

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

How should an organization use the Einstein Trust layer to audit, track, and view masked data?

- A. Utilize the audit trail that captures and stores all LLM submitted prompts in Data Cloud.
- B. In Setup, use Prompt Builder to send a prompt to the LLM requesting for the masked data.
- C. Access the audit trail in Setup and export all user-generated prompts.

Answer: A

Explanation:

The Einstein Trust Layer is designed to ensure transparency, compliance, and security for organizations leveraging Salesforce's AI and generative AI capabilities. Specifically, for auditing, tracking, and viewing masked data, organizations can utilize:

? Audit Trail in Data Cloud: The audit trail captures and stores all prompts submitted

to large language models (LLMs), ensuring that sensitive or masked data interactions are logged. This allows organizations to monitor and audit all AI-generated outputs, ensuring that data handling complies with internal and regulatory guidelines. The Data Cloud provides the infrastructure for managing and accessing this audit data.

? Why not B? Using Prompt Builder in Setup to send prompts to the LLM is for creating and managing prompts, not for auditing or tracking data. It does not interact directly with the audit trail functionality.

? Why not C? Although the audit trail can be accessed in Setup, the user-generated prompts are primarily tracked in the Data Cloud for broader control, auditing, and analysis. Setup is not the primary tool for exporting or managing these audit logs. More information on auditing AI interactions can be found in the Salesforce AI Trust Layer documentation, which outlines how organizations can manage and track generative AI interactions securely.

NEW QUESTION 4

Universal Containers is using Agentforce for Sales to find similar opportunities to help close deals faster. The team wants to understand the criteria used by the Agent to match opportunities. What is one criterion that Agentforce for Sales uses to match similar opportunities?

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

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation:UC uses Agentforce for Sales to identify similar opportunities, aiding deal closure. Let's determine a criterion used by the "Find Similar Opportunities" feature.

? Option A: Matched opportunities have a status of Closed Won from the last 12 months.Agentforce for Sales analyzes historical data to find similar opportunities, prioritizing "Closed Won" deals as successful examples. Documentation specifies a 12-month lookback period for relevance, ensuring recent, applicable matches. This is a key criterion, making it the correct answer.

? Option B: Matched opportunities are limited to the same account.While account context may factor in, Agentforce doesn't restrict matches to the same account—it considers broader patterns across opportunities (e.g., industry, deal size). This is too narrow and incorrect.

? Option C: Matched opportunities were created in the last 12 months.Creation date isn't a primary criterion—status (e.g., Closed Won) and recency of closure matter more. This doesn't align with documented behavior, making it incorrect.

Why Option A is Correct:"Closed Won" status within 12 months is a documented criterion for Agentforce's similarity matching, providing actionable insights for deal closure.

References:

? Salesforce Agentforce Documentation: Agentforce for Sales > Find Similar Opportunities – Specifies Closed Won, 12-month criterion.

? Trailhead: Explore Agentforce Sales Agents – Details opportunity matching logic.

? Salesforce Help: Sales Features in Agentforce – Confirms historical success focus.

NEW QUESTION 5

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 6

What is a valid use case for Data Cloud retrievers?

A. Returning relevant data from the vector database to augment a prompt.

B. Grounding data from external websites to augment a prompt with RAG.

C. Modifying and updating data within the source systems connected to Data Cloud.

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation:Salesforce Data Cloud integrates with Agentforce to provide real-time, unified data access for AI-driven applications. Data Cloud retrievers are specialized components that fetch relevant data from Data Cloud's vector database—a storage system optimized for semantic search and retrieval—to enhance agent responses or actions. A valid use case, as described in Option A, is using these retrievers to return pertinent data (e.g., customer purchase history, support tickets) from the vector database to augment a prompt. This process, often part of Retrieval-Augmented Generation (RAG), allows the LLM to generate more accurate, context-aware responses by grounding its output in structured, searchable data stored in Data Cloud.

? Option B: Grounding data from external websites is not a primary function of Data Cloud retrievers. While RAG can incorporate external data, Data Cloud retrievers specifically work with data within Salesforce's ecosystem (e.g., the vector database or harmonized data lakes), not arbitrary external websites. This makes B incorrect.

? Option C: Data Cloud retrievers are read-only mechanisms designed for data retrieval, not for modifying or updating source systems. Updates to source systems are handled by other Salesforce tools (e.g., Flows or Apex), not retrievers.

Option A is correct because it aligns with the core purpose of Data Cloud retrievers: enhancing prompts with relevant, vectorized data from within Salesforce Data Cloud.

References:

? Salesforce Data Cloud Documentation: "Data Cloud for Agentforce" (Salesforce Help:

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

? Trailhead: "Data Cloud Basics" module (<https://trailhead.salesforce.com/content/learn/modules/data-cloud-basics>)

NEW QUESTION 7

What is the role of the large language model (LLM) in understanding intent and executing an Agent Action?

A. Find similar requested topics and provide the actions that need to be executed.

B. Identify the best matching topic and actions and correct order of execution.

C. Determine a user's topic access and sort actions by priority to be executed.

Answer: B

Explanation:

Comprehensive and Detailed In-Depth Explanation:In Agentforce, the large language model (LLM), powered by the Atlas Reasoning Engine, interprets user

requests and drives Agent Actions. Let??s evaluate its role.

? Option A: Find similar requested topics and provide the actions that need to be executed.While the LLM can identify similar topics, its role extends beyond merely finding them—it matches intents to specific topics and determines execution. This option understates the LLM??s responsibility for ordering actions, making it incomplete and incorrect.

? Option B: Identify the best matching topic and actions and correct order of execution.The LLM analyzes user input to understand intent, matches it to the best-fitting topic (configured in Agent Builder), and selects associated actions. It

also determines the correct sequence of execution based on the agent??s plan (e.g., retrieve data before updating a record). This end-to-end process—from intent recognition to action orchestration—is the LLM??s core role in Agentforce, making this the correct answer.

? Option C: Determine a user??s topic access and sort actions by priority to be executed.Topic access is governed by Salesforce permissions (e.g., user profiles), not the LLM. While the LLM prioritizes actions within its plan, its primary role is intent matching and execution ordering, not access control, making this incorrect.

Why Option B is Correct:The LLM??s role in identifying topics, selecting actions, and ordering execution is central to Agentforce??s autonomous functionality, as detailed in Salesforce documentation.

References:

? Salesforce Agentforce Documentation: Atlas Reasoning Engine – Outlines LLM??s intent and action handling.

? Trailhead: Understand Agentforce Technology – Explains topic matching and execution.

? Salesforce Help: Agentforce Actions – Confirms LLM??s role in orchestrating responses.

NEW QUESTION 8

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 whereisToxicityDetected 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 9

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 10

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 10

An Agentforce is setting up a new org and needs to ensure that users can create and execute prompt templates. The Agentforce Specialist is unsure which roles are necessary for these tasks.

Which permission sets should the Agentforce Specialist assign to users who need to create and execute prompt templates?

- A. Prompt Template Manager for creating templates and Data Cloud Admin for executing templates
- B. Prompt Template Manager for creating templates and Prompt Template User for executing templates
- C. Data Cloud Admin for creating templates and Prompt Template User for executing templates

Answer: B

Explanation:

To effectively manage and use prompt templates, two distinct permission sets are required:

? Prompt Template Manager: This permission set allows users to create prompt templates. It provides the necessary access to define templates, which can be shared and utilized across the organization.

? Prompt Template User: This permission set is designed for users who need to execute the templates. It provides the ability to interact with pre-designed prompts and generate outcomes based on these templates.

The Data Cloud Admin permission set is not directly relevant to creating or executing prompt templates but is more focused on managing the Data Cloud.

Reference:

"Permissions and Access for Prompt Templates | Salesforce Trailhead" .

NEW QUESTION 11

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

- A. Model Builder
- B. Agent Builder
- C. Prompt Builder

Answer: C

Explanation:

Comprehensive and Detailed In-Depth Explanation: The hotel resort team needs an AI-generated guest summary with recommendations, displayed exclusively on the contact record page. Let's assess the options.

? Option A: Model Builder Model Builder in Salesforce creates custom predictive AI models (e.g., for scoring or classification) using Data Cloud or Einstein Platform data. It's not designed for generating text summaries or embedding them on record pages, making it incorrect.

? Option B: Agent Builder Agent Builder in Agentforce Studio creates autonomous AI agents for tasks like lead qualification or customer service. While agents can provide summaries, they operate in conversational interfaces (e.g., chat), not as static content on a record page. This doesn't meet the location-specific requirement, making it incorrect.

? Option C: Prompt Builder Einstein Prompt Builder allows creation of prompt templates that generate text (e.g., summaries, recommendations) using Generative AI. The template can pull data from contact records (e.g., activity preferences) and be embedded as a Lightning component on the contact record page via a Flow or Lightning App Builder. This ensures the summary is available only where specified, meeting the team's needs perfectly and making it the correct answer.

Why Option C is Correct: Prompt Builder's ability to generate contextual summaries and integrate them into specific record pages via Lightning components aligns with the team's requirements, as supported by Salesforce documentation.

References:

? Salesforce Agentforce Documentation: Prompt Builder > Embedding Prompts – Details placement on record pages.

? Trailhead: Build Prompt Templates in Agentforce – Covers summaries from object data.

? Salesforce Help: Customize Record Pages with AI – Confirms Prompt Builder integration.

NEW QUESTION 13

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

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

Answer: A

Explanation:

When using related list merge fields in a prompt template associated with the Account object in Prompt Builder, the Activities related list is not supported due to it

being a polymorphic field. Polymorphic fields can reference multiple different types of objects, which makes them incompatible with some merge field operations in prompt generation.

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

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

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

NEW QUESTION 14

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 18

Which object stores the conversation transcript between the customer and the agent?

- A. Messaging End User
- B. Messaging Session
- C. Case

Answer: B

Explanation:

Why is "Messaging Session" the correct answer?

In Agentforce, the Messaging Session object stores the conversation transcript between the customer and the agent.

Key Features of the Messaging Session Object:

? Stores the Entire Customer-Agent Conversation

? Supports AI-Powered Work Summaries

? Links with Service Cloud for Case Resolution

Why Not the Other Options?

* A. Messaging End User

? Incorrect because this object stores details about the customer (e.g., name, contact details) but not the conversation transcript.

* C. Case

? Incorrect because Cases store structured service requests but do not contain raw conversation transcripts.

? Instead, cases may reference the Messaging Session object.

Agentforce Specialist References

? Salesforce AI Specialist Material confirms that Messaging Sessions store chat conversations and support Einstein Work Summaries.

NEW QUESTION 20

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 25

Universal Containers (UC) wants to use the Draft with Einstein feature in Sales Cloud to create a personalized introduction email. After creating a proposed draft email, which predefined adjustment should UC choose to revise the draft with a more casual tone?

- A. Make Less Formal
- B. Enhance Friendliness

C. Optimize for Clarity

Answer: A

Explanation:

When Universal Containers uses the Draft with Einstein feature in Sales Cloud to create a personalized email, the predefined adjustment to Make Less Formal is the correct option to revise the draft with a more casual tone. This option adjusts the wording of the draft to sound less formal, making the communication more approachable while still maintaining professionalism.

? Enhance Friendliness would make the tone more positive, but not necessarily more casual.

? Optimize for Clarity focuses on making the draft clearer but doesn't adjust the tone. For more details, see Salesforce documentation on Einstein-generated email drafts and tone adjustments.

NEW QUESTION 29

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 34

A customer service representative is looking at a custom object that stores travel information. They recently received a weather alert and now need to cancel flights for the customers that are related to this Itinerary. The representative needs to review the Knowledge articles about canceling and rebooking the customer flights. Which Agentforce capability helps the representative accomplish this?

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

Answer: B

Explanation:

Comprehensive and Detailed In-Depth Explanation:The scenario involves a customer service representative needing to cancel flights due to a weather alert and review existing Knowledge articles for guidance on canceling and rebooking. Agentforce provides capabilities to streamline such tasks. The most suitable option is Option B, which allows the agent to "execute tasks based on available actions" (e.g., canceling flights via a predefined action) while "answering questions using information from accessible Knowledge articles." This capability leverages Agentforce's ability to integrate Knowledge articles into the agent's responses, enabling the representative to ask questions (e.g., "How do I cancel a flight?") and receive AI-generated answers grounded in approved Knowledge content. Simultaneously, the agent can trigger actions (e.g., a Flow to update the custom object) to perform the cancellations, meeting all requirements efficiently.

? Option A: Invoking a Flow to call external data and create a Knowledge article is unnecessary. The representative needs to review existing articles, not create new ones, and there's no indication external data is required for this task.

? Option B: This is correct. It combines task execution (canceling flights) with Knowledge article retrieval, aligning with the representative's need to act and seek guidance from existing content.

? Option C: Generating a new Knowledge article based on prompts is not relevant.

The representative needs to use existing articles, not author new ones, especially in a time-sensitive weather alert scenario.

Option B best supports the representative's workflow in Agentforce.

References:

? Salesforce Agentforce Documentation: "Knowledge Replies and Actions" (Salesforce Help:

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

&type=5)

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

NEW QUESTION 36

Universal Containers' service team wants to customize the standard case summary response from Agentforce. What should the Agentforce Specialist do to achieve this?

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

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation: UC's service team seeks to customize the standard case summary response provided by Agentforce. Let's assess the options for tailoring this output.

? Option A: Create a custom Record Summary prompt template for the Case

object. In Prompt Builder, the standard Record Summary prompt template generates summaries for objects like Case. To customize it, the Agentforce Specialist can create a new custom prompt template, specifying the Case object as the source, and adjust the instructions (e.g., tone, fields included) to meet UC's needs. This new template can then be invoked by an agent or flow, providing a tailored summary. This approach offers full control and aligns with Salesforce's customization process, making it the correct answer.

? Option B: Summarize the Case with a standard Agent action. Standard Agent actions (e.g., "Answer Questions") don't specifically target case summarization—they're broader in scope. There's no out-of-the-box "Summarize Case" action that allows customization of the response format, making this insufficient and incorrect.

? Option C: Customize the standard Record Summary template for the Case object. Standard prompt templates in Prompt Builder (e.g., Record Summary) are read-only and cannot be directly edited. Customization requires cloning or creating a new template, not modifying the standard one, making this incorrect.

Why Option A is Correct: Creating a custom Record Summary prompt template allows full customization of the case summary, leveraging Prompt Builder's flexibility, as per Salesforce best practices.

References:

? Salesforce Agentforce Documentation: Prompt Builder > Custom Templates – Details creating custom summaries.

? Trailhead: Build Prompt Templates in Agentforce – Explains customizing standard outputs.

? Salesforce Help: Record Summaries with AI – Recommends custom templates for tailored results.

NEW QUESTION 40

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-Grounded 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 42

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 44

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 48

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.
- D. 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.

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 50

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 51

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 54

What should Universal Containers consider when deploying an Agentforce Service Agent with multiple topics and Agent Actions to production?

- A. Deploy agent components without a test run in staging, relying on production data for reliable result
- B. Sandbox configuration alone ensures seamless production deployment.
- C. Ensure all dependencies are included, Apex classes meet 75% test coverage, and configuration settings are aligned with production
- D. Plan for version management and post- deployment activation.
- E. Deploy flows or Apex after agents, topics, and Agent Actions to avoid deployment failures and potential production agent issues requiring complete redeployment.

Answer: B

Explanation:

Comprehensive and Detailed In-Depth Explanation:UC is deploying an Agentforce Service Agent with multiple topics and actions to production. Let's assess deployment considerations.

? Option A: Deploy agent components without a test run in staging, relying on production data for reliable results. Sandbox configuration alone ensures seamless production deployment.Skipping staging tests is risky and against best practices. Sandbox configuration doesn't guarantee production success without validation, making this incorrect.

? Option B: Ensure all dependencies are included, Apex classes meet 75% test coverage, and configuration settings are aligned with production. Plan for version management and post-deployment activation.This is a comprehensive approach: dependencies (e.g., flows, Apex) must be deployed, Apex requires 75% coverage, and production settings (e.g., permissions, channels) must align. Version management tracks changes, and post-deployment activation ensures controlled rollout. This aligns with Salesforce deployment best practices for Agentforce, making it the correct answer.

? Option C: Deploy flows or Apex after agents, topics, and Agent Actions to avoid deployment failures and potential production agent issues requiring complete redeployment.Deploying components separately risks failures (e.g., actions needing flows failing). All components should deploy together for consistency, making this incorrect.

Why Option B is Correct:Option B covers all critical deployment considerations for a robust Agentforce rollout, as per Salesforce guidelines.

References:

? Salesforce Agentforce Documentation: Deploy Agents to Production – Lists dependencies and coverage.

? Trailhead: Deploy Agentforce Agents – Emphasizes testing and activation planning.

? Salesforce Help: Agentforce Deployment Best Practices – Confirms comprehensive approach.

NEW QUESTION 59

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 62

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 65

Universal Containers wants to leverage the Record Snapshots grounding feature in a prompt template. What preparations are required?

- A. Configure page layout of the master record type.
- B. Create a field set for all the fields to be grounded.
- C. Enable and configure dynamic form for the object.

Answer: B

Explanation:

Comprehensive and Detailed In-Depth Explanation: Universal Containers (UC) aims to use Record Snapshots grounding in a prompt template to provide context from a specific record. Let's evaluate the preparation steps.

? Option A: Configure page layout of the master record type. While page layouts define field visibility for users, Record Snapshots grounding relies on field accessibility at the object level, not the layout. The AI accesses data based on permissions and configuration, not layout alone, making this insufficient and incorrect.

? Option B: Create a field set for all the fields to be grounded. Record Snapshots in Prompt Builder allow grounding with fields from a record, but you must specify which fields to include. Creating a field set is a recommended preparation step—it groups the fields (e.g., from the object) to be passed to the prompt template, ensuring the AI has the right data. This is a documented best practice for controlling snapshot scope, making it the correct answer.

? Option C: Enable and configure dynamic form for the object. Dynamic Forms enhance UI flexibility but aren't required for Record Snapshots grounding. The feature pulls data directly from the object, not the form configuration, making this irrelevant and incorrect.

Why Option B is Correct: Creating a field set ensures the prompt template uses the intended fields for grounding, a key preparation step per Salesforce documentation.

References:

? Salesforce Agentforce Documentation: Prompt Builder > Record Snapshots – Recommends field sets for grounding.

? Trailhead: Ground Your Agentforce Prompts – Details field set preparation.

? Salesforce Help: Set Up Record Snapshots – Confirms field set usage.

NEW QUESTION 67

Universal Containers (UC) wants to use Generative AI Salesforce functionality to reduce Service Agent handling time by providing recommended replies based on the existing Knowledge articles. On which AI capability should UC train the service agents?

- A. Service Replies
- B. Case Replies
- C. Knowledge Replies

Answer: C

Explanation:

Comprehensive and Detailed In-Depth Explanation: Salesforce Agentforce leverages generative AI to enhance service agent efficiency, particularly through capabilities that generate recommended replies. In this scenario, Universal Containers

aims to reduce handling time by providing replies based on existing Knowledge articles, which are a core component of Salesforce Knowledge. The Knowledge Replies capability is specifically designed for this purpose—it uses generative AI to analyze Knowledge articles, match them to the context of a customer inquiry (e.g., a case or chat), and suggest relevant, pre-formulated responses for service agents to use or adapt. This aligns directly with UC's goal of leveraging existing content to streamline agent workflows.

? Option A (Service Replies): While "Service Replies" might sound plausible, it is not a specific, documented capability in Agentforce. It appears to be a generic distractor and does not tie directly to Knowledge articles.

? Option B (Case Replies): "Case Replies" is not a recognized AI capability in Agentforce either. While replies can be generated for cases, the focus here is on Knowledge article integration, which points to Knowledge Replies.

? Option C (Knowledge Replies): This is the correct capability, as it explicitly connects generative AI with Knowledge articles to produce recommended replies, reducing agent effort and handling time.

Training service agents on Knowledge Replies ensures they can effectively use AI-suggested responses, review them for accuracy, and integrate them into their workflows, fulfilling UC's objective.

References:

? Salesforce Agentforce Documentation: "Knowledge Replies for Service Agents" (Salesforce Help:

[https://help.salesforce.com/s/articleView?id=sf.agentforce_knowledge_replies.htm](https://help.salesforce.com/s/articleView?id=sf.agentforce_knowledge_replies.htm&type=5)

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

NEW QUESTION 70

Universal Containers (UC) wants to build an Agentforce Service Agent that provides the latest, active, and relevant policy and compliance information to customers. The agent must:

- ? Semantically search HR policies, compliance guidelines, and company procedures.
 - ? Ensure responses are grounded on published Knowledge.
 - ? Allow Knowledge updates to be reflected immediately without manual reconfiguration.
- What should UC do to ensure the agent retrieves the right information?

- A. Enable the agent to search all internal records and past customer inquiries.
- B. Set up an Agentforce Data Library to store and index policy documents for AI retrieval.
- C. Manually add policy responses into the AI model to prevent hallucinations.

Answer: B

Explanation:

Comprehensive and Detailed In-Depth Explanation: UC requires an Agentforce Service Agent to deliver accurate, up-to-date policy and compliance info with specific criteria. Let's evaluate.

? Option A: Enable the agent to search all internal records and past customer inquiries. Searching all records and inquiries risks irrelevant or outdated responses, conflicting with the need for published Knowledge grounding and immediate updates. This lacks specificity, making it incorrect.

? Option B: Set up an Agentforce Data Library to store and index policy documents for AI retrieval. The Agentforce Data Library integrates with Salesforce Knowledge, indexing HR policies, compliance guidelines, and procedures for semantic search. It ensures grounding in published Knowledge articles, and updates (e.g., new article versions) are reflected instantly without reconfiguration, as the library syncs with Knowledge automatically. This meets all UC requirements, making it the correct answer.

? Option C: Manually add policy responses into the AI model to prevent hallucinations. Manually embedding responses into the model isn't feasible—Agentforce uses pretrained LLMs, not custom training. It also doesn't support real-time updates, making this incorrect.

Why Option B is Correct: The Data Library meets all criteria—semantic search, Knowledge grounding, and instant updates—per Salesforce's recommended approach.

References:

? Salesforce Agentforce Documentation: Data Library > Knowledge Integration – Details indexing and updates.

? Trailhead: Build Agents with Agentforce – Covers Data Library for accurate responses.

? Salesforce Help: Grounding with Knowledge – Confirms real-time sync.

NEW QUESTION 75

What is the importance of Action Instructions when creating a custom Agent action?

- A. Action Instructions define the expected user experience of an action.
- B. Action Instructions tell the user how to call this action in a conversation.
- C. Action Instructions tell the large language model (LLM) which action to use.

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation: In Salesforce Agentforce, custom Agent actions are designed to enable AI-driven agents to perform specific tasks within a conversational context. Action Instructions are a critical component when creating these actions because they define the expected user experience by outlining how the action should behave, what it should accomplish, and how it interacts with the end user. These instructions act as a blueprint for the action's functionality, ensuring that it aligns with the intended outcome and provides a consistent, intuitive experience for users interacting with the agent. For example, if the action is to "schedule a meeting," the Action Instructions might specify the steps (e.g., gather date and time, confirm with the user) and the tone (e.g., professional, concise), shaping the user experience.

? Option B: While Action Instructions might indirectly influence how a user invokes an action (e.g., by making it clear what inputs are needed), they are not primarily about telling the user how to call the action in a conversation. That's more related to user training or interface design, not the instructions themselves.

? Option C: The large language model (LLM) relies on prompts, parameters, and grounding data to determine which action to execute, not the Action Instructions directly. The instructions guide the action's design, not the LLM's decision-making process at runtime.

Thus, Option A is correct as it emphasizes the role of Action Instructions in defining the user experience, which is foundational to creating effective custom Agent actions in Agentforce.

References:

? Salesforce Agentforce Documentation: "Create Custom Agent Actions" (Salesforce Help:

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

? Trailhead: "Agentforce Basics" module (<https://trailhead.salesforce.com/content/learn/modules/agentforce-basics>)

NEW QUESTION 79

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 81

Universal Containers (UC) is creating a new custom prompt template to populate a field with generated output. UC enabled the Einstein Trust Layer to ensure AI Audit data is captured and monitored for adoption and possible enhancements. Which prompt template type should UC use and which consideration should UC review?

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

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation:Salesforce Agentforce provides various prompt template types to support AI-driven tasks, such as generating text or populating fields. In this case, UC needs a custom prompt template to populate a field with generated output, which directly aligns with the Field Generation prompt template type. This type is designed to use generative AI to create field values (e.g., summaries, descriptions) based on input data or prompts, making it the ideal choice for UC's requirement. Additionally, UC has enabled the Einstein Trust Layer, a governance framework that ensures AI outputs are safe, explainable, and auditable, capturing AI Audit data for monitoring adoption and identifying improvement areas.

The consideration UC should review is whether Dynamic Fields is enabled. Dynamic Fields allow the prompt template to incorporate variable data from Salesforce records (e.g., case details, customer info) into the prompt, ensuring the generated output is contextually relevant to each record. This is critical for field population tasks, as static prompts wouldn't adapt to record-specific needs. The Einstein Trust Layer further benefits from this, as it can track how dynamic inputs influence outputs for audit purposes.

? Option A: Correct. "Field Generation" matches the use case, and "Dynamic Fields" is a key consideration to ensure flexibility and auditability with the Trust Layer.

? Option B: "Field Generation" is correct, but "Dynamic Forms" is unrelated.

Dynamic Forms is a UI feature for customizing page layouts, not a prompt template setting, making this option incorrect.

? Option C: "Flex" templates are more general-purpose and not specifically tailored for field population tasks. While Dynamic Fields could apply, Field Generation is the better fit for UC's stated goal.

Option A is the best choice, as it pairs the appropriate template type (Field Generation) with a relevant consideration (Dynamic Fields) for UC's scenario with the Einstein Trust Layer.

References:

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

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

? Trailhead: "Build Prompt Templates for Agentforce" (<https://trailhead.salesforce.com/content/learn/modules/build-prompt-templates-for-agentforce>)

NEW QUESTION 86

What is the main benefit of using a Knowledge article in an Agentforce Data Library?

- A. Only the retriever for Knowledge articles allows for agents to access Knowledge from both inside the platform and on a customer's website.
- B. It provides a structured, searchable repository of approved documents so the agent can retrieve reliable information for each inquiry..
- C. The retriever for Knowledge articles has better accuracy and performance than the default retriever.

Answer: B

Explanation:

Why is "A structured, searchable repository of approved documents" the correct answer?

Using a Knowledge Article in an Agentforce Data Library ensures that agents can quickly access reliable and pre-approved information during customer interactions.

Key Benefits of Knowledge Articles in an Agentforce Data Library:

- ? Ensures Information Accuracy and Consistency
- ? Improves Searchability and AI-Grounded Responses
- ? Enhances Customer Support and Agent Productivity

Why Not the Other Options?

* A. Only the retriever for Knowledge articles allows for agents to access Knowledge from both inside the platform and on a customer's website.

? Incorrect because other retrievers (e.g., standard Salesforce Data Cloud retrievers) can also provide knowledge access.

? Knowledge articles can be accessed via multiple retrieval mechanisms, not just one specific retriever.

* C. The retriever for Knowledge articles has better accuracy and performance than the default retriever.

? Incorrect because retriever accuracy depends on indexing and search configuration, not the article type.

? The default retriever works just as efficiently when properly configured.

Agentforce Specialist References

? Salesforce AI Specialist Material confirms that Knowledge articles provide structured, searchable, and approved information for AI-grounded responses.

NEW QUESTION 89

An Agentforce implements Einstein Sales Emails for a sales team. The team wants to send personalized follow-up emails to leads based on their interactions and data stored in

Salesforce. The Agentforce Specialist needs to configure the system to use the most accurate and up-to-date information for email generation.

Which grounding technique should the Agentforce Specialist use?

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

Answer: C

Explanation:

For Einstein Sales Emails to generate personalized follow-up emails, it is crucial to ground the email content with the most up-to-date and accurate information.

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

? Record Merge Fields ensure the generated emails are highly personalized using

data like lead name, company, or other Salesforce fields directly from the records.

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

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

References:

? Salesforce Einstein Sales Emails Documentation: https://help.salesforce.com/s/articleView?id=release-notes__einstein_sales_emails.htm

NEW QUESTION 90

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 94

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 96

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 99

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 103

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 107

An Agentforce is tasked to optimize a business process flow by assigning actions to agents within the Salesforce Agentforce Platform. What is the correct method for the Agentforce Specialist to assign actions to an Agent?

- A. Assign the action to a Topic First in Agent Builder.
- B. Assign the action to a Topic first on the Agent Actions detail page.
- C. Assign the action to a Topic first on Action Builder.

Answer: C

Explanation:

? Action Builder is the central place in Salesforce Agentforce where you define and manage actions that your AI agents can perform. This includes connecting actions to various tools and systems.
? Topics in Agentforce represent the different tasks or intents that an AI agent can handle. By assigning an action to a Topic in Action Builder, you're essentially telling the agent, "When you encounter this type of request or situation, perform this action."

NEW QUESTION 110

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 114

Universal Containers wants its AI agent to answer customer questions with precise and up-to-date information. How does an Agentforce Data Library simplify and enable this?

- A. It automates the ingestion, taxonomical classification and storage of knowledge in Data Cloud for precision keyword search retrieval to ground prompts and agents with relevant information.
- B. It automates the ingestion, Indexing of data, and creates a default retriever to be used in prompts and agents for grounding with relevant information.
- C. It automates the ingestion and optical character recognition (OCR) processing of any PDF, and indexes them to enable regular SQL query retrieval to ground prompts and agents with relevant information.

Answer: B

Explanation:

Why is "Automates Ingestion, Indexing, and Default Retriever Creation" the correct answer?
An Agentforce Data Library is a key component in ensuring that an AI agent provides precise and up-to-date responses by:
Automating data ingestion Brings in data from various sources. Indexing the data Organizes it efficiently for AI retrieval. Creating a default retriever Enables the AI to fetch relevant data dynamically when answering customer queries.
Key Features of an Agentforce Data Library:
? Automates Data Ingestion
? Indexes Data for Efficient Retrieval
? Creates a Default Retriever
Why Not the Other Options?
* A. Automates ingestion, taxonomical classification, and precision keyword search retrieval
? Incorrect because Agentforce does not rely on keyword searches but on indexing and AI-driven retrieval.
* C. Automates ingestion and OCR processing of PDFs
? Incorrect because OCR (Optical Character Recognition) is not the primary function of an Agentforce Data Library.
? AI grounding is based on indexed and structured data, not raw OCR-extracted text.
Agentforce Specialist References
? Salesforce AI Specialist Material explains that Agentforce Data Libraries automate data ingestion, indexing, and retriever setup for AI-powered responses.
? Salesforce Instructions for Certification confirm that AI responses are grounded in structured and indexed Data Libraries.

NEW QUESTION 115

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 117

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 usingFlow 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 119

An Agentforce is tasked with analyzing Agent interactions looking into user inputs, requests, and queries to identify patterns and trends.

What functionality allows the AX Specialist to achieve this?

A. User Utterances dashboard

B. Agent Event Logs dashboard

C. AI Audit & Feedback Data dashboard

Answer: A

Explanation:

The User Utterances dashboard (Option A) is the correct functionality for analyzing user inputs, requests, and queries to identify patterns and trends. This dashboard aggregates and categorizes the natural language inputs (utterances) from users, enabling the Agentforce Specialist to:

? Identify Common Queries: Surface frequently asked questions or recurring issues.

? Detect Intent Patterns: Understand how users phrase requests, which helps refine intent detection models.

? Improve Bot Training: Highlight gaps in training data or misclassified utterances that require adjustment.

Why Other Options Are Incorrect:

? B. Agent Event Logs dashboard: Focuses on agent activity (e.g., response times, resolved cases) rather than user input analysis.

? C. AI Audit & Feedback Data dashboard: Tracks AI model performance, audit trails, and user feedback scores but does not directly analyze raw user utterances or queries.

References:

? Salesforce Einstein Agentforce Specialist Certification Guide: Emphasizes the User Utterances dashboard as the primary tool for analyzing user inputs to improve conversational AI.

? Trailhead Module: "Einstein Bots Basics" highlights using the dashboard to refine bot training based on user interaction data.

? Salesforce Help Documentation: Describes the User Utterances dashboard as critical for identifying trends in customer interactions.

NEW QUESTION 121

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 125

Universal Containers (UC) is looking to enhance its operational efficiency. UC has recently adopted Salesforce and is considering implementing Agent to improve its processes.

What is a key reason for implementing Agent?

- A. Improving data entry and data cleansing
- B. Allowing AI to perform tasks without user interaction
- C. Streamlining workflows and automating repetitive tasks

Answer: C

Explanation:

The key reason for implementing Agent is its ability to streamline workflows and automate repetitive tasks. By leveraging AI, Agent can assist users in handling mundane, repetitive processes, such as automatically generating insights, completing actions, and guiding users through complex processes, all of which significantly improve operational efficiency.

? Option A (Improving data entry and cleansing) is not the primary purpose of Agent, as its focus is on guiding and assisting users through workflows.

? Option B (Allowing AI to perform tasks without user interaction) does not accurately describe the role of Agent, which operates interactively to assist users in real time.

Salesforce Agentforce Specialist References: More details can be found in the Salesforce documentation:

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

NEW QUESTION 127

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 130

A support team handles a high volume of chat interactions and needs a solution to provide quick, relevant responses to customer inquiries.

Responses must be grounded in the organization's knowledge base to maintain consistency and accuracy.

Which feature in Einstein for Service should the support team use?

- A. Einstein Service Replies
- B. Einstein Reply Recommendations
- C. Einstein Knowledge Recommendations

Answer: B

Explanation:

The support team should use Einstein Reply Recommendations to provide quick, relevant responses to customer inquiries that are grounded in the organization's

knowledge base. This feature leverages AI to recommend accurate and consistent replies based on historical interactions and the knowledge stored in the system, ensuring that responses are aligned with organizational standards.

? Einstein Service Replies (Option A) is focused on generating replies but doesn't have the same emphasis on grounding responses in the knowledge base.

? Einstein Knowledge Recommendations (Option C) suggests knowledge articles to agents, which is more about assisting the agent in finding relevant articles than providing automated or AI-generated responses to customers.

Salesforce Agentforce Specialist References: For more information on Einstein Reply Recommendations:

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

NEW QUESTION 133

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 136

Universal Containers has grounded a prompt template with a related list. During user acceptance testing (UAT), users are not getting the correct responses. What is causing this issue?

- A. The related list is Read Only.
- B. The related list prompt template option is not enabled.
- C. The related list is not on the parent object??s page layout.

Answer: C

Explanation:

Comprehensive and Detailed In-Depth Explanation:UC has grounded a prompt template with a related list, but the responses are incorrect during UAT. Grounding with related lists in Agentforce allows the AI to access data from child records linked to a parent object. Let??s analyze the options.

? Option A: The related list is Read Only.Read-only status (e.g., via field-level security or sharing rules) might limit user edits, but it doesn??t inherently prevent the AI from accessing related list data for grounding, as long as the running user (or system context) has read access. This is unlikely to cause incorrect responses and is not a primary consideration, making it incorrect.

? Option B: The related list prompt template option is not enabled.There??s no specific "related list prompt template option" toggle in Prompt Builder. When grounding with a Record Snapshot or Flex template, related lists are included if properly configured (e.g., via object relationships). This option seems to be a misphrasing and doesn??t align with documented settings, making it incorrect.

? Option C: The related list is not on the parent object??s page layout.In Agentforce, grounding with related lists relies on the related list being defined and accessible in the parent object??s metadata, often tied to its presence on the page layout. If the related list isn??t on the layout, the AI might not recognize or retrieve its data correctly, leading to incomplete or incorrect responses. Salesforce documentation notes that related list data availability can depend on layout configuration, making this a plausible and common issue during UAT, and thus the correct answer.

Why Option C is Correct:The absence of the related list from the parent object??s page layout can disrupt data retrieval for grounding, leading to incorrect AI responses. This is a known configuration consideration in Agentforce setup and testing, as per official guidance.

References:

? Salesforce Agentforce Documentation: Grounding with Related Lists – Notes dependency on page layout configuration.

? Trailhead: Ground Your Agentforce Prompts – Highlights related list setup for accurate grounding.

? Salesforce Help: Troubleshoot Prompt Responses – Lists layout issues as a common grounding problem.

NEW QUESTION 138

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 143

Universal Containers wants support agents to use Agentforce to ask questions about its product tutorials and product guides.

What should the Agentforce Specialist do to meet this requirement?

- A. Create a prompt template for product tutorials and guides.
- B. Add an Answer Questions custom field in the product object for tutorial instructions.
- C. Publish product tutorials and guides as Knowledge articles.

Answer: C

Explanation:

? Context of the QuestionUniversal Containers (UC) wants its support agents to use Agentforce to ask questions about product tutorials and product guides.

Agentforce typically references knowledge sources to provide accurate and contextual responses.

? Why Knowledge Articles?

? Why Not the Other Options?

? ConclusionTo ensure Agentforce can effectively retrieve and deliver accurate information about products, publishing product tutorials and guides as Knowledge articles is the recommended approach.

Salesforce Agentforce Specialist References & Documents

? Salesforce Documentation: Set Up Salesforce KnowledgeDiscusses how to publish articles for easy access

? by AI-driven assistants and support teams.

? Salesforce Agentforce Specialist Study GuideExplains best practices for feeding knowledge sources to generative AI and Agentforce.

NEW QUESTION 147

Universal Containers (UC) users are complaining that agent answers are not satisfactory. The agent is using PDF files as a knowledge source.

How should UC troubleshoot this issue?

- A. Analyze the data mapping between source fields and Data Cloud object fields.
- B. Check that the agent has the PDF file field permission access for the data library.
- C. Verify the retriever's filter criteria and data source connection.

Answer: C

Explanation:

Why is "Verify the retriever's filter criteria and data source connection" the correct answer?

If agent answers are not satisfactory when using PDF files as a knowledge source, the issue is likely caused by:

- ? Retriever misconfiguration
- ? Incorrect data source connection
- ? Parsing Issues with PDF Files

Why Not the Other Options?

* A. Analyze the data mapping between source fields and Data Cloud object fields.

? Incorrect because data mapping issues primarily affect structured CRM data, not PDF-based knowledge sources.

? The issue likely stems from retrieval settings, not field mapping.

* B. Check that the agent has the PDF file field permission access for the data library.

? Incorrect because permission access issues would prevent AI from accessing PDFs entirely rather than causing poor response quality.

? AI can still generate responses, even if they are inaccurate, which means the issue lies in retriever settings, not permissions.

Agentforce Specialist References

? Salesforce AI Specialist Material details how retriever filters and data sources impact AI-generated answers.

? Salesforce Certification Guide mentions the importance of verifying retriever configurations for accurate knowledge retrieval.

NEW QUESTION 149

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 154

Universal Containers built a Field Generation prompt template that worked for many records, but users are reporting random failures with token limit errors. What is the cause of the random nature of this error?

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

Answer: B

Explanation:

Comprehensive and Detailed In-Depth Explanation: In Salesforce Agentforce, prompt templates are used to generate dynamic responses or field values by leveraging an LLM, often with grounding data from Salesforce records or external sources. The scenario describes a Field Generation prompt template that fails intermittently with token limit errors, indicating that the issue is tied to exceeding the LLM's token capacity (e.g., input + output tokens). The random nature of these failures suggests variability in the token count across different records, which is directly addressed by Option B.

Prompt templates in Agentforce can be dynamic, meaning they pull in record-specific data (e.g., customer names, descriptions, or other fields) to generate output. Since the data varies by record—some records might have short text fields while others have lengthy

ones—the total number of tokens (words, characters, or subword units processed by the LLM) fluctuates. When the token count exceeds the LLM's limit (e.g., 4,096 tokens for some models), the process fails, but this only happens for records with higher token- generating data, explaining the randomness.

? Option A: Switching to a "Flex" template type might sound plausible, but Salesforce documentation does not define "Flex" as a specific template type for handling token variability in this context (there are Flow-based templates, but they're unrelated to token limits). This option is a distractor and not a verified solution.

? Option C: The LLM's token processing capacity is fixed per model (e.g., a set limit like 128,000 tokens for advanced models) and does not vary with user demand. Demand might affect performance or availability, but not the token limit itself.

Option B is the correct answer because it accurately identifies the dynamic nature of the prompt template as the root cause of variable token counts leading to random failures.

References:

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

? Trailhead: "Build Prompt Templates for Agentforce" (<https://trailhead.salesforce.com/content/learn/modules/build-prompt-templates-for-agentforce>)

NEW QUESTION 156

Amid their busy schedules, sales reps at Universal Containers dedicate time to follow up with prospects and existing clients via email regarding renewals or new deals. They spend many hours throughout the week reviewing past communications and details about their customers before performing their outreach. Which standard Agent action helps sales reps draft personalized emails to prospects by generating text based on previous successful communications?

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

Answer: C

Explanation:

Comprehensive and Detailed In-Depth Explanation:UC??s sales reps need an AI action to draft personalized emails based on past successful communications, reducing manual review time. Let??s evaluate the standard Agent actions.

? Option A: Agent Action: Summarize Record"Summarize Record" generates a summary of a record (e.g., Opportunity, Contact), useful for overviews but not for drafting emails or leveraging past communications. This doesn??t meet the requirement, making it incorrect.

? Option B: Agent Action: Find Similar Opportunities"Find Similar Opportunities" identifies past deals to inform strategy, not to draft emails. It provides data, not text generation, making it incorrect.

? Option C: Agent Action: Draft or Revise Sales EmailThe "Draft or Revise Sales Email" action in Agentforce for Sales (sometimes styled as "Draft Sales Email") uses the Atlas Reasoning Engine to generate personalized email content. It can analyze past successful communications (e.g., via Opportunity or Contact history) to tailor emails for renewals or deals, saving reps time. This directly addresses UC??s need, making it the correct answer.

Why Option C is Correct:"Draft or Revise Sales Email" is a standard action designed for personalized email generation based on historical data, aligning with UC??s productivity goal per Salesforce documentation.

References:

? Salesforce Agentforce Documentation: Agentforce for Sales > Draft Sales Email – Details email generation.

? Trailhead: Explore Agentforce Sales Agents – Covers email drafting with past data.

? Salesforce Help: Sales Features in Agentforce – Confirms personalization capabilities.

NEW QUESTION 161

Universal Containers needs to provide insights on the usability of Agents to drive adoption in the organization. What should the Agentforce Specialist recommend?

- A. Agent Analytics
- B. Agentforce Analytics
- C. Agent Studio Analytics

Answer: A

Explanation:

? Agent Analytics: This tool is specifically designed to provide usability insights for Salesforce agents. It tracks metrics like adoption rates, task completion times, and efficiency levels, helping organizations identify areas where agents excel or need additional support.

? Agentforce Analytics: This term does not correspond to a recognized Salesforce feature.

? Agent Studio Analytics: This is unrelated to analyzing agent usability, as it primarily supports customization or development features rather than providing analytics for adoption.

Thus, Agent Analytics is the correct recommendation as it offers actionable insights to drive agent adoption and productivity.

Reference:

"Boost Adoption with Analytics Tools | Salesforce" .

NEW QUESTION 163

In the context of retriever and search indexes, what best describes the data preparation process in Data Cloud?

- A. Data preparation focuses on real-time data ingestion and dynamic indexing to generate dynamic grounding reference data without preprocessing steps.
- B. Data preparation entails aggregating, normalizing, and encoding structured datasets to ensure compliance with data governance and security protocols.
- C. Data preparation Involves loading, chunking, vectorizing, and storing content in a search-optimized manner to support retrieval from the vector database.

Answer: C

Explanation:

Why is "Loading, Chunking, Vectorizing, and Storing" the correct answer? Agentforce AI-powered search and retriever indexing requires data to be structured and optimized for retrieval. The Data Cloud preparation process involves:

Key Steps in the Data Preparation Process for Agentforce:

? Loading Data

? Chunking (Breaking Text into Small Parts)

? Vectorization (Transforming Text for AI Retrieval)

? Storing in a Vector Database

Why Not the Other Options?

* A. Real-time data ingestion and dynamic indexing

? Incorrect because while real-time updates can occur, the primary process involves preprocessing and indexing first.

* B. Aggregating, normalizing, and encoding structured datasets

? Incorrect because this process relates to data compliance and security, not AI retrieval optimization.

Agentforce Specialist References

? Salesforce AI Specialist Material confirms that data preparation includes chunking, vectorizing, and storing for AI retrieval in Data Cloud.

NEW QUESTION 168

When creating a custom retriever in Einstein Studio, which step is considered essential?

- A. Select the search index, specify the associated data model object (DMO) and data space, and optionally define filters to narrow search results.
- B. Define the output configuration by specifying the maximum number of results to return, and map the output fields that will ground the prompt.
- C. Configure the search index, choose vector or hybrid search, choose the fields for filtering, the data space and model, then define the ranking method.

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation: In Salesforce's Einstein Studio (part of the Agentforce ecosystem), creating a custom retriever involves setting up a mechanism to fetch data for AI prompts or responses. The essential step is defining the foundation of the retriever: selecting the search index, specifying the data model object (DMO), and identifying the data space (Option A). These elements establish where and what the retriever searches:

? Search Index: Determines the indexed dataset (e.g., a vector database in Data Cloud) the retriever queries.

? Data Model Object (DMO): Specifies the object (e.g., Knowledge Articles, Custom Objects) containing the data to retrieve.

? Data Space: Defines the scope or environment (e.g., a specific Data Cloud instance) for the data.

Filters are noted as optional in Option A, which is accurate—they enhance precision but aren't mandatory for the retriever to function. This step is foundational because without it, the retriever lacks a target dataset, rendering it unusable.

? Option B: Defining output configuration (e.g., max results, field mapping) is important for shaping the retriever's output, but it's a secondary step. The retriever must first know where to search (A) before output can be configured.

? Option C: This option includes advanced configurations (vector/hybrid search, filtering fields, ranking method), which are valuable but not essential. A basic retriever can operate without specifying search type or ranking, as defaults apply, but it cannot function without a search index, DMO, and data space.

? Option A: This is the minimum required step to create a functional retriever, making it essential.

Option A is the correct answer as it captures the core, mandatory components of retriever setup in Einstein Studio.

References:

? Salesforce Agentforce Documentation: "Custom Retrievers in Einstein Studio" (Salesforce Help:

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

=5)

? Trailhead: "Einstein Studio for Agentforce" (<https://trailhead.salesforce.com/content/learn/modules/einstein-studio-for-agentforce>)

NEW QUESTION 170

A sales manager is using Agent Assistant to streamline their daily tasks. They ask the agent to Show me a list of my open opportunities.

How does the large language model (LLM) in Agentforce identify and execute the action to show the sales manager a list of open opportunities?

- A. The LLM interprets the user's request, generates a plan by identifying the appropriate topics and actions, and executes the actions to retrieve and display the open opportunities
- B. The LLM uses a static set of rules to match the user's request with predefined topics and actions, bypassing the need for dynamic interpretation and planning.
- C. Using a dialog pattern
- D. the LLM matches the user query to the available topic, action and steps then performs the steps for each action, such as retrieving a list of open opportunities.

Answer: A

Explanation:

Agentforce's LLM dynamically interprets natural language requests (e.g., "Show me open opportunities"), generates an execution plan using the planner service, and retrieves data via actions (e.g., querying Salesforce records). This contrasts with static rules (B) or rigid dialog patterns (C), which lack contextual adaptability. Salesforce documentation highlights the planner's role in converting intents into actionable steps while adhering to security and business logic.

Reference:

Salesforce Help Article: Agentforce Planner Service ("Dynamic Request Interpretation" section).

Einstein Agentforce Specialist Trailhead: "How Agentforce Processes User Requests."

NEW QUESTION 173

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

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

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

Answer: A

Explanation:

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

Given that UC has legacy fields and outdated data, Service AI Grounding ensures that only the valid and applicable data is used by Einstein AI to craft responses. This helps improve the relevance of responses and avoids inaccuracies caused by outdated or irrelevant fields. Work Summaries and Service Replies are useful features but do not address the need for grounding AI outputs in specific, current data sources like Service AI Grounding does. For more details, you can refer to Salesforce's Service AI Grounding documentation for managing AI-generated content based on accurate data sources.

NEW QUESTION 176

An Agentforce needs to create a Sales Email with a custom prompt template. They need to ground on the following data.

Opportunity Products Events near the customer Tone and voice examples How should the Agentforce Specialist obtain related items?

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

Answer: A

Explanation:

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

information.

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

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

NEW QUESTION 177

Which use case is best supported by Salesforce Agent's capabilities?

- A. Bring together a conversational interface for interacting with AI for all Salesforce users, such as developers and ecommerce retailers.
- B. Enable Salesforce admin users to create and train custom large language models (LLMs) using CRM data.
- C. Enable data scientists to train predictive AI models with historical CRM data using built- in machine learning capabilities

Answer: A

Explanation:

Salesforce Agent is designed to provide a conversational AI interface that can be utilized by different types of Salesforce users, such as developers, sales agents, and retailers. It acts as an AI-powered assistant that facilitates natural interactions with the system, enabling users to perform tasks and access data easily. This includes tasks like pulling reports, updating records, and generating personalized responses in real time.

? Option A is correct because Agent brings a conversational interface that caters to a wide range of users.

? Option B and Option C are more focused on developing and training AI models, which are not the primary functions of Agent.

References:

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

NEW QUESTION 181

Which element in the Omni-Channel Flow should be used to connect the flow with the agent?

- A. Route Work Action
- B. Assignment
- C. Decision

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation:UC is integrating an Agentforce agent with Omni-Channel Flow to route work. Let??s identify the correct element.

? Option A: Route Work ActionThe "Route Work" action in Omni-Channel Flow assigns work items (e.g., cases, chats) to agents or queues based on routing rules. When connecting to an Agentforce agent, this action links the flow to the agent??s queue or presence, enabling interaction. This is the standard element for agent integration, making it the correct answer.

? Option B: AssignmentThere??s no "Assignment" element in Flow Builder for Omni- Channel. Assignment rules exist separately, but within flows, routing is handled by "Route Work," making this incorrect.

? Option C: DecisionThe "Decision" element branches logic, not connects to agents.

It??s a control structure, not a routing mechanism, making it incorrect.

Why Option A is Correct:"Route Work" is the designated Omni-Channel Flow action for connecting to agents, including Agentforce agents, per Salesforce documentation.

References:

? Salesforce Agentforce Documentation: Omni-Channel Integration – Specifies "Route Work" for agents.

? Trailhead: Omni-Channel Flow Basics – Details routing actions.

? Salesforce Help: Set Up Omni-Channel Flows – Confirms "Route Work" usage.

NEW QUESTION 186

.....

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 FirstTry
- * Agentforce-Specialist Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

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