

Salesforce

AI-201

Salesforce Certified Agentforce Specialist

For More Information – Visit link below:

<https://www.examsempire.com/>

Product Version

1. Up to Date products, reliable and verified.
2. Questions and Answers in PDF Format.



<https://examsempire.com/>

Visit us at: <https://www.examsempire.com/ai-201>

Latest Version: 9.0

Question: 1

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:

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.

Salesforce Agentforce Documentation: "Create Custom Agent Actions" (Salesforce Help: 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>)

Question: 2

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

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

Answer: B

Explanation:

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

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

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

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

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

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

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

Trailhead: "Build Prompt Templates for Agentforce"

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

Question: 3

What is a valid use case for Data Cloud retrievers?

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

Answer: A

Explanation:

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.

Salesforce Data Cloud Documentation: "Data Cloud for Agentforce" (Salesforce Help: https://help.salesforce.com/s/articleView?id=sf.data_cloud_agentforce.htm&type=5)

Trailhead: "Data Cloud Basics" module

(<https://trailhead.salesforce.com/content/learn/modules/data-cloud-basics>)

Question: 4

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

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

Answer: A

Explanation:

Service Replies (specifically Einstein Service Replies) is the Salesforce Generative AI functionality designed to automatically draft responses for service agents in real-time, based on contextual information, including existing knowledge articles. This directly addresses Universal Containers' need to reduce handling time by providing recommended replies grounded in their knowledge base

Question: 5

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:

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

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

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

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

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

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

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

Salesforce Data Cloud Documentation: "Vector Database for AI"

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

Question: 6

Universal Containers (UC) is creating a new custom prompt template to populate a field with generated output. UC enabled the Einstein Trust Layer to ensure AI Audit data is captured and

monitored for adoption and possible enhancements. Which prompt template type should UC use and which consideration should UC review?

- A. Field Generation, and that Dynamic Fields is enabled
- B. Field Generation, and that Dynamic Forms is enabled
- C. Flex, and that Dynamic Fields is enabled

Answer: A

Explanation:

Salesforce Agentforce provides various prompt template types to support AI-driven tasks, such as generating text or populating fields. In this case, UC needs a custom prompt template to populate a field with generated output, which directly aligns with the Field Generation prompt template type. This type is designed to use generative AI to create field values (e.g., summaries, descriptions) based on input data or prompts, making it the ideal choice for UC's requirement. Additionally, UC has enabled the Einstein Trust Layer, a governance framework that ensures AI outputs are safe, explainable, and auditable, capturing AI Audit data for monitoring adoption and identifying improvement areas.

The consideration UC should review is whether Dynamic Fields is enabled. Dynamic Fields allow the prompt template to incorporate variable data from Salesforce records (e.g., case details, customer info) into the prompt, ensuring the generated output is contextually relevant to each record. This is critical for field population tasks, as static prompts wouldn't adapt to record-specific needs. The Einstein Trust Layer further benefits from this, as it can track how dynamic inputs influence outputs for audit purposes.

Option A: Correct. "Field Generation" matches the use case, and "Dynamic Fields" is a key consideration to ensure flexibility and auditability with the Trust Layer.

Option B: "Field Generation" is correct, but "Dynamic Forms" is unrelated. Dynamic Forms is a UI feature for customizing page layouts, not a prompt template setting, making this option incorrect.

Option C: "Flex" templates are more general-purpose and not specifically tailored for field population tasks. While Dynamic Fields could apply, Field Generation is the better fit for UC's stated goal.

Option A is the best choice, as it pairs the appropriate template type (Field Generation) with a relevant consideration (Dynamic Fields) for UC's scenario with the Einstein Trust Layer.

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

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

Trailhead: "Build Prompt Templates for Agentforce"

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

Question: 7

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:

In Salesforce Agentforce, a prompt template designed to populate a custom field (like "Latest Opportunities Summary" on the Account object) requires dynamic data to be fed into the template for AI to generate meaningful output. Here, the task is to gather data from the three most recently opened opportunities related to an account. The most robust and flexible way to achieve this is by using a Flow (Option B). Salesforce Flows allow the Agentforce Specialist to define logic to query the Opportunity object, filter for the three most recent opportunities (e.g., using a Get Records element with a sort by CreatedDate descending and a limit of 3), and pass this data as variables into the prompt template. This approach ensures precise control over the data retrieval process and can handle complex filtering or sorting requirements.

Option A: Selecting the "latest Opportunities related list as a merge field" is not a valid option in Agentforce prompt templates. Merge fields can pull basic field data (e.g., {!Account.Name}), but they don't natively support querying or aggregating related list data like the three most recent opportunities.

Option C: There is no "Account Opportunity object" in Salesforce; this seems to be a misnomer (perhaps implying the Opportunity object or a junction object). Even if interpreted as selecting the Opportunity object as a resource, prompt templates don't directly query related objects without additional logic (e.g., a Flow), making this incorrect.

Option B: Flows integrate seamlessly with prompt templates via dynamic inputs, allowing the Specialist to retrieve and structure the exact data needed (e.g., Opportunity Name, Amount, Close Date) for the AI to summarize.

Thus, Option B is the correct method to gather the necessary data efficiently and accurately. Salesforce Agentforce Documentation: "Integrate Flows with Prompt Templates" (Salesforce Help:

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

Trailhead: "Build Flows for Agentforce"

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

Question: 8

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:

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.

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

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

Trailhead: "Agentforce for Admins"

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

Question: 9

Universal Containers (UC) wants to implement an AI-powered customer service agent that can:
Retrieve proprietary policy documents that are stored as PDFs.

Ensure responses are grounded in approved company data, not generic LLM knowledge.

What should UC do first?

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

Answer: A

Explanation:

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

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

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

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

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

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

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

Question: 10

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:

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

predefined action) while "answering questions using information from accessible Knowledge articles." This capability leverages Agentforce's ability to integrate Knowledge articles into the agent's responses, enabling the representative to ask questions (e.g., "How do I cancel a flight?") and receive AI-generated answers grounded in approved Knowledge content. Simultaneously, the agent can trigger actions (e.g., a Flow to update the custom object) to perform the cancellations, meeting all requirements efficiently.

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

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

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

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

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

Thank You for Trying Our Product
Special 16 USD Discount Coupon: NSZUBG3X

Email: support@examsempire.com

**Check our Customer Testimonials and ratings
available on every product page.**

Visit our website.

<https://examsempire.com/>