

Updated and User Friendly Actual4Labs Agentforce-Specialist Exam PDF Questions File

Thus, Option B is the correct method to gather the necessary data efficiently and accurately.
 Reference: Salesforce Agentforce Documentation: "Integrate Flows with Prompt Templates" (Salesforce Help)
https://help.salesforce.com/articles/view?id=af-agentforce_flow_prompt_integration.htm&type=5
 Trailhead: "Build Flows for Agentforce"
<https://trailhead.salesforce.com/content/learn/modules/flows-for-agentforce/>

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:
 Comprehensive and Detailed In-Depth
 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.
 Reference: Salesforce Agentforce Documentation: "Agent Analytics Overview" (Salesforce Help)
https://help.salesforce.com/articles/view?id=af-agentforce_analytics.htm&type=5
 Trailhead: "Agentforce for Admins"
<https://trailhead.salesforce.com/content/learn/modules/agentforce-for-admins/>

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

What's more, part of that Actual4Labs Agentforce-Specialist dumps now are free: <https://drive.google.com/open?id=1GPYBNqnpKVuMCz9vfDOfhg79Tazm7PpF>

After years of unremitting efforts, our Agentforce-Specialist exam materials and services have received recognition and praises by the vast number of customers. An increasing number of candidates choose our Agentforce-Specialist study braindumps as their exam plan utility. There are a lot of advantages about our Agentforce-Specialist training guide. Not only our Agentforce-Specialist learning questions are always the latest and valid, but also the prices of the different versions are quite favourable.

Salesforce Agentforce-Specialist Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none"> Development Lifecycle: This area addresses testing agents in Testing Center, deploying from sandbox to production, and managing agent adoption and monitoring.
Topic 2	<ul style="list-style-type: none"> AI Agents: This domain covers configuring agent behavior, understanding the reasoning engine, selecting topics and actions for agent types, managing Agent User security, choosing appropriate agent types, and connecting agents to various channels.

Topic 3	<ul style="list-style-type: none"> • Multi-Agent Interoperability: This domain explains Model Context Protocol (MCP), agent-to-agent communication, and when to use Agent API for system interactions.
Topic 4	<ul style="list-style-type: none"> • Data Cloud for Agentforce: This domain covers Agentforce Data Library types, improving responses with unstructured data through chunking and indexing, understanding retrievers, and selecting keyword, vector, or hybrid search types.
Topic 5	<ul style="list-style-type: none"> • Prompt Engineering: This section focuses on using Prompt Builder, managing user roles, creating prompt templates with field generation and flex types, selecting grounding techniques, and applying best practices for effective prompts.

>> **Agentforce-Specialist Best Preparation Materials** <<

Dumps Agentforce-Specialist Download | Certification Agentforce-Specialist Cost

As the rapid development of the world economy and intense competition in the international, the leading status of knowledge-based economy is established progressively. A lot of people are in pursuit of a good job, a Agentforce-Specialist certification, and a higher standard of life. You just need little time to download and install it after you purchase, then you just need spend about 20~30 hours to learn it. We are glad that you are going to spare your precious time to have a look to our Agentforce-Specialist Exam Guide.

Salesforce Certified Agentforce Specialist (AI-201) Sample Questions (Q26-Q31):

NEW QUESTION # 26

Universal Containers has grounded a prompt template with a related list. During user acceptance testing (UAT), users are not getting the correct responses. What is causing this issue?

- **A. The related list is not on the parent object's page layout.**
- B. The related list is Read Only.
- C. The related list prompt template option is not enabled.

Answer: A

Explanation:

UC has grounded a prompt template with a related list, but the responses are incorrect during UAT.

Grounding with related lists in Agentforce allows the AI to access data from child records linked to a parent object. Let's analyze the options.

* Option A: The related list is Read Only. Read-only status (e.g., via field-level security or sharing rules) might limit user edits, but it doesn't inherently prevent the AI from accessing related list data for grounding, as long as the running user (or system context) has read access. This is unlikely to cause incorrect responses and is not a primary consideration, making it incorrect.

* Option B: The related list prompt template option is not enabled. There's no specific "related list prompt template option" toggle in Prompt Builder. When grounding with a Record Snapshot or Flex template, related lists are included if properly configured (e.g., via object relationships). This option seems to be a misphrasing and doesn't align with documented settings, making it incorrect.

* Option C: The related list is not on the parent object's page layout. In Agentforce, grounding with related lists relies on the related list being defined and accessible in the parent object's metadata, often tied to its presence on the page layout. If the related list isn't on the layout, the AI might not recognize or retrieve its data correctly, leading to incomplete or incorrect responses. Salesforce documentation notes that related list data availability can depend on layout configuration, making this a plausible and common issue during UAT, and thus the correct answer.

Why Option C is Correct:

The absence of the related list from the parent object's page layout can disrupt data retrieval for grounding, leading to incorrect AI responses. This is a known configuration consideration in Agentforce setup and testing, as per official guidance.

References:

Salesforce Agentforce Documentation: Grounding with Related Lists - Notes dependency on page layout configuration.

Trailhead: Ground Your Agentforce Prompts - Highlights related list setup for accurate grounding.

Salesforce Help: Troubleshoot Prompt Responses - Lists layout issues as a common grounding problem.

NEW QUESTION # 27

Universal Containers (UC) wants to enable its sales team to get insights into product and competitor names mentioned during calls. How should UC meet this requirement?

- A. Enable Einstein Conversation Insights, connect a recording provider, assign permission sets, and customize insights with up to 25 products.
- B. Enable Einstein Conversation Insights, assign permission sets, define recording managers, and customize insights with up to 50 competitor names.
- C. Enable Einstein Conversation Insights, enable sales recording, assign permission sets, and customize insights with up to 50 products.

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation: UC wants insights into product and competitor mentions during sales calls, leveraging Einstein Conversation Insights. Let's evaluate the options.

* Option A: Enable Einstein Conversation Insights, connect a recording provider, assign permission sets, and customize insights with up to 25 products. Einstein Conversation Insights analyzes call recordings to identify keywords like product and competitor names. Setup requires enabling the feature, connecting an external recording provider (e.g., Zoom, Gong), assigning permission sets (e.g., Einstein Conversation Insights User), and customizing insights by defining up to 25 products or competitors to track. Salesforce documentation confirms the 25-item limit for custom keywords, making this the correct, precise answer aligning with UC's needs.

* Option B: Enable Einstein Conversation Insights, assign permission sets, define recording managers, and customize insights with up to 50 competitor names. There's no "recording managers" role in Einstein Conversation Insights setup-integration is with a provider, not a manager designation.

The limit is 25 keywords (not 50), and the option omits the critical step of connecting a provider, making it incorrect.

* Option C: Enable Einstein Conversation Insights, enable sales recording, assign permission sets, and customize insights with up to 50 products. "Enable sales recording" is vague- Conversation Insights relies on external providers, not a native Salesforce recording feature. The keyword limit is 25, not 50, making this incorrect despite being closer than B.

Why Option A is Correct: Option A accurately reflects the setup process and limits for Einstein Conversation Insights, meeting UC's requirement per Salesforce documentation.

References:

* Salesforce Help: Set Up Einstein Conversation Insights- Details provider connection and 25-keyword limit.

* Trailhead: Einstein Conversation Insights Basics- Covers permissions and customization.

* Salesforce Agentforce Documentation: Sales Features- Confirms integration steps.

NEW QUESTION # 28

A data scientist needs to view and manage models in Einstein Studio, and also needs to create prompt templates in Prompt Builder. Which permission sets should an Agentforce Specialist assign to the data scientist?

- A. Prompt Template User and Data Cloud Admin
- B. Prompt Template Manager and Prompt Template User
- C. Data Cloud Admin and Prompt Template Manager

Answer: C

Explanation:

Comprehensive and Detailed In-Depth Explanation:

The data scientist requires permissions for Einstein Studio (model management) and Prompt Builder (template creation). Note: "Einstein Studio" may be a misnomer for Data Cloud's model management or a related tool, but we'll interpret based on context. Let's evaluate.

* Option A: Prompt Template Manager and Prompt Template User There's no distinct "Prompt Template Manager" or "Prompt Template User" permission set in Salesforce- Prompt Builder access is typically via "Einstein Generative AI User" or similar. This option lacks coverage for Einstein Studio /Data Cloud, making it incorrect.

* Option B: Data Cloud Admin and Prompt Template Manager The "Data Cloud Admin" permission set grants access to manage models in Data Cloud (assumed as Einstein Studio's context), including viewing and editing AI models. "Prompt Template Manager" isn't a real set, but Prompt Builder creation is covered by "Einstein Generative AI Admin" or similar admin-level access (assumed intent).

This combination approximates the needs, making it the closest correct answer despite naming ambiguity.

* Option C: Prompt Template User and Data Cloud Admin "Prompt Template User" isn't a standard set, and user-level access (e.g., Einstein Generative AI User) typically allows execution, not creation.

The data scientist needs to create templates, so this lacks sufficient Prompt Builder rights, making it incorrect.

Why Option B is Correct (with Caveat):

"Data Cloud Admin" covers model management in Data Cloud (likely intended as Einstein Studio), and

"Prompt Template Manager" is interpreted as admin-level Prompt Builder access (e.g., Einstein Generative AI Admin). Despite naming inconsistencies, this fits the requirements per Salesforce permissions structure.

References:

Salesforce Data Cloud Documentation: Permissions- Details Data Cloud Admin for models.

Trailhead: Set Up Einstein Generative AI- Covers Prompt Builder admin access.

Salesforce Help: Agentforce Permission Sets- Aligns with admin-level needs.

NEW QUESTION # 29

An administrator is responsible for ensuring the security and reliability of Universal Containers' (UC) CRM data. UC needs enhanced data protection and up-to-date AI capabilities. UC also needs to include relevant information from a Salesforce record to be merged with the prompt.

Which feature in the Einstein Trust Layer best supports UC's need?

- A. Dynamic grounding with secure data retrieval
- B. Zero-data retention policy
- C. Data masking

Answer: A

Explanation:

Dynamic grounding with secure data retrieval is a key feature in Salesforce's Einstein Trust Layer, which provides enhanced data protection and ensures that AI-generated outputs are both accurate and securely sourced. This feature allows relevant Salesforce data to be merged into the AI-generated responses, ensuring that the AI outputs are contextually aware and aligned with real-time CRM data.

Dynamic grounding means that AI models are dynamically retrieving relevant information from Salesforce records (such as customer records, case data, or custom object data) in a secure manner. This ensures that any sensitive data is protected during AI processing and that the AI model's outputs are trustworthy and reliable for business use.

The other options are less aligned with the requirement:

* Data masking refers to obscuring sensitive data for privacy purposes and is not related to merging Salesforce records into prompts.

* Zero-data retention policy ensures that AI processes do not store any user data after processing, but this does not address the need to merge Salesforce record information into a prompt.

References:

* Salesforce Developer Documentation on Einstein Trust Layer

* Salesforce Security Documentation for AI and Data Privacy

NEW QUESTION # 30

Universal Containers' agent must always look up the customer's account tier and open cases from Salesforce before deciding how to respond.

Based on Agent Script flow of control, what is true about executing deterministic actions at the very start of a subagent?

- A. Actions can only be guaranteed to run by placing them in the config block.
- B. The first instruction in reasoning instructions always runs before the large language model (LLM) is invoked.
- C. Only before_reasoning can guarantee the large language model (LLM) is invoked before an action runs.

Answer: B

Explanation:

The correct answer is C. In Agent Script, reasoning instructions is processed before the final assembled prompt is sent to the LLM. Salesforce explains that procedural logic inside reasoning instructions executes top-to-bottom, prompt text is accumulated, and then the assembled prompt is sent to the LLM. Therefore, if the first instruction deterministically runs the account-tier and open-case lookup actions, those values can be collected before the LLM decides how to respond. Option A is wrong because the config block defines agent-level metadata, not guaranteed action execution. Option B is incorrectly worded and reverses the sequencing issue. For deterministic pre-response data gathering, procedural instructions at the start of reasoning instructions are the correct pattern.

