



# **Salesforce**

## **Exam Questions Agentforce-Specialist**

Salesforce Certified Agentforce Specialist

## About ExamBible

### *Your Partner of IT Exam*

## Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

## Our Advances

### \* 99.9% Uptime

All examinations will be up to date.

### \* 24/7 Quality Support

We will provide service round the clock.

### \* 100% Pass Rate

Our guarantee that you will pass the exam.

### \* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

### NEW QUESTION 1

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

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

**Answer: C**

#### Explanation:

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

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

In contrast:

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

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

Reference:

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

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

### NEW QUESTION 2

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

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

**Answer: A**

#### Explanation:

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

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

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

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

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

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

prompts are primarily tracked in the Data Cloud for broader control, auditing, and analysis. Setup is not the primary tool for exporting or managing these audit logs.

More information on auditing AI interactions can be found in the Salesforce AI Trust Layer documentation, which outlines how organizations can manage and track generative AI interactions securely.

### NEW QUESTION 3

Universal Containers Is Interested In Improving the sales operation efficiency by analyzing their data using AI-powered predictions in Einstein Studio.

Which use case works for this scenario?

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

**Answer: B**

#### Explanation:

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

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

? Option A (customer sentiment) is typically handled through NLP models, while

Option C (product popularity) is more of a marketing analysis use case.

References:

Salesforce Einstein Studio Use Case Overview: [https://help.salesforce.com/s/articleView?id=sf.einstein\\_studio\\_overview](https://help.salesforce.com/s/articleView?id=sf.einstein_studio_overview)

### NEW QUESTION 4

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**

#### Explanation:

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?id=sf.flex\\_prompt\\_builder\\_troubleshoot.htm](https://help.salesforce.com/s/articleView?id=sf.flex_prompt_builder_troubleshoot.htm)

#### NEW QUESTION 5

Universal Container's internal auditing team asks An Agentforce to verify that address information is properly masked in the prompt being generated.

How should the Agentforce Specialist verify the privacy of the masked data in the Einstein Trust Layer?

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

**Answer: C**

#### Explanation:

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

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

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

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

Reference:

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

#### NEW QUESTION 6

What is a valid use case for Data Cloud retrievers?

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

**Answer: A**

#### Explanation:

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

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

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

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

References:

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

[https://help.salesforce.com/s/articleView?id=sf.data\\_cloud\\_agentforce.htm&type=5](https://help.salesforce.com/s/articleView?id=sf.data_cloud_agentforce.htm&type=5))

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

#### NEW QUESTION 7

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

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

**Answer: C**

#### Explanation:

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

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

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

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

References:

- ? Salesforce Agentforce Documentation: "Agent Analytics Overview" (Salesforce Help: [https://help.salesforce.com/s/articleView?id=sf.agentforce\\_analytics.htm&type=5](https://help.salesforce.com/s/articleView?id=sf.agentforce_analytics.htm&type=5))
- ? Trailhead: "Agentforce for Admins" (<https://trailhead.salesforce.com/content/learn/modules/agentforce-for-admins>)

**NEW QUESTION 8**

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

- A. To generate relevant content when engaging with each customer
- B. To analyze past campaign performance
- C. To send automated emails to all customers

**Answer: A**

**Explanation:**

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

- ? Option A is correct as Agent helps generate real-time, personalized content based on comprehensive data about the customer.
- ? Option B refers more to Einstein Analytics or Marketing Cloud Intelligence, and Option C deals with automation, which isn't the primary focus of Agent.

References:

- ? Salesforce Agent Overview: [https://help.salesforce.com/s/articleView?id=einstein\\_copilot\\_overview.htm](https://help.salesforce.com/s/articleView?id=einstein_copilot_overview.htm)

**NEW QUESTION 9**

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

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

**Answer: B**

**Explanation:**

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

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

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

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

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

References:

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

**NEW QUESTION 10**

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

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

**Answer: C**

**Explanation:**

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

- ? Prompt Defense (A) focuses on blocking malicious inputs or prompt injections but does not enhance responses with organizational data.
- ? Data Masking (B) redacts sensitive information but does not contribute to grounding responses in business context.

Reference:

- Salesforce Help Article: Einstein Trust Layer – Dynamic Grounding ("How Dynamic Grounding Works" section).
- Einstein Trust Layer Technical Overview: "Contextual Accuracy with Dynamic Grounding."

**NEW QUESTION 10**

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

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

The data sent cannot be accessed by the LLM provider.

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

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

**Answer: B**

**Explanation:**

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

? No data is used for training or product improvements by third-party LLMs.

? No data is retained outside of the customer's Salesforce organization.

? The LLM provider cannot access any customer data.

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

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

References:

? Salesforce Einstein Trust Layer Documentation: [https://help.salesforce.com/s/articleView?id=sf.einstein\\_trust\\_layer.htm](https://help.salesforce.com/s/articleView?id=sf.einstein_trust_layer.htm)

**NEW QUESTION 14**

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

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

**Answer: A**

**Explanation:**

Why is Outbound Omni-Channel Flow the Correct Answer?

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

Key Features of Outbound Omni-Channel Flow in Agentforce:

? Automates Escalation to a Human Agent Queue

? Seamless Transition from AI to Human Agents

? Ensures Proper Prioritization & Load Balancing

? Integration with Agentforce and Service Cloud

Why Not the Other Options?

\* B. Screen Flow

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

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

\* C. Prompt Flow

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

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

Agentforce Specialist References

The importance of using Omni-Channel Flow for routing AI-generated interactions to human agents is supported in the Agentforce Specialist exam objectives and documentation:

? Salesforce AI Specialist Material: Covers the importance of Omni-Channel routing for managing AI and human agent interactions.

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

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

**NEW QUESTION 16**

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

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

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

**Answer: B**

**Explanation:**

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

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

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

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

Reference:

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

**NEW QUESTION 21**

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

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

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

**Answer:** A

**Explanation:**

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

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

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

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

[https://help.salesforce.com/s/articleView?id=sf.model\\_builder\\_external\\_models.htm](https://help.salesforce.com/s/articleView?id=sf.model_builder_external_models.htm)

**NEW QUESTION 23**

An Agentforce is creating a custom action for Agentforce.

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

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

**Answer:** C

**Explanation:**

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

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

In contrast:

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

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

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

Reference:

Salesforce Help Article: Create Custom Actions for Einstein Bots

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

**NEW QUESTION 28**

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

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

**Answer:** B

**Explanation:**

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

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

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

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

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

References:

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

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

**NEW QUESTION 29**

Universal Containers (UC) plans to automatically populate the Description field on the Account object.

Which type of prompt template should UC use?

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 33**

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

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

**Answer: A**

**Explanation:**

- The use case that necessitates the use of Prompt Builder is creating a draft of a support bulletin post for new product patches. Prompt Builder allows the Agentforce Specialist to create and refine prompts that generate specific, relevant outputs, such as drafting support communication based on product information and patch details.
- ? Option B (agent performance score) would likely involve predictive modeling, not prompt generation.
  - ? Option C (estimating support ticket volume) would require data analysis and predictive tools, not prompt building.
- For more details, refer to Salesforce's Prompt Builder documentation for generative AI content creation.

**NEW QUESTION 38**

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

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

**Answer: B**

**Explanation:**

- Comprehensive and Detailed In-Depth Explanation: UC wants to streamline the use of Generative AI by reducing the time reps spend typing prompts and ensuring consistency, leveraging their existing prompt knowledge. Let's evaluate the options.
- ? Option A: Agent Builder and Action: Query Records. Agent Builder in Agentforce Studio creates autonomous AI agents with actions like "Query Records" to fetch data. While this could retrieve information, it's designed for agent-driven workflows, not for simplifying manual prompt entry or ensuring consistency across user inputs. This doesn't directly address UC's concerns and is incorrect.
  - ? Option B: Einstein Prompt Builder and Prompt Templates. Einstein Prompt Builder, part of Agentforce Studio, allows users to create reusable prompt templates that encapsulate specific instructions and grounding for Generative AI (e.g., using public models via the Atlas Reasoning Engine). UC can predefine prompts based on their known language, saving time for reps by eliminating repetitive typing and ensuring consistency across sales and service teams. Templates can be embedded in flows, Lightning pages, or agent interactions, perfectly addressing UC's needs. This is the correct answer.
  - ? Option C: Einstein Recommendation Builder. Einstein Recommendation Builder generates personalized recommendations (e.g., products, next best actions) using predictive AI, not Generative AI for freeform prompts. It doesn't support custom prompt creation or address time/consistency issues for reps, making it incorrect.
- Why Option B is Correct: Einstein Prompt Builder's prompt templates directly tackle UC's challenges by standardizing prompts and reducing manual effort, leveraging their familiarity with Generative AI language. This is a core feature for such use cases, as per Salesforce documentation.
- References:
- ? Salesforce Agentforce Documentation: Einstein Prompt Builder – Details prompt templates for consistency and efficiency.
  - ? Trailhead: Build Prompt Templates in Agentforce – Explains time-saving benefits of templates.
  - ? Salesforce Help: Generative AI with Prompt Builder – Confirms use for streamlining rep interactions.

**NEW QUESTION 43**

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

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

**Answer: B**

**Explanation:**

- Why is "Messaging Session" the correct answer?
- In Agentforce, the Messaging Session object stores the conversation transcript between the customer and the agent.
- Key Features of the Messaging Session Object:
- ? Stores the Entire Customer-Agent Conversation
  - ? Supports AI-Powered Work Summaries
  - ? Links with Service Cloud for Case Resolution
- Why Not the Other Options?
- \* A. Messaging End User
  - ? Incorrect because this object stores details about the customer (e.g., name, contact details) but not the conversation transcript.

\* C. Case

? Incorrect because Cases store structured service requests but do not contain raw conversation transcripts.

? Instead, cases may reference the Messaging Session object.

Agentforce Specialist References

? Salesforce AI Specialist Material confirms that Messaging Sessions store chat conversations and support Einstein Work Summaries.

#### NEW QUESTION 44

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

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

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

C. Storing this data requires Salesforce big objects.

**Answer: A**

#### Explanation:

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

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

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

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

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

References:

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

? Salesforce Einstein AI Overview:

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

#### NEW QUESTION 45

Universal Containers (UC) wants to use the Draft with Einstein feature in Sales Cloud to create a personalized introduction email. After creating a proposed draft email, which predefined adjustment should UC choose to revise the draft with a more casual tone?

A. Make Less Formal

B. Enhance Friendliness

C. Optimize for Clarity

**Answer: A**

#### Explanation:

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

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

? Optimize for Clarity focuses on making the draft clearer but doesn't adjust the tone. For more details, see Salesforce documentation on Einstein-generated email drafts and tone adjustments.

#### NEW QUESTION 50

A Salesforce Administrator wants to generate personalized, targeted emails that incorporate customer interaction data. The admin wants to leverage large language models (LLMs) to write the emails, and wants to reuse templates for different products and customers. Which solution approach should the admin leverage?

A. Use sales Email standard templates

B. Create a t field Generation prompt template type

C. Create a Sales Email prompt template type.

**Answer: C**

#### Explanation:

To generate personalized emails using LLMs while reusing templates:

? Sales Email Prompt Template Type (Option C): Designed specifically for generating dynamic email content by combining LLMs with structured templates. It allows admins to define placeholders (e.g., customer name, product details) and reuse templates across scenarios.

? Option A: Standard email templates lack LLM integration and dynamic personalization.

? Option B: "t field Generation" is not a valid Salesforce prompt template type.

References:

? Salesforce Help: Sales Email Prompt Templates

? Describes using Sales Email prompt templates to "generate targeted emails using dynamic data and LLMs."

#### NEW QUESTION 51

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

A. Invoke a flow which makes a call to external data to create a Knowledge article.

- B. Execute tasks based on available actions, answering questions using information from accessible Knowledge articles.
- C. Generate Knowledge article based off the prompts that the agent enters to create steps to cancel flights.

**Answer:** B

**Explanation:**

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

- ? Option A: Invoking a Flow to call external data and create a Knowledge article is unnecessary. The representative needs to review existing articles, not create new ones, and there's no indication external data is required for this task.
- ? Option B: This is correct. It combines task execution (canceling flights) with Knowledge article retrieval, aligning with the representative's need to act and seek guidance from existing content.
- ? Option C: Generating a new Knowledge article based on prompts is not relevant. The representative needs to use existing articles, not author new ones, especially in a time-sensitive weather alert scenario.

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

**References:**

- ? Salesforce Agentforce Documentation: "Knowledge Replies and Actions" (Salesforce Help: [https://help.salesforce.com/s/articleView?id=sf.agentforce\\_knowledge\\_replies.htm&type=5](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 55**

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

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

**Answer:** A

**Explanation:**

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

- ? Option A: Create a custom Record Summary prompt template for the Case object. In Prompt Builder, the standard Record Summary prompt template generates summaries for objects like Case. To customize it, the Agentforce Specialist can create a new custom prompt template, specifying the Case object as the source, and adjust the instructions (e.g., tone, fields included) to meet UC's needs. This new template can then be invoked by an agent or flow, providing a tailored summary. This approach offers full control and aligns with Salesforce's customization process, making it the correct answer.
- ? Option B: Summarize the Case with a standard Agent action. Standard Agent actions (e.g., "Answer Questions") don't specifically target case summarization—they're broader in scope. There's no out-of-the-box "Summarize Case" action that allows customization of the response format, making this insufficient and incorrect.
- ? Option C: Customize the standard Record Summary template for the Case object. Standard prompt templates in Prompt Builder (e.g., Record Summary) are read-only and cannot be directly edited. Customization requires cloning or creating a new template, not modifying the standard one, making this incorrect.

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

**References:**

- ? Salesforce Agentforce Documentation: Prompt Builder > Custom Templates – Details creating custom summaries.
- ? Trailhead: Build Prompt Templates in Agentforce – Explains customizing standard outputs.
- ? Salesforce Help: Record Summaries with AI – Recommends custom templates for tailored results.

**NEW QUESTION 60**

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

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

**Answer:** B

**Explanation:**

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

- ? Option B is correct because Einstein Service Replies is responsible for generating AI-driven responses based on knowledge articles.
- ? Option A (Einstein Reply Recommendations) is focused on recommending replies but does not generate them.
- ? Option C (Einstein Grounding) refers to grounding responses in data but is not directly related to drafting replies.

**References:**

- ? Einstein Service Replies Overview: [https://help.salesforce.com/s/articleView?id=sf.einstein\\_service\\_replies.htm](https://help.salesforce.com/s/articleView?id=sf.einstein_service_replies.htm)

**NEW QUESTION 64**

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

- A. The data source can be changed through the Data Cloud settings.
- B. The Data Retriever can be reconfigured to use a different data source.

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

**Answer: C**

**Explanation:**

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

Key Considerations in Agentforce Data Library

? Data Source Lock-In

? Why Can't the Data Source Be Changed?

? Workarounds for Changing Data Sources

Why Not the Other Options?

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

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

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

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

Agentforce Specialist References

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

**NEW QUESTION 65**

Universal Containers implemented Agent for its users.

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

A. Agent Delete Record Action permission is not associated to the user.

B. Agent does not have the permission to delete the user's records.

C. Agent does not support the Delete Record action.

**Answer: C**

**Explanation:**

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

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

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

References:

? Salesforce Agentforce Specialist Documentation - Agent Supported Actions:

? Salesforce Help - Limitations of Agent:

**NEW QUESTION 66**

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

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

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

C. Sales rep users are missing the Use SDR Agent permission set.

**Answer: C**

**Explanation:**

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

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

Key Considerations for This Issue:

? Permission Set Restriction

? Agentforce Role-Based Access Control

? Fixing the Issue

Why Not the Other Options?

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

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

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

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

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

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

Agentforce Specialist References

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

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

**NEW QUESTION 70**

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

A. The text encoding (for example, UTF-8, ASCII) option

B. The maximum number of related list merge fields

C. The choice between XML and JSON rendering formats for the list

**Answer: B**

**Explanation:**

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

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

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

References:

? Salesforce Prompt Builder Documentation: [https://help.salesforce.com/s/articleView?id=sf.prompt\\_builder.htm](https://help.salesforce.com/s/articleView?id=sf.prompt_builder.htm)

**NEW QUESTION 72**

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

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

- A. Assign the Prompt Template User permission set and enable Sales Emails in Setup.
- B. Assign the Prompt Template Manager permission set and enable Sales Emails in setup.
- C. Assign the Data Cloud Admin permission set and enable Sales Emails in Setup.

**Answer:** A

**Explanation:**

To enable Sales Email prompt templates:

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

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

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

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

References:

? Salesforce Help: Prompt Template Permissions

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

? Sales Email Setup outlines enabling the feature in Setup.

**NEW QUESTION 76**

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

What is an important consideration for this requirement?

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

**Answer:** B

**Explanation:**

? Context of the Question

? Why Apex Code for JSON Formatting?

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

Salesforce Agentforce Specialist References & Documents

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

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

**NEW QUESTION 79**

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

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

**Answer:** B

**Explanation:**

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

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

The other options are incorrect because:

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

References:

Salesforce Einstein Documentation on Agent Actions

Salesforce AI Documentation on Large Language Models

**NEW QUESTION 83**

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

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

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

- B. Add the fine-tuned LLM in Einstein Studio Model Builder.
- C. Enable model endpoint on OpenAI and make callouts to the model to generate emails.

**Answer:** B

**Explanation:**

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

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

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

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

**NEW QUESTION 88**

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

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

**Answer:** A

**Explanation:**

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

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

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

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

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

References:

? Salesforce Agentforce Documentation: Agent Builder > Custom Actions > Apex – Notes standard Apex deployment rules apply.

? Salesforce Developer Guide: Apex Testing – Confirms 75% coverage requirement.

? Trailhead: Deploy Apex Code – Emphasizes coverage and dependencies for production.

**NEW QUESTION 93**

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

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

**Answer:** B

**Explanation:**

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

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

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

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

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

References:

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

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

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

**NEW QUESTION 96**

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

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

**Answer: B**

**Explanation:**

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

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

Key Considerations for Einstein Data Libraries:

? Salesforce Knowledge Integration is Optional

? How to Fix This Issue?

Why Not the Other Options?

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

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

\* C. The administrator is not using Lightning Experience

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

Agentforce Specialist References

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

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

**NEW QUESTION 100**

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

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

**Answer: C**

**Explanation:**

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

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

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

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

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

References:

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

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

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

**NEW QUESTION 103**

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

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

**Answer: C**

**Explanation:**

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

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

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

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

References:

? Salesforce Help: Einstein Agent Setup

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

#### NEW QUESTION 107

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

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

**Answer: C**

#### Explanation:

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

References:

- ? Salesforce Agentforce Specialist Documentation - Agent Overview: Details how Agent employs LLMs to interpret user inputs and generate responses within the Salesforce ecosystem.
- ? Salesforce Help - How Agent Works: Explains the underlying mechanisms of how Agent processes user requests using AI technologies.

#### NEW QUESTION 111

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

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

**Answer: C**

#### Explanation:

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

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

? Option B: SOQL SOQL (Salesforce Object Query Language) is a query language,

not an executable action type in Agent Builder. While actions can use queries internally, SOQL isn't a standalone option, making this incorrect.

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

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

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

References:

- ? Salesforce Agentforce Documentation: Agent Builder > Custom Actions – Lists Flows as a supported action type.
- ? Trailhead: Build Agents with Agentforce – Details Flow-based actions.
- ? Salesforce Help: Configure Agent Actions – Confirms Flows integration.

#### NEW QUESTION 116

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

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

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

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

**Answer: B**

#### Explanation:

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

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

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

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

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

References:

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

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

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

#### NEW QUESTION 120

What is the importance of Action Instructions when creating a custom Agent action?

- A. Action Instructions define the expected user experience of an action.
- B. Action Instructions tell the user how to call this action in a conversation.
- C. Action Instructions tell the large language model (LLM) which action to use.

**Answer: A**

#### Explanation:

Comprehensive and Detailed In-Depth Explanation: In Salesforce Agentforce, custom Agent actions are designed to enable AI-driven agents to perform specific tasks within a conversational context. Action Instructions are a critical component when creating these actions because they define the expected user experience by outlining how the action should behave, what it should accomplish, and how it interacts with the end user. These instructions act as a blueprint for the action's functionality, ensuring that it aligns with the intended outcome and provides a consistent, intuitive experience for users interacting with the agent. For example, if the action is to "schedule a meeting," the Action Instructions might specify the steps (e.g., gather date and time, confirm with the user) and the tone (e.g., professional, concise), shaping the user experience.

? Option B: While Action Instructions might indirectly influence how a user invokes an action (e.g., by making it clear what inputs are needed), they are not primarily about telling the user how to call the action in a conversation. That's more related to user training or interface design, not the instructions themselves.

? Option C: The large language model (LLM) relies on prompts, parameters, and grounding data to determine which action to execute, not the Action Instructions directly. The instructions guide the action's design, not the LLM's decision-making process at runtime.

Thus, Option A is correct as it emphasizes the role of Action Instructions in defining the user experience, which is foundational to creating effective custom Agent actions in Agentforce.

References:

? Salesforce Agentforce Documentation: "Create Custom Agent Actions" (Salesforce Help:

[https://help.salesforce.com/s/articleView?id=sf.agentforce\\_custom\\_actions.htm&type=5](https://help.salesforce.com/s/articleView?id=sf.agentforce_custom_actions.htm&type=5))

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

#### NEW QUESTION 121

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

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

**Answer: A**

#### Explanation:

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

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

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

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

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

References:

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

[https://help.salesforce.com/s/articleView?id=sf.agentforce\\_custom\\_retrievers.htm&type=5](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](https://help.salesforce.com/s/articleView?id=sf.data_cloud_retrieval_filters.htm&type=5))

#### NEW QUESTION 122

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

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

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

**Answer: A**

#### Explanation:

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

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

Key Considerations for This Approach:

? Custom Agent Action Triggers the Flow

? Enables AI-Powered Query Execution

? No Need for Manual Data Entry

Why Not the Other Options?

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

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

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

\* C. Create a Flex Prompt Template in Prompt Builder

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

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

Agentforce Specialist References

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

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

### NEW QUESTION 127

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: A**

#### Explanation:

Comprehensive and Detailed In-Depth Explanation: 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 129

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

What should UC consider when using standard Service AI Grounding?

A. Service AI Grounding only works with Case and Knowledge objects.

B. Service AI Grounding only supports String and Text Area type fields.

C. Service AI Grounding visibility works in system mode.

**Answer: B**

#### Explanation:

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

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

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

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

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

References:

? Salesforce Help: Service AI Grounding Requirements

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

### NEW QUESTION 132

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: B**

#### Explanation:

Comprehensive and Detailed In-Depth Explanation: 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.

References:

? Salesforce Agentforce Documentation: "Integrate Flows with Prompt Templates" (Salesforce Help:

[https://help.salesforce.com/s/articleView?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 137

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

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

A. Enable prompt test mode, allocate different prompt variations to a subset of users for evaluation, and standardize the most effective model based on performance feedback.

B. Create concise, clear, and consistent prompt templates with effective grounding, contextual role-playing, clear instructions, and iterative feedback.

C. Submit a prompt review case to Salesforce and conduct thorough testing in the playground to refine outputs until they meet user expectations.

**Answer: B**

#### Explanation:

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

? Effective grounding: Ensuring the prompt pulls data from the correct sources.

? Contextual role-playing: Providing the AI with a clear understanding of its role in generating the summary.

? Clear instructions: Giving unambiguous directions on what to include in the response.

? Iterative feedback: Regularly testing and adjusting prompts based on user feedback.

? Option B is correct because it follows industry best practices for refining prompt design.

? Option A (prompt test mode) is useful but less relevant for refining prompt design itself.

? Option C (prompt review case with Salesforce) would be more appropriate for technical issues or complex prompt errors, not general design refinement.

References:

Salesforce Prompt Design Best Practices: [https://help.salesforce.com/s/articleView?id=sf.prompt\\_design\\_best\\_practices.htm](https://help.salesforce.com/s/articleView?id=sf.prompt_design_best_practices.htm)

### NEW QUESTION 138

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

A. A data pipeline, an indexing engine, and a query processor

B. A data connector, an analytics dashboard, and a workflow rule

C. A data stream, a search index, and a retriever

**Answer: C**

#### Explanation:

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

Key Components Created in Data Cloud:

? Data Stream

? Search Index

? Retriever

Why Not the Other Options?

\* A. A data pipeline, an indexing engine, and a query processor

? Incorrect because Data Cloud does not use a query processor in the same way as traditional databases.

? Instead, retrievers handle AI-powered data searches.

\* B. A data connector, an analytics dashboard, and a workflow rule

? Incorrect because these components are not automatically created when setting up a Data Library.

? Analytics dashboards and workflow rules are separate tools used for reporting and automation.

Agentforce Specialist References

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

### NEW QUESTION 139

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: B**

#### Explanation:

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 143

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: B**

#### Explanation:

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

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

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

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

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

[https://help.salesforce.com/s/articleView?id=sf.einstein\\_copilot\\_event\\_logging.htm](https://help.salesforce.com/s/articleView?id=sf.einstein_copilot_event_logging.htm)

#### NEW QUESTION 144

An Agentforce is tasked to optimize a business process flow by assigning actions to agents within the Salesforce Agentforce Platform.

What is the correct method for the Agentforce Specialist to assign actions to an Agent?

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

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

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

**Answer: C**

#### Explanation:

? Action Builder is the central place in Salesforce Agentforce where you define and manage actions that your AI agents can perform. This includes connecting actions to various tools and systems.

? Topics in Agentforce represent the different tasks or intents that an AI agent can handle. By assigning an action to a Topic in Action Builder, you're essentially telling the agent, "When you encounter this type of request or situation, perform this action."

#### NEW QUESTION 146

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

A. It automates the ingestion, taxonomical classification and storage of knowledge in Data Cloud for precision keyword search retrieval to ground prompts and agents with relevant information.

B. It automates the ingestion, indexing of data, and creates a default retriever to be used in prompts and agents for grounding with relevant information.

C. It automates the ingestion and optical character recognition (OCR) processing of any PDF, and indexes them to enable regular SQL query retrieval to ground

prompts and agents with relevant information.

**Answer:** B

**Explanation:**

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

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

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

Key Features of an Agentforce Data Library:

? Automates Data Ingestion

? Indexes Data for Efficient Retrieval

? Creates a Default Retriever

Why Not the Other Options?

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

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

\* C. Automates ingestion and OCR processing of PDFs

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

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

Agentforce Specialist References

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

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

**NEW QUESTION 150**

An Agentforce turned on Einstein Generative AI in Setup. Now, the Agentforce Specialist would like to create custom prompt templates in Prompt Builder.

However, they cannot access Prompt Builder in the Setup menu.

What is causing the problem?

A. The Prompt Template User permission set was not assigned correctly.

B. The Prompt Template Manager permission set was not assigned correctly.

C. The large language model (LLM) was not configured correctly in Data Cloud.

**Answer:** B

**Explanation:**

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

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

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

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

References:

? Salesforce Prompt Builder Permissions: [https://help.salesforce.com/s/articleView?id=sf.prompt\\_builder\\_permissions.htm](https://help.salesforce.com/s/articleView?id=sf.prompt_builder_permissions.htm)

**NEW QUESTION 155**

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

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

A. Design universal product recommendations.

B. Write a response script for the chatbot.

C. Collect and analyze browsing data.

**Answer:** C

**Explanation:**

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

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

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

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

Options A and D are incorrect because:

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

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

References:

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

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

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

**NEW QUESTION 160**

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

What is a key reason for implementing Agent?

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

**Answer:** C

**Explanation:**

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

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

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

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

[https://help.salesforce.com/s/articleView?id=sf.einstein\\_copilot\\_overview.htm](https://help.salesforce.com/s/articleView?id=sf.einstein_copilot_overview.htm)

**NEW QUESTION 161**

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

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

**Answer:** A

**Explanation:**

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

How It Works:

? Data Masking in the Request Journey:

? Processing by the LLM:

? De-masking in the Response Journey:

Why Option A is Correct:

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

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

Why Options B and C are Incorrect:

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

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

References:

? Salesforce Agentforce Specialist Documentation - Einstein Trust Layer Overview:

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

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

Conclusion:

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

**NEW QUESTION 162**

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

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 166**

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

- A. The User related list should have View All access.
- B. The User related list needs to be included on the record page.
- C. The User related list is not supported in prompt templates.

**Answer:** C

**Explanation:**

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

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

### NEW QUESTION 168

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:** A

#### Explanation:

? 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

? Why Option A is Correct

? Why Not Option B (JSON Formatting)

? Why Not Option C (Default Page Layout)

? 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 173

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

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

**Answer:** C

#### Explanation:

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

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

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

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

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

References:

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

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

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

### NEW QUESTION 177

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

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

**Answer:** C

#### Explanation:

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

Key Steps in the Data Preparation Process for Agentforce:

? Loading Data

? Chunking (Breaking Text into Small Parts)

? Vectorization (Transforming Text for AI Retrieval)

? Storing in a Vector Database

Why Not the Other Options?

\* A. Real-time data ingestion and dynamic indexing

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

\* B. Aggregating, normalizing, and encoding structured datasets

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

Agentforce Specialist References

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

### NEW QUESTION 182

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

What should An Agentforce recommend to generate this email?

- A. Standard email marketing template using Apex or flows for matching interest in products
- B. Custom sales email template which is grounded with interest and product information
- C. Standard email draft with Einstein and choose standard email template

**Answer:** B

**Explanation:**

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

Using grounding ensures that the generative AI pulls the correct data related to customer interests and product matches, making the email more effective. For more information, refer to Salesforce documentation on grounding AI-generated content and email personalization strategies.

**NEW QUESTION 183**

An account manager is preparing for an upcoming customer call and wishes to get a snapshot of key data points from accounts, contacts, leads, and opportunities in Salesforce.

Which feature provides this?

- A. Sales Summaries
- B. Sales Insight Summary
- C. Work Summaries

**Answer:** B

**Explanation:**

Sales Insight Summary aggregates key data points from multiple Salesforce objects (accounts, contacts, leads, opportunities) into a consolidated view, enabling account managers to quickly access relevant information for customer calls.

? Option A (Sales Summaries): Typically refers to Einstein-generated summaries of specific interactions (e.g., emails, calls), not multi-object snapshots.

? Option C (Work Summaries): Focuses on summarizing customer service interactions (e.g., chat transcripts), not sales data.

? Option B (Sales Insight Summary): Directly provides a holistic snapshot of sales-related objects, aligning with the scenario.

References:

? Salesforce Help: Sales Insight Overview

? Describes Sales Insight Summary as "a unified view of account, contact, and opportunity data for sales readiness."

**NEW QUESTION 188**

A sales manager is using Agent Assistant to streamline their daily tasks. They ask the agent to Show me a list of my open opportunities.

How does the large language model (LLM) in Agentforce identify and execute the action to show the sales manager a list of open opportunities?

- A. The LLM interprets the user's request, generates a plan by identifying the appropriate topics and actions, and executes the actions to retrieve and display the open opportunities
- B. The LLM uses a static set of rules to match the user's request with predefined topics and actions, bypassing the need for dynamic interpretation and planning.
- C. Using a dialog pattern
- D. the LLM matches the user query to the available topic, action and steps then performs the steps for each action, such as retrieving a list of open opportunities.

**Answer:** A

**Explanation:**

Agentforce's LLM dynamically interprets natural language requests (e.g., "Show me open opportunities"), generates an execution plan using the planner service, and retrieves data via actions (e.g., querying Salesforce records). This contrasts with static rules (B) or rigid dialog patterns (C), which lack contextual adaptability. Salesforce documentation highlights the planner's role in converting intents into actionable steps while adhering to security and business logic.

Reference:

Salesforce Help Article: Agentforce Planner Service ("Dynamic Request Interpretation" section).

Einstein Agentforce Specialist Trailhead: "How Agentforce Processes User Requests."

**NEW QUESTION 189**

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

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

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

**Answer:** A

**Explanation:**

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

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

? Option C (manual insertion) would not allow for the dynamic and automated grounding of data required for custom prompts.

Refer to Salesforce documentation on flows and grounding for more details on integrating data into custom prompt templates.

**NEW QUESTION 191**

An Agentforce has created a copilot custom action using flow as the reference action type. However, it is not delivering the expected results to the conversation preview, and therefore needs troubleshooting.

What should the Agentforce Specialist do to identify the root cause of the problem?

- A. In Copilot Builder within the Dynamic Panel, turn on dynamic debugging to show the inputs and outputs.

- B. Copilot Builder within the Dynamic Panel, confirm selected action and observe the values in Input and Output sections.
- C. In Copilot Builder, verify the utterance entered by the user and review session event logs for debug information.

**Answer:** A

**Explanation:**

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

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

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

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

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

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

[https://help.salesforce.com/s/articleView?id=sf.copilot\\_custom\\_action\\_debugging.htm](https://help.salesforce.com/s/articleView?id=sf.copilot_custom_action_debugging.htm)

**NEW QUESTION 193**

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

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

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

**Answer:** C

**Explanation:**

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

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

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

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

**NEW QUESTION 194**

Universal Containers (UC) uses a file upload-based data library and custom prompt to support AI-driven training content. However, users report that the AI frequently returns outdated documents. Which corrective action should UC implement to improve content relevancy?

- A. Switch the data library source from file uploads to a Knowledge-based data library, because Salesforce Knowledge bases automatically manage document recency, ensuring current documents are returned.
- B. Configure a custom retriever that includes a filter condition limiting retrieval to documents updated within a defined recent period, ensuring that only current content is used for AI responses.
- C. Continue using the default retriever without filters, because periodic re-uploads will eventually phase out outdated documents without further configuration or the need for custom retrievers.

**Answer:** B

**Explanation:**

Comprehensive and Detailed In-Depth Explanation: UC's issue is that their file upload-based Data Library (where PDFs or documents are uploaded and indexed into Data Cloud's vector database) is returning outdated training content in AI responses. To improve relevancy by ensuring only current documents are retrieved, the most effective solution is to configure a custom retriever with a filter (Option B). In Agentforce, a custom retriever allows UC to define specific conditions—such as a filter on a "Last Modified Date" or similar timestamp field—to limit retrieval to documents updated within a recent period (e.g., last 6 months). This ensures the AI grounds its responses in the most current content, directly addressing the problem of outdated documents without requiring a complete overhaul of the data source.

? Option A: Switching to a Knowledge-based Data Library (using Salesforce Knowledge articles) could work, as Knowledge articles have versioning and expiration features to manage recency. However, this assumes UC's training content is already in Knowledge articles (not PDFs) and requires migrating all uploaded files, which is a significant shift not justified by the question's context. File-based libraries are still viable with proper filtering.

? Option B: This is the best corrective action. A custom retriever with a date filter leverages the existing file-based library, refining retrieval without changing the data source, making it practical and targeted.

? Option C: Relying on periodic re-uploads with the default retriever is passive and inefficient. It doesn't guarantee recency (old files remain indexed until manually removed) and requires ongoing manual effort, failing to proactively solve the issue.

Option B provides a precise, scalable solution to ensure content relevancy in UC's AI-driven training system.

References:

? Salesforce Agentforce Documentation: "Custom Retrievers for Data Libraries" (Salesforce Help:

[https://help.salesforce.com/s/articleView?id=sf.agentforce\\_custom\\_retrievers.htm&type=5](https://help.salesforce.com/s/articleView?id=sf.agentforce_custom_retrievers.htm&type=5))

? Salesforce Data Cloud Documentation: "Filter Retrieval for AI"

([https://help.salesforce.com/s/articleView?id=sf.data\\_cloud\\_retrieval\\_filters.htm&type=5](https://help.salesforce.com/s/articleView?id=sf.data_cloud_retrieval_filters.htm&type=5))

? Trailhead: "Manage Data Libraries in Agentforce" (<https://trailhead.salesforce.com/content/learn/modules/agentforce-data-libraries>)

**NEW QUESTION 197**

.....

## Relate Links

**100% Pass Your Agentforce-Specialist Exam with Exambible Prep Materials**

<https://www.exambible.com/Agentforce-Specialist-exam/>

## Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>