Agentforce-Specialist Dumps

Salesforce Certified Agentforce Specialist

https://www.certleader.com/Agentforce-Specialist-dumps.html

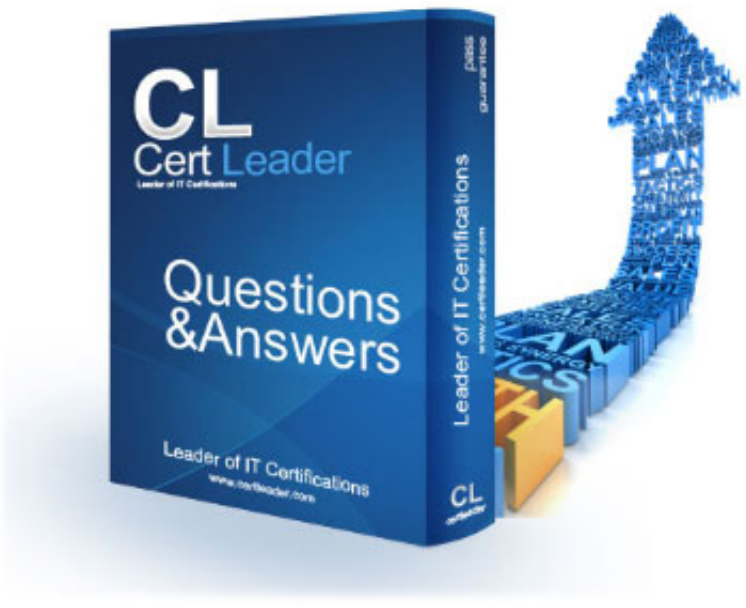

NEW QUESTION 1
Universal Containers (UC) wants to use Flow to bring data from unified Data Cloud objects to prompt templates.
Which type of flow should UC use?

A. Data Cloud-triggered flow
B. Template-triggered prompt flow
C. Unified-object linking flow

Answer: B

Explanation:
In this scenario, Universal Containers wants to bring data from unified Data Cloud objects into prompt templates, and the best way to do that is through a Data Cloud-triggered flow. This type of flow is specifically designed to trigger actions based on data changes within Salesforce Data Cloud objects.
Data Cloud-triggered flows can listen for changes in the unified data model and automatically bring relevant data into the system, making it available for prompt templates. This ensures that the data is both real-time and up-to-date when used in generative AI contexts.
For more detailed guidance, refer to Salesforce documentation on Data Cloud-triggered flows and Data Cloud integrations with generative AI solutions.

NEW QUESTION 2
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 3
An Agentforce Agent has been developed with multiple topics and Agent Actions that use flows and Apex. Which options are available for deploying these to production?

A. Deploy the flows and Apex using normal deployment tools and manually create theagent-related items in production.
B. Use only change sets because the Salesforce CLI does not currently support the deployment of agent-related metadata.
C. Deploy flows, Apex, and all agent-related items using either change sets or the Salesforce CLI/Metadata API.

Answer: C

Explanation:
Why is "Deploy flows, Apex, and all agent-related items using either change sets or the Salesforce CLI/Metadata API" the correct answer?
When deploying an Agentforce Agent with multiple topics and Agent Actions that use flows and Apex, a complete deployment solution is required. Change sets and the Salesforce CLI/Metadata API support the deployment of flows, Apex code, and agent- related metadata.
Key Considerations for Agentforce Deployments:
? Supports Deployment of All Required Components
? Agentforce Metadata Can Be Deployed Using Standard Tools
? Ensures a Complete Migration Without Manual Configuration
Why Not the Other Options?
* A. Deploy the flows and Apex using normal deployment tools and manually create the agent-related items in production.
? Incorrect because manually creating agent-related items in production introduces
risk and inconsistency.
? This approach is error-prone and time-consuming, especially for large Agentforce deployments.
* B. Use only change sets because the Salesforce CLI does not currently support the deployment of agent-related metadata.
? Incorrect because Salesforce CLI and Metadata API fully support Agentforce deployments.
? Change sets are useful but limited in large-scale, automated deployments.
Agentforce Specialist References
? Salesforce AI Specialist Material confirms that Agentforce metadata (flows, actions, and topics) can be deployed using Change Sets or the Metadata API.

NEW QUESTION 4
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 5
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)
? Trailhead: "Agentforce for Admins" (https://trailhead.salesforce.com/content/learn/modules/agentforce-for-admins)

NEW QUESTION 6
Universal Containers (UC) is Implementing Service AI Grounding to enhance its customer service operations. UC wants to ensure that its AI- generated responses are grounded in the most relevant data sources. The team needs to configure the system to include all supported objects for grounding. Which objects should UC select to configure Service AI Grounding?

A. Case, Knowledge, and Case Notes
B. Case and Knowledge
C. Case, Case Emails, and Knowledge

Answer: B

Explanation:
Universal Containers (UC) is implementing Service AI Grounding to enhance its customer service operations. They aim to ensure that AI-generated responses are grounded in the most relevant data sources and need to configure the system to include all supported objects for grounding.
Supported Objects for Service AI Grounding:
? Case
? Knowledge
? Case Object:
? Knowledge Object:
? Exclusion of Other Objects:
Why Options A and C are Incorrect:
? Option A (Case, Knowledge, and Case Notes):
? Option C (Case, Case Emails, and Knowledge):
References:
? Salesforce Agentforce Specialist Documentation - Service AI Grounding Configuration: Details the objects supported for grounding AI responses in Service Cloud.
? Salesforce Help - Implementing Service AI Grounding: Provides guidance on setting up grounding with Case and Knowledge objects.
? Salesforce Trailhead - Enhance Service with AI Grounding: Offers an interactive learning path on using AI grounding in service scenarios.

NEW QUESTION 7
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 8
Which configuration must An Agentforce complete for users to access generative AI- enabled fields in the Salesforce mobile app?

A. Enable Mobile Generative AI.
B. Enable Mobile Prompt Responses.
C. Enable Dynamic Forms on Mobile.

Answer: A

Explanation:
? Context of the Question
? Why Dynamic Forms on Mobile?
? Conclusion
Salesforce Agentforce Specialist References & Documents
? Salesforce Documentation: Dynamic Forms OverviewExplains how to enable Dynamic Forms for both desktop and mobile UIs, allowing newly added fields (including generative AI–enabled ones) to display in the Salesforce Mobile App.
? Salesforce Agentforce Specialist Study GuideReiterates that to expose generative AI fields or components in mobile, you must configure dynamic forms and ensure compatibility on mobile layouts.

NEW QUESTION 9
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 10
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)
? Salesforce Data Cloud Documentation: "Vector Database for AI" (https://help.salesforce.com/s/articleView?id=sf.data_cloud_vector_database.htm& type=5)

NEW QUESTION 10
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 15
What is true of Agentforce Testing Center?

A. Running tests risks modifying CRM data in a production environment.
B. Running tests does not consume Einstein Requests.
C. Agentforce Testing Center can only be used in a production environment.

Answer: B

Explanation:
Comprehensive and Detailed In-Depth Explanation:The Agentforce Testing Center is a tool in Agentforce Studio for validating agent performance. Let??s evaluate the statements.
? Option A: Running tests risks modifying CRM data in a production environment.Agentforce Testing Center runs synthetic interactions in a controlled environment (e.g., sandbox or isolated test space) and doesn??t modify live CRM data. It??s designed for safe pre-deployment testing, making this incorrect.
? Option B: Running tests does not consume Einstein Requests.Einstein Requests are part of the usage quota for Einstein Generative AI features (e.g., prompt executions in production). Testing Center uses synthetic data to simulate interactions without invoking live AI calls that count against this quota. Salesforce documentation confirms tests don??t consume requests, making this the correct answer.
? Option C: Agentforce Testing Center can only be used in a production environment.Testing Center is available in both sandbox and production orgs, but it??s primarily used pre-deployment (e.g., in sandboxes) to validate agents safely. This restriction is false, making it incorrect.
Why Option B is Correct:Not consuming Einstein Requests is a key feature of Testing Center, allowing extensive testing without impacting quotas, as per Salesforce documentation.
References:
? Salesforce Agentforce Documentation: Testing Center > Overview – Confirms no request consumption.
? Trailhead: Test Your Agentforce Agents – Notes quota-free testing.
? Salesforce Help: Agentforce Testing – Details safe, isolated testing.

NEW QUESTION 20
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 24

An Agentforce is considering using a Field Generation prompt template type.
What should the Agentforce Specialist check before creating the Field Generation prompt to ensure it is possible for the field to be enabled for generative AI?

A. That the field chosen must be a rich text field with 255 characters or more.
B. That the org is set to API version 59 or higher
C. That the Lightning page layout where the field will reside has been upgraded to Dynamic Forms

Answer: B

Explanation:
Before creating a Field Generation prompt template, the Agentforce Specialist must ensure that the Salesforce org is set to API version 59 or higher. This version of the API introduces support for advanced generative AI features, such as enabling fields for generative AI outputs. This is a critical technical requirement for the Field Generation prompt template to function correctly.
? Option A (rich text field requirement) is not necessary for generative AI
functionality.
? Option C (Dynamic Forms) does not impact the ability of a field to be generative AI-enabled, although it might enhance the user interface.
For more information, refer to Salesforce documentation on API versioning and Field Generation templates.

NEW QUESTION 29
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 33
How does an Agent respond when it can??t understand the request or find any requested information?

A. With a preconfigured message, based on the action type.
B. With a general message asking the user to rephrase the request.
C. With a generated error message.

Answer: B

Explanation:
Comprehensive and Detailed In-Depth Explanation:Agentforce Agents are designed to handle situations where they cannot interpret a request or retrieve requested data gracefully. Let??s assess the options based on Agentforce behavior.
? Option A: With a preconfigured message, based on the action type.While Agentforce allows customization of responses, there??s no specific mechanism tying preconfigured messages to action types for unhandled requests. Fallback responses are more general, not action-specific, making this incorrect.
? Option B: With a general message asking the user to rephrase the request.When an Agentforce Agent fails to understand a request or find information, it defaults to a general fallback response, typically asking the user to rephrase or clarify their input (e.g., ??I didn??t quite get that—could you try asking again???). This is configurable in Agent Builder but defaults to a user-friendly prompt to encourage retry, aligning with Salesforce??s focus on conversational UX. This is the correct answer per documentation.
? Option C: With a generated error message.Agentforce Agents prioritize user experience over technical error messages. While errors might log internally (e.g., in Event Logs), the user-facing response avoids jargon and focuses on retry prompts, making this incorrect.
Why Option B is Correct:The default behavior of asking users to rephrase aligns with Agentforce??s conversational design principles, ensuring a helpful response when comprehension fails, as noted in official resources.
References:
? Salesforce Agentforce Documentation: Agent Builder > Fallback Responses – Describes general retry messages.
? Trailhead: Build Agents with Agentforce – Covers handling ununderstood requests.
? Salesforce Help: Agentforce Interaction Design – Confirms user-friendly fallback behavior.

NEW QUESTION 38
Universal Containers (UC) configured a new PDF file ingestion in Data Cloud with all the required fields, and also created the mapping and the search Index. UC Is now setting up the retriever and notices a required fleld is missing.
How should UC resolve this?

A. Create a new custom Data Cloud object that includes the desired field.
B. Update the search index to include the desired field.
C. Modify the retriever's configuration to include the desired field..

Answer: B

Explanation:
Why is "Update the search index to include the desired field" the correct answer? When configuring a retriever in Data Cloud for PDF file ingestion, all necessary fields must be included in the search index. If a required field is missing, the correct action is to update the search index to ensure it is available for retrieval.
Key Considerations for Fixing Missing Fields in Data Cloud Retrievers:
? Search Index Controls Which Fields Are Searchable
? Ensures Complete and Accurate Data Retrieval
? Supports AI-Grounded Responses
Why Not the Other Options?
* A. Create a new custom Data Cloud object that includes the desired field.
? Incorrect because the issue is with indexing, not with Data Cloud object structure.
? The field already exists in Data Cloud; it just needs to be indexed.
* C. Modify the retriever's configuration to include the desired field.
? Incorrect because retriever configurations only define query rules; they do not modify the index itself.
? Updating the search index is the required step to ensure the field is retrievable.
Agentforce Specialist References
? Salesforce AI Specialist Material confirms that search indexing is required for retrievers to access specific fields in Data Cloud.

NEW QUESTION 40
What is an appropriate use case for leveraging Agentforce Sales Agent in a sales context?

A. Enable a sates team to use natural language to invoke defined sales tasks grounded in relevant data and be able to ensure company policies are applie
B. conversationally and in the now or work.
C. Enable a sales team by providing them with an interactive step-by-step guide based on business rules to ensure accurate data entry into Salesforce and help close deals fatter.
D. Instantly review and read incoming messages or emails that are then logged to the correct opportunity, contact, and account records to provide a full view of customer interactions and communications.

Answer: A

Explanation:
Agentforce Sales Agent is designed to let sales teams perform tasks via natural language commands, leveraging Salesforce data while adhering to policies. For example, agents can ask the AI to "update the opportunity stage to Closed Won" or "generate a quote," with the system enforcing validations and data security. This use case aligns with Salesforce??s vision of conversational AI streamlining workflows without compromising compliance.
? Step-by-step guides (B) are typically handled by tools like Dynamic Forms or
Guided Selling, not Agentforce.
? Logging messages/emails (C) is managed by Email-to-Case or Service Cloud, not a sales-specific AI agent.
Reference:
Salesforce Help Article: Agentforce for Sales ("Use Cases and Capabilities" section).
Einstein Agentforce Specialist Trailhead: "Sales Automation with Agentforce" (Natural Language Task Execution).

NEW QUESTION 45
What should Universal Containers consider when deploying an Agentforce Service Agent with multiple topics and Agent Actions to production?

A. Deploy agent components without a test run in staging, relying on production data for reliable result
B. Sandbox configuration alone ensures seamless production deployment.
C. Ensure all dependencies are included, Apex classes meet 75% test coverage, and configuration settings are aligned with productio
D. Plan for version management and post- deployment activation.
E. Deploy flows or Apex after agents, topics, and Agent Actions to avoid deployment failures and potential production agent issues requiring complete redeployment.

Answer: B

Explanation:
Comprehensive and Detailed In-Depth Explanation:UC is deploying an Agentforce Service Agent with multiple topics and actions to production. Let??s assess deployment considerations.
? Option A: Deploy agent components without a test run in staging, relying on production data for reliable results. Sandbox configuration alone ensures seamless production deployment.Skipping staging tests is risky and against best practices. Sandbox configuration doesn??t guarantee production success without validation, making this incorrect.
? Option B: Ensure all dependencies are included, Apex classes meet 75% test
coverage, and configuration settings are aligned with production. Plan for version management and post-deployment activation.This is a comprehensive approach: dependencies (e.g., flows, Apex) must be deployed, Apex requires 75% coverage, and production settings (e.g., permissions, channels) must align. Version management tracks changes, and post-deployment activation ensures controlled rollout. This aligns with Salesforce deployment best practices for Agentforce, making it the correct answer.
? Option C: Deploy flows or Apex after agents, topics, and Agent Actions to avoid deployment failures and potential production agent issues requiring complete redeployment.Deploying components separately risks failures (e.g., actions needing flows failing). All components should deploy together for consistency, making this incorrect.
Why Option B is Correct:Option B covers all critical deployment considerations for a robust Agentforce rollout, as per Salesforce guidelines.
References:
? Salesforce Agentforce Documentation: Deploy Agents to Production – Lists dependencies and coverage.
? Trailhead: Deploy Agentforce Agents – Emphasizes testing and activation planning.
? Salesforce Help: Agentforce Deployment Best Practices – Confirms comprehensive approach.

NEW QUESTION 49
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 53
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 54
Universal Containers wants to leverage the Record Snapshots grounding feature in a prompt template. What preparations are required?

A. Configure page layout of the master record type.
B. Create a field set for all the fields to be grounded.
C. Enable and configure dynamic form for the object.

Answer: B

Explanation:
Comprehensive and Detailed In-Depth Explanation:Universal Containers (UC) aims to use Record Snapshots grounding in a prompt template to provide context from a specific record. Let??s evaluate the preparation steps.
? Option A: Configure page layout of the master record type.While page layouts define field visibility for users, Record Snapshots grounding relies on field accessibility at the object level, not the layout. The AI accesses data based on permissions and configuration, not layout alone, making this insufficient and incorrect.
? Option B: Create a field set for all the fields to be grounded.Record Snapshots in Prompt Builder allow grounding with fields from a record, but you must specify which fields to include. Creating a field set is a recommended preparation step—it groups the fields (e.g., from the object) to be passed to the prompt template, ensuring the AI has the right data. This is a documented best practice for controlling snapshot scope, making it the correct answer.
? Option C: Enable and configure dynamic form for the object.Dynamic Forms enhance UI flexibility but aren??t required for Record Snapshots grounding. The feature pulls data directly from the object, not the form configuration, making this irrelevant and incorrect.
Why Option B is Correct:Creating a field set ensures the prompt template uses the intended fields for grounding, a key preparation step per Salesforce documentation.
References:
? Salesforce Agentforce Documentation: Prompt Builder > Record Snapshots – Recommends field sets for grounding.
? Trailhead: Ground Your Agentforce Prompts – Details field set preparation.
? Salesforce Help: Set Up Record Snapshots – Confirms field set usage.

NEW QUESTION 57
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 61
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: C

Explanation:
Comprehensive and Detailed In-Depth Explanation:Salesforce Agentforce leverages generative AI to enhance service agent efficiency, particularly through capabilities that generate recommended replies. In this scenario, Universal Containers
aims to reduce handling time by providing replies based on existing Knowledge articles, which are a core component of Salesforce Knowledge. The Knowledge Replies capability is specifically designed for this purpose—it uses generative AI to analyze Knowledge articles, match them to the context of a customer inquiry (e.g., a case or chat), and suggest relevant, pre-formulated responses for service agents to use or adapt. This aligns directly with UC??s goal of leveraging existing content to streamline agent workflows.
? Option A (Service Replies): While "Service Replies" might sound plausible, it is not a specific, documented capability in Agentforce. It appears to be a generic distractor and does not tie directly to Knowledge articles.
? Option B (Case Replies): "Case Replies" is not a recognized AI capability in Agentforce either. While replies can be generated for cases, the focus here is on Knowledge article integration, which points to Knowledge Replies.
? Option C (Knowledge Replies): This is the correct capability, as it explicitly connects generative AI with Knowledge articles to produce recommended replies, reducing agent effort and handling time.
Training service agents on Knowledge Replies ensures they can effectively use AI- suggested responses, review them for accuracy, and integrate them into their workflows, fulfilling UC??s objective.
References:
? Salesforce Agentforce Documentation: "Knowledge Replies for Service Agents" (Salesforce Help:
https://help.salesforce.com/s/articleView?id=sf.agentforce_knowledge_replies.htm
&type=5)
? Trailhead: "Agentforce for Service" module (https://trailhead.salesforce.com/content/learn/modules/agentforce-for-service)

NEW QUESTION 64
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 68
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 71
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 75
What is the primary function of the reasoning engine in Agentforce?

A. Identifying agent topics and actions to respond to user utterances
B. Offering real-time natural language response during conversations
C. Generating record queries based on conversation history

Answer: A

Explanation:
Why is "Identifying agent topics and actions to respond to user utterances" the correct answer?
In Agentforce, the reasoning engine plays a critical role in interpreting user queries
and determining the appropriate agent response.
Key Functions of the Reasoning Engine in Agentforce:
? Analyzing User Intent
? Selecting the Appropriate Agent Action
? Ensuring AI Accuracy and Context Awareness
Why Not the Other Options?
* B. Offering real-time natural language response during conversations.
? Incorrect because real-time natural language processing (NLP) is handled by the large language model (LLM), not the reasoning engine.
? The reasoning engine focuses on action selection, not linguistic processing.
* C. Generating record queries based on conversation history.
? Incorrect because query generation is handled by Copilot Actions (e.g., Query Records), not the reasoning engine.
? The reasoning engine decides which query should be run, but does not generate queries itself.
Agentforce Specialist References
? Salesforce AI Specialist Material explains that the reasoning engine identifies topics and selects agent actions.
? Salesforce Instructions for the Certification confirm that the reasoning engine determines AI workflow execution.

NEW QUESTION 79
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&ty pe=5)
? Trailhead: "Agentforce Basics" module (https://trailhead.salesforce.com/content/learn/modules/agentforce-basics)

NEW QUESTION 84
What is the main benefit of using a Knowledge article in an Agentforce Data Library?

A. Only the retriever for Knowledge articles allows for agents to access Knowledge from both inside the platform and on a customer's website.
B. It provides a structured, searchable repository of approved documents so the agent can retrieve reliable information for each inquiry..
C. The retriever for Knowledge articles has better accuracy and performance than the default retriever.

Answer: B

Explanation:
Why is "A structured, searchable repository of approved documents" the correct answer?
Using a Knowledge Article in an Agentforce Data Library ensures that agents can quickly access reliable and pre-approved information during customer interactions.
Key Benefits of Knowledge Articles in an Agentforce Data Library:
? Ensures Information Accuracy and Consistency
? Improves Searchability and AI-Grounded Responses
? Enhances Customer Support and Agent Productivity
Why Not the Other Options?
* A. Only the retriever for Knowledge articles allows for agents to access Knowledge from both inside the platform and on a customer's website.
? Incorrect because other retrievers (e.g., standard Salesforce Data Cloud
retrievers) can also provide knowledge access.
? Knowledge articles can be accessed via multiple retrieval mechanisms, not just one specific retriever.
* C. The retriever for Knowledge articles has better accuracy and performance than the default retriever.
? Incorrect because retriever accuracy depends on indexing and search
configuration, not the article type.
? The default retriever works just as efficiently when properly configured.
Agentforce Specialist References
? Salesforce AI Specialist Material confirms that Knowledge articles provide structured, searchable, and approved information for AI-grounded responses.

NEW QUESTION 87
An Agentforce implements Einstein Sales Emails for a sales team. The team wants to send personalized follow-up emails to leads based on their interactions and data stored in
Salesforce. The Agentforce Specialist needs to configure the system to use the most accurate and up-to-date information for email generation.
Which grounding technique should the Agentforce Specialist use?

A. Ground with Apex Merge Fields
B. Ground with Record Merge Fields
C. Automatic grounding using Draft with Einstein feature

Answer: C

Explanation:
For Einstein Sales Emails to generate personalized follow-up emails, it is crucial to ground the email content with the most up-to-date and accurate information. Grounding refers to connecting the AI model with real-time data. The most appropriate technique in this case is Ground with Record Merge Fields. This method ensures that the content in the emails pulls dynamic and accurate data directly from Salesforce records, such as lead or contact information, ensuring the follow-up is relevant and customized based on the specific record.
? Record Merge Fields ensure the generated emails are highly personalized using
data like lead name, company, or other Salesforce fields directly from the records.
? Apex Merge Fields are typically more suited for advanced, custom logic-driven scenarios but are not the most straightforward for this use case.
? Automatic grounding using Draft with Einstein is a different feature where Einstein automatically drafts the email, but it does not specifically ground the content with record-specific data like Record Merge Fields.
References:
? Salesforce Einstein Sales Emails Documentation: https://help.salesforce.com/s/articleView?id=release- notes.rn_einstein_sales_emails.htm

NEW QUESTION 92
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 m 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 95
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_integratio n.htm&type=5)
? Trailhead: "Build Flows for Agentforce" (https://trailhead.salesforce.com/content/learn/modules/flows-for-agentforce)

NEW QUESTION 96
Universal Containers (UC) wants its AI agent to return responses quickly. UC needs to optimize the retriever's configuration to ensure minimal latency when grounding AI responses.
Which configuration aspect should UC prioritize?

A. Configure the retriever to operate in dynamic mode so that it modifies the search Index structure at runtime.
B. Ensure the retriever's filters are defined to limit the scope of each search efficiently.
C. Increase the recency bias setting for the retriever limiting scope to more recent data.

Answer: B

Explanation:
Why is "Ensure the retriever's filters are defined to limit the scope of each search efficiently" the correct answer?
In Agentforce, when optimizing a retriever's configuration to ensure minimal latency in AI-generated responses, the most effective approach is narrowing the scope of searches
by applying specific filters.
Key Considerations for Optimizing Retrievers in Agentforce:
? Defining Effective Filters
? Reducing Query Complexity
? Optimizing the Data Indexing Process
Why Not the Other Options?
* A. Configure the retriever to operate in dynamic mode so that it modifies the search index structure at runtime.
? Incorrect because modifying the search index at runtime increases latency rather
than reducing it.
? Index modifications require restructuring large datasets, which can slow down AI- generated responses.
* C. Increase the recency bias setting for the retriever, limiting scope to more recent data.
? Incorrect because increasing recency bias only prioritizes recent records but does
not necessarily improve overall retrieval speed.
? While it affects relevance, it does not directly address latency issues.

Agentforce Specialist References
? Salesforce AI Specialist Material confirms that retriever efficiency depends on well- defined filtering mechanisms to minimize latency.
? Salesforce Instructions for Certification highlight retriever optimization strategies to improve search response times.

NEW QUESTION 100
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 105
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 setThis 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 setThis 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 setThis 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 PermissionsThe 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 ServiceThis 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 107
Universal Containers (UC) noticed an increase in customer contract cancellations in the last few months. UC is seeking ways to address this issue by implementing a proactive outreach program to
customers before they cancel their contracts and is asking the Salesforce team to provide suggestions.
Which use case functionality of Model Builder aligns with UC's request?

A. Product recommendation prediction
B. Customer churn prediction
C. Contract Renewal Date prediction

Answer: B

Explanation:
Customer churn prediction is the best use case for Model Builder in addressing Universal Containers' concerns about increasing customer contract cancellations. By implementing a model that predicts customer churn, UC can proactively identify customers who are at risk of canceling and take action to retain them before they decide to terminate their contracts. This functionality allows the business to forecast churn probability based on historical data and initiate timely outreach programs.
? Option B is correct because customer churn prediction aligns with UC's need to
reduce cancellations through proactive measures.
? Option A (product recommendation prediction) is unrelated to contract cancellations.
? Option C (contract renewal date prediction) addresses timing but does not focus on predicting potential cancellations.
References:
? Salesforce Model Builder Use Case Overview: https://help.salesforce.com/s/articleView?id=sf.model_builder_use_cases.htm

NEW QUESTION 110
Universal Container (UC) has effectively utilized prompt templates to update summary fields on Lightning record pages. An admin now wishes to incorporate similar functionality into UC's automation process using Flow.
How can the admin get a response from this prompt template from within a flow to use as part of UC's automation?

A. Invocable Apex
B. Flow Action
C. Einstein for Flow

Answer: C

Explanation:
* 1. Context of the Question
oUniversal Container (UC) has used prompt templates to update summary fields on record pages.
oNow, the admin wants to incorporate similar generative AI functionality within a Flow for automation purposes.
* 2. How to Call a Prompt Template Within a Flow
oFlow Action: Salesforce provides a standard way to invoke generative AI templates or prompts within a Flow step. From the Flow Builder, you can add an ??Action?? that references the prompt template you created in Prompt Builder.
oOther Options:
Invocable Apex: Possible fallback if there??s no out-of-the-box Flow Action available. However, Salesforce is releasing native Flow integration for AI prompts, making custom Apex less necessary.
Einstein for Flow: A broad label for Salesforce??s generative AI features within Flow. Under the hood, you typically use a ??Flow Action?? that points to your prompt.
* 3.Conclusion
oThe easiest out-of-the-box solution is to use a Flow Action referencing the prompt template. Hence, Option B is correct.
Salesforce Agentforce Specialist References & Documents
•Salesforce Trailhead: Use Prompt Templates in Flow
Demonstrates how to add an Action in Flow that calls a prompt template.
•Salesforce Documentation: Einstein GPT for Flow

NEW QUESTION 114
How is Data Cloud leveraged by the Answer Questions with Knowledge action in Agentforce?

A. Data Cloud is not required; the articles can be accessed directly from the CRM by the agent.
B. Data Cloud stores and manages the Indexed Knowledge articles.
C. Data Cloud provides the real-time data streams that update the Knowledge articles.

Answer: B

Explanation:
How Does Data Cloud Support "Answer Questions with Knowledge" in Agentforce? The Answer Questions with Knowledge action in Agentforce leverages Salesforce Data Cloud to store, manage, and index Knowledge articles used for AI-powered responses.
? Data Cloud as the Central Storage for Knowledge Articles
? Ensuring Up-to-Date Responses
? Enhancing AI-Driven Customer Service
Why Not the Other Options?
* A. Data Cloud is not required; the articles can be accessed directly from the CRM by the agent.
? Incorrect because Data Cloud is the primary system for storing and indexing
Knowledge articles.
? Without Data Cloud, Einstein AI cannot efficiently retrieve and rank articles dynamically.
* C. Data Cloud provides the real-time data streams that update the Knowledge articles.
? Incorrect because while Data Cloud stores and manages articles, real-time
updates are not its primary function.
? The Knowledge Management system within Salesforce handles article creation and updates.
Agentforce Specialist References
? Salesforce AI Specialist Material highlights that Data Cloud is the core storage system for AI-driven Knowledge management.
? Salesforce Instructions for Certification confirm the central role of Data Cloud in managing indexed Knowledge articles for AI-powered responses.

NEW QUESTION 118
Universal Containers (UC) needs to save agents time with AI-generated case summaries. UC has implemented the Work Summary feature.
What does Einstein consider when generating a summary?

A. Generation is grounded with conversation context, Knowledge articles, and cases.
B. Generation is grounded with existing conversation context only.
C. Generation is grounded with conversation context and Knowledge articles.

Answer: A

Explanation:
When generating a Work Summary, Einstein leverages multiple sources of information to provide a comprehensive and accurate case summary for agents.
? Conversation Context:
? Knowledge Articles:
? Cases:
? Option A is correct as it includes all three: conversation context, Knowledge articles, and cases.
? Option B is incorrect because it limits the grounding to conversation context only, excluding other critical elements.
? Option C is incorrect because it omits case data, which Einstein considers for more accurate and contextually rich summaries.
Reference:
"Einstein Work Summary and AI Case Management | Salesforce Trailhead" .

NEW QUESTION 121
In Model Playground, which hyperparameters of an existing Salesforce-enabled foundational model can An Agentforce change?
A. Temperature, Frequency Penalty, Presence Penalty
B. Temperature, Top-k sampling, Presence Penalty
C. Temperature, Frequency Penalty, Output Tokens

Answer: A

Explanation:
In Model Playground, An Agentforce working with a Salesforce-enabled foundational model has control over specific hyperparameters that can directly affect the behavior of the generative model:
? Temperature: Controls the randomness of predictions. A higher temperature leads
to more diverse outputs, while a lower temperature makes the model's responses more focused and deterministic.
? Frequency Penalty: Reduces the likelihood of the model repeating the same
phrases or outputs frequently.
? Presence Penalty: Encourages the model to introduce new topics in its responses, rather than sticking with familiar, previously mentioned content.
These hyperparameters are adjustable to fine-tune the model??s responses, ensuring that it meets the desired behavior and use case requirements. Salesforce documentation confirms that these three are the key tunable hyperparameters in the Model Playground. For more details, refer to Salesforce AI Model Playground guidance from Salesforce??s official documentation on foundational model adjustments.

NEW QUESTION 122
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 123
Universal Containers (UC) has a legacy system that needs to integrate with Salesforce. UC
wishes to create a digest of account action plans using the generative API feature. Which API service should UC use to meet this requirement?

A. REST API
B. Metadata API
C. SOAP API

Answer: A

Explanation:
To create a digest of account action plans using the generative API feature, Universal Containers should use the REST API. The REST API is ideal for integrating Salesforce with external systems and enabling interaction with Salesforce data, including generative capabilities like creating summaries or digests. It supports modern web standards and is suitable for flexible, lightweight interactions between Salesforce and legacy systems.
? Metadata API is used for retrieving and deploying metadata, not for data
operations like generating summaries.
? SOAP API is an older API used for integration but is less flexible compared to REST for this specific use case.
For more details, refer to Salesforce REST API documentation regarding using REST for data integration and generating content.

NEW QUESTION 124

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

NEW QUESTION 126
Universal Containers aims to streamline the sales team's daily tasks by using AI.
When considering these new workflows, which improvement requires the use of Prompt Builder?

A. Populate an AI-generated time-to close estimation to opportunities
B. Populate an AI generated summary field for sales contracts.
C. Populate an AI generated lead score for new leads.

Answer: B

Explanation:
Prompt Builder is explicitly required to createAI-generated summary fields via prompt templates. These fields use natural language instructions to extract or synthesize information (e.g., summarizing contract terms). Time-to-close estimations (A) and lead scores (C) are typically handled by predictive AI (e.g., Einstein Opportunity Scoring) or analytics tools, which do not require Prompt Builder.
Reference:
Salesforce Help Article: Create AI-Generated Fields with Prompt Builder ("Summary Field Generation" example).
Einstein GPT for Sales Guide: "Automating Contract Summaries."

NEW QUESTION 130
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

NEW QUESTION 135
Universal Containers is evaluating Einstein Generative AI features to improve the productivity of the service center operation.
Which features should the Agentforce Specialist recommend?

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

Answer: A

Explanation:
To improve the productivity of the service center, the Agentforce Specialist should recommend the Service Replies and Case Summaries features.
? Service Replies helps agents by automatically generating suggested responses to customer inquiries, reducing response time and improving efficiency.
? Case Summaries provide a quick overview of case details, allowing agents to get up to speed faster on customer issues.
? Work Summaries are not as relevant for direct customer service operations, and Sales Summaries are focused on sales processes, not service center productivity.
For more information, see Salesforce's Einstein Service Cloud documentation on the use of generative AI to assist customer service teams.

NEW QUESTION 137
Universal Containers (UC) wants to enable its sales team to use AI to suggest recommended products from its catalog. Which type of prompt template should UC use?

A. Record summary prompt template
B. Email generation prompt template
C. Flex prompt template

Answer: C

Explanation:
Comprehensive and Detailed In-Depth Explanation:UC needs an AI solution to suggest products from a catalog for its sales team. Let??s assess the prompt template types in Prompt Builder.
? Option A: Record summary prompt templateRecord summary templates generate concise summaries of records (e.g., Case, Opportunity). They??re not designed for product recommendations, which require dynamic logic beyond summarization, making this incorrect.
? Option B: Email generation prompt templateEmail generation templates craft emails (e.g., customer outreach). While they could mention products, they??re not optimized for standalone recommendations, making this incorrect.
? Option C: Flex prompt templateFlex prompt templates are versatile, allowing custom inputs (e.g., catalog data from objects or Data Cloud) and instructions (e.g., ??Suggest products based on customer preferences??). This flexibility suits UC??s need to recommend products dynamically, making it the correct answer.
Why Option C is Correct:Flex templates offer the customization needed to suggest products from a catalog, aligning with Salesforce??s guidance for tailored AI outputs.
References:
? Salesforce Agentforce Documentation: Prompt Builder > Flex Templates – Details dynamic use cases.
? Trailhead: Build Prompt Templates in Agentforce – Covers Flex for custom scenarios.
? Salesforce Help: Prompt Template Types – Confirms Flex versatility.

NEW QUESTION 138
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 143
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 145
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 149
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 150
Universal Containers (UC) users are complaining that agent answers are not satisfactory. The agent is using PDF files as a knowledge source.
How should UC troubleshoot this issue?

A. Analyze the data mapping between source fields and Data Cloud object fields.
B. Check that the agent has the PDF file field permission access for the data library.
C. Verify the retriever's filter criteria and data source connection.

Answer: C

Explanation:
Why is "Verify the retriever's filter criteria and data source connection" the correct answer?
If agent answers are not satisfactory when using PDF files as a knowledge source, the issue is likely caused by:
? Retriever misconfiguration
? Incorrect data source connection
? Parsing Issues with PDF Files
Why Not the Other Options?
* A. Analyze the data mapping between source fields and Data Cloud object fields.
? Incorrect because data mapping issues primarily affect structured CRM data, not PDF-based knowledge sources.
? The issue likely stems from retrieval settings, not field mapping.
* B. Check that the agent has the PDF file field permission access for the data library.
? Incorrect because permission access issues would prevent AI from accessing PDFs entirely rather than causing poor response quality.
? AI can still generate responses, even if they are inaccurate, which means the issue lies in retriever settings, not permissions.
Agentforce Specialist References
? Salesforce AI Specialist Material details how retriever filters and data sources impact AI-generated answers.
? Salesforce Certification Guide mentions the importance of verifying retriever configurations for accurate knowledge retrieval.

NEW QUESTION 154
Universal Containers has implemented an agent that answers questions based on Knowledge articles. Which topic and Agent Action will be shown in the Agent Builder?

A. General Q&A topic and Knowledge Article Answers action.
B. General CRM topic and Answers Questions with LLM Action.
C. General FAQ topic and Answers Questions with Knowledge Action.

Answer: C

Explanation:
Comprehensive and Detailed In-Depth Explanation:UC??s agent answers questions using Knowledge articles, configured in Agent Builder. Let??s identify the topic and action.

? Option A: General Q&A topic and Knowledge Article Answers action."General Q&A" is not a standard topic name in Agentforce, and "Knowledge Article Answers" isn??t a predefined action. This lacks specificity and doesn??t match documentation, making it incorrect.
? Option B: General CRM topic and Answers Questions with LLM Action."General CRM" isn??t a default topic, and "Answers Questions with LLM" suggests raw LLM responses, not Knowledge-grounded ones. This doesn??t align with the Knowledge focus, making it incorrect.
? Option C: General FAQ topic and Answers Questions with Knowledge Action.In Agent Builder, the "General FAQ" topic is a common default or starting point for question-answering agents. The "Answers Questions with Knowledge" action
(sometimes styled as "Answer with Knowledge") is a prebuilt action that retrieves and grounds responses with Knowledge articles. This matches UC??s implementation and is explicitly supported in documentation, making it the correct answer.
Why Option C is Correct:"General FAQ" and "Answers Questions with Knowledge" are the standard topic-action pair for Knowledge-based question answering in Agentforce, per Salesforce resources.
References:
? Salesforce Agentforce Documentation: Agent Builder > Actions – Lists "Answers Questions with Knowledge."
? Trailhead: Build Agents with Agentforce – Describes FAQ topics with Knowledge actions.
? Salesforce Help: Knowledge in Agentforce – Confirms this configuration.

NEW QUESTION 157
Universal Containers wants to use an external large language model (LLM) in Prompt Builder.
What should An Agentforce recommend?

A. Use Apex to connect to an external LLM and ground the prompt.
B. Use BYO-LLM functionality in Einstein Studio.
C. Use Flow and External Services to bring data from an external LLM.

Answer: B

Explanation:
Bring Your Own Large Language Model (BYO-LLM) functionality in Einstein Studio allows organizations to integrate and use external large language models (LLMs) within the Salesforce ecosystem. Universal Containers can leverage this feature to connect and ground prompts with external LLMs, allowing for custom AI model use cases and seamless integration with Salesforce data.
? Option B is the correct choice as Einstein Studio provides a built-in feature to work
with external models.
? Option A suggests using Apex, but BYO-LLM functionality offers a more streamlined solution.
? Option C focuses on Flow and External Services, which is more about data
integration and isn't ideal for working with LLMs.
References:
Salesforce Einstein Studio BYO-LLM Documentation: https://help.salesforce.com/s/articleView?id=sf.einstein_studio_llm.htm

NEW QUESTION 158
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:
Comprehensive and Detailed In-Depth 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.
References:
? 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)

NEW QUESTION 159
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:
Comprehensive and Detailed In-Depth 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.
? Option B: While "Einstein Reply Recommendations" is correct for the first part, "Case Summaries" generates a summary of the case rather than suggesting specific field values. Summaries are useful for documentation but don??t directly reduce post-chat field entry time.
? Option C: "Einstein Service Replies" is not a distinct, documented feature in Agentforce (possibly a distractor for Reply Recommendations), and "Work Summaries" applies more to summarizing work orders or broader tasks, not case field suggestions in a chat context.
? Option A: This combination precisely targets both in-chat efficiency (Reply Recommendations) and post-chat automation (Case Classification).
Thus, Option A is the correct answer for UC??s needs.
References:
? Salesforce Agentforce Documentation: "Einstein Reply Recommendations" (Salesforce Help: https://help.salesforce.com/s/articleView?id=sf.einstein_reply_recommendations.ht m&type=5)
? Salesforce Agentforce Documentation: "Case Classification" (Salesforce Help: https://help.salesforce.com/s/articleView?id=sf.case_classification.htm&type=5)
? Trailhead: "Agentforce for Service" (https://trailhead.salesforce.com/content/learn/modules/agentforce-for-service)

NEW QUESTION 160
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 CanvasPlan 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 SettingsAgent 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 LogsEvent 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 161
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 163
Universal Containers (UC) wants to make a sales proposal and directly use data from multiple unrelated objects (standard and custom) in a prompt template. How should UC accomplish this?

A. Create a prompt template passing in a special custom object that connects the records temporarily.
B. Create a prompt template-triggered flow to access the data from standard and custom objects.
C. Create a Flex template to add resources with standard and custom objects as inputs.
D. Use a Record Snapshot to combine data from unrelated objects into a single prompt.

Answer: C

Explanation:
Comprehensive and Detailed In-Depth Explanation:UC needs to incorporate data from multiple unrelated objects (standard and custom) into a prompt template for a sales proposal. Let??s evaluate the options based on Agentforce capabilities.
? Option A: Create a prompt template passing in a special custom object that connects the records temporarily.While a custom object could theoretically act as a junction to link unrelated records, this approach requires additional setup (e.g., creating the object, populating it with data via automation), and there??s no direct mechanism in Prompt Builder to "pass in" such an object to a prompt template without grounding or flow support. This is inefficient and not a native feature, making

it incorrect.
? Option B: Create a prompt template-triggered flow to access the data from standard and custom objects.There??s no such thing as a "prompt template- triggered flow" in Salesforce. Flows can invoke prompt templates (e.g., via the "Prompt Template" action), but the reverse—triggering a flow from a prompt template—is not a standard construct. While a flow could gather data from unrelated objects and pass it to a prompt, this option??s terminology is inaccurate, and it??s not the most direct solution, making it incorrect.
? Option C: Create a Flex template to add resources with standard and custom objects as inputs.In Agentforce??s Prompt Builder, a Flex template (short for Flexible Prompt Template) allows users to define dynamic inputs, including data from multiple Salesforce objects (standard or custom), even if they??re unrelated. Resources can be added to the template (e.g., via merge fields or Data Cloud queries), enabling the prompt to pull data directly from specified objects without requiring a junction object or complex flows. This is ideal for generating a sales proposal using disparate data sources and aligns with Salesforce??s documentation on Flex templates, making it the correct answer.
Why Option C is Correct:Flex templates are designed for scenarios requiring flexible data inputs, allowing UC to directly reference multiple unrelated objects in the prompt template. This simplifies the process and leverages Prompt Builder??s native capabilities, as outlined in Salesforce documentation.
References:
? Salesforce Agentforce Documentation: Prompt Builder > Flex Templates – Describes adding multiple object resources as inputs.
? Trailhead: Build Prompt Templates in Agentforce – Highlights Flex templates for dynamic data scenarios.
? Salesforce Help: Create Flexible Prompts – Confirms support for standard and custom object data.

NEW QUESTION 164
How does Secure Data Retrieval ensure that only authorized users can access necessary Salesforce data for dynamic grounding?

A. Retrieves Salesforce data based on the 'Run As" users permissions.
B. Retrieves Salesforce data based on the user??s permissions executing the prompt.
C. Retrieves Salesforces data based on the Prompt template's object permissions.

Answer: B

Explanation:
Secure Data Retrieval enforces Salesforce??s security model by dynamically
grounding data access in the permissions of the user executing the prompt. This ensures compliance with CRUD (Create, Read, Update, Delete) and FLS (Field-Level Security) settings, preventing unauthorized access to sensitive data. For example, if a user lacks access to a specific object or field, the AI model cannot retrieve it for dynamic grounding.
? "Run As" user permissions (A) would bypass user-specific security, posing a
compliance risk.
? Prompt template permissions (C) are not a Salesforce security mechanism; access is always tied to the user??s profile and sharing settings.
Reference:
Salesforce Help Article: Secure Data Retrieval in Einstein Trust Layer ("User Context Enforcement" section).
Einstein Trust Layer Technical Guide: "Dynamic Grounding and Data Security" (User Permissions alignment).

NEW QUESTION 169
An Agentforce at Universal Containers is working on a prompt template to generate personalized emails for product demonstration requests from customers. It is important for the AI-generated email to adhere strictly to the guidelines, using only associated opportunity information, and to encourage the recipient to take the desired action.
How should the Agentforce Specialist include these instructions on a new line in the prompt template?

A. Surround them with triple quotes (""").
B. Make sure merged fields are defined.
C. Use curly brackets {} to encapsulate instructions.

Answer: A

Explanation:
In Salesforce prompt templates, instructions that guide how the Large Language Model (LLM) should generate content (in this case, personalized emails) can be included by surrounding the instruction text with triple quotes (""""). This formatting ensures that the LLM adheres to the specific instructions while generating the email content.
The use of triple quotes allows the AI to understand that the enclosed text is a directive for how to approach the task, such as limiting the content to associated opportunity information or encouraging a specific action from the recipient.
Refer to Salesforce Prompt Builder documentation for detailed instructions on how to structure prompts for generative AI.

NEW QUESTION 170
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 apcMopnete 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 patter
D. the LLM matches the user query to the available topic, action and steps thenperforms the steps for each action, such as retrieving a fast 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 172

Universal Containers (UC) has a mature Salesforce org with a lot of data in cases and Knowledge articles. UC is concerned that there are many legacy fields, with data that might not be applicable for Einstein AI to draft accurate email responses.
Which solution should UC use to ensure Einstein AI can draft responses from a defined data source?

A. Service AI Grounding
B. Work Summaries
C. Service Replies

Answer: A

Explanation:
Service AI Grounding is the solution that Universal Containers should use to ensure Einstein AI drafts responses based on a well-defined data source. Service AI Grounding allows the AI model to be anchored in specific, relevant data sources, ensuring that any AI-generated responses (e.g., email replies) are accurate, relevant, and drawn from up-to-date information, such as Knowledge articles or cases.
Given that UC has legacy fields and outdated data, Service AI Grounding ensures that only the valid and applicable data is used by Einstein AI to craft responses. This helps improve the relevance of responses and avoids inaccuracies caused by outdated or irrelevant fields. Work Summaries and Service Replies are useful features but do not address the need for grounding AI outputs in specific, current data sources like Service AI Grounding does. For more details, you can refer to Salesforce??s Service AI Grounding documentation for managing AI-generated content based on accurate data sources.

NEW QUESTION 175
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

NEW QUESTION 178
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 182
Which element in the Omni-Channel Flow should be used to connect the flow with the agent?

A. Route Work Action
B. Assignment
C. Decision

Answer: A

Explanation:
Comprehensive and Detailed In-Depth Explanation:UC is integrating an Agentforce agent with Omni-Channel Flow to route work. Let??s identify the correct element.
? Option A: Route Work ActionThe "Route Work" action in Omni-Channel Flow assigns work items (e.g., cases, chats) to agents or queues based on routing rules. When connecting to an Agentforce agent, this action links the flow to the agent??s queue or presence, enabling interaction. This is the standard element for agent integration, making it the correct answer.
? Option B: AssignmentThere??s no "Assignment" element in Flow Builder for Omni- Channel. Assignment rules exist separately, but within flows, routing is handled by "Route Work," making this incorrect.
? Option C: DecisionThe "Decision" element branches logic, not connects to agents.
It??s a control structure, not a routing mechanism, making it incorrect.

Why Option A is Correct:"Route Work" is the designated Omni-Channel Flow action for connecting to agents, including Agentforce agents, per Salesforce documentation.
References:
? Salesforce Agentforce Documentation: Omni-Channel Integration – Specifies "Route Work" for agents.
? Trailhead: Omni-Channel Flow Basics – Details routing actions.
? Salesforce Help: Set Up Omni-Channel Flows – Confirms "Route Work" usage.

NEW QUESTION 184
Which business requirement presents a good use case for leveraging Einstein Prompt Builder?

A. Forecast future sales trends based on historical data.
B. Identify potential high-value leads for targeted marketing campaigns.
C. Send reply to a request for proposal via a personalized email.

Answer: C

Explanation:
? Context of the Question
? Einstein Prompt Builder Typical Use Cases
? ConclusionOption C (Send reply to a request for proposal via a personalized email) is the best match for Einstein Prompt Builder??s generative text functionality.
Salesforce Agentforce Specialist References & Documents
? Salesforce Documentation: Einstein Prompt Builder OverviewHighlights how to use Prompt Builder to create and customize text-based responses, especially for email or record fields.
? Salesforce Agentforce Specialist Study GuideExplains that generative AI features in Salesforce are designed for creating or summarizing text, not for advanced predictive use cases (like forecasting or lead scoring).

NEW QUESTION 187
Universal Containers (UC) wants to improve the productivity of its sales team with generative AI technology. However, UC is concerned that public AI virtual assistants lack adequate company data to general useful responses.
Which solution should UC consider?

A. fine-tune the Einstein AI model with CBM data.
B. Build AI model with Einstein discovery and deploy to sales users.
C. Enable Agentforce and deploy to sales users.

Answer: A

Explanation:
? Context of the QUESTION NO: Universal Containers (UC) wants to harness generative AI to boost sales productivity. They are wary of public AI virtual assistants (like generic chatbots) that lack sufficient UC-specific data to generate useful business responses.
? Why Fine-Tune an Einstein AI Model with CRM Data?
? Why Not Build an AI Model with Einstein Discovery (Option B)?
? Why Not Enable Agentforce (Option C)?
? Outcome: Fine-tuning the Einstein AI model with UC??s CRM data (Answer A) is the most direct, Salesforce-native solution to provide generative AI responses that are aligned with UC??s context, driving productivity gains and ensuring data privacy.
Salesforce Agentforce Specialist References & Documents
? Salesforce Official: Einstein GPT Overview
? Salesforce Trailhead: Get Started with Salesforce Einstein
? Salesforce Documentation: Einstein Discovery
? Salesforce Agentforce Specialist Study Guide

NEW QUESTION 192
Where should the Agentforce Specialist go to add/update actions assigned to a copilot?

A. Copilot Actions page, the record page for the copilot action, or the Copilot Action Library tab
B. Copilot Actions page or Global Actions
C. Copilot Detail page, Global Actions, or the record page for the copilot action

Answer: A

Explanation:
To add or update actions assigned to a copilot, An Agentforce can manage this through several areas:
? Copilot Actions Page: This is the central location where copilot actions are
managed and configured.
? Record Page for the Copilot Action: From the record page, individual copilot actions can be updated or modified.
? Copilot Action Library Tab: This tab serves as a repository where predefined or custom actions for Copilot can be accessed and modified.
These areas provide flexibility in managing and updating the actions assigned to Copilot,
ensuring that the AI assistant remains aligned with business requirements and processes. The other options are incorrect:
? B misses the Copilot Action Library, which is crucial for managing actions.
? C includes the Copilot Detail page, which isn't the primary place for action management.
References:
? Salesforce Documentation on Managing Copilot Actions
? Salesforce Agentforce Specialist Guide on Copilot Action Management

NEW QUESTION 193
.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your Agentforce-Specialist Exam with Our Prep Materials Via below:

https://www.certleader.com/Agentforce-Specialist-dumps.html