

Exam Questions Agentforce-Specialist

Salesforce Certified Agentforce Specialist

<https://www.2passeasy.com/dumps/Agentforce-Specialist/>



NEW QUESTION 1

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

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

Answer: A

Explanation:

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

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

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

? Why not B? Using Prompt Builder in Setup to send prompts to the LLM is for

creating and managing prompts, not for auditing or tracking data. It does not interact directly with the audit trail functionality.

? Why not C? Although the audit trail can be accessed in Setup, the user-generated

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

NEW QUESTION 2

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

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

Answer: A

Explanation:

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

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

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

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

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

References:

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

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

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

NEW QUESTION 3

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

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

Answer: C

Explanation:

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

? Service Replies generates responses automatically.

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

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

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

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

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

NEW QUESTION 4

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

- A. Run a report on the Platform Debug Logs.

- B. Query the Agent log data using the Metadata API.
- C. Run Agent Analytics.

Answer: C

Explanation:

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

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

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

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

References:

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

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

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

NEW QUESTION 5

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

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

Answer: C

Explanation:

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

? Why Agentforce Cannot Delete Records

? Why Other Options Are Incorrect

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

Salesforce Agentforce Specialist References & Documents

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

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

NEW QUESTION 6

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

- A. Find similar requested topics and provide the actions that need to be executed.
- B. Identify the best matching topic and actions and correct order of execution.
- C. Determine a user's topic access and sort actions by priority to be executed.

Answer: B

Explanation:

Comprehensive and Detailed In-Depth Explanation: In Agentforce, the large language model (LLM), powered by the Atlas Reasoning Engine, interprets user requests and drives Agent Actions. Let's evaluate its role.

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

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

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

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

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

References:

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

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

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

NEW QUESTION 7

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

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

Answer: C

Explanation:

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

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

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

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

References:

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

NEW QUESTION 8

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

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

Answer: A

Explanation:

Why is Outbound Omni-Channel Flow the Correct Answer?

In Agentforce, when a service agent's conversation needs to be escalated to a human agent queue, Outbound Omni-Channel Flow is the appropriate tool to facilitate this process.

Key Features of Outbound Omni-Channel Flow in Agentforce:

? Automates Escalation to a Human Agent Queue

? Seamless Transition from AI to Human Agents

? Ensures Proper Prioritization & Load Balancing

? Integration with Agentforce and Service Cloud

Why Not the Other Options?

* B. Screen Flow

? Screen Flow is used for interactive guided processes where users manually enter data in predefined steps.

? It does not support automated case routing to human agents in real time.

* C. Prompt Flow

? Prompt Flow is designed to enhance AI-generated responses and workflows rather than routing service agent interactions to human agents.

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

Agentforce Specialist References

The importance of using Omni-Channel Flow for routing AI-generated interactions to

human agents is supported in the Agentforce Specialist exam objectives and documentation:

? Salesforce AI Specialist Material: Covers the importance of Omni-Channel routing

for managing AI and human agent interactions.

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

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

NEW QUESTION 9

Universal Containers?? data science team is hosting a generative large language model (LLM) on Amazon Web Services (AWS).

What should the team use to access externally-hosted models in the Salesforce Platform?

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

Answer: A

Explanation:

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

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

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

Model Builder enables seamless integration with external systems and models, allowing Salesforce users to use external LLMs for generating AI-driven insights and automation. Salesforce Agentforce Specialist References:For more details, check the Model Builder guide here:

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

NEW QUESTION 10

An Agentforce is creating a custom action for Agentforce.

Which setting should the Agentforce Specialist test and iterate on to ensure the action performs as expected?

- A. Action Name
- B. Action Input
- C. Action Instructions

Answer: C

Explanation:

When creating a custom action for Einstein Bots in Salesforce (including Agentforce), Action Instructions are critical for defining how the bot processes and executes the action. These instructions guide the bot on the logic to follow, such as API calls, data transformations, or conditional steps. Testing and iterating on the instructions ensures the bot understands how to handle dynamic inputs, external integrations, and decision-making.

Salesforce documentation emphasizes that Action Instructions directly impact the bot's ability to execute workflows accurately. For example, poorly defined instructions may lead to incorrect API payloads or failure to parse responses. The Einstein Bot Developer Guide highlights that refining instructions is essential for aligning the bot's behavior with business requirements.

In contrast:

? Action Name (A) is a static identifier and does not affect functionality.

? Action Input (B) defines parameters passed to the action but does not dictate execution logic.

Thus, iterating on Action Instructions (C) ensures the action performs as expected.

Reference:

Salesforce Help Article: Create Custom Actions for Einstein Bots

Einstein Bot Developer Guide: "Custom Action Configuration Best Practices" (Section 4.3).

NEW QUESTION 10

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

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

Answer: B

Explanation:

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

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

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

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

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

References:

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

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

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

NEW QUESTION 15

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 18

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

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

Answer: B

Explanation:

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

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

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

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

incorrect.

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

References:

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

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

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

NEW QUESTION 22

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 24

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 25

Universal Containers has a custom Agent action calling a flow to retrieve the real-time status of an order from the order fulfillment system.

For the given flow, what should the Agentforce Specialist consider about the running user's data access?

- A. The flow must have the "with sharing" permission selected in the advanced settings for the permissions, field-level security, and sharing settings to be respected.
- B. The custom action adheres to the permissions, field-level security, and sharing settings configured in the flow.
- C. The Agent will always run flows in system mode so the running user's data access will not affect the data returned.

Answer: B

Explanation:

When a flow is invoked via a custom Agent action, its data access depends on the flow's runtime configuration, not system mode by default. Salesforce flows can be configured to respect the running user's permissions and sharing settings:

? If the flow is set to "Run as the User Who Launched the Flow" (enabled in Flow Settings), it adheres to the user's permissions, field-level security (FLS), and sharing rules.

? Option C is incorrect because flows do not always run in system mode unless explicitly configured to do so.

? Option A is misleading because "with sharing" is an Apex concept, not a flow setting. Flows use runtime settings like FLS and sharing enforcement.

References:

? Salesforce Help: Flow Runtime and Security Context

? Flow Settings: "Run with User Permission and Field-Level Security" ensures data access aligns with the user's permissions.

NEW QUESTION 28

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 32

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 34

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

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

Answer: A

Explanation:

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

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

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

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

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

References:

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

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

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

NEW QUESTION 39

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 43

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

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

Answer: B

Explanation:

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

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

AI-driven responses based on knowledge articles.

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

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

References:

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

NEW QUESTION 48

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

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

Answer: A

Explanation:

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

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

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

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

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

References:

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

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

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

NEW QUESTION 50

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

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

Answer: A

Explanation:

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

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

Guided Selling, not Agentforce.

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

Reference:

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

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

NEW QUESTION 51

Universal Containers implemented Agent for its users.

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

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

Answer: C

Explanation:

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

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

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

References:

? Salesforce Agentforce Specialist Documentation - Agent Supported Actions:

? Salesforce Help - Limitations of Agent:

NEW QUESTION 52

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 56

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

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

Answer: B

Explanation:

? Context of the Question

? Why Apex Code for JSON Formatting?

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

Salesforce Agentforce Specialist References & Documents

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

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

NEW QUESTION 58

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 59

An AI Specialist is tasked with configuring a generative model to create personalized sales emails using customer data stored in Salesforce. The AI Specialist has already fine-tuned a large language model (LLM) on the OpenAI platform. Security and data privacy are critical concerns for the client.

How should the Agentforce Specialist integrate the custom LLM into Salesforce?

- A. Create an application of the custom LLM and embed it in Sales Cloud via iFrame.
- B. Add the fine-tuned LLM in Einstein Studio Model Builder.
- C. Enable model endpoint on OpenAI and make callouts to the model to generate emails.

Answer: B

Explanation:

Since security and data privacy are critical, the best option for the Agentforce Specialist is to integrate the fine-tuned LLM (Large Language Model) into Salesforce by adding it to Einstein Studio Model Builder. Einstein Studio allows organizations to bring their own AI models (BYOM), ensuring the model is securely managed within Salesforce's environment, adhering to data privacy standards.

? Option A (embedding via iFrame) is less secure and doesn't integrate deeply with Salesforce's data and security models.

? Option C (making callouts to OpenAI) raises concerns about data privacy, as sensitive Salesforce data would be sent to an external system.

Einstein Studio provides the most secure and seamless way to integrate custom AI models while maintaining control over data privacy and compliance. More details can be found in Salesforce's Einstein Studio documentation on integrating external models.

NEW QUESTION 64

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 68

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

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

Answer: C

Explanation:

Comprehensive and Detailed In-Depth Explanation:Salesforce Agentforce leverages generative AI to enhance service agent efficiency, particularly through capabilities that generate recommended replies. In this scenario, Universal Containers aims to reduce handling time by providing replies based on existing Knowledge articles, which are a core component of Salesforce Knowledge. The Knowledge Replies capability is specifically designed for this purpose—it uses generative AI to analyze Knowledge articles, match them to the context of a customer inquiry (e.g., a case or chat), and suggest relevant, pre-formulated responses for service agents to use or adapt. This aligns directly with UC's goal of leveraging existing content to streamline agent workflows.

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

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

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

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

References:

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

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

&type=5)

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

NEW QUESTION 69

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

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

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

Answer: B

Explanation:

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

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

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

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

References:

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

NEW QUESTION 72

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 77

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 81

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 82

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

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

A. Use the merge fields to reference a custom related list of opportunities.

B. Use merge fields to reference the default related list of opportunities.

C. Use formula fields to reference the Einstein related list of opportunities.

Answer: B

Explanation:

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

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

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

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

Salesforce References:

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

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

NEW QUESTION 83

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

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

A. Define filters to narrow the search results based on specific conditions.

B. Change the search index to a different data model object (DMO).

C. Increase the maximum number of results returned to capture a broader dataset.

Answer: A

Explanation:

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

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

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

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

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

References:

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

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

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

NEW QUESTION 88

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 89

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

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

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

Answer: B

Explanation:

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

- ? Effective grounding: Ensuring the prompt pulls data from the correct sources.
- ? Contextual role-playing: Providing the AI with a clear understanding of its role in generating the summary.
- ? Clear instructions: Giving unambiguous directions on what to include in the response.
- ? Iterative feedback: Regularly testing and adjusting prompts based on user feedback.
- ? Option B is correct because it follows industry best practices for refining prompt design.
- ? Option A (prompt test mode) is useful but less relevant for refining prompt design itself.
- ? Option C (prompt review case with Salesforce) would be more appropriate for technical issues or complex prompt errors, not general design refinement.

References:

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

NEW QUESTION 90

Universal Containers (UC) noticed an increase in customer contract cancellations in the last few months. UC is seeking ways to address this issue by implementing a proactive outreach program to

customers before they cancel their contracts and is asking the Salesforce team to provide suggestions.

Which use case functionality of Model Builder aligns with UC's request?

- A. Product recommendation prediction
- B. Customer churn prediction
- C. Contract Renewal Date prediction

Answer: B

Explanation:

Customer churn prediction is the best use case for Model Builder in addressing Universal Containers' concerns about increasing customer contract cancellations. By implementing a model that predicts customer churn, UC can proactively identify customers who are at risk of canceling and take action to retain them before they decide to terminate their contracts. This functionality allows the business to forecast churn probability based on historical data and initiate timely outreach programs.

- ? Option B is correct because customer churn prediction aligns with UC's need to reduce cancellations through proactive measures.
- ? Option A (product recommendation prediction) is unrelated to contract cancellations.
- ? Option C (contract renewal date prediction) addresses timing but does not focus on predicting potential cancellations.

References:

? Salesforce Model Builder Use Case Overview: https://help.salesforce.com/s/articleView?id=sf.model_builder_use_cases.htm

NEW QUESTION 93

Universal Container (UC) has effectively utilized prompt templates to update summary fields on Lightning record pages. An admin now wishes to incorporate similar functionality into UC's automation process using Flow.

How can the admin get a response from this prompt template from within a flow to use as part of UC's automation?

- A. Invocable Apex
- B. Flow Action
- C. Einstein for Flow

Answer: C

Explanation:

* 1. Context of the Question

- o Universal Container (UC) has used prompt templates to update summary fields on record pages.
- o Now, the admin wants to incorporate similar generative AI functionality within a Flow for automation purposes.

* 2. How to Call a Prompt Template Within a Flow

- o Flow Action: Salesforce provides a standard way to invoke generative AI templates or prompts within a Flow step. From the Flow Builder, you can add an ??Action?? that references the prompt template you created in Prompt Builder.

Other Options:

Invocable Apex: Possible fallback if there's no out-of-the-box Flow Action available. However, Salesforce is releasing native Flow integration for AI prompts, making custom Apex less necessary.

Einstein for Flow: A broad label for Salesforce's generative AI features within Flow. Under the hood, you typically use a Flow Action that points to your prompt.

* 3. Conclusion

The easiest out-of-the-box solution is to use a Flow Action referencing the prompt template. Hence, Option B is correct.

Salesforce Agentforce Specialist References & Documents

•Salesforce Trailhead: Use Prompt Templates in Flow

Demonstrates how to add an Action in Flow that calls a prompt template.

•Salesforce Documentation: Einstein GPT for Flow

NEW QUESTION 96

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 97

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

A. Temperature, Frequency Penalty, Presence Penalty

B. Temperature, Top-k sampling, Presence Penalty

C. Temperature, Frequency Penalty, Output Tokens

Answer: A

Explanation:

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

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

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

? Frequency Penalty: Reduces the likelihood of the model repeating the same

phrases or outputs frequently.

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

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

NEW QUESTION 102

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 105

Universal Containers (UC) has a legacy system that needs to integrate with Salesforce. UC

wishes to create a digest of account action plans using the generative API feature. Which API service should UC use to meet this requirement?

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

Answer: A

Explanation:

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

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

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

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

NEW QUESTION 108

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

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

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

Answer: B

Explanation:

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

Reference:

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

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

NEW QUESTION 113

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 118

Universal Containers (UC) wants to enable its sales team to use AI to suggest recommended products from its catalog. Which type of prompt template should UC use?

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

Answer: C

Explanation:

Comprehensive and Detailed In-Depth Explanation: UC needs an AI solution to suggest products from a catalog for its sales team. Let's assess the prompt template types in Prompt Builder.

? Option A: Record summary prompt template Record summary templates generate concise summaries of records (e.g., Case, Opportunity). They're not designed for product recommendations, which require dynamic logic beyond summarization, making this incorrect.

? Option B: Email generation prompt template Email generation templates craft emails (e.g., customer outreach). While they could mention products, they're not optimized for standalone recommendations, making this incorrect.

? Option C: Flex prompt template Flex prompt templates are versatile, allowing custom inputs (e.g., catalog data from objects or Data Cloud) and instructions (e.g., "Suggest products based on customer preferences"). This flexibility suits UC's need to recommend products dynamically, making it the correct answer.

Why Option C is Correct: Flex templates offer the customization needed to suggest products from a catalog, aligning with Salesforce's guidance for tailored AI outputs.

References:

- ? Salesforce Agentforce Documentation: Prompt Builder > Flex Templates – Details dynamic use cases.
- ? Trailhead: Build Prompt Templates in Agentforce – Covers Flex for custom scenarios.
- ? Salesforce Help: Prompt Template Types – Confirms Flex versatility.

NEW QUESTION 123

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

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

Answer: A

Explanation:

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

How It Works:

? Data Masking in the Request Journey:

? Processing by the LLM:

? De-masking in the Response Journey:

Why Option A is Correct:

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

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

Why Options B and C are Incorrect:

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

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

References:

? Salesforce Agentforce Specialist Documentation - Einstein Trust Layer Overview:

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

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

Conclusion:

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

NEW QUESTION 125

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 126

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 128

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 132

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

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

Answer: C

Explanation:

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

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

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

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

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

References:

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

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

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

NEW QUESTION 136

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 141

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

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

Answer: B

Explanation:

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

Prompt templates in Agentforce can be dynamic, meaning they pull in record-specific data (e.g., customer names, descriptions, or other fields) to generate output. Since the data varies by record—some records might have short text fields while others have lengthy ones—the total number of tokens (words, characters, or subword units processed by the LLM) fluctuates. When the token count exceeds the LLM's limit (e.g., 4,096 tokens for some models), the process fails, but this only happens for records with higher token-generating data, explaining the randomness.

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

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

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

References:

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

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

NEW QUESTION 143

Universal Containers deploys a new Agentforce Service Agent into the company's website but is getting feedback that the Agentforce Service Agent is not providing answers to customer questions that are found in the company's Salesforce Knowledge articles. What is the likely issue?

- A. The Agentforce Service Agent user is not assigned the correct Agent Type License.
- B. The Agentforce Service Agent user needs to be created under the standard Agent Knowledge profile.
- C. The Agentforce Service Agent user was not given the Allow View Knowledge permission set.

Answer: C

Explanation:

Comprehensive and Detailed In-Depth Explanation: Universal Containers (UC) has deployed an Agentforce Service Agent on its website, but it's failing to provide answers from Salesforce Knowledge articles. Let's troubleshoot the issue.

? Option A: The Agentforce Service Agent user is not assigned the correct Agent Type License. There's no "Agent Type License" in Salesforce—agent functionality is tied to Agentforce licenses (e.g., Service Agent license) and permissions. Licensing affects feature access broadly, but the specific issue of not retrieving Knowledge suggests a permission problem, not a license type, making this incorrect.

? Option B: The Agentforce Service Agent user needs to be created under the standard Agent Knowledge profile. No "standard Agent Knowledge profile" exists. The Agentforce Service Agent runs under a system user (e.g., "Agentforce Agent User") with a custom profile or permission sets. Profile creation isn't the issue—access permissions are, making this incorrect.

? Option C: The Agentforce Service Agent user was not given the Allow View Knowledge permission set. The Agentforce Service Agent user requires read access to Knowledge articles to ground responses. The "Allow View Knowledge" permission (typically via the "Salesforce Knowledge User" license or a permission set like "Agentforce Service Permissions") enables this. If missing, the agent can't access Knowledge, even if articles are indexed, causing the reported failure. This is a common setup oversight and the likely issue, making it the correct answer.

Why Option C is Correct: Lack of Knowledge access permissions for the Agentforce Service Agent user directly prevents retrieval of article content, aligning with the symptoms and Salesforce security requirements.

References:

? Salesforce Agentforce Documentation: Service Agent Setup > Permissions – Requires Knowledge access.

? Trailhead: Set Up Agentforce Service Agents – Lists "Allow View Knowledge" need.

? Salesforce Help: Knowledge in Agentforce – Confirms permission necessity.

NEW QUESTION 148

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 152

Universal Containers needs to provide insights on the usability of Agents to drive adoption in the organization.

What should the Agentforce Specialist recommend?

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

Answer: A

Explanation:

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

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

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

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

Reference:

"Boost Adoption with Analytics Tools | Salesforce" .

NEW QUESTION 157

An Agentforce Specialist wants to troubleshoot their Agent's performance. Where should the Agentforce Specialist go to access all user interactions with the Agent, including Agent errors, incorrectly triggered actions, and incomplete plans?

- A. Plan Canvas
- B. Agent Settings
- C. Event Logs

Answer: C

Explanation:

Comprehensive and Detailed In-Depth Explanation: The Agentforce Specialist needs a comprehensive view of user interactions, errors, and action issues for troubleshooting. Let's evaluate the options.

? Option A: Plan Canvas Plan Canvas in Agent Builder visualizes an agent's execution plan for a single interaction, useful for design but not for aggregated troubleshooting data like errors or all interactions, making it incorrect.

? Option B: Agent Settings Agent Settings configure the agent (e.g., topics, channels), not provide interaction logs or error details. This is for setup, not analysis, making it incorrect.

? Option C: Event Logs Event Logs in Agentforce (accessible via Setup or Agent Analytics) record all user interactions, including errors, incorrectly triggered actions, and incomplete plans. They provide detailed telemetry (e.g., timestamps, action outcomes) for troubleshooting performance issues, making this the correct answer.

Why Option C is Correct: Event Logs offer the full scope of interaction data needed for troubleshooting, as per Salesforce documentation.

References:

? Salesforce Agentforce Documentation: Agent Analytics > Event Logs – Details interaction and error logging.

? Trailhead: Monitor and Optimize Agentforce Agents – Recommends Event Logs for troubleshooting.

? Salesforce Help: Agentforce Performance – Confirms logs for diagnostics.

NEW QUESTION 158

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 161

Universal Containers is planning a marketing email about products that most closely match a customer's expressed interests.

What should An Agentforce recommend to generate this email?

A. Standard email marketing template using Apex or flows for matching interest in products

B. Custom sales email template which is grounded with interest and product information

C. Standard email draft with Einstein and choose standard email template

Answer: B

Explanation:

To generate an email about products that closely match a customer's expressed interests, An Agentforce should recommend using a custom sales email template that is grounded with interest and product information. This ensures that the email content is personalized based on the customer's preferences, increasing the relevance of the marketing message.

Using grounding ensures that the generative AI pulls the correct data related to customer interests and product matches, making the email more effective.

For more information, refer to Salesforce documentation on grounding AI-generated content and email personalization strategies.

NEW QUESTION 164

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

A. Retrieves Salesforce data based on the 'Run As' users permissions.

B. Retrieves Salesforce data based on the user's permissions executing the prompt.

C. Retrieves Salesforce 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 168

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 172

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

Which audit data is available using the Einstein Trust Layer?

A. Response accuracy and offensiveness score

B. Hallucination score and bias score

C. Masked data and toxicity score

Answer: C

Explanation:

Universal Containers is considering the use of the Einstein Trust Layer along with Einstein Generative AI Audit Data. The Einstein Trust Layer provides a secure and compliant way to use AI by offering features like data masking and toxicity assessment.

The audit data available through the Einstein Trust Layer includes information about masked data—which ensures sensitive information is not exposed—and the toxicity score, which evaluates the generated content for inappropriate or harmful language. References:

? Salesforce Agentforce Specialist Documentation - Einstein Trust Layer: Details the auditing capabilities, including logging of masked data and evaluation of generated responses for toxicity to maintain compliance and trust.

NEW QUESTION 175

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

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 179

What considerations should an Agentforce Specialist be aware of when using Record Snapshots grounding in a prompt template?

- A. Activities such as tasks and events are excluded.
- B. Empty data, such as fields without values or sections without limits, is filtered out.
- C. Email addresses associated with the object are excluded.

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation:Record Snapshots

grounding in Agentforce prompt templates allows the AI to access and use data from a specific Salesforce record (e.g., fields and related records) to generate contextually relevant responses. However, there are specific limitations to consider. Let's analyze each option based on official documentation.

? Option A: Activities such as tasks and events are excluded.According to Salesforce Agentforce documentation, when grounding a prompt template with Record Snapshots, the data included is limited to the record's fields and certain related objects accessible via Data Cloud or direct Salesforce relationships. Activities (tasks and events) are not included in the snapshot because they are stored in a separate Activity object hierarchy and are not directly part of the primary record's data structure. This is a key consideration for an Agentforce Specialist, as it means the AI won't have visibility into task or event details unless explicitly provided through other grounding methods (e.g., custom queries). This limitation is accurate and critical to understand.

? Option B: Empty data, such as fields without values or sections without limits, is filtered out.Record Snapshots include all accessible fields on the record, regardless of whether they contain values. Salesforce documentation does not indicate that empty fields are automatically filtered out when grounding a prompt template. The Atlas Reasoning Engine processes the full snapshot, and empty fields are simply treated as having no data rather than being excluded. The phrase "sections without limits" is unclear but likely a typo or misinterpretation; it doesn't align with any known Agentforce behavior. This option is incorrect.

? Option C: Email addresses associated with the object are excluded.There's no specific exclusion of email addresses in Record Snapshots grounding. If an email field (e.g., Contact.Email or a custom email field) is part of the record and accessible to the running user, it is included in the snapshot. Salesforce documentation does not list email addresses as a restricted data type in this context, making this option incorrect.

Why Option A is Correct:The exclusion of activities (tasks and events) is a documented limitation of Record Snapshots grounding in Agentforce. This ensures specialists design prompts with awareness that activity-related context must be sourced differently (e.g., via Data Cloud or custom logic) if needed. Options B and C do not reflect actual Agentforce behavior per official sources.

References:

? Salesforce Agentforce Documentation: Prompt Templates > Grounding with Record Snapshots – Notes that activities are not included in snapshots.

? Trailhead: Ground Your Agentforce Prompts – Clarifies scope of Record Snapshots data inclusion.

? Salesforce Help: Agentforce Limitations – Details exclusions like activities in grounding mechanisms.

NEW QUESTION 180

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual Agentforce-Specialist Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the Agentforce-Specialist Product From:

<https://www.2passeasy.com/dumps/Agentforce-Specialist/>

Money Back Guarantee

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