

---

**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.

Reference:

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>)

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.

Reference:

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

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

### 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.

Reference:

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>)

---

**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.