Vendor: Salesforce

Exam Code: AI-201

Exam Name: Salesforce Certified Agentforce Specialist

Version: DEMO

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.

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.

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.

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.

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.

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.

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.

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.

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.

QUESTION 11
Universal Containers wants to reduce overall customer support handling time by minimizing the time spent typing routine answers for common questions in-chat, and reducing the post-chat analysis by suggesting values for case fields. Which combination of Agentforce for Service features enables this effort?

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

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

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

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

QUESTION 12
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:
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.

QUESTION 13
When creating a custom retriever in Einstein Studio, which step is considered essential?

A. Select the search index, specify the associated data model object (DMO) and data space, and optionally define filters to narrow search results.
B. Define the output configuration by specifying the maximum number of results to return, and map the output fields that will ground the prompt.
C. Configure the search index, choose vector or hybrid search, choose the fields for filtering, the data space and model, then define the ranking method.

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

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

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

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

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

QUESTION 14
When configuring a prompt template, an Agentforce Specialist previews the results of the prompt template they've written. They see two distinct text outputs: Resolution and Response. Which information does the Resolution text provide?

A. It shows the full text that is sent to the Trust Layer.
B. It shows the response from the LLM based on the sample record.
C. It shows which sensitive data is masked before it is sent to the LLM.

Answer: B
Explanation:
In Salesforce Agentforce, when previewing a prompt template, the interface displays two outputs: Resolution and Response. These terms relate to how the prompt is processed and evaluated, particularly in the context of the Einstein Trust Layer, which ensures AI safety, compliance, and auditability. The Resolution text specifically refers to the full text that is sent to the Trust Layer for processing, monitoring, and governance (Option A). This includes the constructed prompt (with grounding data, instructions, and variables) as it's submitted to the large language model (LLM), along with any Trust Layer interventions (e.g., masking, filtering) applied before or after LLM

processing. It's a comprehensive view of the input/output flow that the Trust Layer captures for auditing and compliance purposes.

QUESTION 15
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:
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.

QUESTION 16
Universal Containers (UC) wants to ensure the effectiveness, reliability, and trust of its agents prior to deploying them in production. UC would like to efficiently test a large and repeatable number of utterances. What should the Agentforce Specialist recommend?

A. Leverage the Agent Large Language Model (LLM) UI and test UC's agents with different utterances prior to activating the agent.
B. Deploy the agent in a QA sandbox environment and review the Utterance Analysis reports to review effectiveness.
C. Create a CSV file with UC's test cases in Agentforce Testing Center using the testing template.

Answer: C
Explanation:
The goal of Universal Containers (UC) is to test its Agentforce agents for effectiveness, reliability, and trust before production deployment, with a focus on efficiently handling a large and repeatable number of utterances.

Thank You for Trying Our Product

Lead2pass Certification Exam Features:

★ More than **99,900** Satisfied Customers Worldwide.

★ Average **99.9%** Success Rate.

★ **Free Update** to match latest and real exam scenarios.

★ **Instant Download** Access! No Setup required.

★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.

★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.

★ **100%** Guaranteed Success or **100%** Money Back Guarantee.

★ **Fast**, helpful support **24x7**.

View list of all certification exams: http://www.lead2pass.com/all-products.html

10% Discount Coupon Code: **ASTR14**