

Exam Questions Agentforce-Specialist

Salesforce Certified Agentforce Specialist

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



NEW QUESTION 1

Which feature in the Einstein Trust Layer helps to minimize the risks of jailbreaking and prompt injection attacks?

- A. Secure Data Retrieval and Grounding
- B. Data Masking
- C. Prompt Defense

Answer: C

Explanation:

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

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

In contrast:

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

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

Reference:

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

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

NEW QUESTION 2

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 3

What is a valid use case for Data Cloud retrievers?

- A. Returning relevant data from the vector database to augment a prompt.
- B. Grounding data from external websites to augment a prompt with RAG.
- C. Modifying and updating data within the source systems connected to Data Cloud.

Answer: A

Explanation:

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

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

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

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

References:

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

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

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

NEW QUESTION 4

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

- A. To generate relevant content when engaging with each customer

- B. To analyze past campaign performance
- C. To send automated emails to all customers

Answer: A

Explanation:

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

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

? Option B refers more to Einstein Analytics or

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

References:

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

NEW QUESTION 5

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

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

Answer: B

Explanation:

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

Supported Objects for Service AI Grounding:

? Case

? Knowledge

? Case Object:

? Knowledge Object:

? Exclusion of Other Objects:

Why Options A and C are Incorrect:

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

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

References:

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

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

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

NEW QUESTION 6

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 7

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

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

Answer: B

Explanation:

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

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

incomplete and incorrect.

? Option B: Identify the best matching topic and actions and correct order of execution. The LLM analyzes user input to understand intent, matches it to the best-fitting topic (configured in Agent Builder), and selects associated actions. It also determines the correct sequence of execution based on the agent's plan (e.g., retrieve data before updating a record). This end-to-end process—from intent recognition to action orchestration—is the LLM's core role in Agentforce, making this the correct answer.

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

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

References:

? Salesforce Agentforce Documentation: Atlas Reasoning Engine – Outlines LLM's intent and action handling.
? Trailhead: Understand Agentforce Technology – Explains topic matching and execution.
? Salesforce Help: Agentforce Actions – Confirms LLM's role in orchestrating responses.

NEW QUESTION 8

Universal Containers (UC) wants to assess Salesforce's generative features but has concerns over its company data being exposed to third-party large language models (LLMs). Specifically, UC wants the following capabilities to be part of Einstein's generative AI service.

No data is used for LLM training or product improvements by third-party LLMs. No data is retained outside of UC's Salesforce org.

The data sent cannot be accessed by the LLM provider.

Which property of the Einstein Trust Layer should the Agentforce Specialist highlight to UC that addresses these requirements?

- A. Prompt Defense
- B. Zero-Data Retention Policy
- C. Data Masking

Answer: B

Explanation:

Universal Containers (UC) has concerns about data privacy when using Salesforce's generative AI features, particularly around preventing third-party LLMs from accessing or retaining their data. The Zero-Data Retention Policy in the Einstein Trust Layer is designed to address these concerns by ensuring that:

- ? No data is used for training or product improvements by third-party LLMs.
- ? No data is retained outside of the customer's Salesforce organization.
- ? The LLM provider cannot access any customer data.

This policy aligns perfectly with UC's requirements for keeping their data safe while leveraging generative AI capabilities.

? Prompt Defense and Data Masking are also security features, but they do not directly address the concerns related to third-party data access and retention.

References:

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

NEW QUESTION 9

Universal Containers wants to be able to detect with a high level confidence if content generated by a large language model (LLM) contains toxic language.

Which action should an AI Specialist take in the Trust Layer to confirm toxicity is being appropriately managed?

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

Answer: C

Explanation:

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

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

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

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

References:

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

NEW QUESTION 10

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

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

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

Answer: B

Explanation:

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

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

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

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

Reference:

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

NEW QUESTION 10

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

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

Answer: C

Explanation:

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

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

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

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

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

References:

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

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

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

NEW QUESTION 11

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 12

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 17

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

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

Answer: A

Explanation:

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

NEW QUESTION 21

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

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

Answer: B

Explanation:

Why is "Messaging Session" the correct answer?
In Agentforce, the Messaging Session object stores the conversation transcript between the customer and the agent.
Key Features of the Messaging Session Object:
? Stores the Entire Customer-Agent Conversation
? Supports AI-Powered Work Summaries
? Links with Service Cloud for Case Resolution
Why Not the Other Options?
* A. Messaging End User
? Incorrect because this object stores details about the customer (e.g., name, contact details) but not the conversation transcript.
* C. Case
? Incorrect because Cases store structured service requests but do not contain raw conversation transcripts.
? Instead, cases may reference the Messaging Session object.
Agentforce Specialist References
? Salesforce AI Specialist Material confirms that Messaging Sessions store chat conversations and support Einstein Work Summaries.

NEW QUESTION 26

Universal Containers wants to incorporate the current order fulfillment status into a prompt for a large language model (LLM). The order status is stored in the external enterprise resource planning (ERP) system. Which data grounding technique should the Agentforce Specialist recommend?

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

Answer: A

Explanation:

? Context of the Requirement: Universal Containers wants to pull in real-time order status data from an external ERP system into an LLM prompt.
? Data Grounding in LLM Prompts: Data grounding ensures the Large Language Model has access to the most current and relevant information. In Salesforce, one recommended approach is to use External Objects (via Salesforce Connect) when data resides outside of Salesforce.
? Why External Object Record Merge Fields:
? Why Not External Services Merge Fields or Apex Merge Fields:
? References and Study Resources:

NEW QUESTION 31

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 34

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

How should UC resolve this?

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

Answer: B

Explanation:

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

Key Considerations for Fixing Missing Fields in Data Cloud Retrievers:

? Search Index Controls Which Fields Are Searchable

? Ensures Complete and Accurate Data Retrieval

? Supports AI-Generated Responses

Why Not the Other Options?

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

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

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

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

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

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

Agentforce Specialist References

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

NEW QUESTION 36

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 39

Universal Containers implemented Agent for its users.

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

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

Answer: C

Explanation:

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

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

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

References:

? Salesforce Agentforce Specialist Documentation - Agent Supported Actions:

? Salesforce Help - Limitations of Agent:

NEW QUESTION 41

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

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

Answer: B

Explanation:

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

? Standard Merge Fields in Salesforce Email Templates:

? Why Recipient Account?

? References and Study Resources:

NEW QUESTION 46

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

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

Answer: C

Explanation:

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

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

Key Considerations for This Issue:

? Permission Set Restriction

? Agentforce Role-Based Access Control

? Fixing the Issue

Why Not the Other Options?

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

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

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

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

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

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

Agentforce Specialist References

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

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

NEW QUESTION 47

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

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

Answer: B

Explanation:

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

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

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

References:

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

NEW QUESTION 51

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

- A. Find similar requests and provide actions that need to be executed

- B. Identify the best matching actions and correct order of execution
- C. Determine a user's access and sort actions by priority to be executed

Answer: B

Explanation:

In Agent, the role of the Large Language Model (LLM) is to analyze user inputs and identify the best matching actions that need to be executed. It uses natural language understanding to break down the user's request and determine the correct sequence of actions that should be performed.

By doing so, the LLM ensures that the tasks and actions executed are contextually relevant and are performed in the proper order. This process provides a seamless, AI-enhanced experience for users by matching their requests to predefined Salesforce actions or flows.

The other options are incorrect because:

A mentions finding similar requests, which is not the primary role of the LLM in this context. C focuses on access and sorting by priority, which is handled more by security models and governance than by the LLM.

References:

Salesforce Einstein Documentation on Agent Actions

Salesforce AI Documentation on Large Language Models

NEW QUESTION 56

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 60

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 65

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 68

Universal Containers is rolling out a new generative AI initiative.

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

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

Answer: C

Explanation:

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

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

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

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

Reference:

"Prompt Builder Limitations | Salesforce Documentation" .

NEW QUESTION 73

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 78

A Salesforce Administrator is exploring the capabilities of Agent to enhance user interaction within their organization. They are particularly interested in how Agent processes user requests and the mechanism it employs to deliver responses. The administrator is evaluating whether Agent directly interfaces with a large language model (LLM) to fetch and display responses to user inquiries, facilitating a broad range of requests from users.

How does Agent handle user requests In Salesforce?

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

Answer: C

Explanation:

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

References:

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

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

NEW QUESTION 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 86

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

- ? Semantically search HR policies, compliance guidelines, and company procedures.
- ? Ensure responses are grounded on published Knowledge.
- ? Allow Knowledge updates to be reflected immediately without manual reconfiguration.

What should UC do to ensure the agent retrieves the right information?

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

Answer: B

Explanation:

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

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

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

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

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

References:

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

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

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

NEW QUESTION 91

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 93

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 97

Universal Containers wants to allow its service agents to query the current fulfillment status of an order with natural language. There is an existing autolaunched flow to query the Information from Oracle ERP, which is the system of record for the order fulfillment process. How should an Agentforce Specialist apply the power of conversational AI to this use case?

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

Answer: A

Explanation:

Why is "Create a custom Agent action which calls a flow" the correct answer?
In Agentforce, the best way to allow service agents to query order fulfillment status from an external system (Oracle ERP) using natural language is to create a custom Agent action that invokes an existing autolaunched flow. Key Considerations for This Approach:
? Custom Agent Action Triggers the Flow
? Enables AI-Powered Query Execution
? No Need for Manual Data Entry
Why Not the Other Options?
* B. Configure the Integration Flow Standard Action in Agent Builder
? Incorrect because Integration Flow Standard Actions are for predefined use cases, not custom ERP integrations.
? They do not provide the flexibility needed to connect with Oracle ERP dynamically.
* C. Create a Flex Prompt Template in Prompt Builder
? Incorrect because Flex prompts are used for structuring AI-generated responses, not executing queries on external systems.
? This approach does not enable the AI to retrieve live fulfillment status from Oracle ERP. Agentforce Specialist References
? Salesforce AI Specialist Material confirms that custom Agent actions allow integration with external systems through Salesforce flows.
? Salesforce Instructions for Certification mention that Agentforce supports custom Agent actions for external data retrieval.

NEW QUESTION 99

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 101

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

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

Answer: B

Explanation:

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

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

Key Benefits of Knowledge Articles in an Agentforce Data Library:

? Ensures Information Accuracy and Consistency

? Improves Searchability and AI-Grounded Responses

? Enhances Customer Support and Agent Productivity

Why Not the Other Options?

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

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

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

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

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

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

Agentforce Specialist References

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

NEW QUESTION 104

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

What should UC consider when using standard Service AI Grounding?

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

Answer: B

Explanation:

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

? Field Types: Standard Service AI Grounding supports String and Text Area fields.

Custom rich text fields (e.g., RichTextArea) are not supported, making Option B correct.

? Objects: While Service AI Grounding primarily uses Case and Knowledge objects (Option A), the limitation here is the field type, not the object.

? Visibility: Service AI Grounding respects user permissions and sharing settings unless overridden (Option C is incorrect).

References:

? Salesforce Help: Service AI Grounding Requirements

? Explicitly states support for "Text Area and String fields" only.

NEW QUESTION 106

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

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

Answer: C

Explanation:

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

Key Components Created in Data Cloud:

- ? Data Stream
- ? Search Index
- ? Retriever

Why Not the Other Options?

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

Agentforce Specialist References

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

NEW QUESTION 107

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 **Flow Action** that references the prompt template you created in Prompt Builder.

o 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

- o 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 110

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 114

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 115

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

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

Answer: B

Explanation:

Why is "Automates Ingestion, Indexing, and Default Retriever Creation" the correct answer?

An Agentforce Data Library is a key component in ensuring that an AI agent provides precise and up-to-date responses by:

Automating data ingestion Brings in data from various sources. Indexing the data Organizes it efficiently for AI retrieval. Creating a default retriever Enables the AI to fetch relevant data dynamically when answering customer queries.

Key Features of an Agentforce Data Library:

? Automates Data Ingestion

? Indexes Data for Efficient Retrieval

? Creates a Default Retriever

Why Not the Other Options?

* A. Automates ingestion, taxonomical classification, and precision keyword search retrieval

? Incorrect because Agentforce does not rely on keyword searches but on indexing and AI-driven retrieval.

* C. Automates ingestion and OCR processing of PDFs

? Incorrect because OCR (Optical Character Recognition) is not the primary function of an Agentforce Data Library.

? AI grounding is based on indexed and structured data, not raw OCR-extracted text.

Agentforce Specialist References

? Salesforce AI Specialist Material explains that Agentforce Data Libraries automate data ingestion, indexing, and retriever setup for AI-powered responses.

? Salesforce Instructions for Certification confirm that AI responses are grounded in structured and indexed Data Libraries.

NEW QUESTION 118

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

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

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

Answer: B

Explanation:

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

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

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

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

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

NEW QUESTION 120

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

What is causing the problem?

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

Answer: B

Explanation:

In order to access and create custom prompt templates in Prompt Builder, the Agentforce Specialist must have the Prompt Template Manager permission set assigned. Without this permission, they will not be able to access Prompt Builder in the Setup menu, even though Einstein Generative AI is enabled.

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

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

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

References:

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

NEW QUESTION 124

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

What is a key reason for implementing Agent?

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

Answer: C

Explanation:

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

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

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

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

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

NEW QUESTION 125

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 130

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 131

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 133

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

* 1. Create a new Field Generation prompt template.

* 2. Choose Case as the object type.

* 3. Select the custom field AI_Analysis_c as the target field.

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

Which critical step did the Agentforce Specialist miss?

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

Answer: C

Explanation:

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

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

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

Reference:

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

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

NEW QUESTION 137

Universal Containers (UC) is discussing its AI strategy in an agile Scrum meeting.

Which business requirement would lead An Agentforce to recommend connecting to an external foundational model via Einstein Studio (Model Builder)?

- A. UC wants to fine-tune model temperature.
- B. UC wants a model fine-tuned using company data.
- C. UC wants to change the frequency penalty of the model.

Answer: B

Explanation:

Einstein Studio (Model Builder) allows organizations to connect and utilize external foundational models while fine-tuning them with company-specific data. This capability is particularly suited to businesses like Universal Containers (UC) that require customization of foundational models to better align with their unique data and use cases.

? Option A: Adjusting model temperature is a parameter-level setting for controlling randomness in AI-generated responses but does not necessitate connecting to an external foundational model.

? Option B: This is the correct answer because Einstein Studio supports fine-tuning external models with proprietary company data, enabling a tailored and more accurate AI solution for UC.

? Option C: Changing frequency penalties is another parameter-level adjustment and does not require external foundational models or Einstein Studio.

Reference:

"Using Einstein Studio to Connect Foundational Models | Salesforce Trailhead" .

NEW QUESTION 140

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 145

Universal Containers implements Custom Agent Actions to enhance its customer service operations. The development team needs to understand the core components of a Custom Agent Action to ensure proper configuration and functionality. What should the development team review in the Custom Agent Action configuration to identify one of the core components of a Custom Agent Action?

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

Answer: B

Explanation:

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

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

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

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

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

References:

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

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

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

NEW QUESTION 149

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

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

Answer: C

Explanation:

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

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

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

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

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

References:

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

- ? Trailhead: Build Agents with Agentforce – Describes FAQ topics with Knowledge actions.
- ? Salesforce Help: Knowledge in Agentforce – Confirms this configuration.

NEW QUESTION 154

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

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

Answer: B

Explanation:

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

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

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

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

References:

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

NEW QUESTION 155

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 159

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 161

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

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

Answer: C

Explanation:

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

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

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

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

NEW QUESTION 164

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

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

Answer: C

Explanation:

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

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

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

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

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

References:

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

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

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

NEW QUESTION 168

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 171

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 175

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

A. Create a prompt template passing in a special custom object that connects the records temporarily.

B. Create a prompt template-triggered flow to access the data from standard and custom objects.

C. Create a Flex template to add resources with standard and custom objects as inputs.

D. Use a Record Snapshot to combine data from unrelated objects into a single prompt.

Answer: C

Explanation:

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

? Option A: Create a prompt template passing in a special custom object that connects the records temporarily. While a custom object could theoretically act as a junction to link unrelated records, this approach requires additional setup (e.g., creating the object, populating it with data via automation), and there's no direct mechanism in Prompt Builder to "pass in" such an object to a prompt template without grounding or flow support. This is inefficient and not a native feature, making it incorrect.

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

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

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

References:

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

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

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

NEW QUESTION 178

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

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

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

C. Retrieves Salesforce's data based on the Prompt template's object permissions.

Answer: B

Explanation:

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

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

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

Reference:

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

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

NEW QUESTION 179

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

A. Select the search index, specify the associated data model object (DMO) and data space, and optionally define filters to narrow search results.

B. Define the output configuration by specifying the maximum number of results to return, and map the output fields that will ground the prompt.

C. Configure the search index, choose vector or hybrid search, choose the fields for filtering, the data space and model, then define the ranking method.

Answer: A

Explanation:

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

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

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

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

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

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

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

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

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

References:

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

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

=5)

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

NEW QUESTION 181

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 182

Which scenario best demonstrates when an Agentforce Data Library is most useful for improving an AI agent's response accuracy?

A. When the AI agent must provide answers based on a curated set of policy documents that are stored, regularly updated, and indexed in the data library.

B. When the AI agent needs to combine data from disparate sources based on mutually common data, such as Customer Id and Product Id for grounding.

C. When data is being retrieved from Snowflake using zero-copy for vectorization and retrieval.

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation: The Agentforce Data Library enhances AI accuracy by grounding responses in curated, indexed data. Let's assess the scenarios.

? Option A: When the AI agent must provide answers based on a curated set of policy documents that are stored, regularly updated, and indexed in the data library. The Data Library is designed to store and index structured content (e.g., Knowledge articles, policy documents) for semantic search and grounding. It excels when an agent needs accurate, up-to-date responses from a managed corpus, like policy documents, ensuring relevance and reducing hallucinations. This is a prime use case per Salesforce documentation, making it the correct answer.

? Option B: When the AI agent needs to combine data from disparate sources based on mutually common data, such as Customer Id and Product Id for grounding. Combining disparate sources is more suited to Data Cloud's ingestion and harmonization capabilities, not the Data Library, which focuses on indexed content retrieval. This scenario is less aligned, making it incorrect.

? Option C: When data is being retrieved from Snowflake using zero-copy for vectorization and retrieval. Zero-copy integration with Snowflake is a Data Cloud feature, but the Data Library isn't specifically tied to this process—it's about indexed libraries, not direct external retrieval. This is a different context, making it incorrect.

Why Option A is Correct: The Data Library shines in curated, indexed content scenarios like policy documents, improving agent accuracy, as per Salesforce guidelines.

References:

? Salesforce Agentforce Documentation: Data Library > Use Cases – Highlights

curated content grounding.

? Trailhead: Ground Your Agentforce Prompts – Describes Data Library accuracy benefits.

? Salesforce Help: Agentforce Data Library – Confirms policy document scenario.

NEW QUESTION 187

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 189

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

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

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

Answer: A

Explanation:

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

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

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

NEW QUESTION 191

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

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

Answer: C

Explanation:

? Context of the Question

? Why ??Add Flow Instructions?? Is Needed

? Outcome

Salesforce Agentforce Specialist References & Documents

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

? Salesforce Agentforce Specialist Study GuideOutlines how to configure generative AI prompts that reference real-time Flow data.

NEW QUESTION 194

.....

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