

Salesforce

Exam Questions Agentforce-Specialist

Salesforce Certified Agentforce Specialist



NEW QUESTION 1

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

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

Answer: B

Explanation:

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

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

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

References:

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

NEW QUESTION 2

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

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

Answer: C

Explanation:

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

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

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

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

Reference:

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

NEW QUESTION 3

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 4

Universal Containers recently launched a pilot program to integrate conversational AI into its CRM business operations with Agentforce Agents. How should the Agentforce Specialist monitor Agents' usability and the assignment of actions?

- A. Run a report on the Platform Debug Logs.
- B. Query the Agent log data using the Metadata API.
- C. Run Agent Analytics.

Answer: C

Explanation:

Comprehensive and Detailed In-Depth Explanation: Monitoring the usability and action assignments of Agentforce Agents requires insights into how agents perform, how users interact with them, and how actions are executed within conversations. Salesforce provides Agent Analytics (Option C) as a built-in capability specifically designed for this purpose. Agent Analytics offers dashboards and reports that track metrics such as agent response times, user satisfaction, action invocation frequency, and success rates. This tool allows the Agentforce Specialist to assess usability (e.g., are agents meeting user needs?) and monitor action assignments (e.g., which actions are triggered and how often), providing actionable data to optimize the pilot program.

? Option A: Platform Debug Logs are low-level logs for troubleshooting Apex, Flows, or system processes. They don't provide high-level insights into agent usability or action assignments, making this unsuitable.

? Option B: The Metadata API is used for retrieving or deploying metadata (e.g., object definitions), not runtime log data about agent performance. While Agent log data might exist, querying it via Metadata API is not a standard or documented approach for this use case.

? Option C: Agent Analytics is the dedicated solution, offering a user-friendly way to monitor conversational AI performance without requiring custom development. Option C is the correct choice for effectively monitoring Agentforce Agents in a pilot program.

References:

? Salesforce Agentforce Documentation: "Agent Analytics Overview" (Salesforce Help:

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

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

NEW QUESTION 5

Universal Containers (UC) is Implementing Service AI Grounding to enhance its customer service operations. UC wants to ensure that its AI-generated responses are grounded in the most relevant data sources. The team needs to configure the system to include all supported objects for grounding. Which objects should UC select to configure Service AI Grounding?

- A. Case, Knowledge, and Case Notes
- B. Case and Knowledge
- C. Case, Case Emails, and Knowledge

Answer: B

Explanation:

Universal Containers (UC) is implementing Service AI Grounding to enhance its customer service operations. They aim to ensure that AI-generated responses are grounded in the most relevant data sources and need to configure the system to include all supported objects for grounding.

Supported Objects for Service AI Grounding:

? Case

? Knowledge

? Case Object:

? Knowledge Object:

? Exclusion of Other Objects:

Why Options A and C are Incorrect:

? Option A (Case, Knowledge, and Case Notes):

? Option C (Case, Case Emails, and Knowledge):

References:

? Salesforce Agentforce Specialist Documentation - Service AI Grounding Configuration: Details the objects supported for grounding AI responses in Service Cloud.

? Salesforce Help - Implementing Service AI Grounding: Provides guidance on setting up grounding with Case and Knowledge objects.

? Salesforce Trailhead - Enhance Service with AI Grounding: Offers an interactive learning path on using AI grounding in service scenarios.

NEW QUESTION 6

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 7

Which part of the Einstein Trust Layer architecture leverages an organization's own data within a large language model (LLM) prompt to confidently return relevant and accurate responses?

- A. Prompt Defense
- B. Data Masking
- C. Dynamic Grounding

Answer: C

Explanation:

Dynamic Grounding in the Einstein Trust Layer architecture ensures that large language model (LLM) prompts are enriched with organization-specific data (e.g., Salesforce records, Knowledge articles) to generate accurate and relevant responses. By dynamically injecting contextual data into prompts, it reduces hallucinations and aligns outputs with trusted business data.

? Prompt Defense (A) focuses on blocking malicious inputs or prompt injections but does not enhance responses with organizational data.

? Data Masking (B) redacts sensitive information but does not contribute to grounding responses in business context.

Reference:

Salesforce Help Article: Einstein Trust Layer – Dynamic Grounding ("How Dynamic Grounding Works" section).

Einstein Trust Layer Technical Overview: "Contextual Accuracy with Dynamic Grounding."

NEW QUESTION 8

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 9

For an Agentforce Data Library that contains uploaded files, what occurs once it is created and configured?

- A. Indexes the uploaded files in a location specified by the user
- B. Indexes the uploaded files into Data Cloud
- C. Indexes the uploaded files in Salesforce File Storage

Answer: B

Explanation:

Comprehensive and Detailed In-Depth Explanation: In Salesforce Agentforce, a Data Library is a feature that allows organizations to upload files (e.g., PDFs, documents) to be used as grounding data for AI-driven agents. Once the Data Library is created and configured, the uploaded files are indexed to make their content searchable and usable by the AI (e.g., for retrieval-augmented generation or prompt enhancement). The key question is where this indexing occurs. Salesforce Agentforce integrates tightly with Data Cloud, a unified data platform that includes a vector database optimized for storing and indexing unstructured data like uploaded files. When a Data Library is set up, the files are ingested and indexed into Data Cloud's vector database, enabling the AI to efficiently retrieve relevant information from them during conversations or actions.

? Option A: Indexing files in a "location specified by the user" is not a feature of Agentforce Data Libraries. The indexing process is managed by Salesforce infrastructure, not a user-defined location.

? Option B: This is correct. Data Cloud handles the indexing of uploaded files, storing them in its vector database to support AI capabilities like semantic search and content retrieval.

? Option C: Salesforce File Storage (e.g., where ContentVersion records are stored) is used for general file storage, but it does not inherently index files for AI use. Agentforce relies on Data Cloud for indexing, not basic file storage.

Thus, Option B accurately reflects the process after a Data Library is created and configured in Agentforce.

References:

? Salesforce Agentforce Documentation: "Set Up a Data Library" (Salesforce Help:

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

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

NEW QUESTION 10

Universal Containers (UC) plans to automatically populate the Description field on the Account object. Which type of prompt template should UC use?

- A. Field Generation prompt template
- B. Flex Prompt template
- C. Sales Email prompt template

Answer: A

Explanation:

? Context of the Question Universal Containers (UC) wants to automatically populate the Description field on the Account object. The AI-driven solution must generate textual data and write it directly into a field.
? Field Generation Prompt Template
? Why Not Flex or Sales Email Prompt Templates?
? Conclusion For automatically populating the Description field with AI-generated content, the Field Generation prompt template (Option A) is the correct choice.
Salesforce Agentforce Specialist References & Documents
? Salesforce Documentation: Prompt Template Types Explains various template types (Field Generation, Flex, Email, etc.) and their typical use cases.
? Salesforce Agentforce Specialist Study Guide Highlights Field Generation prompt templates for populating or updating record fields with AI-generated text.

NEW QUESTION 10

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

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

Answer: B

Explanation:

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

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

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

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

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

References:

- ? Salesforce Agentforce Documentation: Einstein Prompt Builder – Details prompt templates for consistency and efficiency.
- ? Trailhead: Build Prompt Templates in Agentforce – Explains time-saving benefits of templates.
- ? Salesforce Help: Generative AI with Prompt Builder – Confirms use for streamlining rep interactions.

NEW QUESTION 15

What is true of Agentforce Testing Center?

- A. Running tests risks modifying CRM data in a production environment.
- B. Running tests does not consume Einstein Requests.
- C. Agentforce Testing Center can only be used in a production environment.

Answer: B

Explanation:

Comprehensive and Detailed In-Depth Explanation: The Agentforce Testing Center is a tool in Agentforce Studio for validating agent performance. Let's evaluate the statements.

? Option A: Running tests risks modifying CRM data in a production environment. Agentforce Testing Center runs synthetic interactions in a controlled environment (e.g., sandbox or isolated test space) and doesn't modify live CRM data. It's designed for safe pre-deployment testing, making this incorrect.

? Option B: Running tests does not consume Einstein Requests. Einstein Requests are part of the usage quota for Einstein Generative AI features (e.g., prompt executions in production). Testing Center uses synthetic data to simulate interactions without invoking live AI calls that count against this quota. Salesforce documentation confirms tests don't consume requests, making this the correct answer.

? Option C: Agentforce Testing Center can only be used in a production environment. Testing Center is available in both sandbox and production orgs, but it's primarily used pre-deployment (e.g., in sandboxes) to validate agents safely. This restriction is false, making it incorrect.

Why Option B is Correct: Not consuming Einstein Requests is a key feature of Testing Center, allowing extensive testing without impacting quotas, as per Salesforce documentation.

References:

- ? Salesforce Agentforce Documentation: Testing Center > Overview – Confirms no request consumption.
- ? Trailhead: Test Your Agentforce Agents – Notes quota-free testing.
- ? Salesforce Help: Agentforce Testing – Details safe, isolated testing.

NEW QUESTION 16

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 19

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

A. Storing this data requires Data Cloud to be provisioned.

B. Storing this data requires a custom object for data to be configured.

C. Storing this data requires Salesforce big objects.

Answer: A

Explanation:

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

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

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

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

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

References:

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

? Salesforce Einstein AI Overview:

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

NEW QUESTION 21

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 22

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

A. With a preconfigured message, based on the action type.

B. With a general message asking the user to rephrase the request.

C. With a generated error message.

Answer: B

Explanation:

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

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

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

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

Why Option B is Correct: The default behavior of asking users to rephrase aligns with Agentforce's conversational design principles, ensuring a helpful response

when comprehension fails, as noted in official resources.

References:

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

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

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

NEW QUESTION 23

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 27

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 32

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

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

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

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

Answer: C

Explanation:

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

Key Considerations in Agentforce Data Library

? Data Source Lock-In

? Why Can't the Data Source Be Changed?

? Workarounds for Changing Data Sources

Why Not the Other Options?

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

? Incorrect because once the data source is linked to an Agentforce Data Library, it cannot be altered, even via Data Cloud settings.

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

? Incorrect as the Data Retriever works within the constraints of the selected data source and does not provide an option to swap data sources post-selection.

Agentforce Specialist References

The Salesforce AI Specialist Material and Salesforce Instructions for the Certification confirm that once a data source is set for an Agentforce Data Library, it cannot be changed.

NEW QUESTION 37

Universal Containers implemented Agent for its users.

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

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

Answer: C

Explanation:

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

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

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

References:

? Salesforce Agentforce Specialist Documentation - Agent Supported Actions:

? Salesforce Help - Limitations of Agent:

NEW QUESTION 41

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 46

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 50

An Agentforce wants to use the related lists from an account in a custom prompt template. What should the Agentforce Specialist consider when configuring the prompt template?

- A. The text encoding (for example, UTF-8, ASCII) option
- B. The maximum number of related list merge fields
- C. The choice between XML and JSON rendering formats for the list

Answer: B

Explanation:

When configuring a custom prompt template to use related lists, the Agentforce Specialist must be aware of the maximum number of related list merge fields that can be included. Salesforce enforces limits to ensure prompt templates perform efficiently and do not overload the system with too much data. As a best practice, it's important to monitor and optimize the number of merge fields used.

? Option B is correct because there is a limit on how many related list merge fields can be included in a prompt template.

? Option A (text encoding) and Option C (XML/JSON rendering) are not key considerations in this context.

References:

- ? Salesforce Prompt Builder Documentation: https://help.salesforce.com/s/articleView?id=sf.prompt_builder.htm

NEW QUESTION 52

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 53

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

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

Answer: A

Explanation:

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

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

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

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

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

References:

- ? Salesforce Agentforce Documentation: Agent Builder > Custom Actions > Apex – Notes standard Apex deployment rules apply.
- ? Salesforce Developer Guide: Apex Testing – Confirms 75% coverage requirement.
- ? Trailhead: Deploy Apex Code – Emphasizes coverage and dependencies for production.

NEW QUESTION 54

A Universal Containers administrator is setting up Einstein Data Libraries. After creating a new library, the administrator notices that only the file upload option is available; there is no option to configure the library using a Salesforce Knowledge base. What is the most likely cause of this Issue?

- A. The current Salesforce org lacks the necessary Einstein for Service permissions that support the Knowledge-based Data Library option, so only the file upload option is presented.
- B. Salesforce Knowledge is not enabled in the organization; without Salesforce Knowledge enabled, the Knowledge-based data source option will not be available in Einstein Data Libraries.
- C. The administrator is not using Lightning Experience, which is required to display all data source options, Including the Knowledge base option, when configuring Einstein Data Libraries.

Answer: B

Explanation:

Why is "Salesforce Knowledge is not enabled" the correct answer?

If an administrator only sees the file upload option in Einstein Data Libraries and cannot configure a Salesforce Knowledge base, the most likely reason is that Salesforce Knowledge is not enabled in the organization.

Key Considerations for Einstein Data Libraries:

? Salesforce Knowledge Integration is Optional

? How to Fix This Issue?

Why Not the Other Options?

* A. The current Salesforce org lacks the necessary Einstein for Service permissions

? Incorrect because even without certain permissions, the Knowledge option would still be visible but greyed out.

* C. The administrator is not using Lightning Experience

? Incorrect because Einstein Data Libraries are accessible in both Classic and Lightning, and Lightning does not control Knowledge base visibility.

Agentforce Specialist References

? Salesforce AI Specialist Material confirms that Salesforce Knowledge must be enabled for Data Libraries to use Knowledge as a data source.

? Salesforce Certification Guide explicitly states that file uploads are the default option if Knowledge is not available.

NEW QUESTION 59

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

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

Answer: C

Explanation:

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

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

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

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

References:

? Salesforce Help: Einstein Agent Setup

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

NEW QUESTION 61

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

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

Answer: C

Explanation:

Comprehensive and Detailed In-Depth Explanation:The Agentforce Specialist is defining a custom action for an Agentforce agent in Agent Builder. Actions determine what the agent does (e.g., retrieve data, update records). Let??s evaluate the options.

? Option A: Apex TriggerApex Triggers are event-driven scripts, not selectable actions in Agent Builder. While Apex can be invoked via other means (e.g., Flows), it??s not a direct option for custom agent actions, making this incorrect.

? Option B: SOQLSOQL (Salesforce Object Query Language) is a query language, not an executable action type in Agent Builder. While actions can use queries internally, SOQL isn??t a standalone option, making this incorrect.

? Option C: FlowsIn Agentforce Studio??s Agent Builder, custom actions can be created using Salesforce Flows. Flows allow complex logic (e.g., data retrieval, updates, or integrations) and are explicitly supported as a custom action type. The specialist can select an existing Flow or create one, making this the correct answer.

? Option D: JavaScriptJavaScript isn??t an option for defining agent actions in Agent Builder. It??s used in Lightning Web Components, not agent configuration, making this incorrect.

Why Option C is Correct:Flows are a native, flexible option for custom actions in Agentforce, enabling tailored functionality for agents, as per official documentation.

References:

? Salesforce Agentforce Documentation: Agent Builder > Custom Actions – Lists Flows as a supported action type.

? Trailhead: Build Agents with Agentforce – Details Flow-based actions.

? Salesforce Help: Configure Agent Actions – Confirms Flows integration.

NEW QUESTION 66

An AI Specialist is tasked with creating a prompt template for a sales team. The template needs to generate a summary of all related opportunities for a given Account.

Which grounding technique should the AI Specialist use to include data from the related list of opportunities in the prompt template?

- A. Use the merge fields to reference a custom related list of opportunities.
- B. Use merge fields to reference the default related list of opportunities.
- C. Use formula fields to reference the Einstein related list of opportunities.

Answer: B

Explanation:

In Salesforce, when creating a prompt template for the sales team, you can include data from related objects such as Opportunities that are linked to an Account. The best method to ground the AI model and provide relevant information from related records, like Opportunities, is by using merge fields.

Merge fields in Salesforce allow you to dynamically reference data from a record or related records, like Opportunities for a given Account. In this scenario, the Agentforce Specialist needs to pull data from the default related list of Opportunities associated with the Account. This is achieved by using merge fields, which pull in data from the standard relationship Salesforce creates between Accounts and Opportunities.

Option A (referencing a custom related list) and Option C (using formula fields with Einstein-related lists) do not align with the standard, practical grounding method for this

task. Custom lists would require additional configurations not typically necessary for a basic use case, and formula fields are typically not used to directly fetch related list data for prompt generation in templates. The standard and straightforward method is using merge fields tied to the default related list of opportunities.

Salesforce References:

? Merge Fields in Templates: <https://help.salesforce.com/s/articleView?id=000387601&type=1>

? Grounding Data in Prompts: https://developer.salesforce.com/docs/atlas.en-us.salesforce_ai.meta/salesforce_ai/grounding_data_prompts

NEW QUESTION 70

Universal Containers (UC) implements a custom retriever to improve the accuracy of AI-generated responses. UC notices that the retriever is returning too many irrelevant results, making the responses less useful. What should UC do to ensure only relevant data is retrieved?

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

Answer: A

Explanation:

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

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

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

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

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

References:

? Salesforce Agentforce Documentation: "Create Custom Retriever" (Salesforce Help:

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

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

NEW QUESTION 75

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 78

How does the AI Retriever function within Data Cloud?

A. It performs contextual searches over an indexed repository to quickly fetch the most relevant documents, enabling grounding AI responses with trustworthy, verifiable information.

B. It monitors and aggregates data quality metrics across various data pipelines to ensure only high-integrity data is used for strategic decision-making.

C. It automatically extracts and reformats raw data from diverse sources into standardized datasets for use in historical trend analysis and forecasting.

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation:The AI Retriever is a key component in Salesforce Data Cloud, designed to support AI-driven processes like Agentforce by retrieving relevant data. Let's evaluate each option based on its documented functionality.

? Option A: It performs contextual searches over an indexed repository to quickly fetch the most relevant documents, enabling grounding AI responses with trustworthy, verifiable information.The AI Retriever in Data Cloud uses vector-based search technology to query an indexed repository (e.g., documents, records, or ingested data) and retrieve the most relevant results based on context. It employs embeddings to match user queries or prompts with stored data, ensuring AI responses (e.g., in Agentforce prompt templates) are grounded in accurate, verifiable information from Data Cloud. This enhances trustworthiness by linking outputs to source data, making it the primary function of the AI Retriever. This aligns with Salesforce documentation and is the correct answer.

? Option B: It monitors and aggregates data quality metrics across various data pipelines to ensure only high-integrity data is used for strategic decision-making.Data quality monitoring is handled by other Data Cloud features, such as Data Quality Analysis or ingestion validation tools, not the AI Retriever. The Retriever's role is retrieval, not quality assessment or pipeline management. This option is incorrect as it misattributes functionality unrelated to the AI Retriever.

? Option C: It automatically extracts and reformats raw data from diverse sources

into standardized datasets for use in historical trend analysis and forecasting.Data extraction and standardization are part of Data Cloud's ingestion and harmonization processes (e.g., via Data Streams or Data Lake), not the AI Retriever's function. The Retriever works with already-indexed data to fetch results, not to process or reformat raw data. This option is incorrect.

Why Option A is Correct:The AI Retriever's core purpose is to perform contextual searches over indexed data, enabling AI grounding with reliable information.

This is critical for Agentforce agents to provide accurate responses, as outlined in Data Cloud and Agentforce documentation.

References:

? Salesforce Data Cloud Documentation: AI Retriever – Describes its role in contextual searches for grounding.

? Trailhead: Data Cloud for Agentforce – Explains how the AI Retriever fetches relevant data for AI responses.

? Salesforce Help: Grounding with Data Cloud – Confirms the Retriever's search functionality over indexed repositories.

NEW QUESTION 82

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 85

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 86

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 89

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 94

What is An Agentforce able to do when the "Enrich event logs with conversation data" setting in Agent is enabled?

A. View the user click path that led to each copilot action.

B. View session data including user input and copilot responses for sessions over the past 7 days.

C. Generate detailed reports on all Copilot conversations over any time period.

Answer: B

Explanation:

When the "Enrich event logs with conversation data" setting is enabled in Agent, it allows An Agentforce or admin to view session data, including both the user input and copilot responses from interactions over the past 7 days. This data is crucial for monitoring how the copilot is being used, analyzing its performance, and improving future interactions based on past inputs.

? This setting enriches the event logs with detailed conversational data for better insights into the interaction history, helping Agentforce Specialists track AI behavior and user engagement.

? Option A, viewing the user click path, focuses on navigation but is not part of the conversation data enrichment functionality.

? Option C, generating detailed reports over any time period, is incorrect because this specific feature is limited to data for the past 7 days.

Salesforce Agentforce Specialist References: You can refer to this documentation for further insights:

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

NEW QUESTION 97

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

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

Answer: A

Explanation:

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

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

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

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

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

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

NEW QUESTION 102

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 107

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

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

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

Answer: B

Explanation:

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

Reference:

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

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

NEW QUESTION 108

An Agentforce configured Data Masking within the Einstein Trust Layer.

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

- A. Use a Flow-based resource in Prompt Builder to debug the fields' merge values using Flow Debugger.
- B. Request the Einstein Generative AI Audit Data from the Security section of the Setup menu.

C. Enable the collection and storage of Einstein Generative AI Audit Data on the Einstein Feedback setup page.

Answer: C

Explanation:

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

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

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

References:

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

NEW QUESTION 112

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

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

A. Design universal product recommendations.

B. Write a response scrip for the chatbot.

C. Collect and analyze browsing data.

Answer: C

Explanation:

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

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

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

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

Options A and D are incorrect because:

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

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

References:

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

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

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

NEW QUESTION 115

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 117

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 120

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 122

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

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

Answer: A

Explanation:

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

How It Works:

? Data Masking in the Request Journey:

? Processing by the LLM:

? De-masking in the Response Journey:

Why Option A is Correct:

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

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

Why Options B and C are Incorrect:

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

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

References:

? Salesforce Agentforce Specialist Documentation - Einstein Trust Layer Overview:

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

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

Conclusion:

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

NEW QUESTION 123

Universal Containers (UC) wants to enable its sales team to get insights into product and competitor names mentioned during calls. How should UC meet this requirement?

- A. Enable Einstein Conversation Insights, connect a recording provider, assign permission sets, and customize insights with up to 25 products.
- B. Enable Einstein Conversation Insights, assign permission sets, define recording managers, and customize insights with up to 50 competitor names.
- C. Enable Einstein Conversation Insights, enable sales recording, assign permission sets, and customize insights with up to 50 products.

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation: UC wants insights into product and competitor mentions during sales calls, leveraging Einstein Conversation Insights. Let's evaluate the options.

? Option A: Enable Einstein Conversation Insights, connect a recording provider, assign permission sets, and customize insights with up to 25 products. Einstein Conversation Insights analyzes call recordings to identify keywords like product and competitor names. Setup requires enabling the feature, connecting an external recording provider (e.g., Zoom, Gong), assigning permission sets (e.g., Einstein Conversation Insights User), and customizing insights by defining up to 25 products or competitors to track. Salesforce documentation confirms the 25-item limit for custom keywords, making this the correct, precise answer aligning with UC??s needs.

? Option B: Enable Einstein Conversation Insights, assign permission sets, define recording managers, and customize insights with up to 50 competitor names. There??s no "recording managers" role in Einstein Conversation Insights setup—integration is with a provider, not a manager designation. The limit is 25 keywords (not 50), and the option omits the critical step of connecting a provider, making it incorrect.

? Option C: Enable Einstein Conversation Insights, enable sales recording, assign permission sets, and customize insights with up to 50 products. "Enable sales recording" is vague—Conversation Insights relies on external providers, not a native Salesforce recording feature. The keyword limit is 25, not 50, making this incorrect despite being closer than B.

Why Option A is Correct: Option A accurately reflects the setup process and limits for Einstein Conversation Insights, meeting UC??s requirement per Salesforce documentation.

References:

? Salesforce Help: Set Up Einstein Conversation Insights – Details provider connection and 25-keyword limit.

? Trailhead: Einstein Conversation Insights Basics – Covers permissions and customization.

? Salesforce Agentforce Documentation: Sales Features – Confirms integration steps.

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 129

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

A. The User related list should have View All access.

B. The User related list needs to be included on the record page.

C. The User related list is not supported in prompt templates.

Answer: C

Explanation:

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

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

NEW QUESTION 131

Universal Containers is very concerned about security compliance and wants to understand:

Which prompt text is sent to the large language model (LLM)

* How it is masked

* The masked response

What should the Agentforce Specialist recommend?

A. Ingest the Einstein Shield Event logs into CRM Analytics.

B. Review the debug logs of the running user.

C. Enable audit trail in the Einstein Trust Layer.

Answer: C

Explanation:

To address security compliance concerns and provide visibility into the prompt text sent to the LLM, how it is masked, and the masked response, the Agentforce Specialist should recommend enabling the audit trail in the Einstein Trust Layer. This feature captures and logs the prompts sent to the large language model (LLM) along with the masking of sensitive information and the AI's response. This audit trail ensures full transparency and compliance with security requirements.

? Option A (Einstein Shield Event logs) is focused on system events rather than specific AI prompt data.

? Option B (debug logs) would not provide the necessary insight into AI prompt masking or responses.

For further details, refer to Salesforce's Einstein Trust Layer documentation about auditing and security measures.

NEW QUESTION 134

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 138

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 141

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 142

Universal Containers wants to use an external large language model (LLM) in Prompt Builder.
What should An Agentforce recommend?

- A. Use Apex to connect to an external LLM and ground the prompt.
- B. Use BYO-LLM functionality in Einstein Studio.
- C. Use Flow and External Services to bring data from an external LLM.

Answer: B

Explanation:

Bring Your Own Large Language Model (BYO-LLM) functionality in Einstein Studio allows organizations to integrate and use external large language models (LLMs) within the Salesforce ecosystem. Universal Containers can leverage this feature to connect and ground prompts with external LLMs, allowing for custom AI model use cases and seamless integration with Salesforce data.

? Option B is the correct choice as Einstein Studio provides a built-in feature to work with external models.

? Option A suggests using Apex, but BYO-LLM functionality offers a more streamlined solution.

? Option C focuses on Flow and External Services, which is more about data integration and isn't ideal for working with LLMs.

References:

Salesforce Einstein Studio BYO-LLM Documentation: https://help.salesforce.com/s/articleView?id=sf.einstein_studio_llm.htm

NEW QUESTION 146

Universal Containers (UC) wants to enable its sales team with automatic post-call visibility into mention of competitors, products, and other custom phrases. Which feature should the Agentforce Specialist set up to enable UC's sales team?

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

Answer: C

Explanation:

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

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

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

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

NEW QUESTION 149

Universal Containers recently added a custom flow for processing returns and created a new Agent Action. Which action should the company take to ensure the Agentforce Service Agent can run this new flow as part of the new Agent Action?

- A. Recreate the flow using the Agentforce agent user.
- B. Assign the Manage Users permission to the Agentforce Agent user.
- C. Assign the Run Flows permission to the Agentforce Agent user.

Answer: C

Explanation:

Comprehensive and Detailed In-Depth Explanation:UC has created a custom flow for processing returns and linked it to a new Agent Action for the Agentforce Service Agent, an AI-driven agent for customer service tasks. The agent must have the ability to execute this flow. Let's assess the options.

? Option A: Recreate the flow using the Agentforce agent user.Flows are authored by admins or developers, not "recreated" by specific users like the Agentforce agent user (a system user for agent operations). The issue isn't the flow's creation context but its execution permissions. This option is impractical and incorrect.

? Option B: Assign the Manage Users permission to the Agentforce Agent user.The "Manage Users" permission allows user management (e.g., creating or editing users), which is unrelated to running flows. This permission is excessive and irrelevant for the Service Agent's needs, making it incorrect.

? Option C: Assign the Run Flows permission to the Agentforce Agent user.The Agentforce Service Agent operates under a dedicated system user (e.g., "Agentforce Agent User") with a specific profile or permission set. To execute a flow as part of an Agent Action, this user must have the "Run Flows" permission, either via its profile or a permission set (e.g., Agentforce Service Permissions). This ensures the agent can invoke the custom flow for processing returns, aligning with Salesforce's security model and Agentforce setup requirements. This is the correct answer.

Why Option C is Correct:Granting the "Run Flows" permission to the Agentforce Agent user is the standard, documented step to enable flow execution in Agent Actions, ensuring the Service Agent can process returns as intended.

References:

? Salesforce Agentforce Documentation: Agent Builder > Custom Actions – Requires "Run Flows" for flow-based actions.

? Trailhead: Set Up Agentforce Service Agents – Lists "Run Flows" in agent user permissions.

? Salesforce Help: Agentforce Security > Permissions – Confirms flow execution needs.

NEW QUESTION 153

Universal Containers (UC) wants to make a sales proposal and directly use data from multiple unrelated objects (standard and custom) in a prompt template. How should UC accomplish this?

- A. Create a prompt template passing in a special custom object that connects the records temporarily.
- B. Create a prompt template-triggered flow to access the data from standard and custom objects.
- C. Create a Flex template to add resources with standard and custom objects as inputs.
- D. Use a Record Snapshot to combine data from unrelated objects into a single prompt.

Answer: C

Explanation:

Comprehensive and Detailed In-Depth Explanation:UC needs to incorporate data from multiple unrelated objects (standard and custom) into a prompt template for a sales proposal. Let's evaluate the options based on Agentforce capabilities.

? Option A: Create a prompt template passing in a special custom object that connects the records temporarily.While a custom object could theoretically act as a junction to link unrelated records, this approach requires additional setup (e.g., creating the object, populating it with data via automation), and there's no direct

mechanism in Prompt Builder to "pass in" such an object to a prompt template without grounding or flow support. This is inefficient and not a native feature, making it incorrect.

? Option B: Create a prompt template-triggered flow to access the data from standard and custom objects. There's no such thing as a "prompt template-triggered flow" in Salesforce. Flows can invoke prompt templates (e.g., via the "Prompt Template" action), but the reverse—triggering a flow from a prompt template—is not a standard construct. While a flow could gather data from unrelated objects and pass it to a prompt, this option's terminology is inaccurate, and it's not the most direct solution, making it incorrect.

? Option C: Create a Flex template to add resources with standard and custom objects as inputs. In Agentforce's Prompt Builder, a Flex template (short for Flexible Prompt Template) allows users to define dynamic inputs, including data from multiple Salesforce objects (standard or custom), even if they're unrelated. Resources can be added to the template (e.g., via merge fields or Data Cloud queries), enabling the prompt to pull data directly from specified objects without requiring a junction object or complex flows. This is ideal for generating a sales proposal using disparate data sources and aligns with Salesforce's documentation on Flex templates, making it the correct answer.

Why Option C is Correct: Flex templates are designed for scenarios requiring flexible data inputs, allowing UC to directly reference multiple unrelated objects in the prompt template. This simplifies the process and leverages Prompt Builder's native capabilities, as outlined in Salesforce documentation.

References:

? Salesforce Agentforce Documentation: Prompt Builder > Flex Templates – Describes adding multiple object resources as inputs.

? Trailhead: Build Prompt Templates in Agentforce – Highlights Flex templates for dynamic data scenarios.

? Salesforce Help: Create Flexible Prompts – Confirms support for standard and custom object data.

NEW QUESTION 158

How does Secure Data Retrieval ensure that only authorized users can access necessary Salesforce data for dynamic grounding?

- A. Retrieves Salesforce data based on the 'Run As' users permissions.
- B. Retrieves Salesforce data based on the user's permissions executing the prompt.
- C. Retrieves Salesforce's data based on the Prompt template's object permissions.

Answer: B

Explanation:

Secure Data Retrieval enforces Salesforce's security model by dynamically grounding data access in the permissions of the user executing the prompt. This ensures compliance with CRUD (Create, Read, Update, Delete) and FLS (Field-Level Security) settings, preventing unauthorized access to sensitive data. For example, if a user lacks access to a specific object or field, the AI model cannot retrieve it for dynamic grounding.

? "Run As" user permissions (A) would bypass user-specific security, posing a compliance risk.

? Prompt template permissions (C) are not a Salesforce security mechanism; access is always tied to the user's profile and sharing settings.

Reference:

Salesforce Help Article: Secure Data Retrieval in Einstein Trust Layer ("User Context Enforcement" section).

Einstein Trust Layer Technical Guide: "Dynamic Grounding and Data Security" (User Permissions alignment).

NEW QUESTION 160

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 162

Universal Containers is considering leveraging the Einstein Trust Layer in conjunction with Einstein Generative AI Audit Data.

Which audit data is available using the Einstein Trust Layer?

- A. Response accuracy and offensiveness score
- B. Hallucination score and bias score
- C. Masked data and toxicity score

Answer: C

Explanation:

Universal Containers is considering the use of the Einstein Trust Layer along with Einstein Generative AI Audit Data. The Einstein Trust Layer provides a secure and compliant way to use AI by offering features like data masking and toxicity assessment. The audit data available through the Einstein Trust Layer includes information about masked data—which ensures sensitive information is not exposed—and the toxicity score, which evaluates the generated content for inappropriate or harmful language. References:
? Salesforce Agentforce Specialist Documentation - Einstein Trust Layer: Details the auditing capabilities, including logging of masked data and evaluation of generated responses for toxicity to maintain compliance and trust.

NEW QUESTION 167

An Agentforce at Universal Containers (UC) is building with no-code tools only. They have many small accounts that are only touched periodically by a specialized sales team, and UC wants to maximize the sales operations team's time. UC wants to help prep the sales team for the calls by summarizing past purchases, interests in products shown by the Contact captured via Data Cloud, and a recap of past email and phone conversations for which there are transcripts. Which approach should the Agentforce Specialist recommend to achieve this use case?

- A. Use a prompt template grounded on CRH and Data Cloud data using standard foundation model.
- B. Fine-Tune the standard foundational model due to the complexity of the data.
- C. Deploy UC's own custom foundational model on this data first.

Answer: A

Explanation:

For no-code implementations, Prompt Builder allows Agentforce Specialists to create prompt templates that dynamically ground responses in Salesforce CRM data (e.g., past purchases) and Data Cloud insights (e.g., product interests) without custom coding. The standard foundation model (e.g., Einstein GPT) can synthesize this data into summaries, leveraging structured and unstructured sources (e.g., email/phone transcripts). Fine-tuning (B) or custom models (C) require code and are unnecessary here, as the use case does not involve unique data patterns requiring model retraining.

Reference:

Salesforce Help Article: Prompt Builder for No-Code AI ("Grounding in CRM and Data Cloud" section).

Einstein GPT Implementation Guide: "Generating Summaries with Pre-Built Models."

NEW QUESTION 170

An Agentforce has created a copilot custom action using flow as the reference action type. However, it is not delivering the expected results to the conversation preview, and therefore needs troubleshooting. What should the Agentforce Specialist do to identify the root cause of the problem?

- A. In Copilot Builder within the Dynamic Panel, turn on dynamic debugging to show the inputs and outputs.
- B. Copilot Builder within the Dynamic Panel, confirm selected action and observe the values in Input and Output sections.
- C. In Copilot Builder, verify the utterance entered by the user and review session event logs for debug information.

Answer: A

Explanation:

When troubleshooting a copilot custom action using flow as the reference action type, enabling dynamic debugging within Copilot Builder's Dynamic Panel is the most effective way to identify the root cause. By turning on dynamic debugging, the Agentforce Specialist can see detailed logs showing both the inputs and outputs of the flow, which helps identify where the action might be failing or not delivering the expected results.

? Option B, confirming selected actions and observing the Input and Output

sections, is useful for monitoring flow configuration but does not provide the deep diagnostic details available with dynamic debugging.

? Option C, verifying the user utterance and reviewing session event logs, could

provide helpful context, but dynamic debugging is the primary tool for identifying issues with inputs and outputs in real time.

Salesforce Agentforce Specialist References: To explore more about dynamic debugging in Copilot Builder, see:

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

NEW QUESTION 173

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

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

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

Answer: C

Explanation:

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

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

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

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

NEW QUESTION 177

A sales rep at Universal Containers is extremely busy and sometimes will have very long sales calls on voice and video calls and might miss key details. They are just starting to adopt new generative AI features.

Which Einstein Generative AI feature should An Agentforce recommend to help the rep get the details they might have missed during a conversation?

- A. Call Summary
- B. Call Explorer
- C. Sales Summary

Answer: A

Explanation:

For a sales rep who may miss key details during long sales calls, the Agentforce Specialist should recommend the Call Summary feature. Call Summary uses Einstein Generative AI to automatically generate a concise summary of important points discussed during the call, helping the rep quickly review the key information they might have missed.

? Call Explorer is designed for manually searching through call data but doesn't summarize.

? Sales Summary is focused more on summarizing overall sales activity, not call- specific content.

For more details, refer to Salesforce's Call Summary documentation on how AI- generated summaries can improve sales rep productivity.

NEW QUESTION 180

Universal Containers (UC) is building a Flex prompt template. UC needs to use data returned by the flow in the prompt template.

Which flow element should UC use?

- A. Add Flex Instructions
- B. Add Prompt Instructions
- C. Add Flow Instructions

Answer: C

Explanation:

? Context of the Question

? Why ??Add Flow Instructions?? Is Needed

? Outcome

Salesforce Agentforce Specialist References & Documents

? Salesforce Help & Training: Using Prompt Templates with FlowExplains how to pass Flow variables into a prompt template via a specialized step (e.g., ??Flow Instructions??).

? Salesforce Agentforce Specialist Study GuideOutlines how to configure generative

AI prompts that reference real-time Flow data.

NEW QUESTION 183

Universal Containers (UC) wants to improve the productivity of its sales team with generative AI technology. However, UC is concerned that public AI virtual assistants lack adequate company data to general useful responses.

Which solution should UC consider?

- A. fine-tune the Einstein AI model with CBM data.
- B. Build AI model with Einstein discovery and deploy to sales users.
- C. Enable Agentforce and deploy to sales users.

Answer: A

Explanation:

? Context of the QUESTION NO: Universal Containers (UC) wants to harness generative AI to boost sales productivity. They are wary of public AI virtual assistants (like generic chatbots) that lack sufficient UC-specific data to generate useful business responses.

? Why Fine-Tune an Einstein AI Model with CRM Data?

? Why Not Build an AI Model with Einstein Discovery (Option B)?

? Why Not Enable Agentforce (Option C)?

? Outcome: Fine-tuning the Einstein AI model with UC??s CRM data (Answer A) is the most direct, Salesforce-native solution to provide generative AI responses that are aligned with UC??s context, driving productivity gains and ensuring data privacy.

Salesforce Agentforce Specialist References & Documents

? Salesforce Official: Einstein GPT Overview

? Salesforce Trailhead: Get Started with Salesforce Einstein

? Salesforce Documentation: Einstein Discovery

? Salesforce Agentforce Specialist Study Guide

NEW QUESTION 188

Universal Containers (UC) currently tracks Leads with a custom object. UC is preparing to implement the Sales Development Representative (SDR) Agent. Which consideration should UC keep in mind?

- A. Agentforce SDR only works with the standard Lead object.
- B. Agentforce SDR only works on Opportunities.
- C. Agentforce SDR only supports custom objects associated with Accounts.

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation:Universal Containers (UC) uses a custom object for Leads and plans to implement the Agentforce Sales Development Representative (SDR) Agent. The SDR Agent is a prebuilt, configurable AI agent designed to assist sales teams by qualifying leads and scheduling meetings. Let??s evaluate the options based on its functionality and limitations.

? Option A: Agentforce SDR only works with the standard Lead object.Per Salesforce documentation, the Agentforce SDR Agent is specifically designed to interact with the standard Lead object in Salesforce. It includes preconfigured logic to qualify leads, update lead statuses, and schedule meetings, all of which rely on standard Lead fields (e.g., Lead Status, Email, Phone). Since UC tracks leads in a custom object, this is a critical consideration—they would need to migrate data to the standard Lead object or create a workaround (e.g., mapping custom object data to Leads) to leverage the SDR Agent effectively. This limitation is accurate and aligns with the SDR Agent??s out-of-the-box capabilities.

? Option B: Agentforce SDR only works on Opportunities.The SDR Agent??s primary focus is lead qualification and initial engagement, not opportunity management. Opportunities are handled by other roles (e.g., Account Executives) and potentially other Agentforce agents (e.g., Sales Agent), not the SDR Agent. This option is incorrect, as it misaligns with the SDR Agent??s purpose.

? Option C: Agentforce SDR only supports custom objects associated with Accounts.There??s no evidence in Salesforce documentation that the SDR Agent supports custom objects, even those related to Accounts. The SDR Agent is tightly coupled with the standard Lead object and does not natively extend to custom

objects, regardless of their relationships. This option is incorrect.

Why Option A is Correct:The Agentforce SDR Agent??s reliance on the standard Lead object is a documented constraint. UC must consider this when planning implementation, potentially requiring data migration or process adjustments to align their custom object with the SDR Agent??s capabilities. This ensures the agent can perform its intended functions, such as lead qualification and meeting scheduling.

References:

? Salesforce Agentforce Documentation: SDR Agent Setup – Specifies the SDR Agent??s dependency on the standard Lead object.

? Trailhead: Explore Agentforce Sales Agents – Describes SDR Agent functionality tied to Leads.

? Salesforce Help: Agentforce Prebuilt Agents – Confirms Lead object requirement for SDR Agent.

NEW QUESTION 191

Universal Containers tests out a new Einstein Generative AI feature for its sales team to create personalized and contextualized emails for its customers.

Sometimes, users find that the draft email contains placeholders for attributes that could have been derived from the recipient??s contact record. What is the most likely explanation for why the draft email shows these placeholders?

- A. The user does not have permission to access the fields.
- B. The user??s locale language is not supported by Prompt Builder.
- C. The user does not have Einstein Sales Emails permission assigned.

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation:UC is using an Einstein Generative AI feature (likely Einstein Sales Emails) to draft personalized emails, but placeholders (e.g., {!Contact.FirstName}) appear instead of actual data from the contact record. Let??s analyze the options.

? Option A: The user does not have permission to access the fields.Einstein Sales Emails, built on Prompt Builder, pulls data from contact records to populate email drafts. If the user lacks field-level security (FLS) or object-level permissions to access relevant fields (e.g., FirstName, Email), the system cannot retrieve the data, leaving placeholders unresolved. This is a common issue in Salesforce when permissions restrict data access, making it the most likely explanation and the correct answer.

? Option B: The user??s locale language is not supported by Prompt Builder.Prompt Builder and Einstein Sales Emails support multiple languages, and locale mismatches typically affect formatting or translation, not data retrieval. Placeholders appearing instead of data isn??t a documented symptom of language support issues, making this unlikely and incorrect.

? Option C: The user does not have Einstein Sales Emails permission assigned.The Einstein Sales Emails permission (part of the Einstein Generative AI license) enables the feature itself. If missing, users couldn??t generate drafts at all—not just see placeholders. Since drafts are being created, this permission is likely assigned, making this incorrect.

Why Option A is Correct:Permission restrictions are a frequent cause of unresolved placeholders in Salesforce AI features, as the system respects FLS and sharing rules. This is well-documented in troubleshooting guides for Einstein Generative AI.

References:

? Salesforce Help: Einstein Sales Emails > Troubleshooting – Lists permissions as a cause of data issues.

? Trailhead: Set Up Einstein Generative AI – Emphasizes field access for personalization.

? Agentforce Documentation: Prompt Builder > Data Access – Notes dependency on user permissions.

NEW QUESTION 193

Universal Containers (UC) is using Einstein Generative AI to generate an account summary. UC aims to ensure the content is safe and inclusive, utilizing the Einstein Trust Layer's toxicity scoring to assess the content's safety level.

What does a safety category score of 1 indicate in the Einstein Generative Toxicity Score?

- A. Not safe
- B. Safe
- C. Moderately safe

Answer: B

Explanation:

In the Einstein Trust Layer, the toxicity scoring system is used to evaluate the safety level of content generated by AI, particularly to ensure that it is non-toxic, inclusive, and appropriate for business contexts. A toxicity score of 1 indicates that the content is deemed safe.

The scoring system ranges from 0 (unsafe) to 1 (safe), with intermediate values indicating varying degrees of safety. In this case, a score of 1 means that the generated content is fully safe and meets the trust and compliance guidelines set by the Einstein Trust Layer. For further reference, check Salesforce??s official Einstein Trust Layer documentation regarding toxicity scoring for AI-generated content.

NEW QUESTION 195

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons 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](#)