

Salesforce.Agentforce-Specialist.v2026-04-15.q124

Exam Code:	Agentforce-Specialist
Exam Name:	Salesforce Certified Agentforce Specialist
Certification Provider:	Salesforce
Free Question Number:	124
Version:	v2026-04-15
# of views:	107
# of Questions views:	1706
https://www.freecram.net/torrent/Salesforce.Agentforce-Specialist.v2026-04-15.q124.html	

NEW QUESTION: 1

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: (SHOW ANSWER)

1.Context of the Question

oUniversal Container (UC) has used prompt templates to update summary fields on record pages.

oNow, 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

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

oOther 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

oThe 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: 2

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

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

A. Create an application of the custom LLM and embed it in Sales Cloud via iFrame.

B. Add the fine-tuned LLM in Einstein Studio Model Builder.

C. Enable model endpoint on OpenAI and make callouts to the model to generate emails.

Answer: (SHOW ANSWER)

Since security and data privacy are critical, the best option for the Agentforce Specialist is to integrate the fine-tuned LLM (Large Language Model) into Salesforce by adding it to Einstein Studio Model Builder.

Einstein Studio allows organizations to bring their own AI models (BYOM), ensuring the model is securely managed within Salesforce's environment, adhering to data privacy standards.

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

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

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

NEW QUESTION: 3

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: (SHOW ANSWER)

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: 4

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: (SHOW ANSWER)

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

- * The search index defines which fields are indexed and accessible to the retriever.
- * If a field is missing, it must be added to the index before it can be queried.
- * Ensures Complete and Accurate Data Retrieval
- * Without indexing, the retriever cannot reference the missing field in AI responses.
- * Updating the index makes the field available for AI-powered retrieval.
- * Supports AI-Grounded Responses
- * Agentforce relies on Retriever-Augmented Generation (RAG) to ground AI responses in searchable Data Cloud content.
- * Ensuring all relevant fields are indexed improves AI-generated answer accuracy.

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: 5

Universal Containers (UC) needs to save agents time with AI-generated case summaries. UC has implemented the Work Summary feature.

What does Einstein consider when generating a summary?

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

Answer: (SHOW ANSWER)

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

* Conversation Context:

- * Einstein analyzes the details of the customer interaction, including chat or email threads, to extract relevant information for the summary.

* Knowledge Articles:

- * It considers linked Knowledge Articles or articles referred to during the case resolution process, ensuring the summary incorporates accurate resolutions or additional resources provided to the customer.

* Cases:

- * Einstein also examines historical cases and related case records to ground the summary in context from past resolutions or interactions.

* Option A is correct as it includes all three: conversation context, Knowledge articles, and cases.

* Option B is incorrect because it limits the grounding to conversation context only, excluding other critical elements.

* Option C is incorrect because it omits case data, which Einstein considers for more accurate and contextually rich summaries.

Reference:

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

NEW QUESTION: 6

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: (SHOW ANSWER)

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: 7

An Agentforce Specialist builds a new Service Agent that uses a custom action built on a flow. The agent has been tested in a sandbox and is now ready to deploy.

What is a key consideration regarding the activation status of the agent in the production environment?

- A. The agent will be activated automatically only if the flow is also active.
- B. The agent must be manually activated in production, regardless of its status in the sandbox.
- C. The agent will automatically be activated upon successful deployment.

Answer: (SHOW ANSWER)

According to the AgentForce Deployment and Lifecycle Management Guide, when an agent is deployed from a sandbox to a production environment, activation does not carry over automatically. The documentation clarifies: "Each environment maintains its own activation state. Agents must be manually activated in production after deployment to ensure controlled rollout and compliance validation." This ensures that only verified configurations are activated intentionally. Option A is incorrect because activation is not dependent solely on the flow's active status. Option C is also incorrect, as automatic activation upon deployment is explicitly prevented by design to maintain environment safety.

Therefore, Option B correctly reflects the deployment requirement for manual activation in production.

References (AgentForce Documents / Study Guide):

- * AgentForce Deployment Guide: "Activating Agents in Production Environments"
- * AgentForce Implementation Handbook: "Environment Lifecycle and Activation Controls"
- * Salesforce Release Management Study Notes: "Post-Deployment Activation Steps"

NEW QUESTION: 8

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

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

Answer: (SHOW ANSWER)

In Salesforce Agentforce, prompt templates are used to generate dynamic responses or field values by leveraging an LLM, often with grounding data from Salesforce records or external sources. The scenario describes a Field Generation prompt template that fails intermittently with token limit errors, indicating that the issue is tied to exceeding the LLM's

token capacity (e.g., input + output tokens). The random nature of these failures suggests variability in the token count across different records, which is directly addressed by Option B.

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

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

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

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

:

Salesforce Agentforce Documentation: "Prompt Templates" (Salesforce Help:

<https://help.salesforce.com/s>

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

Trailhead: "Build Prompt Templates for

Agentforce" (<https://trailhead.salesforce.com/content/learn/modules>

[/build-prompt-templates-for-agentforce](https://trailhead.salesforce.com/content/learn/modules/build-prompt-templates-for-agentforce))

NEW QUESTION: 9

Universal Containers (UC) has a library of custom-built personalized investment portfolio APIs, and is planning to extend it to agents.

Which method should UC's agent choose to dynamically use the best API service?

- A.** Agent-to-Agent (A2A) protocol support
- B.** Model Context Protocol (MCP) server support
- C.** MuleSoft connector for custom hosted processes

Answer: (SHOW ANSWER)

The most appropriate and advanced method for an Agentforce agent to dynamically select and use the best API service from a library of custom-built APIs is through Model Context Protocol (MCP) server support (B).

The Model Context Protocol (MCP) is an open standard specifically designed to standardize how AI agents and Large Language Models (LLMs) interact with external tools, systems, and data sources (like custom APIs). An external system, such as a server

hosting UC's custom portfolio APIs, can be exposed as an MCP Server. This server provides rich, standardized, human-readable metadata about its "tools" (the APIs it offers). The Agentforce Atlas Reasoning Engine can interpret this metadata to understand the function of each API, the required inputs, and the expected outputs. This allows the agent to dynamically discover, reason over, and select the most appropriate API to execute based on a user's request (e.g., "Show me the best-performing portfolio" vs. "Adjust my risk tolerance").

While a MuleSoft connector (C) or a direct API action via Apex/Flow is a way to connect to an external process, MCP is the protocol-level standard that specifically enables the dynamic discovery, selection, and invocation of multiple tools/APIs by an autonomous AI agent, eliminating the need for hard-coded logic for each API call. Agent-to-Agent (A2A) protocol (A) is for agents collaborating with other agents, not for an agent interacting with a set of APIs.

Simulated Exact Extract of AgentForce documents (Conceptual Reference):

"For Agentforce to intelligently and autonomously interact with external, custom-built API services, the system must be configured to utilize Model Context Protocol (MCP). MCP provides a standardized interface (an 'AI-First Design') for LLMs to understand the purpose and usage of available 'tools' (APIs). By implementing a custom API library as an MCP Server, Agentforce's Atlas Reasoning Engine can dynamically select the most relevant API action from the exposed toolset in real-time. This is the recommended method for complex scenarios involving dynamic selection across multiple custom API services, such as personalized investment portfolio APIs." Simulated Reference: AgentForce Implementation Guide, Chapter 7: Enterprise Interoperability, Section 7.3: Model Context Protocol (MCP), p. 185.

NEW QUESTION: 10

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

Answer: (SHOW ANSWER)

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: 11

Choose 1 option.

An administrator at Universal Containers has successfully deployed a new agent from a sandbox to production using a change set.

The agent uses a prompt template that invokes a Salesforce flow to perform a complex calculation. In production, when users interact with the agent, it fails with an error message every time the flow is supposed to run. The flow was included in the change set and is present in production.

What is the most likely cause of this issue?

- A.** The flow was not manually activated in the production org after the deployment.
- B.** The user in production does not have permission to run the flow.
- C.** The change set did not include the dependent Apex classes for the flow.

Answer: (SHOW ANSWER)

Per the AgentForce Deployment and Flow Integration Guide, when deploying flows via change sets, the flows arrive in the production org in "inactive" status by default. The administrator must manually activate the flow post-deployment before it can be executed by agents or users.

This explains why the agent encounters an error when attempting to run the flow-the system recognizes the flow but cannot invoke it because it remains inactive.

Option B is incorrect since permission errors would display an "insufficient privileges" message, not a runtime failure. Option C is unlikely because dependent Apex classes would be automatically handled if properly included in the deployment.

Therefore, the most likely cause is Option A - The flow was not manually activated in production after deployment.

Reference: AgentForce Deployment Guide - "Post-Deployment Flow Activation Requirements."

NEW QUESTION: 12

After creating a foundation model in Einstein Studio, which hyperparameter should An Agentforce use to adjust the balance between consistency and randomness of a response?

- A.** Presence Penally
- B.** Variability
- C.** Temperature

Answer: (SHOW ANSWER)

The Temperature hyperparameter controls the randomness of model outputs:

- * Low Temperature (e.g., 0.2): More deterministic, consistent responses.
- * High Temperature (e.g., 1.0): More creative, varied responses.
- * Presence Penalty (Option A): Discourages repetition of tokens, unrelated to randomness.
- * Variability (Option B): Not a standard hyperparameter in Einstein Studio.

References:

- * Einstein Studio Documentation: Model Hyperparameters
- * Explicitly states "Temperature adjusts the balance between predictable and random outputs."

NEW QUESTION: 13

Sales reps at Universal Containers should not be able to create or edit prompt templates. Which permission set should an AgentForce Specialist assign to the sales reps?

- A.** Prompt Execute User
- B.** Prompt Template Manager
- C.** Prompt Template User

Answer: (SHOW ANSWER)

Comprehensive and Detailed Explanation From Exact Extract of AgentForce Documents: According to the AgentForce Permissions and Role Management Guide, permission sets control the level of access users have to Prompt Templates - which define how the reasoning engine interprets and generates responses. AgentForce distinguishes between users who create or manage templates and those who only execute or use them during interactions.

* The Prompt Execute User permission set is designed for users who can invoke and run prompts but cannot create, modify, or delete prompt templates. This role is intended for front-line users such as sales reps, service agents, or support staff who utilize preconfigured templates in daily workflows without altering them.

* The Prompt Template Manager permission set (Option B) grants full administrative access - allowing users to create, edit, and delete templates. This is reserved for system administrators or AgentForce specialists responsible for managing prompt configurations.

* The Prompt Template User permission set (Option C) provides limited management capabilities, enabling viewing and cloning of templates but still allowing minor modifications, which does not meet the stated restriction.

Therefore, to ensure sales reps can only use but not edit or create prompt templates, the correct permission set is Option A - Prompt Execute User.

Reference: AgentForce Security and Permissions Documentation - "Prompt Template Access Levels and Role Assignment."

NEW QUESTION: 14

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: (SHOW ANSWER)

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: 15

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: (SHOW ANSWER)

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.

:

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

NEW QUESTION: 16

What is one key purpose of action instructions when creating a custom agent action in Agentforce?

A. Action instructions help the reasoning engine decide which action to use.

B. Action instructions define the temperature of the large language model (LLM) powering the Reasoning Engine.

C. Action instructions tell the user how to call this action in a conversation.

Answer: A (LEAVE A REPLY)

According to the AgentForce Action Design and Configuration Guide, action instructions serve as the directive text that informs the Reasoning Engine about when and how to invoke a specific action. The documentation states: "Action instructions guide the reasoning engine by describing the action's purpose, inputs, and when it should be selected. Clear, concise instructions improve the LLM's accuracy in mapping user intent to the correct action." Option A is correct because it reflects this documented purpose - connecting user intent with the appropriate system action.

Option B is incorrect since the temperature parameter for the LLM is defined at the system or configuration level, not within action instructions.

Option C is incorrect because action instructions are meant for the reasoning engine, not for end users.

References (AgentForce Documents / Study Guide):

* AgentForce Action Design Guide: "How Action Instructions Support Reasoning"

* AgentForce Builder Handbook: "Optimizing Action Selection through Clear Instructions"

* AgentForce Study Guide: "Purpose of Action Instructions in Custom Actions"

Valid Agentforce-Specialist Dumps shared by ExamDiscuss.com for Helping Passing Agentforce-Specialist Exam! ExamDiscuss.com now offer the **newest Agentforce-**

Specialist exam dumps, the ExamDiscuss.com Agentforce-Specialist exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com Agentforce-Specialist dumps with Test Engine here:
<https://www.examdiscuss.com/Salesforce/exam/Agentforce-Specialist/premium/> (300 Q&As Dumps, **35%OFF Special Discount Code: freecram**)

NEW QUESTION: 17

Choose 1 option.

Universal Containers (UC) plans to answer questions based on similar cases that have been successfully resolved in the past.

What should UC consider when implementing this approach?

- A. No action is needed, as past cases are used to answer the question.
- B. Create a data model object (DMO) based on Case object and create an index on it.
- C. Create an unstructured data model object (UDMO) based on Case object and create an index on it.

Answer: ([SHOW ANSWER](#))

According to the AgentForce Data Configuration and Retrieval Guide, when an organization like Universal Containers wants to enable its AI agent to answer questions using historical case data, the correct implementation is to create an Unstructured Data Model Object (UDMO) based on the Case object, then index that data for retrieval.

The documentation clearly explains:

"When using previous case records to power AI-driven Q&A or similarity-based retrieval, create a UDMO mapped to the Case object. UDMOs allow the system to process and semantically index unstructured text fields such as Case Description, Resolution, and Comments, enabling the LLM to surface contextually similar resolved cases." This allows the AgentForce retrieval engine to perform semantic searches across historical support data, returning cases that are most contextually relevant to the user's query.

Option A is incorrect because past cases cannot be used automatically without indexing them.

Option B is incorrect because a DMO is for structured data (tables, numeric fields) and doesn't support semantic text retrieval.

Therefore, Option C is correct and aligns fully with Salesforce's documented best practices.

References (AgentForce Documents / Study Guide):

* AgentForce Data Configuration Guide: "Using UDMOs for Case-Based Reasoning"

* AgentForce Implementation Handbook: "Indexing Historical Case Records for Semantic Search"

* AgentForce Study Guide: "Creating Unstructured Data Model Objects from Case Objects"

NEW QUESTION: 18

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: (SHOW ANSWER)

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: 19

An administrator wants to check the response of the Flex prompt template they've built, but the preview button is greyed out.

What is the reason for this?

- A.** The records related to the prompt have not been selected.
- B.** The prompt has not been saved and activated,
- C.** A merge field has not been inserted in the prompt.

Answer: A (LEAVE A REPLY)

When the preview button is greyed out in a Flex prompt template, it is often because the records related to the prompt have not been selected. Flex prompt templates pull data dynamically from Salesforce records, and if there are no records specified for the prompt, it can't be previewed since there is no content to generate based on the template.

- * Option B, not saving or activating the prompt, would not necessarily cause the preview button to be greyed out, but it could prevent proper functionality.
- * Option C, missing a merge field, would cause issues with the output but would not directly grey out the preview button.

Ensuring that the related records are correctly linked is crucial for testing and previewing how the prompt will function in real use cases.

Salesforce Agentforce Specialist References:

Refer to the documentation on troubleshooting Flex templates

here: [https://help.salesforce.com/s/articleView?](https://help.salesforce.com/s/articleView?id=sf.flex_prompt_builder_troubleshoot.htm)

[id=sf.flex_prompt_builder_troubleshoot.htm](https://help.salesforce.com/s/articleView?id=sf.flex_prompt_builder_troubleshoot.htm)

NEW QUESTION: 20

- 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 (LEAVE A REPLY)

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
- * Users must have the "Use SDR Agent" permission set to access Agentforce SDR in their Salesforce environment.
- * If they lack this permission, the SDR Agent will not appear in their interface.
- * Agentforce Role-Based Access Control
- * Agentforce assigns specific permissions based on user roles.
- * Sales reps require explicit permission to access the SDR Agent.
- * Fixing the Issue
- * The Salesforce Admin should assign the "Use SDR Agent" permission set to all relevant sales reps.
- * This is done in Setup # Permission Sets # Assign to Users.

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: 21

A customer service representative is looking at a custom object that stores travel information. They recently received a weather alert and now need to cancel flights for the

customers that are related to this Itinerary. The representative needs to review the Knowledge articles about canceling and rebooking the customer flights.

Which Agentforce capability helps the representative accomplish this?

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

Answer: (SHOW ANSWER)

The scenario involves a customer service representative needing to cancel flights due to a weather alert and review existing Knowledge articles for guidance on canceling and rebooking. Agentforce provides capabilities to streamline such tasks. The most suitable option is Option B, which allows the agent to

"execute tasks based on available actions" (e.g., canceling flights via a predefined action) while "answering questions using information from accessible Knowledge articles." This capability leverages Agentforce's ability to integrate Knowledge articles into the agent's responses, enabling the representative to ask questions (e.g., "How do I cancel a flight?") and receive AI-generated answers grounded in approved Knowledge content.

Simultaneously, the agent can trigger actions (e.g., a Flow to update the custom object) to perform the cancellations, meeting all requirements efficiently.

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

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

* Option C: Generating a new Knowledge article based on prompts is not relevant. The representative needs to use existing articles, not author new ones, especially in a time-sensitive weather alert scenario.

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

:

Salesforce Agentforce Documentation: "Knowledge Replies and Actions" (Salesforce Help: <https://help>.

https://help.salesforce.com/s/articleView?id=sf.agentforce_knowledge_replies.htm&type=5) Trailhead: "Agentforce for Service" (<https://trailhead.salesforce.com/content/learn/modules/agentforce-for-service>)

NEW QUESTION: 22

Universal Containers (UC) is implementing Agentforce Service Agent on Email, UC made an email template and now needs to connect it to a Service Agent.

What should an Agentforce Specialist recommend?

- A. Create an Email Configuration for the Service Agent.
- B. Create an Omni-Channel flow to point to an email template.
- C. No action needed; the Service Agent connects automatically.

Answer: (SHOW ANSWER)

According to the AgentForce for Service Configuration Guide, when implementing Service Agents on Email, administrators must create an Email Configuration to connect the agent with the appropriate email channel and templates. The documentation specifies: "To enable Service Agents to handle emails, create an Email Configuration that links the agent to the email address, template, and routing parameters. This configuration allows the Service Agent to read, interpret, and respond using the defined template." Option B (creating an Omni-Channel flow) applies to routing live messages or chats, not configuring email agents.

Option C is incorrect because Service Agents do not automatically connect to email templates - a manual configuration is required.

Thus, Option A is correct as it aligns with Salesforce's documented process for connecting email templates to AgentForce Service Agents.

References (AgentForce Documents / Study Guide):

- * AgentForce for Service Setup Guide: "Creating Email Configurations for Service Agents"
- * Salesforce Service Cloud Email Configuration Overview
- * AgentForce Study Guide: "Deploying Service Agents on Email Channels"

NEW QUESTION: 23

An Agentforce Specialist created a Field Generation prompt template.

What should the Agentforce Specialist do to expose the template to the user?

- A. Use a screen flow to associate the Field Generation prompt template.
- B. Associate the template with the form field on the Lightning page.
- C. Call a template using an autolaunched flow.

Answer: (SHOW ANSWER)

The Field Generation prompt template type is specifically designed to enable generative AI within the context of a Salesforce record field. To expose this functionality to an end-user, the Agentforce Specialist must associate the template with the form field on the Lightning page (B). This is accomplished using the Lightning App Builder:

- * The Agentforce Specialist first creates a custom field (often a Long Text Area or Rich Text Area) on the desired object to store the AI-generated output.
- * In the Lightning App Builder for the object's Record Page, the Specialist selects the field component.
- * In the properties panel for that field component, there is a setting (often a dropdown) to select an active Field Generation Prompt Template.
- * Once associated, an Einstein icon (or "Generate" button) appears next to the field on the record page, allowing the user to click it to run the prompt, review the AI-generated content, and then decide to use it to populate the field.

Options A and C (using Flows) are methods for calling prompt templates to automate the generation of content or to ground the prompt with more complex data (like related list information). However, for the Field Generation prompt template to be exposed directly to the user for on-demand generation and manual review (the intended user experience for this template type), it must be bound to the field itself on the Lightning Record Page.

Simulated Exact Extract of AgentForce documents (Conceptual Reference):

"The Field Generation prompt template is surfaced to the user via the Lightning Record Page. After the prompt template is created and activated in Prompt Builder, the Agentforce Specialist must edit the Lightning Record Page in the Lightning App Builder. The key step is to select the target field component and, within its property panel, assign the Field Generation Prompt Template from the available dropdown menu. This action binds the generative AI capability directly to the field, displaying the 'Generate' button to the user to trigger the AI-assisted content creation upon the record." Simulated Reference:

AgentForce Study Guide, Chapter 3: Prompt Builder, Section 3.2: Field Generation Deployment, p. 55.

NEW QUESTION: 24

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 (LEAVE A REPLY)

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: 25

An Agentforce Specialist is creating a custom agent action. The topic is selected correctly, but the action is not.

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

A. Action Scope

B. Action Instructions

C. Classification Description

Answer: (SHOW ANSWER)

Per the AgentForce Custom Action Development Guide, if a topic is correctly triggered but the wrong action is executed, the issue typically lies in action instructions. The documentation notes: "Action instructions provide the LLM with explicit guidance on when and how to use a given action. Poorly written or ambiguous instructions can cause the reasoning engine to select an incorrect action, even within the right topic." Option A (Action Scope) defines data inputs/outputs, not reasoning behavior.

Option C (Classification Description) pertains to topic-level intent, not action execution. Thus, Option B - refining and testing the action instructions - ensures accurate behavior and action selection.

References (AgentForce Documents / Study Guide):

* AgentForce Action Creation Guide: "Testing and Refining Action Instructions"

* AgentForce Builder User Guide: "Ensuring Correct Action Selection"

* AgentForce Study Guide: "Troubleshooting Incorrect Action Mapping"

NEW QUESTION: 26

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: (SHOW ANSWER)

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: 27

Universal Containers (UC) wants to use the Draft with Einstein feature in Sales Cloud to create a personalized introduction email.

After creating a proposed draft email, which predefined adjustment should UC choose to revise the draft with a more casual tone?

A. Make Less Formal

B. Enhance Friendliness

C. Optimize for Clarity

Answer: ([SHOW ANSWER](#))

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

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

* Optimize for Clarity focuses on making the draft clearer but doesn't adjust the tone.

For more details, see Salesforce documentation on Einstein-generated email drafts and tone adjustments.

NEW QUESTION: 28

Universal Containers has an active standard email prompt template that does not fully deliver on the business requirements. Which steps should an Agentforce Specialist take to use the content of the standard prompt email template in question and customize it to fully meet the business requirements?

- A. Save as New Template and edit as needed.
- B. Clone the existing template and modify as needed.
- C. Save as New Version and edit as needed.

Answer: (SHOW ANSWER)

Universal Containers (UC) has a standard email prompt template (likely a prebuilt template provided by Salesforce) that isn't meeting their needs, and they want to customize it while retaining its original content as a starting point. Let's assess the options based on Agentforce prompt template management practices.

* Option A: Save as New Template and edit as needed. In Agentforce Studio's Prompt Builder, there's no explicit "Save as New Template" option for standard templates. This phrasing suggests creating a new template from scratch, but the question specifies using the content of the existing standard template

. Without a direct "save as" feature for standards, this option is imprecise and less applicable than cloning.

* Option B: Clone the existing template and modify as needed. Salesforce documentation confirms that standard prompt templates (e.g., for email drafting or summarization) can be cloned in Prompt Builder. Cloning creates a custom copy of the standard template, preserving its original content and structure while allowing modifications. The Agentforce Specialist can then edit the cloned template- adjusting instructions, grounding, or output format- to meet UC's specific business requirements. This is the recommended approach for customizing standard templates without altering the original, making it the correct answer.

* Option C: Save as New Version and edit as needed. Prompt Builder supports versioning for custom templates, allowing users to save new versions of an existing template to track changes. However, standard templates are typically read-only and cannot be versioned directly- versioning applies to custom templates after cloning. The question implies starting with the standard template's content, so cloning precedes versioning. This option is a secondary step, not the initial action, making it incorrect.

Why Option B is Correct:

Cloning is the documented method to repurpose a standard prompt template's content while enabling customization. After cloning, the specialist can modify the new custom template (e.g., tweak the email prompt's tone, structure, or grounding) to align with UC's requirements. This preserves the original standard template and follows Salesforce best practices.

References:

Salesforce Agentforce Documentation: Prompt Builder > Managing Templates - Details cloning standard templates for customization.

Trailhead: Build Prompt Templates in Agentforce - Explains how to clone standard templates to create editable copies.

Salesforce Help: Customize Standard Prompt Templates - Recommends cloning as the first step for modifying prebuilt templates.

NEW QUESTION: 29

A Service Agent at Universal Containers (UC) is designed to help customers resolve issues by searching against knowledge articles.

Knowledge articles have PDF attachments that add critical details. UC reports that the agent provides excellent summaries of the knowledge articles, but seems completely unaware of the PDF attachments.

How should an Agentforce Specialist configure the Data Cloud search index to include the content of these attached files?

- A.** Increase article chunk size and token limits for Knowledge indexing so larger contexts capture attachment references.
- B.** Enable 'Include Related Attachments' for Knowledge-- kav and map the ContentDocumentLink unstructured data model object (UDMO).
- C.** Use Data Cloud's 'Include Attachments' option and select the ContentDocumentVersion unstructured data model object (UDMO).

Answer: ([SHOW ANSWER](#))

The AgentForce Data Cloud Indexing Guide clearly states that to include content from attached files such as PDFs in Knowledge articles, the correct configuration is to enable "Include Attachments" and map the ContentDocumentVersion unstructured data model object (UDMO). The documentation specifies:

"When indexing Knowledge or Case data, enabling the 'Include Attachments' option allows the Data Cloud index to extract and embed content from linked ContentDocumentVersion records, ensuring the agent retrieves relevant information from attachments." Option A (increasing chunk size) does not enable attachment ingestion. Option B (ContentDocumentLink mapping) only establishes a relationship, not content extraction. Therefore, Option C ensures attachment content becomes searchable within AgentForce retrieval.

References (AgentForce Documents / Study Guide):

- * AgentForce Data Cloud Indexing and Retrieval Guide: "Including Attachments in Knowledge Indexes"
- * AgentForce Implementation Handbook: "Mapping UDMO for ContentDocumentVersion"
- * AgentForce Study Guide: "Attachment Content Inclusion for Knowledge Articles"

NEW QUESTION: 30

A. Filters narrow the search results by applying up to 10 conditions based on fields defined in the search index, thereby enhancing the relevancy of the content returned.

B. Filters automatically encrypt and mask sensitive fields in the search index to ensure that only non- confidential information is retrieved for public queries.

C. Filters reformat and aggregate multiple documents into a single summary output to streamline and unify retriever output for more efficient and accurate AI grounding.

Answer: ([SHOW ANSWER](#))

The AgentForce Retriever Configuration Guide specifies that filters are used to refine and constrain search results within a retriever setup. Filters operate by applying conditions (up to 10) on indexed fields such as document type, category, region, or update date. This targeted filtering ensures that the retrieved data is highly relevant to the current user query or context.

For instance, an AgentForce retriever could be configured to include only documents tagged as "Active" or within a specific product line, reducing noise and improving grounding accuracy for the LLM. This mechanism supports precision retrieval, which directly improves both the accuracy and reliability of generated responses.

Option B is incorrect because filters do not handle encryption or masking of sensitive data - those functions are managed through Data Cloud security and access controls. Option C is incorrect because retrievers do not aggregate or summarize documents; they retrieve data for grounding, leaving summarization to the LLM reasoning layer.

Therefore, the correct answer is Option A - Filters narrow search results using field-based conditions to improve relevancy and retrieval precision.

Reference: AgentForce Implementation Guide - "Configuring Filters for Targeted Retrieval in Custom Retriever Settings."

NEW QUESTION: 31

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

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

- * How it is masked
- * The masked response

What should the Agentforce Specialist recommend?

- A.** Ingest the Einstein Shield Event logs into CRM Analytics.
- B.** Review the debug logs of the running user.
- C.** Enable audit trail in the Einstein Trust Layer.

Answer: C (LEAVE A REPLY)

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

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

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

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

Valid Agentforce-Specialist Dumps shared by ExamDiscuss.com for Helping Passing Agentforce-Specialist Exam! ExamDiscuss.com now offer the **newest Agentforce-Specialist exam dumps**, the ExamDiscuss.com Agentforce-Specialist exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com Agentforce-Specialist dumps with Test Engine here:
<https://www.examd Discuss.com/Salesforce/exam/Agentforce-Specialist/premium/> (300 Q&As Dumps, **35%OFF Special Discount Code: freecram**)

NEW QUESTION: 32

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

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

Answer: (SHOW ANSWER)

When using related list merge fields in a prompt template associated with the Account object in Prompt Builder, the Activities related list is not supported due to it being a polymorphic field. Polymorphic fields can reference multiple different types of objects, which makes them incompatible with some merge field operations in prompt generation.

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

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

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

NEW QUESTION: 33

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

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

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

Answer: (SHOW ANSWER)

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

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

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

Model Builder enables seamless integration with external systems and models, allowing Salesforce users to use external LLMs for generating AI-driven insights and automation.

Salesforce Agentforce Specialist References:

For more details, check the Model Builder guide here:

[https://help.salesforce.com/s/articleView?id=sf.](https://help.salesforce.com/s/articleView?id=sf.model_builder_external_models.htm)

[model_builder_external_models.htm](https://help.salesforce.com/s/articleView?id=sf.model_builder_external_models.htm)

NEW QUESTION: 34

What is true of Agentforce Testing Center?

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

Answer: (SHOW ANSWER)

The Agentforce Testing Center is a tool in Agentforce Studio for validating agent performance. Let's evaluate the statements.

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

* Option B: Running tests does not consume Einstein Requests. Einstein Requests are part of the usage quota for Einstein Generative AI features (e.g., prompt executions in production). Testing Center uses synthetic data to simulate interactions without invoking live AI calls that count against this quota.

Salesforce documentation confirms tests don't consume requests, making this the correct answer.

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

Why Option B is Correct:

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

References:

Salesforce Agentforce Documentation: Testing Center > Overview - Confirms no request consumption.

Trailhead: Test Your Agentforce Agents - Notes quota-free testing.

Salesforce Help: Agentforce Testing - Details safe, isolated testing.

NEW QUESTION: 35

Choose 1 option.

Universal Containers (UC) recently attended a major trade show and received thousands of new leads from event badge scans. UC is struggling to follow up with each lead in a timely, personalized way. Leadership wants to:

Qualify and nurture leads 24/7.

* Provide accurate answers to prospect questions.

* Automatically book meetings with qualified prospects.

* Free up reps to focus on building relationships and closing deals.

Which Agentforce capability should UC implement to meet these goals?

A. SDR Agent

B. Sales Coach

C. Commerce Agent

Answer: (SHOW ANSWER)

Universal Containers (UC) needs a solution that can automatically qualify and nurture thousands of new leads

24/7, provide accurate and consistent responses to prospects, schedule meetings for qualified leads, and allow sales representatives to focus on relationship building and closing deals. These needs align precisely with the Agentforce SDR Agent.

According to official AgentForce documentation, "Agentforce SDR helps sales teams qualify and nurture leads at scale, around the clock. It acts as a digital sales development representative capable of engaging new leads instantly, asking the right qualifying questions, answering inquiries accurately using connected Salesforce data, and automatically scheduling meetings on behalf of the sales team." The documentation further explains that the SDR Agent is designed to "personalize outreach, manage follow-up sequences, and book meetings directly from your website or campaign pages." This automation "frees your human reps to focus on high-value interactions and closing opportunities rather than manual lead qualification." By contrast, the Sales Coach capability focuses on guiding and coaching sales representatives internally rather than interacting with prospects, and the Commerce Agent is designed for e-commerce use cases such as assisting shoppers with product discovery and order management-not lead nurturing.

References (AgentForce Documents / Study Guide):

* AgentForce SDR Overview - Salesforce AgentForce Documentation

* AgentForce for Sales - SDR Agent Use Cases

- * AgentForce Study Guide: "Qualify and Nurture Leads at Scale with SDR Agents"
- * Salesforce Trailhead: "Get to Know AgentForce SDR"

NEW QUESTION: 36

What does it mean when a prompt template version is described as immutable?

- A.** Only the latest version of a template can be activated.
- B.** Every modification on a template will be saved as a new version automatically.
- C.** Prompt template version is activated; no further changes can be saved to that version.

Answer: C (LEAVE A REPLY)

When a prompt template version is immutable, it means that once the version is activated, it cannot be edited or modified. This ensures consistency in production environments where changes could disrupt workflows.

- * Option A is incorrect: Any version (not just the latest) can be activated, depending on the use case.
- * Option D is incorrect: Modifications require manually creating a new version; automatic versioning is not enforced.
- * Option C is correct: Activation locks the version, enforcing immutability.

References:

- * Salesforce Help: Prompt Template Versioning
- * States that "activated prompt template versions are immutable and cannot be edited."

NEW QUESTION: 37

Universal Containers implements three custom actions to get three distinct types of sales summaries for its users. Users are complaining that they are not getting the right summary based on their utterances. What should the Agentforce Specialist investigate as the root cause?

- A.** Review that the custom action is assigned to an Agent.
- B.** Review the action instructions to ensure they are unique.
- C.** Ensure the input and output types are correctly chosen.

Answer: (SHOW ANSWER)

The root cause of users receiving incorrect sales summaries lies in non-unique action instructions (Option B). In Einstein Bots, custom actions are triggered based on how well user utterances align with the action instructions defined for each action. If the instructions for the three custom actions overlap or lack specificity, the bot's natural language processing (NLP) cannot reliably distinguish between them, leading to mismatched responses.

Steps to Investigate:

- * Review Action Instructions: Ensure each custom action has distinct, context-specific instructions. For example:
- * Action 1: "Summarize quarterly sales by region."
- * Action 2: "Generate a product-wise sales breakdown for the current fiscal year."

* Action 3: "Provide a comparison of sales performance between online and in-store channels." Ambiguous or overlapping instructions (e.g., "Get sales summary") cause confusion.

* Test Utterance Matching: Use Einstein Bot's training tools to validate if user utterances map to the correct action.Overlap indicates instruction ambiguity.

* Refine Instructions: Incorporate keywords or phrases unique to each sales summary type to improve intent detection.

Why Other Options Are Incorrect:

* A. Assigning actions to an agent is irrelevant, as custom actions are automated bot components.

* C. Input/output types relate to data formatting, not intent routing.While important for execution, they don't resolve utterance mismatches.

:

Einstein Bot Developer Guide: Stresses the need for unique action instructions to avoid intent conflicts.

Trailhead Module: "Build AI-Powered Bots with Einstein" highlights instruction specificity for accurate action triggering.

Salesforce Help Documentation: Recommends testing and refining action instructions to ensure clarity in utterance mapping.

NEW QUESTION: 38

Universal Containers is indexing millions of product manuals where users may ask both structured queries (model numbers) and natural language questions (for example, "How do I reset my device?"), Which retrieval approach should the company use?

A. Use keyword search only, since model numbers dominate queries.

B. Use semantic search only, as natural language is always preferred.

C. Use hybrid search to combine keyword precision with semantic flexibility.

Answer: (SHOW ANSWER)

According to the AgentForce Retrieval Optimization Guide, when users ask both structured (exact) and unstructured (natural language) questions, the best practice is to use hybrid search. The documentation states:

"Hybrid search combines the precision of keyword retrieval for structured terms, such as IDs or model numbers, with the semantic flexibility of vector search for natural language queries. This approach ensures both deterministic accuracy and contextual understanding." Option A (keyword search only) fails for natural language queries, which require semantic understanding.

Option B (semantic search only) may misinterpret or overlook structured identifiers like product or model numbers.

Therefore, Option C-hybrid search-provides the ideal balance between exact match precision and contextual recall.

References (AgentForce Documents / Study Guide):

* AgentForce Retrieval and Indexing Guide: "Choosing Between Keyword, Semantic, and Hybrid Search"

* AgentForce Data Cloud Handbook: "Optimizing Multi-Modal Retrieval Strategies"

* AgentForce Study Guide: "Hybrid Search for Structured and Unstructured Queries"

NEW QUESTION: 39

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

What should An Agentforce recommend to generate this email?

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

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

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

Answer: ([SHOW ANSWER](#))

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

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

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

NEW QUESTION: 40

A. The user does not have the Prompt Template User permission set assigned.

B. The prompt template associated with summary field is not activated for that user.

C. The user does not have the field Generative AI User permission set assigned.

Answer: ([SHOW ANSWER](#))

In Salesforce, Generative AI capabilities are controlled by specific permission sets. To use features such as generating summaries with AI, users need to have the correct permission sets that allow access to these functionalities.

* **Generative AI User Permission Set:** This is a key permission set required to enable the generative AI capabilities for a user. In this case, the missing Generative AI User permission set prevents the user from seeing the generative AI-enabled field icon. Without this permission, the generative AI feature in the Competitor Analysis custom object won't be accessible.

* **Why not A?** The Prompt Template User permission set relates specifically to users who need access to prompt templates for interacting with Einstein GPT, but it's not directly related to the visibility of AI-enabled field icons.

* **Why not B?** While a prompt template might need to be activated, this is not the primary issue here. The question states that other users with the same profile can see the icon, so the problem is more likely to be permissions-based for this particular user.

For more detailed information, you can review Salesforce documentation on permission sets related to AI capabilities at [Salesforce AI Documentation](#) and [Einstein GPT permissioning guidelines](#).

NEW QUESTION: 41

Which configuration must An Agentforce complete for users to access generative AI-enabled fields in the Salesforce mobile app?

- A.** Enable Mobile Generative AI.
- B.** Enable Mobile Prompt Responses.
- C.** Enable Dynamic Forms on Mobile.

Answer: ([SHOW ANSWER](#))

* Context of the Question

* Universal Containers (UC) has generative AI-enabled fields that users can access in the desktop experience.

* The Agentforce Specialist needs these same fields to be visible and usable in the Salesforce Mobile App.

* Why Dynamic Forms on Mobile?

* Dynamic Forms allow you to configure record pages so that fields and sections can appear or be hidden based on certain criteria.

* When you enable "Dynamic Forms for Mobile," any generative AI-enabled fields placed on the dynamic layout become accessible in the Salesforce mobile experience.

* There is no standard Setup option labeled "Enable Mobile Generative AI" or "Enable Mobile Prompt Responses" as a universal toggle; the existing official approach is to ensure dynamic forms (and the relevant fields) are supported on mobile.

* Conclusion

* Ensuring that these AI-driven fields are visible on mobile is accomplished by turning on Dynamic Forms on Mobile and adding those fields to the dynamic layout. Therefore, Option C is correct.

Salesforce Agentforce Specialist References & Documents

* [Salesforce Documentation: Dynamic Forms Overview](#) Explains how to enable Dynamic Forms for both desktop and mobile UIs, allowing newly added fields (including generative AI-enabled ones) to display in the Salesforce Mobile App.

* [Salesforce Agentforce Specialist Study Guide](#) Reiterates that to expose generative AI fields or components in mobile, you must configure dynamic forms and ensure compatibility on mobile layouts.

NEW QUESTION: 42

An Agentforce Specialist wants to ensure their custom agent action performs as expected in conversations.

What should the Agentforce Specialist focus on when creating action instructions?

- A.** Write concise agent action instructions and test in Agentforce Builder.

- B.** Ensure the agent action label matches the utterance's intent.
- C.** Include comprehensive detailed descriptions and perform smoke testing.

Answer: ([SHOW ANSWER](#))

The AgentForce Action Design Guide emphasizes that concise and clear action instructions are essential for predictable and reliable agent behavior. The documentation states: "Action instructions should be concise, clearly define the purpose of the action, and specify its expected inputs and outputs. Each action must be validated through testing in AgentForce Builder to confirm the LLM interprets and executes it as intended." This aligns with Option A, which focuses on clarity and testing.

Option B incorrectly focuses on label-intent matching, which is a naming best practice but not sufficient for accuracy.

Option C promotes overly detailed instructions, which can lead to LLM confusion and inconsistent results.

Therefore, Option A best aligns with Salesforce's official best practices for creating and validating custom actions.

References (AgentForce Documents / Study Guide):

- * AgentForce Action Design and Testing Guide
- * AgentForce Builder Documentation: "Validating Custom Actions"
- * AgentForce Study Guide: "Creating Clear and Testable Action Instructions"

NEW QUESTION: 43

What is best practice when refining Agent custom action instructions?

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

Answer: ([SHOW ANSWER](#))

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

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

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

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

NEW QUESTION: 44

Universal Containers' sales team engages in numerous video sales calls with prospects across the nation. Sales management wants an easy way to understand key information such as deal terms or customer sentiments.

Which Einstein Generative AI feature should An Agentforce recommend for this request?

- A. Einstein Call Summaries
- B. Einstein Conversation Insights
- C. Einstein Video KPI

Answer: (SHOW ANSWER)

Einstein Conversation Insights analyzes sales calls (including video and voice) to capture key moments, deal terms, competitor mentions, and customer sentiment. It helps sales management quickly understand what was discussed without manually reviewing the entire call. This feature ensures actionable insights are delivered directly into the Salesforce CRM, allowing sales managers to gain a concise overview without manually reviewing long recordings.

Reference:

"Boost Sales with Automated AI Strategies | Salesforce Trailhead" .

"Introduction to Einstein Discovery | Salesforce" .

NEW QUESTION: 45

Universal Containers would like to route SMS text messages to a service rep from an Agentforce Service Agent. Which Service Channel should the company use in the flow to ensure it's routed properly?

- A. Messaging
- B. Route Work Action
- C. Live Agent

Answer: A (LEAVE A REPLY)

UC wants to route SMS text messages from an Agentforce Service Agent to a service rep using a flow. Let's identify the correct Service Channel.

* Option A: Messaging In Salesforce, the "Messaging" Service Channel (part of Messaging for In-App and Web or SMS) handles text-based interactions, including SMS. When integrated with Omni-Channel Flow, the "Route Work" action uses this channel to route SMS messages to agents. This aligns with UC' s requirement for SMS routing, making it the correct answer.

* Option B: Route Work Action "Route Work" is an action in Omni-Channel Flow, not a Service Channel. It uses a channel (e.g., Messaging) to route work, so this is a component, not the channel itself, making it incorrect.

* Option C: Live Agent "Live Agent" refers to an older chat feature, not the current Messaging framework for SMS. It's outdated and unrelated to SMS routing, making it incorrect.

* Option D: SMS Channel There's no standalone "SMS Channel" in Salesforce Service Channels-SMS is encompassed within the "Messaging" channel. This is a misnomer, making it incorrect.

Why Option A is Correct:

The "Messaging" Service Channel supports SMS routing in Omni-Channel Flow, ensuring proper handoff from the Agentforce Service Agent to a rep, per Salesforce documentation.

References:

Salesforce Agentforce Documentation: Omni-Channel Integration > Messaging - Details SMS in Messaging channel.

Trailhead: Omni-Channel Flow Basics - Confirms Messaging for SMS.

Salesforce Help: Service Channels - Lists Messaging for text-based routing.

NEW QUESTION: 46

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: ([SHOW ANSWER](#))

* Context of the Question

* Universal Containers (UC) wants to build a Flex prompt template that uses data returned by a Flow.

* "Flex Prompt Templates" allow admins and Agentforce Specialists to incorporate external or dynamic data into generative AI prompts.

* Why "Add Flow Instructions" Is Needed

* Passing Flow Data into Prompt Templates: When configuring the prompt, you must specify how data from the running Flow is passed into the Flex template. The designated element for that is typically "Flow Instructions," which map the Flow outputs to the prompt.

* Other Options:

* Add Flex Instructions: Typically controls how the AI responds or structures the output, not how to bring Flow data into the template.

* Add Prompt Instructions: Usually for static or manual instructions that shape the AI's response, rather than referencing dynamic data from the Flow.

* Outcome

* "Add Flow Instructions" ensures the prompt can dynamically use the data that the Flow returns- making Option C correct.

Salesforce Agentforce Specialist References & Documents

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

* Salesforce Agentforce Specialist Study Guide Outlines how to configure generative AI prompts that reference real-time Flow data.

Valid Agentforce-Specialist Dumps shared by ExamDiscuss.com for Helping Passing Agentforce-Specialist Exam! ExamDiscuss.com now offer the **newest Agentforce-Specialist exam dumps**, the ExamDiscuss.com Agentforce-Specialist exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com Agentforce-Specialist dumps with Test Engine here:
<https://www.examdiscuss.com/Salesforce/exam/Agentforce-Specialist/premium/> (300 Q&As Dumps, **35%OFF Special Discount Code: freecram**)

NEW QUESTION: 47

An Agentforce implements Einstein Sales Emails for a sales team. The team wants to send personalized follow-up emails to leads based on their interactions and data stored in Salesforce. The Agentforce Specialist needs to configure the system to use the most accurate and up-to-date information for email generation.

Which grounding technique should the Agentforce Specialist use?

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

Answer: C (LEAVE A REPLY)

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

* Record Merge Fields ensure the generated emails are highly personalized using data like lead name, company, or other Salesforce fields directly from the records.

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

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

:

Salesforce Einstein Sales Emails Documentation:

https://help.salesforce.com/s/articleView?id=release-notes.rn_einstein_sales_emails.htm

NEW QUESTION: 48

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

In the score of 1 indicate?

- A.** The response is the least toxic Einstein Generative AI Toxicity Scoring system, what does a toxicity category.
- B.** The response is not toxic.
- C.** The response is the most toxic.

Answer: (SHOW ANSWER)

Einstein Trust Layer's Toxicity Scoring categorizes content on a scale of 0 to 1, where 1 indicates the highest level of toxicity (e.g., harmful, biased, or inappropriate language). This scoring helps organizations filter unsafe AI-generated content. A score of 1 triggers mitigation actions, such as blocking the response or alerting administrators.

* A score of 0 would indicate no toxicity (B is incorrect).

* The scoring system does not use "least toxic" as a category (A is misleading).

Reference:

Salesforce Help Article: Einstein Trust Layer - Toxicity Scoring ("Interpreting Toxicity Scores" section).

Einstein GPT Safety Overview: "Mitigating Harmful Content with Toxicity Detection."

NEW QUESTION: 49

A service agent is looking at a custom object that stores travel information. They recently received a weather alert and now need to cancel flights for the customers that are related with this itinerary. The service agent needs to review the Knowledge articles about canceling and rebooking the customer flights.

Which Agent capability helps the agent accomplish this?

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

Answer: (SHOW ANSWER)

In this scenario, the Agent capability that best helps the agent is its ability to execute tasks based on available actions and answer questions using data from Knowledge articles. Agent can assist the service agent by providing relevant Knowledge articles on canceling and rebooking flights, ensuring that the agent has access to the correct steps and procedures directly within the workflow.

This feature leverages the agent's existing context (the travel itinerary) and provides actionable insights or next steps from the relevant Knowledge articles to help the agent quickly resolve the customer's needs.

The other options are incorrect:

* B refers to invoking a flow to create a Knowledge article, which is unrelated to the task of retrieving existing Knowledge articles.

* C focuses on generating Knowledge articles, which is not the immediate need for this situation where the agent requires guidance on existing procedures.

:

Salesforce Documentation on Agent

Trailhead Module on Einstein for Service

NEW QUESTION: 50

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: (SHOW ANSWER)

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: 51

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

A. Service Replies

B. Case Replies

C. Knowledge Replies

Answer: (SHOW ANSWER)

Service Replies (specifically Einstein Service Replies) is the Salesforce Generative AI functionality designed to automatically draft responses for service agents in real-time, based on contextual information, including existing knowledge articles. This directly

addresses Universal Containers' need to reduce handling time by providing recommended replies grounded in their knowledge base

NEW QUESTION: 52

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: ([SHOW ANSWER](#))

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: 53

A service manager wants to use Salesforce Prompt Builder to help agents summarize customer case notes after a support call.

The summary should:

- * Capture the customer's issue, troubleshooting steps taken, and next actions.
- * Be no longer than five sentences.
- * Use plain language (no technical jargon).

If no next action is identified, the summary should explicitly state "No next action required."

Which prompt template follows Salesforce prompt design best practices?
required."

Format: Use numbered sentences for clarity.

A. Role: You are an experienced support agent.

Task: Summarize the case notes,

Context: Include customer issue, troubleshooting steps, and next actions.

Constraints: Limit to 5 sentences, use plain language, and if no next action is found, state "No next action"

B. Role: You are a support agent writing a case summary.

Task: Provide a professional summary of the issue and troubleshooting steps.

Context: Include customer issue, steps taken, and next actions if available.

Constraints: No strict sentence limit, but use plain language. If no next action is found, leave it out.

Format: Use paragraphs for readability.

C. Role: You are a case documentation assistant,

Task: Write a summary of the support call.

Context: Always describe the customer issue, troubleshooting, and resolution details.

Constraints: The summary should be comprehensive and professional, but there is no limit on length or language style.

Format: Use complete sentences in a narrative style.

Answer: A (LEAVE A REPLY)

According to the Salesforce Prompt Builder Best Practices Guide, an effective prompt must include Role, Task, Context, Constraints, and Format clearly defined - a structure known as the RTCCF model. The documentation explains: "Prompts should specify the assistant's role, define a clear task, include context and constraints, and provide output format instructions to ensure predictable and high-quality responses." Option A follows this framework precisely. It defines:

- * Role: The assistant's identity ("experienced support agent").

- * Task: Summarizing case notes.

- * Context: Customer issue, troubleshooting steps, next actions.

- * Constraints: Limit of 5 sentences, plain language, include "No next action required" if applicable.

- * Format: Numbered sentences for clarity.

Options B and C omit critical prompt design elements such as strict constraints or output formatting and therefore do not align with Salesforce's prompt design standard.

References (AgentForce Documents / Study Guide):

- * Salesforce Prompt Builder Guide: "Prompt Structure Using RTCCF Model"

- * AgentForce Prompt Template Design Guide: "Best Practices for Summarization Prompts"

- * Salesforce AI Prompt Engineering Study Guide

NEW QUESTION: 54

Support agents at Universal Containers are using Agentforce to find troubleshooting information. They've reported that the agent frequently provides knowledge articles that are outdated, even when newer versions of the articles are available. The administrator has confirmed that all articles are correctly chunked and indexed.

Which configuration change in the Data Cloud hybrid search index best addresses this problem?

- A. Disable the keyword index to rely solely on the vector index.
- B. Switch the chunking strategy from section-aware to fixed-size.
- C. Add a ranking factor for recency based on the LastModifiedDate field.

Answer: (SHOW ANSWER)

The AgentForce Data Cloud Retrieval and Ranking Guide highlights that when outdated Knowledge articles appear before newer ones, administrators should configure ranking factors that prioritize content based on recency. The documentation specifies: "Adding a recency ranking factor using the LastModifiedDate or LastPublishedDate fields ensures the retrieval prioritizes the most up-to-date documents, improving response relevance." Option A (disabling keyword index) would remove precision in retrieval and does not address recency.

Option B (changing chunking strategy) affects data segmentation, not ranking order. Therefore, Option C - adding a ranking factor for recency - is the correct way to ensure updated articles are prioritized.

References (AgentForce Documents / Study Guide):

- * AgentForce Data Cloud Hybrid Search Configuration Guide: "Applying Recency Ranking"
- * AgentForce Knowledge Management Handbook: "Prioritizing Updated Articles in Search"
- * AgentForce Study Guide: "Ranking and Weighting Strategies for Knowledge Retrieval"

NEW QUESTION: 55

Universal Containers wants to incorporate CRM data as well-formatted JSON in a prompt to a large language model (LLM).

What is an important consideration for this requirement?

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

Answer: (SHOW ANSWER)

* Context of the Question

* Universal Containers (UC) wants to send well-formatted JSON data in a prompt to a large language model (LLM).

* The question is about an important technical or design consideration for including CRM data as JSON in that prompt.

* Why Apex Code for JSON Formatting?

* Apex to Generate JSON: Salesforce does not have a simple "checkbox" or single setting to

"convert CRM data to JSON." Typically, to structure data as JSON in a template, you either:

* Use an Apex class that queries or processes the data, then returns a JSON string.

* Use a Flow or formula approach (though complex data structures often require Apex).

* No Built-In "Enable JSON Format in Prompt Builder": Prompt Builder doesn't have a toggle that automatically transforms data into JSON.

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

Salesforce Agentforce Specialist References & Documents

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

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

NEW QUESTION: 56

Universal Containers has a new AI project.

What should An Agentforce consider when adding a related list on the Account object to be used in the prompt template?

A. After selecting a related list from the Account, use the field picker to choose merge fields in Prompt Builder.

B. Prompt Builder must be used to assign the fields from the related list as a JSON format.

C. The fields for the related list are based on the default page layout of the Account for the current user.

Answer: (SHOW ANSWER)

* Context of the Question Universal Containers (UC) wants to include details from a related list on the Account object in a prompt template. This is typically done via Prompt Builder in Salesforce's generative AI setup.

* Prompt Builder Behavior

* Selecting a Related List: Within Prompt Builder, you can navigate to the object (Account) and choose which related list (e.g., Contacts, Opportunities) you want to reference.

* Field Picker: Once a related list is chosen, Prompt Builder provides a field picker interface, allowing you to select specific fields from that related list. These fields then become available for merge fields or dynamic insertion within your prompt.

* Why Option A is Correct

* Direct Alignment with the Standard Process: The recommended approach in Salesforce's documentation is to select a related list and then use the field picker to add the necessary fields into your AI prompt. This ensures the prompt has exactly the data you need from that related list.

* Why Not Option B (JSON Formatting)

* No Mandatory JSON Requirement: Although you can structure data as JSON if you desire advanced formatting, Prompt Builder does not require you to manually assign the fields from the related list in JSON. The platform automatically handles how the data is passed along in the background.

* Why Not Option C (Default Page Layout)

* Independent of Page Layout: Prompt Builder does not rely strictly on the default page layout for fields. You can configure the fields you want from the related list, independent of how the user's page layout is set up in the UI.

* Conclusion Since the official Salesforce approach involves selecting a related list and then using the field picker to insert merge fields, Option A is the correct and verified answer. Salesforce Agentforce Specialist References & Documents

* Salesforce Official Documentation: Prompt Builder Basics Explains how to reference objects and related lists when building AI prompts.

* Salesforce Trailhead: Get Started with Prompt Builder Provides hands-on exercises demonstrating how to pick fields from related objects or lists.

* Salesforce Agentforce Specialist Study Guide Outlines best practices for referencing related records and fields in generative AI prompts.

NEW QUESTION: 57

Choose 1 option.

Before activating a custom agent action, an AgentForce Specialist would like to evaluate multiple real-world user utterances to ensure the action is being selected appropriately.

Which tool should the AgentForce Specialist recommend?

A. Testing Center

B. AgentForce Builder

C. Prompt Builder

Answer: (SHOW ANSWER)

Comprehensive and Detailed Explanation From Exact Extract of AgentForce Documents: The AgentForce Testing Center is the recommended tool for validating and refining agent behavior before activation. According to the AgentForce Quality Assurance and Testing Framework, the Testing Center enables specialists to simulate and analyze how the reasoning engine classifies user utterances, selects topics, and triggers actions across multiple test cases.

This testing environment provides detailed diagnostics, showing classification confidence, prompt resolution paths, and grounding behavior. It allows for adjustments before production deployment to ensure accuracy and relevance in real-world scenarios.

Option B, AgentForce Builder, is used for creating and configuring agents, topics, and actions but not for performing comparative or scenario-based testing. Option C, Prompt Builder, is specifically for designing and testing individual prompt templates, not full agent action selection across utterances.

Thus, the correct answer is Option A - Testing Center, as it provides the purpose-built environment for multi-utterance validation and end-to-end agent testing.

Reference: AgentForce Specialist Handbook - "Using Testing Center for Action Selection Validation and Performance Evaluation."

NEW QUESTION: 58

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

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

A. The flow must have the "with sharing" permission selected in the advanced settings for the permissions, field-level security, and sharing settings to be respected.

B. The custom action adheres to the permissions, field-level security, and sharing settings configured in the flow.

C. The Agent will always run flows in system mode so the running user's data access will not affect the data returned.

Answer: (SHOW ANSWER)

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

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

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

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

References:

* Salesforce Help: Flow Runtime and Security Context

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

NEW QUESTION: 59

Choose 1 option.

Coral Cloud Resorts is uploading thousands of new HTML knowledge articles files for a resort launch.

To ensure Agentforce retrieves accurate responses quickly, which chunking strategy should be used when creating a new index?

A. Semantic-based passage extraction

B. Conversation-based chunking

C. Section-aware chunking

Answer: (SHOW ANSWER)

In AgentForce documentation on Knowledge Indexing and Chunking Strategies, Salesforce emphasizes that when uploading large volumes of structured content such as HTML or documentation files, the system should use section-aware chunking. The guide states: "Section-aware chunking preserves the logical boundaries of headings, paragraphs, and sub-sections in structured documents like HTML or PDF files, allowing the agent to retrieve contextually accurate and relevant responses quickly." This method

ensures that the agent does not split content mid-section or lose contextual relationships between headings and body text. It enhances both retrieval speed and answer precision. Option A, semantic-based passage extraction, is better suited for free-text knowledge bases, where meaning needs to be inferred. Option B, conversation-based chunking, applies only to chat logs or dialogue histories.

For HTML documentation and structured articles, section-aware chunking ensures optimized retrieval and minimal latency in AgentForce responses.

References (AgentForce Documents / Study Guide):

- * AgentForce Knowledge Management Guide: "Choosing the Right Chunking Strategy"
- * AgentForce Indexing and Retrieval Optimization Study Notes
- * AgentForce Developer Handbook: "Implementing Section-Aware Chunking for Structured Files"

NEW QUESTION: 60

A. Create a prompt template for product tutorials and guides.

B. Add an Answer Questions custom field in the product object for tutorial instructions.

C. Publish product tutorials and guides as Knowledge articles.

Answer: ([SHOW ANSWER](#))

* Context of the Question Universal Containers (UC) wants its support agents to use Agentforce to ask questions about product tutorials and product guides. Agentforce typically references knowledge sources to provide accurate and contextual responses.

* Why Knowledge Articles?

* Centralized Repository: Publishing product tutorials and guides as Knowledge articles in Salesforce ensures that the information is readily available and searchable by Agentforce.

* AI Integration: Salesforce's AI solutions, including Agentforce, can often be configured to pull content directly from Salesforce Knowledge articles, giving users on-demand answers without manual data duplication.

* Maintenance & Updates: Storing content in Salesforce Knowledge simplifies content updates, versioning, and user permissions.

* Why Not the Other Options?

* Option A (Create a Prompt Template): Creating a prompt template alone does not solve how the underlying content (tutorials, guides) is stored or accessed by Agentforce. Prompt templates shape the queries/responses but do not provide the knowledge base.

* Option B (Add an Answer Questions Custom Field): A single field on the product object is insufficient for the depth of information found in tutorials and guides. It also lacks the robust search and user-friendly interface that Knowledge articles provide.

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

Salesforce Agentforce Specialist References & Documents

* Salesforce Documentation: Set Up Salesforce Knowledge Discusses how to publish articles for easy access

* by AI-driven assistants and support teams.

* Salesforce Agentforce Specialist Study Guide Explains best practices for feeding knowledge sources to generative AI and Agentforce.

NEW QUESTION: 61

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: (SHOW ANSWER)

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).

Valid Agentforce-Specialist Dumps shared by ExamDiscuss.com for Helping Passing Agentforce-Specialist Exam! ExamDiscuss.com now offer the **newest Agentforce-Specialist exam dumps**, the ExamDiscuss.com Agentforce-Specialist exam **questions have been updated** and **answers have been corrected** get the **newest**

ExamDiscuss.com Agentforce-Specialist dumps with Test Engine here:

<https://www.examdisscuss.com/Salesforce/exam/Agentforce-Specialist/premium/> (300

Q&As Dumps, **35%OFF Special Discount Code: freecram**)

NEW QUESTION: 62

Universal Containers (UC) has configured a data library and wants to restrict indexing of knowledge articles to articles which are only publicly available in their knowledge base, UC also wants the agent to link sources that the large language model (LLM) grounded its response on.

Which settings should help UC with this?

A. Use Data Categories to categorize publicly available articles to index. Sources are automatically displayed when knowledge articles are categorized as Public.

B. In the data library setting window, under Knowledge Settings, enable Use Public Knowledge Article. It is not possible to display articles that the LLM grounded its response in.

C. In the data library setting window, under Knowledge Settings, enable Use Public Knowledge Article and select Show sources,

Answer: ([SHOW ANSWER](#))

According to the AgentForce Data Library Configuration Guide, administrators can restrict indexing and retrieval of Knowledge articles to publicly available ones and enable source visibility for LLM-grounded responses. The documentation states:

"Within the data library settings, under Knowledge Settings, enable 'Use Public Knowledge Articles' to ensure only publicly visible content is indexed. To display citations, enable 'Show Sources' so the agent links the specific articles or data records used to ground its response." Option A correctly reflects these two documented configuration steps.

Option B is incorrect because Salesforce explicitly supports source display for transparency through the

"Show Sources" setting.

Option C incorrectly assumes that Data Categories control indexing visibility and source linking, which is handled by explicit Knowledge Settings, not categorization.

References (AgentForce Documents / Study Guide):

* AgentForce Data Library Configuration Guide: "Knowledge Settings for Public Article Indexing"

* AgentForce Transparency and Source Attribution Notes: "Show Sources Option"

* AgentForce Study Guide: "Configuring Knowledge Visibility and Source Display"

NEW QUESTION: 63

An administrator at Universal Containers is setting up a new Sales Development Representative (SDR) Agent.

The agent's purpose is to nurture cold leads before connecting them to the assigned sales rep.

To ensure the agent has all the necessary access to the leads in the North America sales region, what should the administrator data?

A. Assign the user in the highest-level role within the North America role hierarchy as the SDR Agent User.

B. Grant View All record permission of the Lead object to the 'Einstein Agent User' profile,
C. Create a criteria-based sharing rule to grant access to targeted lead records to SDR Agent User.

Answer: (SHOW ANSWER)

According to the AgentForce Security and Data Access Configuration Guide, the best practice for ensuring an agent (such as an SDR Agent) can access specific records while maintaining security is to use criteria-based sharing rules. The documentation states: "When an AI Agent needs access to a subset of records (for example, regional leads), create a sharing rule granting access to that agent's user context based on defined criteria such as region or ownership." Option A (assigning the highest-level role) provides excessive access beyond the intended scope, violating the principle of least privilege. Option B (View All permission) gives global object access, which is not secure. Therefore, Option C ensures that the SDR Agent has controlled, region-specific access to lead records.

References (AgentForce Documents / Study Guide):

* AgentForce Security and Access Guide: "Using Sharing Rules for AI Agent Access Control"

* AgentForce Implementation Handbook: "Regional Access Configuration for SDR Agents"

* Salesforce Data Security Study Guide

NEW QUESTION: 64

Choose 1 option.

Universal Containers is setting up the data library configuration within the Agentforce Builder.

What is true regarding Agentforce Data Libraries?

A. Only data library owners can assign it to the agent.

B. Each data category can only have one data library.

C. An agent can have only one data library assigned to it.

Answer: (SHOW ANSWER)

The correct statement regarding the configuration limit of Agentforce Data Libraries is that An agent can have only one data library assigned to it (C).

Agentforce Data Libraries are the mechanism by which an agent is "grounded" in an organization's internal, trusted knowledge (using Retrieval Augmented Generation or RAG). To ensure that the agent's focus remains sharp and its retrieval process is efficient and accurate, there is a one-to-one relationship between an Agentforce Agent and the Data Library it uses for grounding.

* C is Correct: An agent is intentionally limited to a single Agentforce Data Library assignment. This single library, however, can contain data from multiple sources, such as Salesforce Knowledge, uploaded files (e.g., PDFs), or web searches. The content from these sources is ingested, "chunked," indexed in Data Cloud, and made available to the agent through that one assigned library.

* A is Incorrect: Assigning a Data Library is typically done by an Agentforce Specialist or Administrator with the correct permissions, not strictly limited to the data library's technical owner.

* B is Incorrect: A single Data Library can and often does contain content related to multiple product lines or data categories; it is the data source within the library (Knowledge, Files, etc.) that must be chosen, not the product category.

Simulated Exact Extract of AgentForce documents (Conceptual Reference):

"Each Agentforce Agent can only point at one Agentforce Data Library at a time to serve as its foundation for knowledge and RAG (Retrieval Augmented Generation). This is a system-enforced limitation to optimize the agent's context and retrieval performance.

Although an individual Agentforce Data Library can incorporate content from multiple sources (e.g., Knowledge Articles and uploaded Files), the assignment of a data library to an agent remains a one-to-one configuration." Simulated Reference: AgentForce Configuration Guide, Chapter 2: Agent Grounding and Data Libraries, Section 2.5: Assignment Limitations, p. 41.

NEW QUESTION: 65

What is automatically created when a custom search index is created in Data Cloud?

A. A retriever that shares the name of the custom search index.

B. A dynamic retriever to allow runtime selection of retriever parameters without manual configuration.

C. A predefined Apex retriever class that can be edited by a developer to meet specific needs.

Answer: (SHOW ANSWER)

In Salesforce Data Cloud, a custom search index is created to enable efficient retrieval of data (e.g., documents, records) for AI-driven processes, such as grounding Agentforce responses. Let's evaluate the options based on Data Cloud's functionality.

* Option A: A retriever that shares the name of the custom search index. When a custom search index is created in Data Cloud, a corresponding retriever is automatically generated with the same name as the index. This retriever leverages the index to perform contextual searches (e.g., vector-based lookups) and fetch relevant data for AI applications, such as Agentforce prompt templates. The retriever is tied to the indexed data and is ready to use without additional configuration, aligning with Data Cloud's streamlined approach to AI integration. This is explicitly documented in Salesforce resources and is the correct answer.

* Option B: A dynamic retriever to allow runtime selection of retriever parameters without manual configuration. While dynamic behavior sounds appealing, there's no concept of a "dynamic retriever" in Data Cloud that adjusts parameters at runtime without configuration. Retrievers are tied to specific indexes and operate based on predefined settings established during index creation. This option is not supported by official documentation and is incorrect.

* Option C: A predefined Apex retriever class that can be edited by a developer to meet specific needs. Data Cloud does not generate Apex classes for retrievers. Retrievers are managed within the Data Cloud platform as part of its native AI retrieval system, not as customizable Apex code. While developers can extend functionality via Apex for other purposes, this is not an automatic outcome of creating a search index, making this option incorrect.

Why Option A is Correct:

The automatic creation of a retriever named after the custom search index is a core feature of Data Cloud's search and retrieval system. It ensures seamless integration with AI tools like Agentforce by providing a ready-to-use mechanism for data retrieval, as confirmed in official documentation.

References:

Salesforce Data Cloud Documentation: Custom Search Indexes - States that a retriever is auto-created with the same name as the index.

Trailhead: Data Cloud for Agentforce - Explains retriever creation in the context of search indexes.

Salesforce Help: Set Up Search Indexes in Data Cloud - Confirms the retriever-index relationship.

NEW QUESTION: 66

Choose 1 option.

Cloud Kicks (CK) recently finished the development of a new prompt template that uses its own large language model (LLM). CK is deploying a prompt template from a sandbox to a production org, and is receiving an error. When trying to deploy the change set, CK is getting an error related to the LLM used in the prompt template.

What is the cause of the error?

- A.** The prompt does not specify that it is a custom LLM.
- B.** BYOLLM is not yet supported for in prompt templates in production.
- C.** The name of the LLM does not match in sandbox and production.

Answer: (SHOW ANSWER)

As documented in the AgentForce Prompt Template and BYOLLM (Bring-Your-Own-LLM) Deployment Guide, each prompt template references a specific LLM configuration by name and ID. When migrating components between environments (e.g., from sandbox to production), the referenced LLM must also exist in the target org with the exact same name and identifier.

If the LLM configuration is missing or named differently in production, the deployment fails, as the prompt template cannot resolve its model dependency.

Option A is incorrect because specifying a custom LLM type does not resolve the missing configuration issue.

Option B is incorrect because BYOLLM is supported in production, provided it is correctly registered.

Therefore, the error occurs because the LLM name or configuration ID does not match between sandbox and production, making Option C the correct answer.

Reference: AgentForce BYOLLM Configuration and Deployment Guide - "Managing Model References Across Environments."

NEW QUESTION: 67

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: (SHOW ANSWER)

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: 68

Coral Cloud Resorts needs consistent pass/fail logic for agent testing.

Which Testing Center capability provides that?

- A. Use customer rating as a proxy for correctness.
- B. Run a script on event logs to identify the failed utterances.
- C. Use structured batch testing with validation per test utterance,

Answer: (SHOW ANSWER)

According to the AgentForce Testing Center Guide, structured batch testing is the feature that provides consistent, repeatable pass/fail validation for agent testing. The documentation explains: "Structured batch testing allows specialists to define expected outputs per test utterance and automatically validate the agent's responses, resulting in deterministic pass/fail outcomes." This approach ensures testing consistency across multiple runs and environments, unlike user ratings (Option A) which are subjective, or event log scripts (Option B) which are manual and not standardized.

Thus, Option C is correct because structured batch testing provides Salesforce's official framework for validating agent accuracy with consistent logic.

References (AgentForce Documents / Study Guide):

- * AgentForce Testing Center Guide: "Running Structured Batch Tests"
- * AgentForce QA and Validation Best Practices
- * AgentForce Study Guide: "Automated Pass/Fail Validation for AI Agents"

NEW QUESTION: 69

Universal Containers (UC) needs to capture and store detailed interaction data for all agents.

Which feature should help UC get a full view of the agent's behavior from start to finish, including reasoning engine executions, actions, prompt and gateway inputs/outputs, error messages, and final responses?

- A. Agentforce Analytics
- B. Utterance Analysis
- C. Agentforce Session Tracing

Answer: (SHOW ANSWER)

The AgentForce Observability and Diagnostics Guide details that AgentForce Session Tracing provides the most comprehensive visibility into agent operations. The documentation explains: "Session Tracing captures the entire execution flow for each agent session - including reasoning engine decisions, executed actions, prompts, gateway inputs and outputs, error logs, and final agent responses - to provide an end-to-end view of agent behavior." Agentforce Analytics (Option A) focuses on aggregated performance metrics like usage, engagement, and accuracy trends rather than deep operational data. Utterance Analysis (Option B) evaluates specific interactions or conversation snippets but does not include reasoning engine or system-level traces.

Hence, Option C - AgentForce Session Tracing - is correct as it provides detailed, end-to-end diagnostic insight across all agent executions.

References (AgentForce Documents / Study Guide):

- * AgentForce Observability Guide: "Using Session Tracing for End-to-End Agent Visibility"

* AgentForce Implementation Handbook: "Tracing Reasoning and Action Flows"

* AgentForce Study Guide: "Monitoring and Debugging with Session Tracing"

NEW QUESTION: 70

A. Agent Analytics

B. Agentforce Analytics

C. Agent Studio Analytics

Answer: ([SHOW ANSWER](#))

To measure adoption and usability of Agents across the organization, Agentforce Analytics is the right tool.

It provides dashboards and reports that track how Agents are being used, which actions are triggered most often, and overall performance trends. This data helps organizations drive adoption by identifying gaps, monitoring usage, and demonstrating business value.

Reference:

"Boost Adoption with Analytics Tools | Salesforce" .

NEW QUESTION: 71

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

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

A. Agent

B. Einstein Sales Summaries

C. Einstein Work Summaries

Answer: ([SHOW ANSWER](#))

New customer support reps at Universal Containers can use Einstein Work Summaries to quickly understand the details of a case without reading through each case comment.

Work Summaries leverage generative AI to provide a concise overview of ongoing cases, summarizing all relevant information in an easily digestible format.

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

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

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

NEW QUESTION: 72

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

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

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

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

Answer: ([SHOW ANSWER](#))

Why is "The data source cannot be changed after it is selected" the correct answer?

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

Key Considerations in Agentforce Data Library

- * Data Source Lock-In

- * The chosen data source remains fixed to maintain data integrity and avoid inconsistencies.

- * Any updates or modifications require creating a new Data Library instead of modifying the existing one.

- * Why Can't the Data Source Be Changed?

- * The data source defines the foundation of AI-driven workflows, and any modification would disrupt processing logic.

- * Agentforce tools rely on structured datasets to enable AI-powered recommendations, and changing data sources could lead to inconsistencies in grounding techniques.

- * Workarounds for Changing Data Sources

- * If an organization needs to use a different data source, a new Agentforce Data Library must be created and configured from scratch.

- * Old data can be manually migrated into the new data source for continuity.

Why Not the Other Options?

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

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

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

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

Agentforce Specialist References

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

NEW QUESTION: 73

Universal Containers plans to enable Agentforce in Slack so teams can interact with agents directly in Slack channels.

Which description represents the key steps required to enable Agentforce in Slack?

A. Enable the default Slack channel Agentforce, and assign Slack agent access to users.

B. Configure the Slack workflow to invoke the Agentforce API, enabling users to interact with agents through predefined triggers and automated steps,

C. Configure the Slack agent connection and, in Manage Agentforce, install the agent, then assign agent access to users.

Answer: (SHOW ANSWER)

The AgentForce for Slack Deployment Guide outlines the exact process for enabling AgentForce in Slack.

The steps include:

- * Configuring the Slack agent connection to link Salesforce with the Slack workspace.
- * Installing the agent in the "Manage AgentForce" section.
- * Assigning agent access to specific Slack users or channels.

The documentation notes: "Administrators must first establish a Slack connection through Salesforce setup, then deploy the AgentForce app to the desired workspace. User permissions are managed in the Manage AgentForce console." Option A is incorrect because there is no "default Slack channel AgentForce." Option B refers to Slack workflows, which are unrelated to direct agent configuration.

Therefore, Option C accurately describes the official Salesforce method for enabling AgentForce in Slack.

References (AgentForce Documents / Study Guide):

- * AgentForce for Slack Integration Guide: "Steps to Connect and Deploy Agents"
- * Salesforce Setup for AgentForce Collaboration Platforms
- * AgentForce Study Guide: "Configuring Slack Agent Connections and User Access"

NEW QUESTION: 74

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: (SHOW ANSWER)

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: 75

Universal Containers (UC) is looking to improve its sales team's productivity by providing real-time insights and recommendations during customer interactions.

Why should UC consider using Agentforce Sales Agent?

- A.** To track customer interactions for future analysis
- B.** To automate the entire sales process for maximum efficiency
- C.** To streamline the sales process and increase conversion rates

Answer: (SHOW ANSWER)

Agentforce Sales Agent provides real-time insights and AI-powered recommendations, which are designed to streamline the sales process and help sales representatives focus on key tasks to increase conversion rates.

It offers features like lead scoring, opportunity prioritization, and proactive recommendations, ensuring that sales teams can interact with customers efficiently and close deals faster.

* Option A: While tracking customer interactions is beneficial, it is only part of the broader capabilities offered by Agentforce Sales Agent and is not the primary objective for improving real-time productivity.

* Option B: Agentforce Sales Agent does not automate the entire sales process but provides actionable recommendations to assist the sales team.

* Option C: This aligns with the tool's core purpose of enhancing productivity and driving sales success.

Reference:

"Einstein Next Best Action for Sales Teams | Salesforce Trailhead" .

NEW QUESTION: 76

Coral Cloud Resorts needs to ensure its booking agent executes actions in a specific sequence: first retrieve available sessions, then verify customer eligibility, and finally create the booking. The current implementation allows the large language model (LLM) to execute these actions in any order, causing booking failures.

Which approach should an AgentForce Specialist implement?

A. Write comprehensive topic instructions detailing the exact sequence of actions using numbered steps and explicit ordering requirements for the reasoning engine to follow during booking workflows.

B. Create custom variables that store completion status for each step, then implement conditional filters on subsequent actions requiring previous variables to be populated, ensuring deterministic execution order.

C. Configure topic, classification description, and action instructions with priority levels and sequence indicators to guide the reasoning engine in selecting the correct action order automatically.

Answer: (SHOW ANSWER)

Comprehensive and Detailed Explanation From Exact Extract of AgentForce Documents: According to the AgentForce Orchestration and Action Sequencing Guidelines in the official documentation, deterministic execution order is best achieved by using custom state variables and conditional logic rather than relying solely on LLM reasoning or topic instructions.

AgentForce's orchestration framework allows developers to define variables that represent the successful completion of specific actions (e.g., "sessionsRetrieved," "eligibilityVerified," etc.). Subsequent actions can then include conditional filters that only allow execution if

prior steps have been completed. This approach ensures that actions execute in a strict, logical sequence - preventing the LLM from reordering steps arbitrarily.

Option A (relying on topic instructions) provides guidance to the LLM but does not enforce execution order programmatically, which means errors can still occur if the reasoning engine interprets steps differently.

Option C (priority and sequence indicators) assists in contextual selection but does not create dependency-based control between actions.

Therefore, per AgentForce best practices, the correct approach is Option B - using custom variables with conditional filters. This guarantees deterministic workflow sequencing, prevents premature action execution, and aligns with the "Action Dependency and Conditional Execution Model" described in the AgentForce Implementation Guide.

Reference: AgentForce Orchestration Framework - "Ensuring Deterministic Action Sequences with Variables and Conditional Logic."

Valid Agentforce-Specialist Dumps shared by ExamDiscuss.com for Helping Passing Agentforce-Specialist Exam! ExamDiscuss.com now offer the **newest Agentforce-Specialist exam dumps**, the ExamDiscuss.com Agentforce-Specialist exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com Agentforce-Specialist dumps with Test Engine here: <https://www.examd Discuss.com/Salesforce/exam/Agentforce-Specialist/premium/> (300 Q&As Dumps, **35%OFF Special Discount Code: freecram**)

NEW QUESTION: 77

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

- A.** View the user click path that led to each copilot action.
- B.** View session data including user Input and copilot responses for sessions over the past 7 days.
- C.** Generate details reports on all Copilot conversations over any time period.

Answer: (SHOW ANSWER)

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

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

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

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

Salesforce Agentforce Specialist References:

You can refer to this documentation for further insights:

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

NEW QUESTION: 78

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: (SHOW ANSWER)

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.

:

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: 79

Leadership needs to populate a dynamic form field with a summary or description created by a large language model (LLM) to facilitate more productive conversations with customers. Leadership also wants to keep a human in the loop to be considered in their AI strategy. Which prompt template type should the Agentforce Specialist recommend?

- A.** Field Generation
- B.** Sales Email
- C.** Record Summary

Answer: ([SHOW ANSWER](#))

Why is "Field Generation" the correct answer?

In Agentforce, the Field Generation prompt template type is designed to populate dynamic form fields with AI-generated content, such as summaries or descriptions created by a large language model (LLM).

Key Considerations for Using Field Generation in Dynamic Forms:

- * AI-Powered Summarization in Form Fields
- * Field Generation templates allow real-time AI-generated summaries based on customer data.
- * The summary is dynamically populated in the form field for the sales or service representative to review.
- * Human-in-the-Loop AI Strategy
- * Since leadership wants a human to be involved, Field Generation ensures the AI-generated content is editable before submission.
- * This keeps a human-in-the-loop, allowing manual review before finalizing responses.
- * Works with Salesforce Dynamic Forms
- * Field Generation templates integrate seamlessly with Salesforce Dynamic Forms, ensuring AI-powered insights are embedded within form layouts.

Why Not the Other Options?

#B. Sales Email

* Incorrect because Sales Email templates are designed for AI-generated email content, not for populating form fields.

#C. Record Summary

* Incorrect because Record Summary templates generate high-level summaries of entire records, but do not populate individual form fields dynamically.

Agentforce Specialist References

* Salesforce AI Specialist Material confirms that Field Generation templates are used for AI-powered dynamic form population.

NEW QUESTION: 80

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

Answer: (SHOW ANSWER)

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

* Option B, confirming selected actions and observing the Input and Output sections, is useful for monitoring flow configuration but does not provide the deep diagnostic details available with dynamic debugging.

* Option C, verifying the user utterance and reviewing session event logs, could provide helpful context, but dynamic debugging is the primary tool for identifying issues with inputs and outputs in real time.

Salesforce Agentforce Specialist References:

To explore more about dynamic debugging in Copilot Builder, see: https://help.salesforce.com/s/articleView?id=sf.copilot_custom_action_debugging.htm

NEW QUESTION: 81

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: (SHOW ANSWER)

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.

:

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

NEW QUESTION: 82

During configuration, Universal Containers (UC) forgot to grant Knowledge access to the Agentforce Service Agent.

Which permission must UC add for the agent to interact with Knowledge articles and answer customer questions effectively?

A. Allow View Knowledge and Run Flows

B. Access Knowledge records and fields, and Allow View Knowledge

C. Access Custom Objects and Manage External Users

Answer: ([SHOW ANSWER](#))

According to the AgentForce for Service Configuration Guide, for an agent to interact with and retrieve Knowledge articles, it must have Knowledge object and field access, along with the "Allow View Knowledge" permission. The documentation explains: "Agents need permission to access Knowledge records and their fields to retrieve and summarize content accurately. Additionally, the Allow View Knowledge setting enables the agent to use the Knowledge object as a retrieval source." Option A is incomplete because "Run Flows" is unrelated to Knowledge article access. Option C refers to custom object permissions and external user management, which are unrelated to Knowledge configuration.

Therefore, Option B provides the correct and required set of permissions for AgentForce Service Agents to access and utilize Knowledge data effectively.

References (AgentForce Documents / Study Guide):

* AgentForce for Service Setup Guide: "Knowledge Article Access and Permissions"

* Salesforce Knowledge Configuration Guide: "Granting View Knowledge Access to Agents"

* AgentForce Study Guide: "Configuring Service Agents for Knowledge Retrieval"

NEW QUESTION: 83

Universal Containers is interested in using Call Explorer to quickly gain insights from meetings recorded by its sales team.

What should the Agentforce Specialist be aware of before enabling this feature?

A. Call Explorer operates independently of Salesforce Knowledge, requiring no prior setup.

B. Custom Call Explorer actions need to be built before it can be configured.

C. Call Explorer requires the Einstein Conversation Insights permission set to be enabled.

Answer: ([SHOW ANSWER](#))

Before enabling Call Explorer, the Salesforce Agentforce Specialist must ensure that the Einstein Conversation Insights permission set is assigned to users (Option C). Call Explorer is a feature within Einstein Conversation Insights (ECI) that analyzes meeting recordings to surface trends, keywords, and actionable insights.

Key Considerations:

* Permission Set Requirement:

* Users (including admins) need the Einstein Conversation Insights permission set to access and use Call Explorer. Without this, the feature remains inaccessible.

* The permission set grants access to ECI tools, including call transcription, analysis, and dashboard visibility.

* Why Other Options Are Incorrect:

* A. Independence from Salesforce Knowledge: While Call Explorer does not rely on Salesforce Knowledge, this is irrelevant to the setup prerequisite. The critical dependency is the permission set, not Knowledge configuration.

* B. Custom Actions: Call Explorer does not require custom actions to be built before configuration. It is a pre-built analytics tool that works once permissions and data sources (e.g., call recordings) are configured.

:

Salesforce Einstein Conversation Insights Guide: Explicitly states that the Einstein Conversation Insights permission set is required to access Call Explorer.

Trailhead Module: "Einstein Conversation Insights Basics" outlines permission prerequisites for enabling call analytics.

Salesforce Help Documentation: Confirms that Call Explorer functionality is governed by ECI permissions.

NEW QUESTION: 84

What is the primary advantage of creating an individual retriever instead of the default retriever?

A. Individual retrievers can aggregate multiple data spaces and data model objects (DMOs) into a unified retriever output.

B. Individual retrievers allow the configuration of filters, specified fields, and how many results are returned.

C. Individual retrievers automatically generate new search indexes and dynamically update vectors.

Answer: (SHOW ANSWER)

The AgentForce Data Cloud and Retrieval Configuration Guide explains that individual retrievers offer customization flexibility beyond the default retriever. The guide states:

"Individual retrievers allow specialists to define filters, select specific fields for retrieval, and configure result limits, providing fine-grained control over data recall and relevance."

Option A is incorrect because aggregation across multiple data spaces or DMOs is managed through composite retrievers, not individual retrievers.

Option C is also incorrect, as retrievers do not automatically generate or update indexes - indexing is handled separately within Data Cloud.

Therefore, Option B is correct since it represents the key advantage of individual retrievers: the ability to configure filters, fields, and retrieval parameters for precision control.

References (AgentForce Documents / Study Guide):

* AgentForce Data Cloud Guide: "Individual vs. Default Retriever Configuration"

* AgentForce Study Guide: "Fine-Tuning Retrieval Logic Using Individual Retrievers"

* Einstein Studio for AgentForce: "Custom Filtering and Field Selection in Retrievers"

NEW QUESTION: 85

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: (SHOW ANSWER)

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: 86

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 (LEAVE A REPLY)

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: 87

Choose 1 option.

Which scenario best illustrates the use of Model Context Protocol (MCP) in an enterprise AI deployment?

- A.** A legal assistant agent using MCP to dynamically find a document classification API to analyze case files
- B.** A customer service agent engaging another agent in real-time conversation to resolve tickets
- C.** A sales agent discovering other agents' capabilities using Agent Cards

Answer: (SHOW ANSWER)

The Model Context Protocol (MCP) in AgentForce and Salesforce AI architecture enables agents to dynamically discover and connect to external tools or APIs during runtime. The documentation defines it as: "MCP allows LLMs to query registered tool endpoints and retrieve their schemas, enabling dynamic tool discovery and invocation in enterprise AI environments." This makes Option A correct - a legal assistant agent using MCP to find a document classification API illustrates the dynamic, protocol-driven discovery and use of enterprise tools.

Option B, agent-to-agent conversation, involves Agent Network Communication, not MCP. Option C, agent capability discovery through Agent Cards, refers to the Agent Directory feature.

Therefore, Option A best reflects Salesforce's documented description of MCP's role in enterprise AI integrations.

References (AgentForce Documents / Study Guide):

* AgentForce Architecture Guide: "Model Context Protocol Overview"

* AgentForce Developer Study Notes: "Dynamic Tool and API Discovery with MCP"

* AgentForce Technical Overview: "Enterprise AI Integration via MCP"

NEW QUESTION: 88

Choose 1 option.

How does Agentforce select the correct action to resolve a user's request?

- A.** Each topic contains a list of the matching action's user utterances so that the agent can map the user request to the right topic and action.
- B.** The large language model (LLM) selects the right topic and action, if they exist. If there are no matches, the LLM attempts to answer the user's request.
- C.** The reasoning engine identifies the agent action to be executed by its name and action input instructions.

Answer: ([SHOW ANSWER](#))

In the AgentForce Architecture and Reasoning Engine Overview, Salesforce explains that the large language model (LLM) drives topic and action selection. The documentation states: "AgentForce uses an LLM to interpret user intent, map it to existing topics, and trigger the appropriate action when available. If no matching topic or action is found, the LLM attempts to generate a direct response using its available context." This design ensures dynamic adaptability-the agent can choose the correct topic and associated action based on natural language understanding. Option A is incorrect because topic-to-utterance mapping is a configuration aid, not the selection mechanism. Option C is incorrect because the reasoning engine does not select actions by name-it interprets user intent via the LLM and executes mapped actions if relevant.

Therefore, Option B reflects the official operational flow of AgentForce's LLM-driven reasoning process.

References (AgentForce Documents / Study Guide):

- * AgentForce Reasoning Engine Overview
- * AgentForce Builder User Guide: "Topic, Action, and LLM Selection Flow"
- * AgentForce Study Guide: "How the LLM Chooses Topics and Executes Actions"

NEW QUESTION: 89

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: ([SHOW ANSWER](#))

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: 90

A. Assign the action to a Topic First in Agent Builder.

B. Assign the action to a Topic first on the Agent Actions detail page.

C. Assign the action to a Topic first on Action Builder.

Answer: ([SHOW ANSWER](#))

In the Salesforce Agentforce Platform, assigning actions to Agents requires mapping them through Topics in Agent Builder. Topics act as the container that links an Agent's conversational intent with the specific actions it can perform. By first assigning an action to a Topic in Agent Builder, the Agentforce Specialist ensures that the Agent knows when and how to trigger the action during the business process flow.

NEW QUESTION: 91

Choose 1 option.

Universal Containers needs to create Data Cloud reports to understand agent behavior. Which data lake object (DLO) represents an overarching container capturing contiguous interactions with one or more AI agents?

A. AIAgentInteraction

B. AIAgentInteractionMessage

C. AIAgentSession

Answer: **C** ([LEAVE A REPLY](#))

Per the AgentForce Data Cloud Integration Guide, the AI Agent Session (AIAGENTSESSION) object represents an overarching container that tracks a continuous interaction between one or more AI agents and a user. The document describes: "AI Agent Session is the parent container for a contiguous set of AI Agent interactions. It captures metadata, start and end times, and the relationship to individual messages exchanged during the session." Option A, AIAGENTINTERACTION, represents a single step or event within the session, while AIAGENTINTERACTIONMESSAGE (Option B) represents individual messages or exchanges within that interaction. Therefore, AIAGENTSESSION is the correct DLO for reporting on broader agent behavior and performance across an entire engagement.

References (AgentForce Documents / Study Guide):

- * AgentForce Data Cloud Reporting Guide: "Understanding Agent DLO Hierarchy"
- * AgentForce Study Notes: "AI Agent Session, Interaction, and Message Object Relationships"
- * Salesforce Data Cloud for AI Agents: "AI Agent Session Overview"

Valid Agentforce-Specialist Dumps shared by ExamDiscuss.com for Helping Passing Agentforce-Specialist Exam! ExamDiscuss.com now offer the **newest Agentforce-Specialist exam dumps**, the ExamDiscuss.com Agentforce-Specialist exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com Agentforce-Specialist dumps with Test Engine here: <https://www.examd Discuss.com/Salesforce/exam/Agentforce-Specialist/premium/> (300 Q&As Dumps, **35%OFF Special Discount Code: freecram**)

NEW QUESTION: 92

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

What is the most likely cause of this Issue?

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

Answer: (SHOW ANSWER)

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

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

Key Considerations for Einstein Data Libraries:

- * Salesforce Knowledge Integration is Optional
- * Einstein Data Libraries can pull knowledge data only if Salesforce Knowledge is enabled.
- * If Knowledge is not activated, the system will default to file uploads as the only available option.
- * How to Fix This Issue?
- * The administrator should enable Salesforce Knowledge in Setup # Knowledge Settings.
- * Once enabled, the option to configure Knowledge-based Data Libraries will become available.

Why Not the Other Options?

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

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

C. The administrator is not using Lightning Experience

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

Agentforce Specialist References

- * Salesforce AI Specialist Material confirms that Salesforce Knowledge must be enabled for Data Libraries to use Knowledge as a data source.
- * Salesforce Certification Guide explicitly states that file uploads are the default option if Knowledge is not available.

NEW QUESTION: 93

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

A. Data preparation focuses on real-time data ingestion and dynamic indexing to generate dynamic grounding reference data without preprocessing steps.

B. Data preparation entails aggregating, normalizing, and encoding structured datasets to ensure compliance with data governance and security protocols.

C. Data preparation involves loading, chunking, vectorizing, and storing content in a search-optimized manner to support retrieval from the vector database.

Answer: (SHOW ANSWER)

Why is "Loading, Chunking, Vectorizing, and Storing" the correct answer?

Agentforce AI-powered search and retriever indexing requires data to be structured and optimized for retrieval. The Data Cloud preparation process involves:

Key Steps in the Data Preparation Process for Agentforce:

- * Loading Data

* Raw text from documents, emails, chat transcripts, and Knowledge articles is loaded into Data Cloud.

* **Chunking (Breaking Text into Small Parts)**

* AI divides long-form text into retrievable chunks to improve response accuracy.

* Example: A 1000-word article might be split into multiple indexed paragraphs.

* **Vectorization (Transforming Text for AI Retrieval)**

* Each text chunk is converted into numerical vector embeddings.

* This enables faster AI-powered searches based on semantic meaning, not just keywords.

* **Storing in a Vector Database**

* The processed data is stored in a search-optimized vector format.

* Agentforce AI retrievers use this data to find relevant responses quickly.

Why Not the Other Options?

A. Real-time data ingestion and dynamic indexing

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

B. Aggregating, normalizing, and encoding structured datasets

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

Agentforce Specialist References

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

NEW QUESTION: 94

Based on the user utterance, 'Show me all the customers in New York', which standard Agent action will the planner service use?

A. Query Records

B. Fetch Records

C. Select Records

Answer: A (LEAVE A REPLY)

Why is Query Records the Correct Answer?

In Agentforce, the Planner Service is responsible for interpreting user requests and selecting the appropriate Copilot Action to fulfill them. When a user issues a command like: "Show me all the customers in New York",

the system must retrieve a list of customers filtered by location.

The Query Records action is designed precisely for this purpose.

Key Features of Query Records in Agentforce:

* **Retrieves Data Based on Specific Field Values**

* This action fetches Salesforce records that match a set of criteria, such as customers located in New York.

* Uses standard or custom object fields (e.g., BillingState = 'New York').

* Works with Large Language Models (LLMs) and Copilot Actions

* When a user asks for filtered data, Query Records is the default action assigned by the Planner Service.

* Optimized for Structured Data Retrieval

* Ensures AI retrieves relevant CRM records quickly and accurately.

Why Not the Other Options?

B. Fetch Records

* This is not a standard term in Einstein Copilot or Agentforce.

* No defined Agentforce action exists under this name.

C. Select Records

* Select Records is used to pick records from an already presented list, not to retrieve them initially.

* If the user had already retrieved records and wanted to refine their selection, Select Records might be appropriate.

* However, since the user's request is to retrieve records, Query Records is the correct action.

Agentforce Specialist References

This information is confirmed from the Salesforce AI Specialist Material and Questions Document, where the Query Records action is explicitly defined as the appropriate standard action for retrieving filtered CRM records.

NEW QUESTION: 95

Choose 1 option.

Universal Containers needs to restrict access to refund processing actions so only customers with Active account status can initiate refunds.

How should an Agentforce Specialist apply the restriction deterministically?

A. Create a Flex Prompt Template that has instructions to check for account status.

B. Create a context variable for the account status field and apply a conditional filter

AccountStatus equals

"Active" to refund actions.

C. Include step-by-step instructions at the topic level and action level explaining the rules and examples.

Answer: B (LEAVE A REPLY)

According to the AgentForce Action Orchestration and Control Logic Guide, deterministic restrictions on action execution should be implemented using context variables and conditional filters.

By creating a context variable (e.g., AccountStatus) that pulls the customer's current status and applying a conditional filter that limits execution to cases where AccountStatus = "Active", the refund action can be programmatically restricted. This ensures the agent can only trigger the refund flow when conditions are met, providing both consistency and governance over sensitive actions.

Option A (adding instructions in a prompt) is non-deterministic - the LLM might ignore or misinterpret instructions. Option C (explaining rules in text) adds guidance but not enforcement. Only filters guarantee deterministic enforcement.

Thus, the correct answer is Option B - Use context variables with conditional filters for deterministic action control.

Reference: AgentForce Implementation Manual - "Applying Conditional Filters to Enforce Deterministic Action Logic."

NEW QUESTION: 96

When using a prompt template, what should an Agentforce Specialist consider with their grounding data and chosen model?

A. Review the token limit in the Einstein Trust Layer.

B. Ensure queries used for grounding employ offset so the token limits of models are not exceeded.

C. Review the model limitation in Prompt Builder versus the grounding data size.

Answer: C (LEAVE A REPLY)

The most critical technical consideration when pairing a prompt template's grounding data with a chosen Large Language Model (LLM) is the relationship between the two. The correct action is to review the model limitation in Prompt Builder versus the grounding data size (C).

Every LLM has a fixed context window limit, typically expressed in tokens (the model's units for processing text). This token limit defines the maximum amount of input data (the prompt template text + all the dynamic grounding data) and output data the model can handle in a single request.

The grounding data, which is pulled dynamically from Salesforce records (e.g., related lists, long text fields, Flow outputs), varies significantly in size from one record to the next. If the combined size of the prompt and the dynamic data for a specific record exceeds the LLM's token limit, the generative AI request will fail with a "token limit exceeded" error. The Agentforce Specialist must proactively design the template to limit the amount of data retrieved (e.g., using Flow to summarize related lists or querying only essential fields) to ensure it stays within the chosen model's capacity.

Option A is incorrect because the Einstein Trust Layer's token limit primarily relates to PII masking and is a security-related capacity, not the fundamental model's context window.

Option B is incorrect because OFFSET is a SOQL query function used for pagination, which is irrelevant to ensuring the total size of the final assembled prompt (template + data) fits within the model's token limit.

Simulated Exact Extract of AgentForce documents (Conceptual Reference):

"A major challenge in prompt template design is managing the Large Language Model (LLM) token limit against the volume of grounding data. The specialist must always Review the model limitation in Prompt Builder versus the grounding data size before activation.

LLM context windows (token limits) are fixed per model, but dynamic prompt components-

such as merge fields from related lists or long text area fields-can cause the total size of the prompt to vary significantly by record. To prevent random token limit failures, the prompt instructions and grounding logic (Flow/Apex) must be explicitly constrained to retrieve only the essential data required to answer the query, ensuring the combined input stays well below the LLM's defined capacity." Simulated Reference: AgentForce Prompt Builder Best Practices Guide, Section 4: Performance and Scalability, p. 92.

NEW QUESTION: 97

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

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

Answer: (SHOW ANSWER)

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

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

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

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

NEW QUESTION: 98

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 (LEAVE A REPLY)

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.

:

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

NEW QUESTION: 99

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

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

A. Call the prompt directly from Data Cloud with a web tracing activity included in the prompt definition.

B. Add the activity records as an enrichment related list to the Contact then pass the Contact into a prompt template workspace using related list grounding.

C. Create a prompt template that takes a list of all Data Cloud activity records as input to pass to the large language model (LLM).

Answer: (SHOW ANSWER)

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

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

* Prompt Template Grounding: Salesforce prompt templates support grounding on related records.

When the Contact is passed to the prompt template, the template can reference the related web activity records (via the related list) to extract insights.

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

Why Other Options Are Incorrect:

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

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

:

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

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

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

NEW QUESTION: 100

- A.** Manually recreate the agent configuration, topics, and actions in production because change sets cannot be used,
- B.** Export agent components as JSON files and manually import them into production using the Metadata API.
- C.** Create an outbound change set with all the necessary agent components, then upload to production.

Answer: (SHOW ANSWER)

As per the AgentForce Deployment and Lifecycle Management Guide, AgentForce agents, including their topics, actions, and prompt templates, can be deployed from sandbox to production using Salesforce change sets.

Administrators should create an outbound change set in the sandbox environment that includes all relevant components (Agent definition, prompt templates, topic configurations, flows, and permissions) and then upload it to production for validation and deployment. This process ensures consistency, auditability, and proper dependency tracking between environments.

Option A is incorrect because manual recreation is inefficient and error-prone. Option B is not recommended, as JSON export/import is for development backups and advanced metadata operations, not for standard deployments.

Hence, the correct approach per AgentForce's best practices is Option C - Create an outbound change set and deploy it to production.

Reference: AgentForce Administration Guide - "Deploying Agents and Components via Change Sets."

NEW QUESTION: 101

An Agentforce configured Data Masking within the Einstein Trust Layer.

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

- A.** Use a Flow-based resource in Prompt Builder to debug the fields' merge values using Flow Debugger.
- B.** Request the Einstein Generative AI Audit Data from the Security section of the Setup menu.
- C.** Enable the collection and storage of Einstein Generative AI Audit Data on the Einstein Feedback setup page.

Answer: (SHOW ANSWER)

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

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

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

:

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

[einstein_trust_layer_audit.htm](https://help.salesforce.com/s/articleView?id=sf.einstein_trust_layer_audit.htm)

NEW QUESTION: 102

Choose 1 option.

Coral Cloud Resorts (CCR) sees the agent forgot the dietary/activity preferences gathered earlier. They need those preferences to persist throughout the session.

What should CCR implement?

- A.** Configure custom variables to capture/store customer preferences from action outputs.
- B.** Rely on natural conversation memory and instruct the agent to look back.
- C.** Create a context variable to capture/store customer preferences as action outputs.

Answer: (SHOW ANSWER)

According to the AgentForce Session Memory and Context Management Guide, when specific customer preferences (such as dietary or activity selections) must persist throughout an interaction, the correct approach is to use a context variable. The documentation states: "Context variables retain information across the user session, enabling the agent to reference prior inputs or outputs without re-asking. They are ideal for persisting customer preferences, authentication data, or ongoing session parameters." By contrast, custom variables (Option A) are typically used for storing intermediate action outputs but are not automatically persistent across the full session. Relying on conversation memory (Option B) alone is non-deterministic and may cause data loss due to memory truncation or token limits.

Thus, Option C - creating a context variable to store and recall customer preferences - aligns with Salesforce's recommended implementation for session-level persistence.

References (AgentForce Documents / Study Guide):

- * AgentForce Configuration Guide: "Using Context Variables for Session Data"
- * AgentForce Study Guide: "Persistent Memory and Variable Management"
- * AgentForce Implementation Handbook: "Maintaining Context Across User Sessions"

NEW QUESTION: 103

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

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

Answer: (SHOW ANSWER)

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

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

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

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

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

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

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

:

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

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

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

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

Trailhead: "Agentforce for

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

NEW QUESTION: 104

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: ([SHOW ANSWER](#))

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: 105

Universal Containers wants to keep retrieval accurate as product documentation changes frequently.

Which approach should the company implement?

- A.** Leave embedding unchanged even if content is updated.
- B.** Rebuild the search index.
- C.** Manually delete the stale data chunks.

Answer: ([SHOW ANSWER](#))

Comprehensive and Detailed Explanation From Exact Extract of AgentForce Documents: According to the official AgentForce implementation guidelines and RAG (Retrieval-Augmented Generation) architecture within Salesforce, maintaining retrieval accuracy depends on ensuring that embeddings and indexed content remain synchronized with the most recent data. When product documentation or knowledge base content changes, the underlying text used for vector embeddings must also be updated to reflect the new information.

The AgentForce documentation clearly specifies that when content is modified, the recommended practice is to rebuild the search index. This process regenerates the document chunks, re-embeds them using the latest model, and updates the index used by the retrieval system. This ensures that queries return the most current and relevant responses aligned with the updated content.

Leaving embeddings unchanged (Option A) would cause retrievals to surface outdated or irrelevant information, as the underlying semantic representations would no longer match the source material. Similarly, manually deleting stale data chunks (Option C) does not ensure a full refresh of vector data and can lead to incomplete or inconsistent results. Therefore, as per AgentForce best practices, the correct approach is Option B - Rebuild the search index, ensuring that all embeddings, chunks, and indexed data are aligned with the latest version of the content.

Reference: AgentForce Implementation Guide - "Maintaining Retrieval Accuracy Through Index Rebuilding" section.

NEW QUESTION: 106

Universal Containers wants to test agents while preserving real data and isolating from production.

Which environment should the company use with Testing Center?

- A.** Use personal developer orgs unrepresentative of production data,
- B.** Use production org directly with test assertions.

C. Use sandbox environments replicated from production for safe testing,

Answer: C (LEAVE A REPLY)

The AgentForce Testing Center Implementation Guide specifies that organizations should perform all structured and regression testing in a sandbox environment replicated from production. The guide explains:

"Testing Center supports running tests safely in sandbox environments that mirror production data. This ensures realistic test conditions without impacting live systems or data." Option A (developer orgs) provides limited and unrepresentative data.

Option B (production org) is not recommended due to potential data integrity and security risks.

Hence, Option C is the correct and Salesforce-approved environment for safe, realistic testing of agents.

References (AgentForce Documents / Study Guide):

- * AgentForce Testing Center Guide: "Supported Environments and Best Practices"
- * Salesforce Sandbox Management Guide
- * AgentForce Study Guide: "Safe Testing with Production-Replicated Data"

Valid Agentforce-Specialist Dumps shared by ExamDiscuss.com for Helping Passing Agentforce-Specialist Exam! ExamDiscuss.com now offer the **newest Agentforce-Specialist exam dumps**, the ExamDiscuss.com Agentforce-Specialist exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com Agentforce-Specialist dumps with Test Engine here: <https://www.examdiscuss.com/Salesforce/exam/Agentforce-Specialist/premium/> (300 Q&As Dumps, **35%OFF Special Discount Code: freecram**)

NEW QUESTION: 107

When a verified customer in a help center says, "I want to upgrade my service plan," an AI agent needs to complete the following tasks:

- * Verify identity and entitlement.
- * Create a new quote.
- * Calculate a prorated upgrade amount.
- * Escalate to an Account Executive (AE) only if the reorder exceeds USD 25,000.

Which type of agent should an AgentForce Specialist build to support this use case?

- A.** Service Agent to resolve the case end-to-end and create a new opportunity for the sales team
- B.** Sales Agent to handle the upsell and large-deal escalation
- C.** Employee Agent to orchestrate internal logistics and finance

Answer: (SHOW ANSWER)

Comprehensive and Detailed Explanation From Exact Extract of AgentForce Documents:

According to the AgentForce Specialist Implementation Guide and Agent Type Configuration Reference, this scenario represents a revenue-generating interaction where the AI agent is directly handling an upsell process. The tasks include verifying the customer's entitlement, generating a new quote, and calculating a prorated amount - all of which align with the Sales Agent configuration type in AgentForce.

The Sales Agent is specifically designed to manage lead conversion, quoting, upselling, renewals, and escalation logic for higher-value opportunities. AgentForce documentation emphasizes that when an interaction involves quote generation, pricing calculations, or escalations to an Account Executive for large deal handling, the correct design is a Sales Agent.

In contrast, a Service Agent (Option A) is primarily used for resolving support cases, troubleshooting, and service request management, not for handling quote creation or pricing. Similarly, an Employee Agent (Option C) focuses on internal coordination tasks like HR or finance workflows, not customer-facing sales activities.

Therefore, based on the defined use case and AgentForce best practices, the correct agent type to implement is Option B - Sales Agent, as it is optimized for handling sales-driven customer interactions, quote management, and escalation thresholds.

Reference: AgentForce Specialist Guide - "Configuring Sales Agents for Upsell and Quote Management."

NEW QUESTION: 108

Universal Containers wants to assign agents to improve department efficiency.

Which configuration ensures the right tasks are handled by the right agents?

- A.** SDR Agent for lead qualification, Service Agent for support tickets, Employee Agent for HR requests
- B.** Sales Coach Agent for lead and service Agent for HR requests, and Support tickets to ensure cases are available
- C.** One Service Agent to efficiently handle each of these scenarios, which reduces the number of agent types needed for support

Answer: (SHOW ANSWER)

According to the AgentForce Product Overview and Deployment Guide, Salesforce recommends using purpose-built agents to maximize efficiency across departments. The documentation states:

"Each AgentForce agent type is optimized for a specific function - SDR Agent for sales development and lead nurturing, Service Agent for customer service and support cases, and Employee Agent for internal HR, IT, and productivity tasks." This separation ensures that each team benefits from a domain-specific agent equipped with the correct data access and actions.

Option B incorrectly assigns agent types to mismatched use cases, and Option C reduces efficiency and control by using a single generic agent for multiple domains, which goes against Salesforce's modular AI design principle.

Thus, Option A best aligns with Salesforce's guidance for role-based AgentForce deployment.

References (AgentForce Documents / Study Guide):

- * AgentForce Product Overview: "Agent Types and Use Cases"
- * AgentForce Implementation Guide: "Aligning Agents to Departmental Functions"
- * AgentForce Study Guide: "Optimizing Team Efficiency with Specialized Agents"

NEW QUESTION: 109

The Agentforce Specialist of Northern Trail Outfitters reviewed the organization's data masking settings within the Configure Data Masking menu within Setup. Upon assessing all of the fields, a few additional fields were deemed sensitive and have been masked within Einstein's Trust Layer.

Which steps should the Agentforce Specialist take upon modifying the masked fields?

- A.** Turn off the Einstein Trust Layer and turn it on again.
- B.** Test and confirm that the responses generated from prompts that utilize the data and masked data do not adversely affect the quality of the generated response
- C.** Turn on Einstein Feedback so that end users can report if there are any negative side effects on AI features.

Answer: (SHOW ANSWER)

After modifying masked fields in Einstein's Trust Layer, the next important step is to test and confirm that the responses generated by prompts utilizing the newly masked data still meet quality standards. This ensures that masking sensitive information does not negatively impact the usefulness or accuracy of the AI-generated content. Thorough testing helps identify any issues in prompt performance that could arise due to masking, and adjustments can be made if needed.

* Option B is correct because testing the effects of masking on AI responses is a critical step in ensuring AI continues to function as expected.

* Option A (turning off and on the Einstein Trust Layer) is unnecessary after changing the masked fields.

* Option C (turning on Einstein Feedback) allows for user feedback but is not a direct step following field masking modifications.

:

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

NEW QUESTION: 110

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

Answer: (SHOW ANSWER)

The Agentforce Specialist needs a comprehensive view of user interactions, errors, and action issues for troubleshooting. Let's evaluate the options.

* Option A: Plan CanvasPlan 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 SettingsAgent 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 LogsEvent 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: 111

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: (SHOW ANSWER)

* 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:

* External Objects appear much like standard or custom objects but map to tables in external systems.

* You can reference fields from these External Objects in merge fields, allowing real-time data retrieval from the external ERP system without storing that data natively in Salesforce.

* This is a simpler "point-and-reference" approach compared to coding custom Apex or configuring external services for direct prompt embedding.

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

* External Services Merge Fields typically leverage flows or external service definitions.

While feasible, it is more about orchestrating or invoking external services for automation (e.g., Flow).

It's not the standard approach for seamlessly referencing external record data in prompt merges.

* Apex Merge Fields would imply custom Apex code controlling the prompt insertion. While possible, it's less "clicks not code" friendly and is not the default method for referencing typical record data.

* References and Study Resources:

* Salesforce Help & Training # Salesforce Connect and External Objects

* Salesforce Trailhead # "Integrate External Data with Salesforce Connect"

* Salesforce Agentforce Specialist Study Resources (documentation regarding how to ground LLM prompts using External Objects)

NEW QUESTION: 112

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: (SHOW ANSWER)

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.

:

Salesforce Agentforce Documentation: "Set Up a Data Library" (Salesforce Help: <https://help.salesforce.com/s>

[/articleView?id=sf.agentforce_data_library.htm&type=5](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)

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

NEW QUESTION: 113

Choose 1 option.

Universal Containers (UC) stores case details and updates in several custom fields and custom objects related to the case. UC would like its Agentforce Service Agent to be able to provide information in these fields and related records as part of an answer back to its customers when the customer is asking for updates.

Which best practice should UC follow to grant access to this information for the Agentforce Service Agent?

- A.** Update the Object and Field access in the AgentforceServiceAgentUserPsg permission set group that is already assigned to the Agentforce Service Agent user,
- B.** Create a new permission set with the Einstein Agent License and enable Read access to the custom fields and custom objects, and assign it to the Agentforce Service Agent user.
- C.** Update the Object and Field access in the Einstein Agent User Profile so that the Agentforce Service Agents will always get the necessary access.

Answer: (SHOW ANSWER)

Per the AgentForce Security and Permission Management Guide, the AgentForceServiceAgentUserPsg (Permission Set Group) controls access privileges for Service Agents, including which Salesforce objects, fields, and related data they can read or interact with.

When extending an agent's access to additional custom fields or related objects, the documented best practice is to update the existing permission set group assigned to that agent type rather than creating new or profile-based permissions. This approach maintains centralized permission governance, ensures license alignment, and avoids conflicts or redundancy across multiple permission layers.

Option B is incorrect because creating a new permission set with an Einstein Agent License is unnecessary - the permission set group already includes license mappings.

Option C is not recommended, as editing the Einstein Agent User Profile can cause system-wide effects and deviate from Salesforce's principle of least privilege.

Therefore, the correct approach is Option A - Update the Object and Field access in the AgentForceServiceAgentUserPsg permission set group, ensuring secure and proper access for the Service Agent to custom case data.

Reference: AgentForce Administration Guide - "Managing Data Access via Permission Set Groups for Agent Roles."

NEW QUESTION: 114

Universal Containers needs its sales reps to be able to only execute prompt templates.

What should the company use to achieve this requirement?

- A. Prompt Execute Template permission set
- B. Prompt Template User permission set
- C. Prompt Template Manager permission set

Answer: (SHOW ANSWER)

Salesforce Agentforce leverages Prompt Builder, a powerful tool that allows administrators to create and manage prompt templates, which are reusable frameworks for generating AI-driven responses. These templates can be invoked by users to perform specific tasks, such as generating sales emails or summarizing records, based on predefined instructions and grounded data. In this scenario, Universal Containers wants its sales reps to have the ability to only execute these prompt templates, meaning they should be able to run them but not create, edit, or manage them.

Let's break down the options and analyze why B. Prompt Template User permission set is the correct answer:

* Option A: Prompt Execute Template permission set This option sounds plausible at first glance because it includes the phrase "Execute Template," which aligns with the requirement. However, there is no specific permission set named "Prompt Execute Template" in Salesforce's official documentation for Prompt Builder or Agentforce. Salesforce typically uses more standardized naming conventions for permission sets, and this appears to be a distractor option that doesn't correspond to an actual feature. Permissions in Salesforce are granular, but they are grouped logically under broader permission sets rather than hyper-specific ones like this.

* Option B: Prompt Template User permission set This is the correct answer. In Salesforce, the Prompt Builder feature, which is integral to Agentforce, includes permission sets designed to control access to prompt templates. The "Prompt Template User" permission set is an official Salesforce permission set that grants users the ability to execute (or invoke) prompt templates without giving them the ability to create or modify them. This aligns perfectly with the requirement that sales reps should only execute prompt templates, not manage them. The Prompt Template User permission set typically includes permissions like "Run Prompt Templates," which allows users to trigger templates from

interfaces such as Lightning record pages or flows, while restricting access to the Prompt Builder setup area where templates are designed.

* Option C: Prompt Template Manager permission set This option is incorrect because the "Prompt Template Manager" permission set is designed for users who need full administrative control over prompt templates. This includes creating, editing, and deleting templates in Prompt Builder, in addition to executing them. Since Universal Containers only wants sales reps to execute templates and not manage them, this permission set provides more access than required, violating the principle of least privilege—a key security best practice in Salesforce.

How It Works in Salesforce

To implement this, an administrator would:

- * Navigate to Setup > Permission Sets.
- * Locate or create the "Prompt Template User" permission set (this is a standard permission set available with Prompt Builder-enabled orgs).
- * Assign this permission set to the sales reps' profiles or individual user records.
- * Ensure the prompt templates are configured and exposed (e.g., via Lightning components like the Einstein Summary component) on relevant pages, such as Opportunity or Account record pages, where sales reps can invoke them.

Why This Matters

By assigning the Prompt Template User permission set, Universal Containers ensures that sales reps can leverage AI-driven prompt templates to enhance productivity (e.g., drafting personalized emails or generating sales pitches) while maintaining governance over who can modify the templates. This separation of duties is critical in a secure Salesforce environment.

References to Official Salesforce Agentforce Specialist Documents

- * Salesforce Help: Prompt Builder Permissions The official Salesforce documentation outlines permission sets for Prompt Builder, including "Prompt Template User" for execution-only access and "Prompt Template Manager" for full control.
- * Trailhead: Configure Agentforce for Service This module discusses how permissions are assigned to control Agentforce features, including prompt-related capabilities.
- * Salesforce Ben: Why Prompt Builder Is Vital in an Agentforce World (November 25, 2024) This resource explains how Prompt Builder integrates with Agentforce and highlights the use of permission sets like Prompt Template User to enable end-user functionality.

NEW QUESTION: 115

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

What functionality allows the AX Specialist to achieve this?

- A.** User Utterances dashboard
- B.** Agent Event Logs dashboard
- C.** AI Audit & Feedback Data dashboard

Answer: (SHOW ANSWER)

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

- * Identify Common Queries: Surface frequently asked questions or recurring issues.
- * Detect Intent Patterns: Understand how users phrase requests, which helps refine intent detection models.
- * Improve Bot Training: Highlight gaps in training data or misclassified utterances that require adjustment.

Why Other Options Are Incorrect:

- * B. Agent Event Logs dashboard: Focuses on agent activity (e.g., response times, resolved cases) rather than user input analysis.
- * C. AI Audit & Feedback Data dashboard: Tracks AI model performance, audit trails, and user feedback scores but does not directly analyze raw user utterances or queries.

:

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

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

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

NEW QUESTION: 116

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

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

Answer: (SHOW ANSWER)

UC uses Agentforce for Sales to identify similar opportunities, aiding deal closure. Let's determine a criterion used by the "Find Similar Opportunities" feature.

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

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

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

Why Option A is Correct:

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

References:

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

Trailhead: Explore Agentforce Sales Agents - Details opportunity matching logic.

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

NEW QUESTION: 117

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: (SHOW ANSWER)

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: 118

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

- A.** Bring together a conversational interface for interacting with AI for all Salesforce users, such as developers and ecommerce retailers.
- B.** Enable Salesforce admin users to create and train custom large language models (LLMs) using CRM data.

C. Enable data scientists to train predictive AI models with historical CRM data using built-in machine learning capabilities

Answer: ([SHOW ANSWER](#))

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

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

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

:

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

NEW QUESTION: 119

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** ([LEAVE A REPLY](#))

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: 120

A. After a prompt template version is activated, no further changes can be saved to that version.

B. Only the latest version of a template can be activated.

C. Every modification on a template will be saved as a new version automatically.

Answer: (SHOW ANSWER)

According to the AgentForce Prompt Template Versioning Guide, when a prompt template version is marked as immutable, it means that no edits or modifications can be made to that version after it has been activated. This ensures that the logic, wording, and grounding parameters tied to that version remain locked and consistent for auditing, reproducibility, and compliance purposes.

If further changes are required, the system automatically creates a new draft version of the template. The old version remains immutable and preserved for traceability.

Option B is incorrect because multiple versions can exist, and older versions can remain active under specific testing or rollback scenarios. Option C is partially true but incomplete—the immutability refers to the frozen nature of an active version, not just version creation. Thus, the correct explanation is Option A - Once a prompt template version is activated, it becomes immutable and cannot be changed.

Reference: AgentForce Prompt Management Documentation - "Immutable Version Control in Prompt Templates."

NEW QUESTION: 121

Choose 1 option.

What is an Agentforce Specialist able to do when the 'Enrich event logs with conversation data' setting in the Agentforce configuration is enabled?

- A. View the user click path that led to each agent action.
- B. View session data including user input and agent responses for sessions.
- C. Generate details reports on all agent conversations over any time period.

Answer: (SHOW ANSWER)

The AgentForce Event and Logging Configuration Guide states that enabling "Enrich event logs with conversation data" allows administrators to capture session-level details, including both user inputs and agent responses. The documentation explains: "When this setting is enabled, conversation transcripts, user messages, and agent responses are appended to the event logs for improved visibility and troubleshooting." This provides a comprehensive record for analytics, training, and quality review. It does not, however, track user click paths (Option A) or generate aggregated historical reports across all time periods automatically (Option C).

Therefore, Option B is correct, as it directly reflects the documented functionality of the conversation data enrichment feature within AgentForce configuration.

References (AgentForce Documents / Study Guide):

- * AgentForce Configuration and Monitoring Guide: "Enrich Event Logs with Conversation Data"
- * AgentForce Data and Analytics Study Notes
- * AgentForce Implementation Handbook: "Session and Conversation Log Management"

Valid Agentforce-Specialist Dumps shared by ExamDiscuss.com for Helping Passing Agentforce-Specialist Exam! ExamDiscuss.com now offer the **newest Agentforce-Specialist exam dumps**, the ExamDiscuss.com Agentforce-Specialist exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com Agentforce-Specialist dumps with Test Engine here:
<https://www.examd Discuss.com/Salesforce/exam/Agentforce-Specialist/premium/> (300 Q&As Dumps, **35%OFF Special Discount Code: freecram**)

NEW QUESTION: 122

An Agentforce Specialist needs to create a prompt template to fill a custom field named Latest Opportunities Summary on the Account object with information from the three most recently opened opportunities. How should the Agentforce Specialist gather the necessary data for the prompt template?

- A. Select the latest Opportunities related list as a merge field.
- B. Create a flow to retrieve the opportunity information.
- C. Select the Account Opportunity object as a resource when creating the prompt template.

Answer: (SHOW ANSWER)

In Salesforce Agentforce, a prompt template designed to populate a custom field (like "Latest Opportunities Summary" on the Account object) requires dynamic data to be fed into the template for AI to generate meaningful output. Here, the task is to gather data from the three most recently opened opportunities related to an account. The most robust and flexible way to achieve this is by using a Flow (Option B). Salesforce Flows allow the Agentforce Specialist to define logic to query the Opportunity object, filter for the three most recent opportunities (e.g., using a Get Records element with a sort by CreatedDate descending and a limit of

3), and pass this data as variables into the prompt template. This approach ensures precise control over the data retrieval process and can handle complex filtering or sorting requirements.

* Option A: Selecting the "latest Opportunities related list as a merge field" is not a valid option in Agentforce prompt templates. Merge fields can pull basic field data (e.g., {! Account.Name}), but they don't natively support querying or aggregating related list data like the three most recent opportunities.

* Option C: There is no "Account Opportunity object" in Salesforce; this seems to be a misnomer (perhaps implying the Opportunity object or a junction object). Even if interpreted as selecting the Opportunity object as a resource, prompt templates don't directly query related objects without additional logic (e.g., a Flow), making this incorrect.

* Option B: Flows integrate seamlessly with prompt templates via dynamic inputs, allowing the Specialist to retrieve and structure the exact data needed (e.g., Opportunity Name, Amount, Close Date) for the AI to summarize.

Thus, Option B is the correct method to gather the necessary data efficiently and accurately.

:

Salesforce Agentforce Documentation: "Integrate Flows with Prompt Templates" (Salesforce Help:

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

[id=sf.agentforce_flow_prompt_integration.htm&type=5](https://help.salesforce.com/s/articleView?id=sf.agentforce_flow_prompt_integration.htm&type=5)) Trailhead: "Build Flows for Agentforce" (<https://trailhead.salesforce.com/content/learn/modules/flows-for-agentforce>)

NEW QUESTION: 123

An Agentforce Specialist at Universal Containers (UC) is building with no-code tools only. They have many small accounts that are only touched periodically by a specialized sales team, and UC wants to maximize the sales operations team's time, UC wants to help prep the sales team for calls by:

- * Summarizing past purchases
 - * Displaying products the contact has shown interest in (with data captured via Data Cloud)
 - * Providing a recap of past email and phone conversations that have transcripts
- Which approach should the Agentforce Specialist recommend to achieve this goal?

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

Answer: (SHOW ANSWER)

The AgentForce No-Code Builder and Prompt Template Guide specifies that when using Salesforce's standard foundation models, users can ground prompts on CRM and Data Cloud data without requiring fine-tuning or model deployment. The documentation notes: "No-code users can leverage standard foundation models with prompt templates grounded in Salesforce data, including CRM and Data Cloud, to summarize records, highlight insights, and prepare agents for customer engagements." This approach allows the sales operations team to automatically summarize customer data (purchases, interests, and communications) without needing model customization.

Option A (deploying a custom foundational model) requires data science expertise, not available in a no-code setup. Option B (fine-tuning) is unnecessary because standard models with grounding are optimized for contextual data use.

Thus, Option C aligns with Salesforce's recommended low-code/no-code configuration for contextual AI assistance.

References (AgentForce Documents / Study Guide):

- * AgentForce Builder Guide: "No-Code Prompt Templates Grounded in CRM and Data Cloud"
- * Salesforce Einstein Studio Overview: "Standard vs. Fine-Tuned Models"
- * AgentForce Study Guide: "Empowering Sales Teams with Grounded Prompts"

NEW QUESTION: 124

A business stakeholder wants to use AI to generate a summary based on Data Cloud data. Which method(s) should the stakeholder use to access Data Cloud data from Prompt Builder?

- A. Accessing data model objects (DMOs) directly in Flex templates, using Data Cloud related lists, and fetching Data Cloud data using prompt-initiated flows
- B. Using only external APIs to import Data Cloud data into Prompt Builder
- C. Using Data Cloud related lists and fetching Data Cloud data using prompt-initiated flows

Answer: (SHOW ANSWER)

The Prompt Builder and Data Cloud Integration Guide explains that Data Cloud information can be accessed directly through Data Cloud related lists or prompt-initiated flows, which fetch relevant data dynamically. The documentation states: "Prompt Builder supports retrieving Data Cloud data using related lists for contextual grounding or invoking flows that query Data Cloud objects at runtime. This enables AI prompts to generate summaries, recommendations, or insights directly from unified customer profiles." Option A is incorrect because direct access to data model objects (DMOs) in Flex templates is not supported in Prompt Builder. Option C (external APIs) is unnecessary, as Prompt Builder has native integration with Data Cloud.

Thus, Option B is the correct and Salesforce-documented method to access Data Cloud data from Prompt Builder.

References (AgentForce Documents / Study Guide):

- * Salesforce Prompt Builder Guide: "Integrating with Data Cloud"
- * AgentForce Study Guide: "Fetching Data Cloud Data with Prompt-Initiated Flows"
- * Salesforce Data Cloud Documentation: "Using Related Lists in Prompt Templates"

Valid Agentforce-Specialist Dumps shared by ExamDiscuss.com for Helping Passing Agentforce-Specialist Exam! ExamDiscuss.com now offer the **newest Agentforce-Specialist exam dumps**, the ExamDiscuss.com Agentforce-Specialist exam **questions have been updated** and **answers have been corrected** get the **newest**

ExamDiscuss.com Agentforce-Specialist dumps with Test Engine here:

<https://www.examdiscuss.com/Salesforce/exam/Agentforce-Specialist/premium/> (300

Q&As Dumps, **35%OFF Special Discount Code: freecram**)