filters.

Key Considerations for Optimizing Retrievers in Agentforce:

- ? Defining Effective Filters
- ? Reducing Query Complexity
- ? Optimizing the Data Indexing Process

Why Not the Other Options?

* A. Configure the retriever to operate in dynamic mode so that it modifies the search index structure at runtime.

? Incorrect because modifying the search index at runtime increases latency rather than reducing it.

? Index modifications require restructuring large datasets, which can slow down AI-generated responses.

* C. Increase the recency bias setting for the retriever, limiting scope to more recent data.

? Incorrect because increasing recency bias only prioritizes recent records but does not necessarily improve overall retrieval speed.

? While it affects relevance, it does not directly address latency issues.

Agentforce Specialist References

? Salesforce AI Specialist Material confirms that retriever efficiency depends on well-defined filtering mechanisms to minimize latency.

? Salesforce Instructions for Certification highlight retriever optimization strategies to improve search response times.

NEW QUESTION 126

Universal Containers (UC) has implemented Generative AI within Salesforce to enable summarization of a custom object called Guest. Users have reported mismatches in the generated information.

In refining its prompt design strategy, which key practices should UC prioritize?

- A. Enable prompt test mode, allocate different prompt variations to a subset of users for evaluation, and standardize the most effective model based on performance feedback.
- B. Create concise, clear, and consistent prompt templates with effective grounding, contextual role-playing, clear instructions, and iterative feedback.
- C. Submit a prompt review case to Salesforce and conduct thorough testing in the playground to refine outputs until they meet user expectations.

Answer: B

Explanation:

For Universal Containers (UC) to refine its Generative AI prompt design strategy and improve the accuracy of the generated summaries for the custom object Guest, the best practice is to focus on crafting concise, clear, and consistent prompt templates. This includes:

? Effective grounding: Ensuring the prompt pulls data from the correct sources.

? Contextual role-playing: Providing the AI with a clear understanding of its role in generating the summary.

? Clear instructions: Giving unambiguous directions on what to include in the response.

? Iterative feedback: Regularly testing and adjusting prompts based on user feedback.

? Option B is correct because it follows industry best practices for refining prompt design.

? Option A (prompt test mode) is useful but less relevant for refining prompt design itself.

? Option C (prompt review case with Salesforce) would be more appropriate for technical issues or complex prompt errors, not general design refinement.

References:

Salesforce Prompt Design Best Practices: https://help.salesforce.com/s/articleView?id=sf.prompt_design_best_practices.htm

NEW QUESTION 129

After configuring and saving a Salesforce Agentforce Data Library (regardless of the data source), which components are automatically created and available in Data Cloud?

- A. A data pipeline, an indexing engine, and a query processor
- B. A data connector, an analytics dashboard, and a workflow rule
- C. A data stream, a search index, and a retriever

Answer: C

Explanation:

Why is "A data stream, a search index, and a retriever" the correct answer? When a Salesforce Agentforce Data Library is configured and saved, it automatically creates three essential components in Data Cloud to facilitate AI-driven search and retrieval.

Key Components Created in Data Cloud:

- ? Data Stream
- ? Search Index
- ? Retriever

Why Not the Other Options?

- * A. A data pipeline, an indexing engine, and a query processor
- ? Incorrect because Data Cloud does not use a query processor in the same way as traditional databases.
- ? Instead, retrievers handle AI-powered data searches.
- * B. A data connector, an analytics dashboard, and a workflow rule
- ? Incorrect because these components are not automatically created when setting up a Data Library.
- ? Analytics dashboards and workflow rules are separate tools used for reporting and automation.

Agentforce Specialist References

? Salesforce AI Specialist Material confirms that a Data Stream, Search Index, and Retriever are created automatically in Data Cloud when configuring a Data Library.

NEW QUESTION 134

Universal Containers needs its sales reps to be able to only execute prompt templates. What should the company use to achieve this requirement?

- A. Prompt Execute Template permission set
- B. Prompt Template User permission set
- C. Prompt Template Manager permission set

Answer: B

Explanation:

Salesforce Agentforce leverages Prompt Builder, a powerful tool that allows administrators to create and manage prompt templates, which are reusable frameworks for generating AI-driven responses. These templates can be invoked by users to perform specific tasks, such as generating sales emails or summarizing records, based on predefined instructions and grounded data. In this scenario, Universal Containers wants its sales reps to have the ability to only execute these prompt templates, meaning they should be able to run them but not create, edit, or manage them.

Let's break down the options and analyze why B. Prompt Template User permission set is the correct answer

? Option A: Prompt Execute Template permission set This option sounds plausible at first glance because it includes the phrase "Execute Template," which aligns with the requirement. However, there is no specific permission set named "Prompt Execute Template" in Salesforce's official documentation for Prompt Builder or Agentforce. Salesforce typically uses more standardized naming conventions for permission sets, and this appears to be a distractor option that doesn't correspond to an actual feature. Permissions in Salesforce are granular, but they are grouped logically under broader permission sets rather than hyper-specific ones like this.

? Option B: Prompt Template User permission set This is the correct answer. In Salesforce, the Prompt Builder feature, which is integral to Agentforce, includes permission sets designed to control access to prompt templates. The "Prompt Template User" permission set is an official Salesforce permission set that grants users the ability to execute (or invoke) prompt templates without giving them the ability to create or modify them. This aligns perfectly with the requirement that sales reps should only execute prompt templates, not manage them. The Prompt Template User permission set typically includes permissions like "Run Prompt Templates," which allows users to trigger templates from interfaces such as Lightning record pages or flows, while restricting access to the Prompt Builder setup area where templates are designed.

? Option C: Prompt Template Manager permission set This option is incorrect because the "Prompt Template Manager" permission set is designed for users who need full administrative control over prompt templates. This includes creating, editing, and deleting templates in Prompt Builder, in addition to executing them. Since Universal Containers only wants sales reps to execute templates and not manage them, this permission set provides more access than required, violating the principle of least privilege—a key security best practice in Salesforce.

How It Works in Salesforce

To implement this, an administrator would:

- ? Navigate to Setup > Permission Sets.
- ? Locate or create the "Prompt Template User" permission set (this is a standard permission set available with Prompt Builder-enabled orgs).
- ? Assign this permission set to the sales reps' profiles or individual user records.
- ? Ensure the prompt templates are configured and exposed (e.g., via Lightning components like the Einstein Summary component) on relevant pages, such as Opportunity or Account record pages, where sales reps can invoke them.

Why This Matters

By assigning the Prompt Template User permission set, Universal Containers ensures that sales reps can leverage AI-driven prompt templates to enhance productivity (e.g., drafting personalized emails or generating sales pitches) while maintaining governance over who can modify the templates. This separation of duties is critical in a secure Salesforce environment.

References to Official Salesforce Agentforce Specialist Documents

? Salesforce Help: Prompt Builder Permissions The official Salesforce documentation outlines permission sets for Prompt Builder, including "Prompt Template User" for execution-only access and "Prompt Template Manager" for full control.

? Trailhead: Configure Agentforce for Service This module discusses how permissions are assigned to control Agentforce features, including prompt-related capabilities.

? Salesforce Ben: Why Prompt Builder Is Vital in an Agentforce World (November 25, 2024) This resource explains how Prompt Builder integrates with Agentforce and highlights the use of permission sets like Prompt Template User to enable end-user functionality.

NEW QUESTION 137

What is An Agentforce able to do when the "Enrich event logs with conversation data" setting in Agent 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 Agent, it allows An Agentforce 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 Agentforce 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 Agentforce Specialist References: You can refer to this documentation for further insights:
https://help.salesforce.com/s/articleView?id=sf.einstein_copilot_event_logging.htm

NEW QUESTION 141

In Model Playground, which hyperparameters of an existing Salesforce-enabled foundational model can An Agentforce 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 Agentforce 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 143

Universal Containers (UC) has a legacy system that needs to integrate with Salesforce. UC wishes to create a digest of account action plans using the generative API feature. Which API service should UC use to meet this requirement?

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

Answer: A

Explanation:

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

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

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

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

NEW QUESTION 145

Universal Containers (UC) wants to create a new Sales Email prompt template in Prompt Builder using the "Save As" function. However, UC notices that the new template produces different results compared to the standard Sales Email prompt due to missing hyperparameters. What should UC do to ensure the new prompt template produces results comparable to the standard Sales Email prompts?

- A. Use Model Playground to create a model configuration with the specified parameters.
- B. Manually add the hyperparameters to the new template.
- C. Revert to using the standard template without modifications.

Answer: B

Explanation:

When Universal Containers creates a new Sales Email prompt template using the "Save As" function, missing hyperparameters can result in different outputs. To ensure the new prompt produces comparable results to the standard Sales Email prompt, the Agentforce Specialist should manually add the necessary hyperparameters to the new template.

? Hyperparameters like Temperature, Frequency Penalty, and Presence Penalty directly affect how the AI generates responses. Ensuring that these are consistent with the standard template will result in similar outputs.

? Option A (Model Playground) is not necessary here, as it focuses on fine-tuning models, not adjusting templates directly.

? Option C (Reverting to the standard template) does not solve the issue of customizing the prompt template.

For more information, refer to Prompt Builder documentation on configuring hyperparameters in custom templates.

NEW QUESTION 147

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

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

Answer: B

Explanation:

In order to access and create custom prompt templates in Prompt Builder, the Agentforce Specialist must have the Prompt Template Manager permission set

assigned. Without this permission, they will not be able to access Prompt Builder in the Setup menu, even though Einstein Generative AI is enabled.

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

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

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

References:

? Salesforce Prompt Builder Permissions: https://help.salesforce.com/s/articleView?id=sf.prompt_builder_permissions.htm

NEW QUESTION 152

Universal Containers aims to streamline the sales team's daily tasks by using AI.

When considering these new workflows, which improvement requires the use of Prompt Builder?

- A. Populate an AI-generated time-to close estimation to opportunities
- B. Populate an AI generated summary field for sales contracts.
- C. Populate an AI generated lead score for new leads.

Answer: B

Explanation:

Prompt Builder is explicitly required to create AI-generated summary fields via prompt templates. These fields use natural language instructions to extract or synthesize information (e.g., summarizing contract terms). Time-to-close estimations (A) and lead scores (C) are typically handled by predictive AI (e.g., Einstein Opportunity Scoring) or analytics tools, which do not require Prompt Builder.

Reference:

Salesforce Help Article: Create AI-Generated Fields with Prompt Builder ("Summary Field Generation" example).

Einstein GPT for Sales Guide: "Automating Contract Summaries."

NEW QUESTION 157

An Agentforce configured Data Masking within the Einstein Trust Layer.

How should the Agentforce Specialist begin validating that the correct fields are being masked?

- A. Use a Flow-based resource in Prompt Builder to debug the fields?? merge values using Flow Debugger.
- B. Request the Einstein Generative AI Audit Data from the Security section of the Setup menu.
- C. Enable the collection and storage of Einstein Generative AI Audit Data on the Einstein Feedback setup page.

Answer: C

Explanation:

To begin validating that the correct fields are being masked in Einstein Trust Layer, the Agentforce Specialist should request the Einstein Generative AI Audit Data from the Security section of the Salesforce Setup menu. This audit data allows the Agentforce Specialist to see how data is being processed, including which fields are being masked, providing transparency and validation that the configuration is working as expected.

? Option B is correct because it allows for the retrieval of audit data that can be used to validate data masking.

? Option A (Flow Debugger) and Option C (Einstein Feedback) do not relate to validating field masking in the context of the Einstein Trust Layer.

References:

? Salesforce Einstein Trust Layer Documentation: https://help.salesforce.com/s/articleView?id=sf.einstein_trust_layer_audit.htm

NEW QUESTION 159

Universal Containers (UC) is implementing generative AI and wants to leverage a prompt template to provide responses to customers that gives personalized product recommendations to website visitors based on their browsing history.

Which initial step should UC take to ensure the chatbot can deliver accurate recommendations?

- A. Design universal product recommendations.
- B. Write a response script for the chatbot.
- C. Collect and analyze browsing data.

Answer: C

Explanation:

To enable personalized product recommendations using generative AI, the foundational step for Universal Containers (UC) is collecting and analyzing browsing data (Option C). Personalized recommendations depend on understanding user behavior, which requires structured data about their browsing history. Without this data, the AI model lacks the context needed to generate relevant suggestions.

? Data Collection: UC must first aggregate browsing data (e.g., pages visited, products viewed, session duration) to build a dataset that reflects user preferences.

? Data Analysis: Analyzing this data identifies patterns (e.g., frequently viewed categories) that inform how prompts should be structured to retrieve relevant recommendations.

? Grounding in Data: Salesforce??s Prompt Templates rely on grounding data to generate accurate outputs. Without analyzing browsing data, the prompt template cannot reference meaningful insights for personalization.

Options A and D are incorrect because:

? Universal recommendations (A) ignore personalization, which is the core requirement.

? Writing a response script (D) addresses chatbot interaction design, not the accuracy of recommendations.

References:

? Salesforce Agentforce Specialist Certification Guide: Highlights the importance of grounding prompts in relevant data sources to ensure accuracy.

? Trailhead Module: "Einstein for Developers" emphasizes data preparation as a prerequisite for effective AI-driven personalization.

? Salesforce Help Documentation: Recommends analyzing user behavior data to tailor generative AI outputs in commerce use cases.

NEW QUESTION 164

Universal Containers is interested in using Call Explorer to quickly gain insights from meetings recorded by its sales team.

What should the Agentforce Specialist be aware of before enabling this feature?

- A. Call Explorer operates independently of Salesforce Knowledge, requiring no prior setup.
- B. Custom Call Explorer actions need to be built before it can be configured.
- C. Call Explorer requires the Einstein Conversation Insights permission set to be enabled.

Answer: C

Explanation:

Before enabling Call Explorer, the Salesforce Agentforce Specialist must ensure that the Einstein Conversation Insights permission set is assigned to users (Option C). Call Explorer is a feature within Einstein Conversation Insights (ECI) that analyzes meeting recordings to surface trends, keywords, and actionable insights.

Key Considerations:

? Permission Set Requirement:

? Why Other Options Are Incorrect:

References:

? Salesforce Einstein Conversation Insights Guide: Explicitly states that the Einstein Conversation Insights permission set is required to access Call Explorer.

? Trailhead Module: "Einstein Conversation Insights Basics" outlines permission prerequisites for enabling call analytics.

? Salesforce Help Documentation: Confirms that Call Explorer functionality is governed by ECI permissions.

NEW QUESTION 166

The Agentforce 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 Agentforce 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

NEW QUESTION 168

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

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

Answer: C

Explanation:

Salesforce has restrictions on which objects and related lists can be used for grounding prompt templates. This is likely due to security and privacy concerns related to user data. While it might seem intuitive to use the User related list to provide context to the LLM, Salesforce prevents this to ensure that sensitive user information is not inadvertently exposed or misused.

Therefore, the Agentforce Specialist needs to explore alternative ways to incorporate the necessary user information into the prompt template, perhaps by using other related objects or fields that are supported.

NEW QUESTION 172

Universal Containers (UC) wants to enable its sales reps to explore opportunities that are similar to previously won opportunities by entering the utterance, "Show me other opportunities like this one."

How should UC achieve this with Agents?

- A. Use the standard Agent action.
- B. Create a custom Agent action calling a flow.
- C. Create a custom Agent action calling an Apex class.

Answer: A

Explanation:

Universal Containers can achieve the request to explore similar opportunities by using the standard Copilot action. Agent has built-in actions to handle natural language queries, such as "Show me other opportunities like this one." The standard action will process the query and return results based on predefined matching criteria like opportunity details and past Closed Won deals.

This approach avoids the need to create custom flows or Apex classes, leveraging out-of-the-box functionality.

For further details, refer to Agent for Sales documentation regarding standard actions and natural language processing.

NEW QUESTION 177

Universal Containers implements Custom Agent Actions to enhance its customer service operations. The development team needs to understand the core

components of a Custom Agent Action to ensure proper configuration and functionality. What should the development team review in the Custom Agent Action configuration to identify one of the core components of a Custom Agent Action?

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

Answer: B

Explanation:

Comprehensive and Detailed In-Depth Explanation:UC??s development team needs to identify a core component of a Custom Agent Action in Agent Builder. Let??s assess the options.

? Option A: Action Triggers"Action Triggers" isn??t a term used in Agentforce Custom Agent Action configuration. Actions are invoked by topics or plans, not standalone triggers, making this incorrect.

? Option B: InstructionsInstructions are a core component of a Custom Agent Action in Agentforce. Defined in Agent Builder, they guide the Atlas Reasoning Engine on how to execute the action (e.g., what to do with inputs, how to process data). Reviewing the instructions helps the team understand the action??s purpose and logic, making this the correct answer.

? Option C: Output TypesWhile outputs are part of an action??s result, "Output Types" isn??t a distinct configuration element in Agent Builder. Outputs are determined by the action??s execution (e.g., Flow or Apex), not a separate setting, making this less core and incorrect.

Why Option B is Correct:Instructions are a fundamental component of Custom Agent Actions, providing the AI??s execution directives, as per Salesforce documentation.

References:

? Salesforce Agentforce Documentation: Agent Builder > Custom Actions – Highlights instructions as key.

? Trailhead: Build Agents with Agentforce – Details configuring actions with instructions.

? Salesforce Help: Create Custom Actions – Confirms instructions?? role.

NEW QUESTION 179

What is best practice when refining Agent 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 Agent 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 Agent documentation on building and refining custom actions.

NEW QUESTION 181

Universal Containers has implemented an agent that answers questions based on Knowledge articles. Which topic and Agent Action will be shown in the Agent Builder?

- A. General Q&A topic and Knowledge Article Answers action.
- B. General CRM topic and Answers Questions with LLM Action.
- C. General FAQ topic and Answers Questions with Knowledge Action.

Answer: C

Explanation:

Comprehensive and Detailed In-Depth Explanation:UC??s agent answers questions using Knowledge articles, configured in Agent Builder. Let??s identify the topic and action.

? Option A: General Q&A topic and Knowledge Article Answers action."General Q&A" is not a standard topic name in Agentforce, and "Knowledge Article Answers" isn??t a predefined action. This lacks specificity and doesn??t match documentation, making it incorrect.

? Option B: General CRM topic and Answers Questions with LLM Action."General CRM" isn??t a default topic, and "Answers Questions with LLM" suggests raw LLM responses, not Knowledge-grounded ones. This doesn??t align with the Knowledge focus, making it incorrect.

? Option C: General FAQ topic and Answers Questions with Knowledge Action.In Agent Builder, the "General FAQ" topic is a common default or starting point for question-answering agents. The "Answers Questions with Knowledge" action (sometimes styled as "Answer with Knowledge") is a prebuilt action that retrieves and grounds responses with Knowledge articles. This matches UC??s implementation and is explicitly supported in documentation, making it the correct answer.

Why Option C is Correct:"General FAQ" and "Answers Questions with Knowledge" are the standard topic-action pair for Knowledge-based question answering in Agentforce, per Salesforce resources.

References:

? Salesforce Agentforce Documentation: Agent Builder > Actions – Lists "Answers Questions with Knowledge."

? Trailhead: Build Agents with Agentforce – Describes FAQ topics with Knowledge actions.

? Salesforce Help: Knowledge in Agentforce – Confirms this configuration.

NEW QUESTION 185

When configuring a prompt template, an Agentforce Specialist previews the results of the prompt template they've written. They see two distinct text outputs: Resolution and Response. Which information does the Resolution text provide?

- A. It shows the full text that is sent to the Trust Layer.
- B. It shows the response from the LLM based on the sample record.
- C. It shows which sensitive data is masked before it is sent to the LLM.

Answer: B

Explanation:

Comprehensive and Detailed In-Depth Explanation: In Salesforce Agentforce, when previewing a prompt template, the interface displays two outputs: Resolution and Response. These terms relate to how the prompt is processed and evaluated, particularly in the context of the Einstein Trust Layer, which ensures AI safety, compliance, and auditability. The Resolution text specifically refers to the full text that is sent to the Trust Layer for processing, monitoring, and governance (Option A). This includes the constructed prompt (with grounding data, instructions, and variables) as it's submitted to the large language model (LLM), along with any Trust Layer interventions (e.g., masking, filtering) applied before or after LLM processing. It's a comprehensive view of the input/output flow that the Trust Layer captures for auditing and compliance purposes.

? Option B: The "Response" output in the preview shows the LLM's generated text based on the sample record, not the Resolution. Resolution encompasses more than just the LLM response—it includes the entire payload sent to the Trust Layer.

? Option C: While the Trust Layer does mask sensitive data (e.g., PII) as part of its guardrails, the Resolution text doesn't specifically isolate "which sensitive data is masked." Instead, it shows the full text, including any masked portions, as processed by the Trust Layer—not a separate masking log.

? Option A: This is correct, as Resolution provides a holistic view of the text sent to the Trust Layer, aligning with its role in monitoring and auditing the AI interaction.

Thus, Option A accurately describes the purpose of the Resolution text in the prompt template preview.

References:

? Salesforce Agentforce Documentation: "Preview Prompt Templates" (Salesforce Help: https://help.salesforce.com/s/articleView?id=sf.agentforce_prompt_preview.htm&type=5)

? Salesforce Einstein Trust Layer Documentation: "Trust Layer Outputs" (https://help.salesforce.com/s/articleView?id=sf.einstein_trust_layer.htm&type=5)

NEW QUESTION 189

Universal Containers wants to reduce overall customer support handling time by minimizing the time spent typing routine answers for common questions in-chat, and reducing the post-chat analysis by suggesting values for case fields. Which combination of Agentforce for Service features enables this effort?

- A. Einstein Reply Recommendations and Case Classification
- B. Einstein Reply Recommendations and Case Summaries
- C. Einstein Service Replies and Work Summaries

Answer: B

Explanation:

Comprehensive and Detailed In-Depth Explanation: Universal Containers (UC) aims to streamline customer support by addressing two goals: reducing in-chat typing time for routine answers and minimizing post-chat analysis by auto-suggesting case field values. In Salesforce Agentforce for Service, Einstein Reply Recommendations and Case Classification (Option A) are the ideal combination to achieve this.

? Einstein Reply Recommendations: This feature uses AI to suggest pre-formulated responses based on chat context, historical data, and Knowledge articles. By providing agents with ready-to-use replies for common questions, it significantly reduces the time spent typing routine answers, directly addressing UC's first goal.

? Case Classification: This capability leverages AI to analyze case details (e.g., chat transcripts) and suggest values for case fields (e.g., Subject, Priority, Resolution) during or after the interaction. By automating field population, it reduces post-chat analysis time, fulfilling UC's second goal.

? Option B: While "Einstein Reply Recommendations" is correct for the first part, "Case Summaries" generates a summary of the case rather than suggesting specific field values. Summaries are useful for documentation but don't directly reduce post-chat field entry time.

? Option C: "Einstein Service Replies" is not a distinct, documented feature in Agentforce (possibly a distractor for Reply Recommendations), and "Work Summaries" applies more to summarizing work orders or broader tasks, not case field suggestions in a chat context.

? Option A: This combination precisely targets both in-chat efficiency (Reply Recommendations) and post-chat automation (Case Classification).

Thus, Option A is the correct answer for UC's needs.

References:

? Salesforce Agentforce Documentation: "Einstein Reply Recommendations" (Salesforce Help: https://help.salesforce.com/s/articleView?id=sf.einstein_reply_recommendations.htm&type=5)

? Salesforce Agentforce Documentation: "Case Classification" (Salesforce Help: https://help.salesforce.com/s/articleView?id=sf.case_classification.htm&type=5)

? Trailhead: "Agentforce for Service" (<https://trailhead.salesforce.com/content/learn/modules/agentforce-for-service>)

NEW QUESTION 192

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

Agentforce-Specialist Practice Exam Features:

- * Agentforce-Specialist Questions and Answers Updated Frequently
- * Agentforce-Specialist Practice Questions Verified by Expert Senior Certified Staff
- * Agentforce-Specialist Most Realistic Questions that Guarantee you a Pass on Your First Try
- * Agentforce-Specialist Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The Agentforce-Specialist Practice Test Here](#)