

Salesforce

Exam Questions Agentforce-Specialist

Salesforce Certified Agentforce Specialist



NEW QUESTION 1

An Agentforce Service Agent, who has been successfully assisting customers with service requests in Salesforce, is now unable to help customers with issues related to a new product replacement process. The company recently implemented a custom Product Replacement object in Salesforce to track and manage these replacements. Which Agentforce Agent User change must be implemented to address this issue?

- A. The permission set group assigned to the Agent User needs to grant access to the Product Replacement flow.
- B. The permission set assigned to the Agent User needs Read access to the custom Product Replacement object.
- C. The profile assigned to the Agentforce Agent User needs AI training permission to the custom Product Replacement object.

Answer: B

Explanation:

Why is "Permission Set Read Access" the correct answer?

If an Agentforce Service Agent is unable to assist customers with the new Product Replacement process, it is likely due to missing object permissions.

Key Considerations for Object Access in Agentforce:

? Custom Objects Require Permission Set Access

? Ensuring Full Data Access for Agents

? Aligning AI and Agent Workflows

Why Not the Other Options?

* A. The permission set group assigned to the Agent User needs to grant access to the Product Replacement flow.

? Incorrect because flow permissions only control automation access, not direct object access.

? If an agent cannot view the object, the flow will not be visible or usable.

* C. The profile assigned to the Agentforce Agent User needs AI training permission to the custom Product Replacement object.

? Incorrect because AI training permissions relate to model learning and improvement, not object visibility.

Agentforce Specialist References

? Salesforce AI Specialist Material confirms that permission sets control object-level access for Agentforce users.

NEW QUESTION 2

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

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

Answer: C

Explanation:

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

? Service Replies generates responses automatically.

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

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

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

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

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

NEW QUESTION 3

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 4

An Agentforce Agent has been developed with multiple topics and Agent Actions that use flows and Apex. Which options are available for deploying these to production?

- A. Deploy the flows and Apex using normal deployment tools and manually create the agent-related items in production.

- B. Use only change sets because the Salesforce CLI does not currently support the deployment of agent-related metadata.
C. Deploy flows, Apex, and all agent-related items using either change sets or the Salesforce CLI/Metadata API.

Answer: C

Explanation:

Why is "Deploy flows, Apex, and all agent-related items using either change sets or the Salesforce CLI/Metadata API" the correct answer?

When deploying an Agentforce Agent with multiple topics and Agent Actions that use flows and Apex, a complete deployment solution is required. Change sets and the Salesforce CLI/Metadata API support the deployment of flows, Apex code, and agent-related metadata.

Key Considerations for Agentforce Deployments:

? Supports Deployment of All Required Components

? Agentforce Metadata Can Be Deployed Using Standard Tools

? Ensures a Complete Migration Without Manual Configuration

Why Not the Other Options?

* A. Deploy the flows and Apex using normal deployment tools and manually create the agent-related items in production.

? Incorrect because manually creating agent-related items in production introduces risk and inconsistency.

? This approach is error-prone and time-consuming, especially for large Agentforce deployments.

* B. Use only change sets because the Salesforce CLI does not currently support the deployment of agent-related metadata.

? Incorrect because Salesforce CLI and Metadata API fully support Agentforce deployments.

? Change sets are useful but limited in large-scale, automated deployments.

Agentforce Specialist References

? Salesforce AI Specialist Material confirms that Agentforce metadata (flows, actions, and topics) can be deployed using Change Sets or the Metadata API.

NEW QUESTION 5

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 6

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 7

Universal Containers (UC) uses Salesforce Service Cloud to support its customers and agents handling cases. UC is considering implementing Agent and

extending Service Cloud to mobile users.
When would Agent implementation be most advantageous?

- A. When the goal is to streamline customer support processes and improve response times
- B. When the main objective is to enhance data security and compliance measures
- C. When the focus is on optimizing marketing campaigns and strategies

Answer: A

Explanation:

Agent implementation would be most advantageous in Salesforce Service Cloud when the goal is to streamline customer support processes and improve response times. Agent can assist agents by providing real-time suggestions, automating repetitive tasks, and generating contextual responses, thus enhancing service efficiency.

? Option B (data security) is not the primary focus of Agent, which is more about improving operational efficiency.

? Option C (marketing campaigns) falls outside the scope of Service Cloud and Agent's primary benefits, which are aimed at improving customer service and case management.

For further reading, refer to Salesforce documentation on Agent for Service Cloud and how it improves support processes.

NEW QUESTION 8

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 9

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

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

Answer: C

Explanation:

? Context of the QuestionUniversal Containers (UC) uses Agentforce, a specialized AI-driven assistant for Salesforce. A user reports that an Agent is unable to delete recent activities.

? Why Agentforce Cannot Delete Records

? Why Other Options Are Incorrect

? ConclusionThe core reason for the issue is that Agentforce does not support a standard Delete Record action (Choice C).

Salesforce Agentforce Specialist References & Documents

? Salesforce Official Documentation – Agentforce(Note: Agentforce may be a pilot or specialized feature; check pilot release notes or official docs for standard actions.)

? Salesforce Agentforce Specialist Study GuideCovers the limitations of certain AI- enabled features regarding record operations.

NEW QUESTION 10

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 10

Universal Containers wants to be able to detect with a high level confidence if content generated by a large language model (LLM) contains toxic language. Which action should an AI Specialist take in the Trust Layer to confirm toxicity is being appropriately managed?

- A. Access the Toxicity Detection log in Setup and export all entries where `isToxicityDetected` is true.
- B. Create a flow that sends an email to a specified address each time the toxicity score from the response exceeds a predefined threshold.
- C. Create a Trust Layer audit report within Data Cloud that uses a toxicity detector type filter to display toxic responses and their respective scores.

Answer: C

Explanation:

To ensure that content generated by a large language model (LLM) is appropriately screened for toxic language, the Agentforce Specialist should create a Trust Layer audit report within Data Cloud. By using the toxicity detector type filter, the report can display toxic responses along with their respective toxicity scores, allowing Universal Containers to monitor and manage any toxic content generated with a high level of confidence.

? Option C is correct because it enables visibility into toxic language detection within the Trust Layer and allows for auditing responses for toxicity.

? Option A suggests checking a toxicity detection log, but Salesforce provides more comprehensive options via the audit report.

? Option B involves creating a flow, which is unnecessary for toxicity detection monitoring.

References:

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

NEW QUESTION 13

An Agentforce is creating a custom action in Agent.

Which option is available for the Agentforce Specialist to choose for the custom copilot action?

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

Answer: C

Explanation:

When creating a custom action in Agent, one of the available options is to use Flows. Flows are a powerful automation tool in Salesforce, allowing the Agentforce Specialist to define custom logic and actions within the Copilot system. This makes it easy to extend Copilot's functionality without needing custom code.

While Apex triggers and SOQL are important Salesforce tools, Flows are the recommended method for creating custom actions within Agent because they are declarative and highly adaptable.

For further guidance, refer to Salesforce Flow documentation and Agent customization resources.

NEW QUESTION 15

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

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

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

Answer: B

Explanation:

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

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

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

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

Reference:

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

NEW QUESTION 17

Universal Containers plans to enhance the customer support team's productivity using AI. Which specific use case necessitates the use of Prompt Builder?

- A. Creating a draft of a support bulletin post for new product patches
- B. Creating an AI-generated customer support agent performance score
- C. Estimating support ticket volume based on historical data and seasonal trends

Answer: A

Explanation:

The use case that necessitates the use of Prompt Builder is creating a draft of a support bulletin post for new product patches. Prompt Builder allows the Agentforce Specialist to create and refine prompts that generate specific, relevant outputs, such as drafting support communication based on product information and patch details.

? Option B (agent performance score) would likely involve predictive modeling, not prompt generation.

? Option C (estimating support ticket volume) would require data analysis and predictive tools, not prompt building.

For more details, refer to Salesforce's Prompt Builder documentation for generative AI content creation.

NEW QUESTION 21

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

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

A. That the field chosen must be a rich text field with 255 characters or more.

B. That the org is set to API version 59 or higher

C. That the Lightning page layout where the field will reside has been upgraded to Dynamic Forms

Answer: B

Explanation:

Before creating a Field Generation prompt template, the Agentforce Specialist must ensure that the Salesforce org is set to API version 59 or higher. This version of the API introduces support for advanced generative AI features, such as enabling fields for generative AI outputs. This is a critical technical requirement for the Field Generation prompt template to function correctly.

? Option A (rich text field requirement) is not necessary for generative AI functionality.

? Option C (Dynamic Forms) does not impact the ability of a field to be generative AI-enabled, although it might enhance the user interface.

For more information, refer to Salesforce documentation on API versioning and Field Generation templates.

NEW QUESTION 26

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

Which data grounding technique should the Agentforce Specialist recommend?

A. External Object Record Merge Fields

B. External Services Merge Fields

C. Apex Merge Fields

Answer: A

Explanation:

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

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

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

? Why External Object Record Merge Fields:

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

? References and Study Resources:

NEW QUESTION 27

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 30

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 33

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 35

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

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

Answer: B

Explanation:

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

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

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

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

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

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

References:

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

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

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

NEW QUESTION 40

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

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

Answer: A

Explanation:

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

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

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

? Option B: Summarize the Case with a standard Agent action.Standard Agent actions (e.g., "Answer Questions") don??t specifically target case summarization—they??re broader in scope. There??s no out-of-the-box "Summarize Case" action that allows customization of the response format, making this insufficient and incorrect.

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

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

References:

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

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

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

NEW QUESTION 45

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 49

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

A. Enable a sales team to use natural language to invoke defined sales tasks grounded in relevant data and be able to ensure company policies are applied

B. conversationally and in the now or work.

C. Enable a sales team by providing them with an interactive step-by-step guide based on business rules to ensure accurate data entry into Salesforce and help close deals faster.

D. Instantly review and read incoming messages or emails that are then logged to the correct opportunity, contact, and account records to provide a full view of customer interactions and communications.

Answer: A

Explanation:

Agentforce Sales Agent is designed to let sales teams perform tasks via natural language commands, leveraging Salesforce data while adhering to policies. For example, agents can ask the AI to "update the opportunity stage to Closed Won" or "generate a quote," with the system enforcing validations and data security.

This use case aligns with Salesforce??s vision of conversational AI streamlining workflows without compromising compliance.

? Step-by-step guides (B) are typically handled by tools like Dynamic Forms or

Guided Selling, not Agentforce.

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

Reference:

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

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

NEW QUESTION 52

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 56

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 59

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

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

A. Assign the Prompt Template User permission set and enable Sales Emails in Setup.

B. Assign the Prompt Template Manager permission set and enable Sales Emails in setup.

C. Assign the Data Cloud Admin permission set and enable Sales Emails in Setup.

Answer: A

Explanation:

To enable Sales Email prompt templates:

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

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

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

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

References:

? Salesforce Help: Prompt Template Permissions

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

? Sales Email Setup outlines enabling the feature in Setup.

NEW QUESTION 64

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

A. Leverage the Agent Large Language Model (LLM) UI and test UC's agents with different utterances prior to activating the agent.

B. Deploy the agent in a QA sandbox environment and review the Utterance Analysis reports to review effectiveness.

C. Create a CSV file with UC's test cases in Agentforce Testing Center using the testing template.

Answer: C

Explanation:

Comprehensive and Detailed In-Depth Explanation: The goal of Universal Containers (UC) is to test its Agentforce agents for effectiveness, reliability, and trust before production deployment, with a focus on efficiently handling a large and repeatable number of utterances. Let's evaluate each option against this requirement and Salesforce's official Agentforce tools and best practices.

? Option A: Leverage the Agent Large Language Model (LLM) UI and test UC's agents with different utterances prior to activating the agent. While Agentforce leverages advanced reasoning capabilities (powered by the Atlas Reasoning Engine), there's no specific "Agent Large Language Model (LLM) UI" referenced in Salesforce documentation for testing agents. Testing utterances directly within an LLM interface might imply manual experimentation, but this approach lacks

scalability and repeatability for a large number of utterances. It's better suited for ad-hoc testing of individual responses rather than systematic evaluation, making it inefficient for UC's needs.

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

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

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

References:

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

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

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

NEW QUESTION 65

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 66

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

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

Answer: B

Explanation:

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

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

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

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

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

References:

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

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

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

NEW QUESTION 71

Universal Containers (UC) recently rolled out Einstein Generative AI capabilities and has created a custom prompt to summarize case records. Users have reported that the case summaries generated are not returning the appropriate information. What is a possible explanation for the poor prompt performance?

- A. The prompt template version is incompatible with the chosen LLM.
- B. The data being used for grounding is incorrect or incomplete.
- C. The Einstein Trust Layer is incorrectly configured.

Answer: B

Explanation:

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

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

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

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

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

References:

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

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

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

NEW QUESTION 74

Universal Containers is rolling out a new generative AI initiative.

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

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

Answer: C

Explanation:

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

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

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

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

Reference:

"Prompt Builder Limitations | Salesforce Documentation" .

NEW QUESTION 78

Universal Containers (UC) has configured an Agentforce Data Library using Knowledge articles. When testing in Agent Builder and the Experience Cloud site, the agent is not responding with grounded Knowledge article information. However, when tested in Prompt Builder, the response returns correctly. What should UC do to troubleshoot the issue?

- A. Create a new permission set that assigns "Manage Knowledge" and assign it to the Agentforce Service Agent User.
- B. Ensure the assigned User permission set includes access to the prompt template used to access the Knowledge articles.
- C. Ensure the Data Cloud User permission set has been assigned to the Agentforce Service Agent User.

Answer: C

Explanation:

Comprehensive and Detailed In-Depth Explanation:UC has set up an Agentforce Data Library with Knowledge articles, and while Prompt Builder retrieves the data correctly, the agent fails to do so in Agent Builder and Experience Cloud. Let's troubleshoot the issue.

? Option A: Create a new permission set that assigns "Manage Knowledge" and assign it to the Agentforce Service Agent User.The "Manage Knowledge" permission is for authoring and managing Knowledge articles, not for reading or retrieving them in an agent context. The Agentforce Service Agent User (a system user) needs read access to Knowledge, not management rights. This option is excessive and irrelevant to the grounding issue, making it incorrect.

? Option B: Ensure the assigned User permission set includes access to the prompt template used to access the Knowledge articles.Prompt templates in Prompt Builder don't require specific permissions beyond general Einstein Generative AI access. Since the Prompt Builder test works, the template and its grounding are accessible to the testing user. The issue lies with the agent's runtime access, not the template itself, making this incorrect.

? Option C: Ensure the Data Cloud User permission set has been assigned to the Agentforce Service Agent User.When Knowledge articles are grounded via an Agentforce Data Library, they are often ingested into Data Cloud for indexing and retrieval. The Agentforce Service Agent User, which runs the agent, needs the "Data Cloud User" permission set (or equivalent) to access Data Cloud resources, including the Data Library. If this permission is missing, the agent cannot retrieve Knowledge article data during runtime (e.g., in Agent Builder or Experience Cloud), even though Prompt Builder (running under a different user context) succeeds. This is a common setup oversight and aligns with the symptoms, making it the correct answer.

Why Option C is Correct:The Agentforce Service Agent User's lack of Data Cloud access explains the failure in agent-driven contexts while Prompt Builder (likely run by an admin with broader permissions) succeeds. Assigning the "Data Cloud User" permission set resolves this, per Salesforce documentation.

References:

? Salesforce Agentforce Documentation: Data Library Setup > Permissions – Requires Data Cloud access for agents.

? Trailhead: Ground Your Agentforce Prompts – Notes Data Cloud User permission for Knowledge grounding.

? Salesforce Help: Agentforce Security > Agent User Setup – Lists required permission sets.

NEW QUESTION 79

A Salesforce Administrator is exploring the capabilities of Agent to enhance user interaction within their organization. They are particularly interested in how Agent processes user requests and the mechanism it employs to deliver responses. The administrator is evaluating whether Agent directly interfaces with a large

language model (LLM) to fetch and display responses to user inquiries, facilitating a broad range of requests from users.
How does Agent handle user requests In Salesforce?

- A. Agent will trigger a flow that utilizes a prompt template to generate the message.
- B. Agent will perform an HTTP callout to an LLM provider.
- C. Agent analyzes the user's request and LLM technology is used to generate and display the appropriate response.

Answer: C

Explanation:

Agent is designed to enhance user interaction within Salesforce by leveraging Large Language Models (LLMs) to process and respond to user inquiries. When a user submits a request, Agent analyzes the input using natural language processing techniques. It then utilizes LLM technology to generate an appropriate and contextually relevant response, which is displayed directly to the user within the Salesforce interface. Option C accurately describes this process. Agent does not necessarily trigger a flow (Option A) or perform an HTTP callout to an LLM provider (Option B) for each user request. Instead, it integrates LLM capabilities to provide immediate and intelligent responses, facilitating a broad range of user requests.

References:

? Salesforce Agentforce Specialist Documentation - Agent Overview: Details how Agent employs LLMs to interpret user inputs and generate responses within the Salesforce ecosystem.

? Salesforce Help - How Agent Works: Explains the underlying mechanisms of how Agent processes user requests using AI technologies.

NEW QUESTION 84

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

- ? Retrieve proprietary policy documents that are stored as PDFs.
 - ? Ensure responses are grounded in approved company data, not generic LLM knowledge.
- What should UC do first?

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

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation: To implement an AI-powered customer service agent that retrieves proprietary policy documents (stored as PDFs) and ensures responses are grounded in approved company data, UC must first establish a foundation for the AI to access and use this data. The Agentforce Data Library (Option A) is the correct starting point. A Data Library allows UC to upload PDFs containing policy documents, index them into Salesforce Data Cloud's vector database, and make them available for AI retrieval. This setup ensures the agent can perform Retrieval-Augmented Generation (RAG), grounding its responses in the specific, approved content from the PDFs rather than relying on generic LLM knowledge, directly meeting UC's requirements. ? Option B: Expanding the AI agent's scope to search all Salesforce records is too broad and unnecessary at this stage. The requirement focuses on PDFs with policy documents, not all Salesforce data (e.g., cases, accounts), making this premature and irrelevant as a first step. ? Option C: "Add the files to the content, and then select the data library option" is vague and not a precise process in Agentforce. While uploading files is part of setting up a Data Library, the phrasing suggests adding files to Salesforce Content (e.g., ContentDocument) without indexing, which doesn't enable AI retrieval. Setting up the Data Library (A) encompasses the full process correctly.

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

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

References:

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

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

? Salesforce Data Cloud Documentation: "Ground AI Responses with Data Cloud"

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

NEW QUESTION 87

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 Trigger Apex 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: SOQL SOQL (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: Flows In 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: JavaScript JavaScript 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 90

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 92

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

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

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

Answer: A

Explanation:

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

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

Key Considerations for This Approach:

? Custom Agent Action Triggers the Flow

? Enables AI-Powered Query Execution

? No Need for Manual Data Entry

Why Not the Other Options?

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

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

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

* C. Create a Flex Prompt Template in Prompt Builder

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

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

Agentforce Specialist References

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

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

NEW QUESTION 96

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 97

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 99

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

Which configuration aspect should UC prioritize?

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

B. Ensure the retriever's filters are defined to limit the scope of each search efficiently.

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

Answer: B

Explanation:

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

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

by applying specific filters.

Key Considerations for Optimizing Retrievers in Agentforce:

? Defining Effective Filters

? Reducing Query Complexity

? Optimizing the Data Indexing Process

Why Not the Other Options?

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

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

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

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

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

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

Agentforce Specialist References

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

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

NEW QUESTION 103

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 106

How is Data Cloud leveraged by the Answer Questions with Knowledge action in Agentforce?

- A. Data Cloud is not required; the articles can be accessed directly from the CRM by the agent.
- B. Data Cloud stores and manages the Indexed Knowledge articles.
- C. Data Cloud provides the real-time data streams that update the Knowledge articles.

Answer: B

Explanation:

How Does Data Cloud Support "Answer Questions with Knowledge" in Agentforce? The Answer Questions with Knowledge action in Agentforce leverages Salesforce Data Cloud to store, manage, and index Knowledge articles used for AI-powered responses.

- ? Data Cloud as the Central Storage for Knowledge Articles
- ? Ensuring Up-to-Date Responses
- ? Enhancing AI-Driven Customer Service

Why Not the Other Options?

- * A. Data Cloud is not required; the articles can be accessed directly from the CRM by the agent.
 - ? Incorrect because Data Cloud is the primary system for storing and indexing Knowledge articles.
 - ? Without Data Cloud, Einstein AI cannot efficiently retrieve and rank articles dynamically.
- * C. Data Cloud provides the real-time data streams that update the Knowledge articles.
 - ? Incorrect because while Data Cloud stores and manages articles, real-time updates are not its primary function.
 - ? The Knowledge Management system within Salesforce handles article creation and updates.

Agentforce Specialist References

- ? Salesforce AI Specialist Material highlights that Data Cloud is the core storage system for AI-driven Knowledge management.
- ? Salesforce Instructions for Certification confirm the central role of Data Cloud in managing indexed Knowledge articles for AI-powered responses.

NEW QUESTION 108

Universal Containers (UC) needs to save agents time with AI-generated case summaries. UC has implemented the Work Summary feature. What does Einstein consider when generating a summary?

- A. Generation is grounded with conversation context, Knowledge articles, and cases.
- B. Generation is grounded with existing conversation context only.
- C. Generation is grounded with conversation context and Knowledge articles.

Answer: A

Explanation:

When generating a Work Summary, Einstein leverages multiple sources of information to provide a comprehensive and accurate case summary for agents.

- ? Conversation Context:
- ? Knowledge Articles:
- ? Cases:
- ? Option A is correct as it includes all three: conversation context, Knowledge articles, and cases.
- ? Option B is incorrect because it limits the grounding to conversation context only, excluding other critical elements.
- ? Option C is incorrect because it omits case data, which Einstein considers for more accurate and contextually rich summaries.

Reference:

"Einstein Work Summary and AI Case Management | Salesforce Trailhead" .

NEW QUESTION 110

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

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

Answer: C

Explanation:

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

NEW QUESTION 115

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 117

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

What should UC do to ensure the new prompt template produces results comparable to the standard Sales Email prompts?

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

Answer: B

Explanation:

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

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

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

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

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

NEW QUESTION 120

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 122

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 126

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

Which AI capability should UC use?

A. Einstein Email Replies

B. Einstein Service Replies for Email

C. Einstein Generative Service Replies for Email

Answer: B

Explanation:

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

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

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

References:

Einstein Service Replies Overview:

NEW QUESTION 128

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 129

What is the primary function of the planner service in the Agent system?

- A. Generating record queries based on conversation history
- B. Offering real-time language translation during conversations
- C. Identifying copilot actions to respond to user utterances

Answer: C

Explanation:

The primary function of the planner service in the Agent system is to identify copilot actions that should be taken in response to user utterances. This service is responsible for analyzing the conversation and determining the appropriate actions (such as querying records, generating a response, or taking another action) that the Agent should perform based on user input.

NEW QUESTION 132

An Agentforce at Universal Containers is trying to set up a new Field Generation prompt template. They take the following steps.

- * 1. Create a new Field Generation prompt template.
- * 2. Choose Case as the object type.
- * 3. Select the custom field AI_Analysis_c as the target field.

After creating the prompt template, the Agentforce Specialist saves, tests, and activates it. However, when they go to a case record, the AI Analysis field does not show the (Sparkle) icon on the Edit pencil. When the Agentforce Specialist was editing the field, it was behaving as a normal field.

Which critical step did the Agentforce Specialist miss?

- A. They forgot to reactivate the Lightning page layout for the Case object after activating their Field Generation prompt template.
- B. They forgot that the Case Object is not supported for Add generation as Einstein Service Replies should be used instead.
- C. They forgot to edit the Lightning page layout and associate the field to a prompt template

Answer: C

Explanation:

For Field Generation prompt templates to display the Sparkle icon (indicating AI-generated content), the target field must be explicitly associated with the prompt template on the Lightning page layout. Even if the prompt template is activated, failing to add the field to the page layout and link it to the template will result in the field behaving as a standard field. Salesforce documentation emphasizes that page layout configuration is mandatory to enable AI-driven field interactions.

? Reactivating the layout (A) is unnecessary unless the layout itself was modified after activation.

? Case objects are supported for Field Generation (B is incorrect).

Reference:

Salesforce Help Article: Configure Field Generation Prompt Templates ("Associating Fields with Page Layouts" section).

Einstein GPT Implementation Guide: "Enabling AI-Generated Fields in Lightning Pages."

NEW QUESTION 135

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 140

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 144

Universal Containers (UC) wants to limit an agent's access to Knowledge articles while deploying the "Answer Questions with Knowledge" action. How should UC achieve this?

- A. Define scope instructions to the agent specifying a list of allowed article titles or IDs.
- B. Update the Data Library Retriever to filter on a custom field on the Knowledge article.
- C. Assign Data Categories to Knowledge articles, and define Data Category filters in the Agentforce Data Library.

Answer: C

Explanation:

Comprehensive and Detailed In-Depth Explanation: UC wants to restrict the "Answer Questions with Knowledge" action to a subset of Knowledge articles. Let's evaluate the options for scoping agent access.

? Option A: Define scope instructions to the agent specifying a list of allowed article titles or IDs. Agent instructions in Agent Builder guide behavior but cannot enforce granular data access restrictions like a specific list of article titles or IDs. This approach is impractical and bypasses Salesforce's security model, making it incorrect.

? Option B: Update the Data Library Retriever to filter on a custom field on the Knowledge article. While Data Library Retrievers in Data Cloud can filter data, this requires custom development (e.g., modifying indexing logic) and assumes articles are ingested with a custom field for filtering. This is less straightforward than native Knowledge features and not a standard option, making it incorrect.

? Option C: Assign Data Categories to Knowledge articles, and define Data Category filters in the Agentforce Data Library. Salesforce Knowledge uses Data Categories to organize articles (e.g., by topic or type). In Agentforce, when configuring a Data Library with Knowledge, you can apply Data Category filters to limit which articles the agent accesses. For the "Answer Questions with Knowledge" action, this ensures the agent only retrieves articles within the specified categories, aligning with UC's goal. This is a native, documented solution, making it the correct answer.

Why Option C is Correct: Using Data Categories and filters in the Data Library is the recommended, scalable way to limit Knowledge article access for agent actions, as per Salesforce documentation.

References:

? Salesforce Agentforce Documentation: Data Library > Knowledge Filters – Describes Data Category filtering.

? Trailhead: Ground Your Agentforce Prompts – Covers limiting Knowledge scope.

? Salesforce Help: Knowledge in Agentforce – Recommends categories for access control.

NEW QUESTION 148

What is best practice when refining Agent custom action instructions?

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

Answer: A

Explanation:

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

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

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

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

NEW QUESTION 149

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 154

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 158

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

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

Answer: B

Explanation:

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

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

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

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

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

References:

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

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

NEW QUESTION 160

Universal Containers (UC) is tracking web activities in Data Cloud for a unified contact, and wants to use that in a prompt template to help extract insights from the data.

Assuming that the Contact object is one of the objects associated with the prompt template, what is a valid way for DC to do this?

- A. Call the prompt directly from Data Cloud with a web tracing activity included in the prompt definition.
- B. Add the activity records as an enrichment related list to the Contact then pass the Contact into a prompt template workspace using related list grounding.
- C. Create a prompt template that takes a list of all Data Cloud activity records as input to pass to the large language model (LLM).

Answer: B

Explanation:

To integrate web activity data from Data Cloud into a prompt template, the correct approach is to enrich the Contact object with the activity records as a related list and use related list grounding (Option B). Here's why:

? Data Cloud Integration: Data Cloud unifies web activity data and associates it with the unified Contact record. By adding these activities as a related list to the Contact, the data becomes accessible to the prompt template.

? Prompt Template Grounding: Salesforce prompt templates support grounding on related records. When the Contact is passed to the prompt template, the template can reference the related web activity records (via the related list) to extract insights.

? Structured Data Handling: This method aligns with Salesforce best practices for grounding, ensuring the large language model (LLM) receives structured, context-rich data without overwhelming it with raw activity lists.

Why Other Options Are Incorrect:

? A. Calling the prompt directly from Data Cloud: Prompt templates are invoked within Salesforce, not directly from Data Cloud. Grounding requires associating data with Salesforce objects, not ad-hoc web activity inclusion.

? C. Passing a list of activity records as input: While technically possible, this bypasses Salesforce's grounding framework, which relies on object relationships. It also risks exceeding LLM input limits and lacks scalability.

References:

? Salesforce Data Cloud Implementation Guide: Explains how to enrich standard/custom objects with related data for AI use cases.

? Prompt Template Documentation: Highlights grounding on related lists to leverage contextual data for LLM prompts.

? Trailhead Module: "Einstein Prompt Builder Basics" demonstrates grounding techniques using related records.

NEW QUESTION 164

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 168

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

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

Answer: B

Explanation:

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

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

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

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

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

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

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

References:

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

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

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

NEW QUESTION 170

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

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

Answer: C

Explanation:

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

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

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

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

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

References:

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

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

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

NEW QUESTION 172

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 173

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

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

Answer: C

Explanation:

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

Key Steps in the Data Preparation Process for Agentforce:

? Loading Data

? Chunking (Breaking Text into Small Parts)

? Vectorization (Transforming Text for AI Retrieval)

? Storing in a Vector Database

Why Not the Other Options?

* A. Real-time data ingestion and dynamic indexing

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

* B. Aggregating, normalizing, and encoding structured datasets

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

Agentforce Specialist References

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

NEW QUESTION 175

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 Salesforces 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 178

An Agentforce at Universal Containers is working on a prompt template to generate personalized emails for product demonstration requests from customers. It is important for the AI-generated email to adhere strictly to the guidelines, using only associated opportunity information, and to encourage the recipient to take the

desired action.

How should the Agentforce Specialist include these instructions on a new line in the prompt template?

- A. Surround them with triple quotes (""").
- B. Make sure merged fields are defined.
- C. Use curly brackets {} to encapsulate instructions.

Answer: A

Explanation:

In Salesforce prompt templates, instructions that guide how the Large Language Model (LLM) should generate content (in this case, personalized emails) can be included by surrounding the instruction text with triple quotes ("""). This formatting ensures that the LLM adheres to the specific instructions while generating the email content.

The use of triple quotes allows the AI to understand that the enclosed text is a directive for how to approach the task, such as limiting the content to associated opportunity information or encouraging a specific action from the recipient.

Refer to Salesforce Prompt Builder documentation for detailed instructions on how to structure prompts for generative AI.

NEW QUESTION 179

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

Agentforce-Specialist Practice Exam Features:

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

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