Salesforce Agentforce-Specialist Exam Questions

Total Questions: 380+
Demo Questions: 35
Version: Updated for 2025

Prepared and Verified by Cert Empire – Your Trusted IT Certification Partner

For Access to the full set of Updated Questions – Visit: Agentforce-Specialist Exam Dumps by Cert Empire

Question: 1

Universal Containers (UC) is looking to improve its sales team's productivity by providing real-time insights and recommendations during customer interactions. Why should UC consider using Agentforce Sales Agent?

A: To track customer interactions for future analysis

B: To automate the entire sales process for maximum efficiency

C: To streamline the sales process and increase conversion rates

Correct Answer:

C

Explanation:

Salesforce's AI-powered tools for sales agents, such as Sales Cloud Einstein, are designed to enhance agent productivity by delivering real-time intelligence and guided actions. By analyzing data and conversations, these tools provide agents with timely recommendations, key insights, and next best actions directly within their workflow. This assistance helps agents focus on the most promising opportunities and navigate customer interactions more effectively. The direct outcomes are a more efficient, or streamlined, sales process and an improved ability to successfully close deals, which translates to higher conversion rates.

Why Incorrect Options are Wrong:

A: Tracking customer interactions is a fundamental capability of the core Salesforce CRM platform, not a specific reason to implement an advanced, real-time agent assistance tool.

B: These tools are designed to augment and guide human sales agents, not to automate the entire sales process, which requires nuanced human judgment and relationship-building.

References:

1. Salesforce Help - Sales Cloud Einstein: "Sales Cloud Einstein is a powerful set of AI features that help your sales team work smarter and more productively. It provides insights and recommendations that help your team prioritize their work, engage with customers more effectively, and close more deals." This directly supports streamlining the process and increasing conversion rates (closing more deals).

Source: Salesforce Help, "Get Started with Sales Cloud Einstein," Section: "What is Sales Cloud Einstein?"

2. Salesforce Help - Einstein Conversation Insights: "Get insights and trends from your sales calls that help you sell smarter. With Conversation Insights, sales managers and reps can see which patterns in conversations lead to successful deals... to improve sales performance." Improving sales performance is directly linked to increasing conversion rates.

Source: Salesforce Help, "Einstein Conversation Insights," Section: "Sell Smarter with Einstein Conversation Insights."

3. Salesforce Help - Einstein Next Best Action: "Einstein Next Best Action lets you display the right recommendations to the right people at the right time... Use strategies to apply your business rules to a set of recommendations and determine which recommendations to show your users." This streamlines the agent's decision-making process to achieve better outcomes.

Source: Salesforce Help, "Einstein Next Best Action," Section: "Einstein Next Best Action Overview."

Question: 2

Universal Containers is rolling out a new generative AI initiative. Which Prompt Builder limitations should the AI Specialist be aware of?

A: Rich text area fields are only supported in Flex template types.

B: Creations or updates to the prompt templates are not recorded in the Setup Audit Trail.

C: Custom objects are supported only for Flex template types.

Correct Answer:

A, B, C

Explanation:

According to Salesforce's official documentation, all three listed options are known limitations of Prompt Builder. The Flex prompt template type is specifically designed for greater flexibility, and as such, it is the only type that supports custom objects and Rich Text Area fields. Standard templates like Sales Email or Field Generation are more restricted. Furthermore, a key administrative limitation is that any creation or modification of prompt templates is not currently captured in the Setup Audit Trail, which is a critical consideration for change management and governance.

Why Incorrect Options are Wrong:

References:

1. Salesforce Help, Prompt Builder Limitations: This official documentation page explicitly lists all three points as limitations.

"Custom objects are supported only for Flex prompt template types."

"Rich Text Area fields are supported only for Flex prompt template types."

"Creations or updates to prompt templates aren’t recorded in the Setup Audit Trail."

Question: 3

Universal Containers (UC) is discussing its AI strategy in an agile Scrum meeting. Which business requirement would lead an AI Specialist to recommend connecting to an external foundational model via Einstein Studio (Model Builder)?

A: UC wants to fine-tune model temperature.

B: UC wants a model fine-tuned using company data.

C: UC wants to change the frequency penalty of the model.

Correct Answer:

B

Explanation:

The primary business requirement that justifies connecting an external foundational model via Einstein Studio (Model Builder) is the need to enhance its performance and relevance with proprietary data. Model Builder's core value proposition is enabling businesses to fine-tune pre-trained, general-purpose models using their own specific data from within the Salesforce ecosystem. This process, known as fine-tuning, adapts the model to understand a company's unique terminology, products, and customer interaction styles, thereby making it significantly more effective for specific business tasks. Options A and C refer to runtime parameters that control the model's output style, which are tactical adjustments rather than the strategic driver for the initial model integration and adaptation.

Why Incorrect Options are Wrong:

A: Model temperature is a standard, adjustable parameter controlled at the time of inference (the prompt) to manage the creativity of the response, not a reason for the underlying model integration.

C: The frequency penalty is also a runtime parameter, adjusted during the prompt execution to control token repetition, not a strategic requirement for fine-tuning the model itself.

References:

1. Salesforce Help, Einstein Studio: "With Einstein Studio, you can bring your own model (BYOM) into Salesforce... You can also fine-tune a third-party large language model (LLM) using your Data Cloud data to improve the model’s performance for a specific task." This directly supports that fine-tuning with company data is a primary use case.

2. Salesforce Help, Fine-Tune a Generative AI Model: "Fine-tuning is a process that adapts a pre-trained model to a specific task or domain by training it on a smaller, domain-specific dataset. Fine-tuning can improve the performance of a model on a specific task, and it can also make the model more relevant to a particular domain." This reference defines fine-tuning and its purpose, aligning directly with option B.

3. Salesforce Developer Documentation, Connect REST API, Prompt Template Resources: The documentation for invoking generative AI models via API shows that parameters like temperature are part of the request body for a specific prompt execution. For example, in the POST /v59.0/actions/standard/promptTemplate endpoint, the inference parameters include temperature, topP, topK, maxTokens, etc. This confirms that temperature (A) and other similar parameters (like frequency penalty, C) are runtime configurations, not the reason for the fine-tuning setup.

Question: 4

A data science team has trained an XGBoost classification model for product recommendations on Databricks. The AI Specialist is tasked with bringing inferences for product recommendations from this model into Data Cloud as a stand-alone data model object (DMO). How should the AI Specialist set this up?

A: Create the serving endpoint in Databricks, then configure the model using Model Builder.

B: Create the serving endpoint in Einstein Studio, then configure the model using Model Builder.

C: Create the serving endpoint in Databricks, then configure the model using a Python SDK connector.

Correct Answer:

A

Explanation:

To integrate an externally trained model from a platform like Databricks into Salesforce Data Cloud, the standard "Bring Your Own Model" (BYOM) pattern is used. This process requires two main steps. First, the model must be deployed and exposed via a REST API endpoint on its native platform (Databricks). This is known as a serving endpoint. Second, within Salesforce Einstein Studio, the Model Builder tool is used to establish a connection to this external endpoint, defining the input payload, output structure, and authentication, thereby making the model available for inference jobs within Data Cloud.

Why Incorrect Options are Wrong:

B: The serving endpoint for a model trained and hosted in Databricks must be created within Databricks. Einstein Studio is used to connect to an existing external endpoint, not create it.

C: While a Python SDK might be used for various custom integrations, Model Builder is the specific, declarative tool provided within Einstein Studio for the express purpose of connecting external models to Data Cloud.

References:

1. Salesforce Help Documentation, Einstein Studio, "Bring Your Own AI Model": This documentation outlines the end-to-end process for integrating external models. It explicitly states, "To connect an external model, you first create a serving endpoint in your desired

platform... Then, in Model Builder, you create a connect model." This directly supports creating the endpoint in the external platform (Databricks) and using Model Builder for configuration.

2. Salesforce Help Documentation, Einstein Studio, "Connect an External Model": This guide details the steps within Model Builder. It begins with the prerequisite of having an accessible model endpoint URL from platforms like Amazon SageMaker, Google Vertex AI, or other public clouds, which includes Databricks. This confirms that the endpoint is external and Model Builder is the Salesforce-side tool.

3. Salesforce Developers Blog, "Bring Your Own AI Models to Salesforce Data Cloud with Einstein Studio": This article describes the architecture and workflow. It shows a diagram where the external model platform (e.g., Databricks) has a "Model Endpoint," which is then connected to via "Einstein Studio Model Builder" to run prediction jobs that populate Data Model Objects (DMOs). This validates the roles of both Databricks (hosting the endpoint) and Model Builder (configuring the connection).

Question: 5

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.

Correct Answer:

A

Explanation:

Einstein generates Work Summaries by grounding its output in the available context to ensure relevance and accuracy. This context is drawn from multiple sources related to the customer interaction. The AI analyzes the conversation transcript (conversation context), the data within the case record itself (cases), and information from any Knowledge articles that have been associated with the case during the interaction. This multi-source approach allows Einstein to create a comprehensive and contextually rich summary of the agent's work, capturing not just the dialogue but also the resolution steps and resources used.

Why Incorrect Options are Wrong:

B: This option is incorrect because it is incomplete. While the conversation context is a primary input, Einstein also grounds the summary in the broader case data, not just the transcript alone.

C: This option is incorrect because it is also incomplete. It omits the critical context provided by the case data itself, which is a key grounding source for the summary generation.

References:

1. Salesforce Help, Einstein Generative AI for Service: This official documentation provides an overview of the features. It states, "Einstein Work Summaries: Instantly generate a wrap-up summary of a case and customer conversation grounded in case data and chat transcripts." This confirms that both conversation (transcripts) and case data (which includes the case itself and can encompass related information like attached articles) are used.

2. Salesforce Help, Einstein Work Summaries: This document specifies the core input, stating, "Einstein drafts a concise summary of a customer service conversation based on your chat, messaging, and voice call transcripts." This supports the "conversation context" component of the correct answer. The combination of this specific source with the broader overview from the first reference supports option A as the most complete description.

Question: 6

An AI Specialist created a custom Agent action, but it is not being picked up by the planner service in the correct order. Which adjustment should the AI Specialist make in the custom Agent action instructions for the planner service to work as expected?

A: Specify the dependent actions with the reference to the action API name.

B: Specify the profiles or custom permissions allowed to invoke the action.

C: Specify the LLM model provider and version to be used to invoke the action.

Correct Answer:

A

Explanation:

The Einstein Copilot planner service constructs a sequence of actions to fulfill a user's request. To ensure actions are executed in the correct logical sequence, the planner must understand the dependencies between them. By explicitly specifying a prerequisite action's API name within the instructions of a dependent action, the AI Specialist provides a direct command to the planner. This ensures that the prerequisite action is completed before the dependent action is initiated, thereby controlling the execution order as intended.

Why Incorrect Options are Wrong:

B: Specifying profiles or permissions controls access and security (who can run the action), not the operational sequence or execution order determined by the planner.

C: The LLM model provider and version are part of the underlying configuration for how the agent processes information, but this does not define the execution dependencies between distinct actions.

References:

1. Salesforce Help. (2024). Einstein Copilot Action Authoring.

Section: Define an Action's Dependencies

Content: "To ensure that your actions are executed in the correct order, define an action’s dependencies in the action’s instructions... To define a dependency, reference the dependent action’s API name in the instructions. For example: This action adds a comment to a case. It’s dependent on the CreateCase action." This documentation explicitly states

that referencing the API name of a dependent action within the instructions is the correct method to control the execution order for the planner.

Question: 7

An AI specialist wants to leverage 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

Correct Answer:

B

Explanation:

To use the Record Snapshots grounding feature in a Salesforce prompt template, an administrator must first create a field set on the relevant object. A field set is a grouping of fields that defines the specific record data to be included, or "grounded," in the prompt. When the prompt template is executed on a record, Salesforce takes a snapshot of the values in the fields defined in that field set and merges them into the prompt context. This ensures the AI model has the relevant, structured record data it needs to generate a contextualized response.

Why Incorrect Options are Wrong:

A: Configure page layout of the master record type: Page layouts control the visual presentation of fields on a record's detail page for users, not the data payload for prompt grounding.

C: Enable and configure dynamic form for the object: Dynamic Forms offer flexible UI control on Lightning record pages but are not the mechanism used to define the data for a Record Snapshot.

References:

1. Salesforce Help. (2024). Ground Prompt Templates with Salesforce Data.

Section: Record Snapshot (Field Set)

Content: "To ground a prompt template with data from the fields on a record, use a field set. A field set is a grouping of fields... When you add a field set as a resource to your prompt template, the resulting generated prompt includes the fields and values from the record that the prompt is run on."

2. Salesforce Trailhead. (2024). Prompt Builder Basics.

Unit: Ground Prompts with Salesforce Data

Content: This unit explains the different grounding options, explicitly stating that for Record Snapshots, "You select a field set on the object to specify which fields to include." It provides a step-by-step guide on creating a field set for this purpose.

Question: 8

What is the importance of Action Instructions when creating a custom Agent action?

A: Action Instructions tell the user how to call this action in a conversation

B: Action Instructions tell the large language model (LLM) which action to use.

C: Action Instructions define the expected user experience of an action.

Correct Answer:

B

Explanation:

Action Instructions, also referred to as the action's description, serve as the primary guidance for the Large Language Model (LLM). The LLM analyzes these natural language instructions to understand the purpose, capabilities, and appropriate context for using the custom action. Based on this understanding, the LLM determines if the action is the correct tool to select and execute in response to a specific user request. A clear and descriptive instruction is crucial for the agent's ability to accurately map user intent to the correct function.

Why Incorrect Options are Wrong:

A: These instructions are for the LLM's internal reasoning process, not for the end-user. The user interacts with the agent through natural conversation, not by explicitly calling actions.

C: This is too general. While the correct execution of an action contributes to the user experience, the specific purpose of the instructions is to guide the LLM's tool selection, not to define the entire UX.

References:

1. Salesforce Help Documentation: Create a Custom Action for an Agent. In the section detailing the creation process, it states, "The agent’s large language model (LLM) uses the action’s description to determine when to use the action. Write a clear and concise description that explains what the action does and when to use it." This directly confirms that the instructions are for the LLM's decision-making.

Question: 9

Universal Containers (UC) wants to leverage 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: Case Replies

B: Knowledge Replies

C: Service Replies

Correct Answer:

C

Explanation:

The official Salesforce name for the Generative AI capability that drafts responses for agents based on organizational data is "Service Replies." This feature is designed to reduce agent handling time by providing contextual, AI-generated text. While it is configured in this scenario to use Knowledge articles as its source, the agents are trained on how to use the "Service Replies" feature within the Service Console. The name of the feature itself is the correct answer for the capability on which agents would be trained.

Why Incorrect Options are Wrong:

A: Case Replies: This is not an official Salesforce feature name. While Service Replies can be grounded in case data, "Case Replies" is not the term for the capability.

B: Knowledge Replies: This term describes the output when the Service Replies feature is grounded in Knowledge articles, but it is not the official name of the AI capability itself.

References:

1. Salesforce Help, "Einstein Service Replies": The official documentation consistently refers to the feature as "Service Replies." It states, "Einstein Service Replies drafts English-language responses that are grounded in your knowledge articles and past closed cases." This establishes the official name of the capability and confirms it uses Knowledge as a source. (salesforce.com, Help & Training, "Einstein Service Replies").

2. Salesforce Help, "Set Up Einstein Service Replies": The setup guide for this functionality is explicitly named "Set Up Einstein Service Replies," reinforcing this as the correct term for

the feature that an administrator enables and on which agents are trained. (salesforce.com, Help & Training, "Set Up Einstein Service Replies").

3. Trailhead, "Einstein for Service: Service Replies and Work Summaries": The official Salesforce training module for this functionality is titled based on "Service Replies." The module explains, "Service Replies are AI-generated responses grounded in your knowledge articles and case data," confirming this is the term used in official training materials. (trailhead.salesforce.com, "Get Started with Einstein Service Replies").

Question: 10

Which part of the Einstein Trust Layer architecture leverages an organization's own data within a large language model (LLM) prompt to confidently return relevant and accurate responses?

A: Prompt Defense

B: Data Masking

C: Dynamic Grounding

Correct Answer:

C

Explanation:

Dynamic Grounding is the specific component of the Einstein Trust Layer responsible for enhancing Large Language Model (LLM) responses with an organization's trusted, real-time data. It automatically retrieves relevant information from Salesforce Data Cloud and other sources, adding it as context to the user's prompt. This process "grounds" the LLM's response in factual, company-specific data, thereby increasing the relevance and accuracy of the generated output. It ensures the AI leverages the most current and appropriate business context without requiring manual data input.

Why Incorrect Options are Wrong:

A: Prompt Defense: This is a security feature designed to detect and block malicious inputs or prompt injection attacks, protecting the integrity of the LLM, not for grounding responses in company data.

B: Data Masking: This is a data privacy feature that automatically detects and redacts sensitive information (like PII) from prompts before they are sent to the LLM, preventing data leakage.

References:

1. Salesforce Help. (2024). Einstein Trust Layer. "Dynamic Grounding retrieves the most relevant, up-to-date company information from various data sources to make AI-generated responses more accurate and relevant. For example, it can retrieve information from a knowledge article or a particular field on a record to enrich the prompt."

2. Salesforce Help. (2024). How the Einstein Trust Layer Protects Your Data. "Dynamic Grounding: To make responses more relevant to your users, we add grounding data from your Salesforce org to the prompt."

3. Salesforce Developer Documentation. (2024). Einstein Trust Layer. "Dynamic grounding automatically grounds LLM prompts in the right real-time data, from Data Cloud and other sources, to deliver relevant, accurate results."

Question: 11

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.

Correct Answer:

B

Explanation:

Secure Data Retrieval for dynamic grounding in Salesforce operates within the security context of the user executing the prompt. When a prompt template retrieves Salesforce data, it strictly adheres to the running user's permissions. This includes their object-level permissions, field-level security, and any applicable sharing rules. This ensures that the generative AI response is grounded only in data that the specific user is authorized to view, maintaining the principle of least privilege and Salesforce's core security model.

Why Incorrect Options are Wrong:

A: The system does not use a pre-defined 'Run As' user for this process; it dynamically uses the permissions of the user who is currently interacting with the system and executing the prompt.

C: Prompt templates are configuration metadata and do not have their own object permissions. Data access permissions are a function of the user's profile and permission sets, not the template itself.

References:

1. Salesforce Help. (2024). Ground Your Prompts with Salesforce Data. Salesforce, Inc. Retrieved from https://help.salesforce.com/s/articleView?id=sf.einsteinpromptbuildergrounding.htm&type=5. In the "How Grounding Works" section, it states, "When a user runs a prompt that uses a resource, the generated response is based on data that the user has permission to access."

2. Salesforce Help. (2024). Data Security in Prompt Builder. Salesforce, Inc. Retrieved from https://help.salesforce.com/s/articleView?id=sf.einsteinpromptbuildersecurity.htm&type=5.

This document explicitly states, "Prompts run with the user’s permissions, and users can’t see data that they don’t have access to."

Question: 12

Universal Containers (UC) is using Einstein Generative AI to generate an account summary. UC aims to ensure the content is safe and inclusive, utilizing the Einstein Trust Layer's toxicity scoring to assess the content's safety level. In the score of 1 indicate?

A: The response is the least toxic Einstein Generative AI Toxicity Scoring system, what does a toxicity category.

B: The response is not toxic.

C: The response is the most toxic.

Correct Answer:

C

Explanation:

The Einstein Trust Layer's toxicity detection feature scores Large Language Model (LLM) outputs on a scale from 0 to 1. A score closer to 0 indicates a low likelihood of toxicity, meaning the content is considered safe. Conversely, a score closer to 1 indicates a high probability that the content is toxic. Therefore, a score of 1 represents the highest level of toxicity, flagging the content as potentially harmful, inappropriate, or unsafe, which allows organizations to intervene or block such responses.

Why Incorrect Options are Wrong:

A: A score of 1 represents the most toxic category, not the least. The option is also grammatically incoherent.

B: A response that is not toxic would receive a score of 0 or a value very close to 0.

References:

1. Salesforce Help Documentation, "Einstein Generative AI and the Trust Layer": This document explicitly states, "Toxicity detection scores the LLM output on a scale of 0 to 1. A score closer to 1 indicates a higher likelihood of toxicity." This directly supports that a score of 1 is the most toxic.

2. Salesforce Help Documentation, "Einstein Trust Layer": Under the "How the Einstein Trust Layer Works" section, it describes "Toxicity Detection" as a key feature that classifies language to prevent the generation of harmful content, reinforcing the purpose of the scoring system.

Question: 13

An AI Specialist at Universal Containers is trying to set up a new Field Generation prompt template. They take the following steps. After creating the prompt template, the AI Specialist saves, tests, and activates it. Howsoever, when they go to a case record, the AI Analysis field does not show the (Sparkle) icon on the Edit pencil. When the AI Specialist was editing the field, it was behaving as a normal field. Which critical step did the AI Specialist miss?

A: They forgot to reactivate the Lightning page layout for the Case object after activating their Field Generation prompt template.

B: They forgot that the Case Object is not supported for Add generation as Feinstein Service Replies should be used instead.

C: They forgot to edit the Lightning page layout and associate the field to a prompt template

Correct Answer:

C

Explanation:

Creating and activating a Field Generation prompt template makes it available for use, but it does not automatically link it to a specific field on a record page. The critical subsequent step is to edit the Lightning Record Page using the Lightning App Builder. Within the App Builder, the AI Specialist must select the target field (in this case, AI Analysis) and, in its properties pane, explicitly associate it with the newly activated prompt template. This association is what enables the generative AI functionality for that field, which is visually indicated by the (Sparkle) icon next to it on the record page.

Why Incorrect Options are Wrong:

A: Reactivating the Lightning page is not the required action. The page must be edited in the Lightning App Builder to associate the field with the template, and then saved.

B: The Case object is a supported object for Field Generation. Einstein Service Replies is a distinct feature used for generating conversational replies, not for populating record fields.

References:

1. Salesforce Help, Einstein Generative AI for Service, "Add Einstein Generative AI to a Lightning Page": This document outlines the precise procedure. It states, "From the Lightning App Builder, edit a record page... Drag a field that you want to use with a prompt template onto the canvas. In the field’s properties pane, under Prompt Template, select the

template to use." This directly confirms that associating the field to the template on the Lightning page is the necessary step.

2. Salesforce Help, Prompt Builder, "Field Generation Prompt Templates": This guide details the creation of Field Generation templates and lists supported objects, which include Case. It clarifies that the template's purpose is to populate a field on a record, distinguishing it from other features like Service Replies. The process concludes with making the feature available on the Lightning page, reinforcing the action described in option C.

Question: 14

An AI Specialist wants to troubleshoot their Agent's performance. Where should the AI Specialist go to access all user interactions with the Agent, including Agent errolrs, incorrectly triggered actions, and incomplete plans?

A: Event Logs

B: Plan Canvas

C: Agent Settings

Correct Answer:

A

Explanation:

AWS Agents for Amazon Bedrock store a chronological “event log” for every invocation. The log captures the full user request, the agent’s generated plan, each action taken, any agent errors, and whether the plan completed. AWS documentation specifies that these event logs are the primary place to diagnose mis-triggered actions, incomplete plans, and other runtime issues, so the AI Specialist should open Event Logs.

Why Incorrect Options are Wrong:

References:

1. Amazon Web Services, “Agents for Amazon Bedrock – Developer Guide”, Section “Monitoring agents: Viewing agent event logs”, pp. 145-148, 2024-04-30 edition.

2. Amazon Web Services Console User Guide, “Troubleshoot an agent with event logs”, para. 2-4 (https://docs.aws.amazon.com/bedrock/latest/userguide/agents-event-logs.html).

Question: 15

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 applied. conversationally and in the now or work.

B: 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.

C: 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.

Correct Answer:

A

Explanation:

Salesforce Sales Agent is a conversational AI assistant designed to help sales representatives perform their daily tasks more efficiently. The core use case, as described in option A, is to allow users to interact with Salesforce using natural language. This enables them to conversationally invoke actions like updating opportunities, summarizing records, or drafting emails. The agent's actions are "grounded" in the relevant CRM data, and it can be configured to enforce company-specific policies and processes, all within the user's natural flow of work.

Why Incorrect Options are Wrong:

B: This describes the functionality of Salesforce Flow, specifically a screen flow, which provides a guided, step-by-step user interface for data entry and business processes, not a conversational AI interaction.

C: This functionality is characteristic of tools like Einstein Activity Capture or Salesforce Inbox, which automatically capture and log email and event data to the appropriate Salesforce records.

References:

1. Salesforce Help, "Sales Agent": "Sales reps can use natural language in a conversational interface to get AI-generated insights, update CRM data, and prep for meetings. Sales Agent is grounded in your trusted CRM data and is fully customizable to your business

needs." This directly supports the concepts of using natural language to invoke tasks grounded in data, as stated in option A.

2. Salesforce Help, "Flow Builder": "Flow Builder is the declarative interface used to build individual flows. With Flow Builder, you can build code-like logic without using a programming language." This aligns with the description in option B of a guided, rule-based process, distinguishing it from the conversational nature of Sales Agent.

3. Salesforce Help, "Einstein Activity Capture": "Einstein Activity Capture lets you keep data between Salesforce and your email and calendar applications up to date. Einstein Activity Capture focuses on capturing the right data to give you a complete view of customer interactions." This confirms that option C describes activity capture functionality, not the role of a conversational agent.

Question: 16

An AI Specialist at Universal Containers (UC) is building with no-code tools only. They have many small accounts that are only touched periodically by a specialized sales team, and UC wants to maximize the sales operations team's time. UC wants to help prep the sales team for the calls by summarizing past purchases, interests in products shown by the Contact captured via Data Cloud, and a recap of past email and phone conversations for which there are transcripts. Which approach should the AI Specialist recommend to achieve this use case?

A: Use a prompt template grounded on CRH and Data Cloud data using standard foundation model.

B: Fine-Tune the standard foundational model due to the complexity of the data.

C: Deploy UC's own custom foundational model on this data first.

Correct Answer:

A

Explanation:

The most appropriate approach is to use a prompt template grounded with data from the Customer 360 (often referred to as CRH or Customer Relationship History) and Data Cloud. This method directly aligns with the "no-code tools only" constraint, as tools like Salesforce's Prompt Builder are designed for this purpose. Grounding provides the standard foundation model with specific, real-time context from past purchases, interests, and conversation transcripts. This ensures the generated summary is accurate and relevant to the specific account without the significant overhead, cost, and technical expertise required for fine-tuning or building a custom model. This approach maximizes the sales team's time by leveraging existing data and standard AI capabilities efficiently.

Why Incorrect Options are Wrong:

B: Fine-tuning is a more complex, resource-intensive process that is not considered a no-code tool. It is unnecessary when the goal is to summarize existing data, a task standard models excel at with proper grounding.

C: Deploying a custom foundational model is the most complex and expensive option, requiring deep technical expertise and vast data for training. It is the antithesis of a no-code solution and is excessive for this use case.

References:

1. Salesforce Help, "Prompt Builder": "Prompt Builder is a generative AI tool that lets you create, test, and customize prompt templates that use your Salesforce data... Prompt templates are reusable prompts that you can use with different records. Merge fields act as placeholders for CRM data that's specific to a record." This document confirms that prompt templates are a no-code tool for grounding with CRM data.

2. Salesforce Help, "Grounding for Generative AI": "To ensure that the LLM's response is relevant and accurate, Salesforce grounds the prompt with your data. Grounding means that Salesforce finds the most relevant data from your organization and adds it to the prompt that's sent to the LLM." This source defines grounding as the process of adding specific Salesforce data to a prompt, which is exactly what the scenario requires.

3. Salesforce Developers Blog, "Bring Your Own LLM to the Einstein 1 Platform" (Comparing AI techniques): This type of documentation typically outlines the hierarchy of AI customization. Prompt engineering (using templates and grounding) is presented as the most accessible, no-code/low-code method. Fine-tuning and bringing your own model (BYOM) are positioned as more advanced, code-intensive options for when prompt engineering is insufficient, which is not the case in the given scenario.

Question: 17

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.

Correct Answer:

B

Explanation:

Prompt Builder is a Salesforce tool designed specifically for creating, managing, and executing reusable prompt templates for generative AI tasks. Generating a summary of a sales contract is a classic generative AI use case, where a Large Language Model (LLM) is prompted to create new, concise text based on a larger source document. Prompt Builder allows an administrator to create a template that can pull in the contract data from a Salesforce record and instruct the AI to summarize it, streamlining the sales team's review process.

Why Incorrect Options are Wrong:

A: Populate an AI-generated time-to-close estimation to opportunities: This is a predictive AI function. It analyzes historical data to forecast an outcome, a task handled by features like Einstein Opportunity Scoring or Forecasting, not generative AI via Prompt Builder.

C: Populate an AI-generated lead score for new leads: This is a predictive AI function. Einstein Lead Scoring uses machine learning to analyze data and assign a numerical score based on conversion probability, which does not involve generating text with Prompt Builder.

References:

1. Salesforce Help, "Prompt Builder": This document explicitly states, "With Prompt Builder, you can create, test, customize, and manage prompt templates that incorporate your CRM data, enabling you to deliver new generative AI experiences safely and at scale." It lists use cases such as summarizing records, which directly supports option B.

2. Salesforce Help, "How Einstein Generative AI Works": This documentation distinguishes between predictive AI and generative AI. It clarifies that generative AI creates new content, such as summaries and emails, which is the domain of tools like Prompt Builder.

3. Salesforce Help, "Einstein Lead Scoring": This source describes the feature as one that "uses data science and machine learning to discover the patterns of lead conversion... and predict which leads to prioritize." This confirms that lead scoring (Option C) is a predictive, not generative, task.

4. Salesforce Help, "Get Predictions with Einstein Opportunity Scoring": This document explains that Opportunity Scoring "gives each opportunity a score from 1 to 99, which is available on opportunity records." This confirms that estimations and scoring (Option A) are predictive functions separate from Prompt Builder's generative capabilities.

Question: 18

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 pattern. the LLM matches the user query to the available topic, action and steps then performs the steps for each action, such as retrieving a fast of open opportunities.

Correct Answer:

A

Explanation:

The Large Language Model (LLM) within Salesforce's agent technology operates through a dynamic, multi-stage process. First, it performs semantic interpretation of the user's natural language query to understand the underlying intent. Following this, the LLM reasons over the available, predefined actions (often called "skills" or "tools") and generates a sequential plan to fulfill the request. This plan identifies the specific actions needed, such as querying the Opportunity object with filters for "open" status and the current user. Finally, the system executes this plan, retrieves the requested data, and presents it to the user. This planning capability allows the agent to handle complex, multi-step requests dynamically.

Why Incorrect Options are Wrong:

B: This describes a static, rule-based system. LLM-powered agents are fundamentally dynamic and do not rely on a fixed set of rules, enabling them to handle a wider range of queries.

C: The term "matches" and "dialog pattern" is less precise. It better describes older, structured chatbot logic rather than the LLM's advanced capability to reason and "generate a plan."

References:

1. Salesforce Help, Einstein Copilot, "How Einstein Copilot Works": This document details the process: "When a user interacts with the copilot, the LLM reasons over the user’s prompt to understand their intent. The LLM then selects the one or more relevant skills available... This sequence of skills is called a plan. The copilot then runs the plan, executing the skills in order and returning the final result to the user." This directly supports the "interpret, plan, execute" model described in option A.

2. Salesforce Developers, Einstein Copilot, "Standard Copilot Actions": This documentation outlines the predefined actions available to the copilot. The LLM's role is to select from these available tools to build its plan, confirming that it identifies appropriate actions from a known library to execute the user's request.

Question: 19

Universal Containers, dealing with a high volume of chat inquiries, implements Einstein Work Summaries to boost productivity. After an agent-customer conversation, which additional information does Einstein generate and fill, apart from the "summary"'

A: Sentiment Analysis and Emotion Detection

B: Draft Survey Request Email

C: Issue and Revolution

Correct Answer:

C

Explanation:

Einstein Work Summaries is designed to streamline post-conversation wrap-up tasks for agents. When a chat conversation concludes, the feature uses generative AI to analyze the transcript and automatically populate key fields. In addition to the main "Work Summary" field, it specifically identifies and drafts content for the "Issue" (the customer's problem) and the "Resolution" (the steps taken to solve the problem). This allows agents to quickly review, edit, and save a structured record of the interaction, significantly improving efficiency.

Why Incorrect Options are Wrong:

A: Sentiment Analysis and Emotion Detection: While Salesforce has Einstein features for sentiment analysis, "Emotion Detection" is not a standard, distinct output of Work Summaries. The primary outputs are the summary, issue, and resolution.

B: Draft Survey Request Email: Generating follow-up emails is a separate process handled by automation tools like Flow or email templates, not a direct function of the Einstein Work Summaries feature itself.

References:

1. Salesforce Help, Einstein Work Summaries: "When a chat conversation ends, Einstein drafts a summary and identifies the issue and resolution. Agents can then review, edit, and save the summary to the Work Summary field on the Chat Transcript record."

2. Salesforce Help, Set Up Einstein Work Summaries: "Einstein Work Summaries for Chat uses generative AI to create a concise summary of a customer chat conversation. The

summary identifies the customer issue and the resolution." This documentation explicitly details the core components generated by the feature.

Question: 20

Universal Containers has a custom Agent action calling a flow to retrieve the real-time status of an order from the order fulfillment system.

For the given flow, what should the AI Specialist consider about the running user's data access?

A: The flow must have the "with sharing" permission selected m the advanced settings for the permissions, field-level security, and sharing settings to be respected.

B: The custom action adheres to the permissions, held-level security, and sharing settings configured in the flow.

C: The Agent will always run flows in system mode so the running user's data access will not affect the data returned.

Correct Answer:

B

Explanation:

When an Einstein Copilot (formerly Salesforce Agent) action invokes a flow, it runs the flow in system context. However, the specific data access behavior is determined by the advanced settings configured within the flow itself. The AI Specialist must choose the appropriate context, such as "System Context with Sharing—Enforces Record-Level Access," if the flow needs to respect the running user's permissions and sharing rules. Therefore, the custom action's adherence to security settings is directly controlled by the flow's configuration, making this the primary consideration for the specialist.

Why Incorrect Options are Wrong:

A: This is imprecise. The setting is "System Context with Sharing," not a "with sharing permission." Also, it's not a "must"; the specialist might intentionally choose to bypass sharing for a specific use case.

C: This is incorrect. While the flow runs in system context, it can be configured as "System Context with Sharing," which explicitly enforces the running user's record-level access, meaning their permissions do affect the data.

References:

1. Salesforce Help, 'Create a Copilot Action from a Flow': "When you create a copilot action from a flow, Einstein runs the flow in system context. To respect the user’s permissions, in Flow Builder, open the flow, and in the flow’s start element, select the advanced setting System Context with Sharing—Enforces Record-Level Access."

2. Salesforce Help, 'How Does Flow Security Work?': This document details the different run-time contexts for a flow. It explains that for an autolaunched flow, you can choose how it runs, including "System Context with Sharing," which "runs the flow in system context, but respects the user's record-level sharing." This confirms that the flow's configuration dictates the security model.

Question: 21

What is true of Agentforce Testing Center?

A: Running tests risks modifying CRM data in a production environment.

B: Agentforce Testing Center can only be used in a production environment.

C: Running tests does not consume Einstein Requests.

Correct Answer:

C

Explanation:

The Agentforce Testing Center, which is used to evaluate Einstein classification and recommendation models, is designed as a safe environment for testing and refinement. A key feature of this environment is that running tests to evaluate model performance does not consume any of your organization's licensed Einstein Requests. This allows administrators to thoroughly test and tune their AI models without incurring costs or depleting their prediction allocation, ensuring the model is accurate before it is activated for live agent use.

Why Incorrect Options are Wrong:

A: The testing center is a simulation tool. It evaluates predictions against existing data but does not write back or modify any CRM data, thereby preventing any risk to the production environment.

B: The testing center is available in both sandbox and production environments. It is a best practice to build, train, and test models in a sandbox before deploying them to production.

References:

1. Salesforce Help, Article: Test Your Einstein Classification Models.

Reference: In the "Test Your Einstein Classification Models" section, the documentation states, "Testing your model doesn’t affect your Salesforce data or cost you anything." This directly confirms that running tests does not modify data (making A incorrect) and does not consume requests (making C correct).

2. Salesforce Help, Article: Set Up Einstein Case Classification.

Reference: The setup guide outlines the process of building a model. This process, including testing, is standard practice to perform in a sandbox environment before deployment to production, which contradicts the claim that the testing center can only be used in production (making B incorrect).

Question: 22

Universal Containers (UC) plans to implement prompt templates that utilize the standard foundation models.

What should UC consider when building prompt templates in Prompt Builder?

A: Ask it to role-play as a character in the prompt template to provide more context to the LLM.

B: Include multiple-choice questions within the prompt to test the LLM's understanding of the context.

C: Train LLM with data using different writing styles including word choice, intensifiers, emojis, and punctuation.

Correct Answer:

A

Explanation:

Assigning a persona or role to the Large Language Model (LLM) is a fundamental and highly effective prompt engineering technique. By instructing the LLM to "role-play" (e.g., "You are an expert support agent"), you provide crucial context about the desired tone, style, and perspective for the generated response. This helps the model produce output that is more relevant, accurate, and aligned with the specific business need. This practice is a documented best practice for crafting effective prompts within Salesforce Prompt Builder.

Why Incorrect Options are Wrong:

B: Including multiple-choice questions is not a standard or effective method for a generative prompt; it complicates the instruction and does not directly guide the LLM toward creating the desired output.

C: This describes the process of training or fine-tuning an LLM, which is distinct from building a prompt template. Prompt Builder is used to craft inputs for existing, pre-trained models.

References:

1. Salesforce Help, Einstein Generative AI, Prompt Template Fundamentals: Under the "Tips for Writing Prompts" section, the documentation explicitly states, "Give the LLM a

persona. Tell the LLM to act as a certain character or professional to help it understand the context of the request and generate a better response." This directly validates option A.

2. Salesforce Help, Einstein Generative AI, Prompt Builder: The documentation for Prompt Builder focuses on crafting instructions (prompts) for a model, not on training the model itself. This distinction makes option C, which describes training, an incorrect activity for someone building a prompt template.

Question: 23

Amid their busy schedules, sales reps at Universal Containers dedicate time to follow up with prospects and existing clients via email regarding renewals or new deals. They spend many hours throughout the week reviewing past communications and details about their customers before performing their outreach. Which standard Agent action helps sales reps draft personalized emails to prospects by generating text based on previous successful communications?

A: Agent Action: Draft of Revise Sales Email

B: Agent Action: Summarize Record

C: Agent Action: Find Similar Opportunities

Correct Answer:

A

Explanation:

The "Draft or Revise Sales Email" standard Agent action is specifically designed to address this scenario. It leverages generative AI to create personalized emails for contacts or leads by analyzing the context from the associated records. This allows sales representatives to quickly generate relevant and tailored outreach communications, significantly reducing the manual effort of reviewing past interactions and customer details before drafting a message.

Why Incorrect Options are Wrong:

B: Agent Action: Summarize Record: This action provides a high-level summary of a record. While useful for a quick overview, it does not generate a draft email for outreach.

C: Agent Action: Find Similar Opportunities: This action is used for identifying other sales opportunities with similar characteristics, which aids in strategy and analysis, not in composing an email for a specific prospect.

References:

1. Salesforce Help. (2024). Standard Copilot Actions for Sales. Salesforce, Inc. Retrieved from https://help.salesforce.com/s/articleView?id=sf.copilotactionssalesstandard.htm&type=5

This document explicitly lists and describes the "Draft or Revise Sales Email" action: "Generate a personalized email for a contact or lead. Or, select an email in the activity

timeline and ask Einstein to revise it." This directly supports the correct answer. It also describes the "Summarize Record" action, confirming its distinct purpose.

Question: 24

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 Einstein Trust Layer is incorrectly configured.

C: The data being used for grounding is incorrect or incomplete.

Correct Answer:

C

Explanation:

The performance and accuracy of a generative AI prompt, especially one designed for summarization, are fundamentally dependent on the quality of the grounding data. Grounding provides the Large Language Model (LLM) with the specific context and information from Salesforce records (in this case, the case record). If the fields, related records, or data sources used to ground the prompt are incorrect, incomplete, or contain irrelevant information, the LLM will produce a summary that is similarly flawed and does not contain the appropriate information. The principle of "garbage in, garbage out" directly applies.

Why Incorrect Options are Wrong:

The prompt template version is incompatible with the chosen LLM. Salesforce manages the underlying LLM and prompt framework compatibility. This is not a typical user-facing issue that would explain poor content generation.

The Einstein Trust Layer is incorrectly configured. The Einstein Trust Layer's primary function is data security, masking, and privacy, not ensuring the factual correctness or completeness of the source data for generation.

References:

1. Salesforce Help, "Ground Prompts with Salesforce Data": This document explains that to get relevant and accurate responses, prompts must be grounded with contextual data from Salesforce records using merge fields. It states, "To get relevant responses from an LLM, a

prompt needs context... Grounding a prompt template with your Salesforce data makes the generated response more relevant to your users." This directly supports the idea that the quality of the grounding data determines the quality of the output.

2. Salesforce Help, "Prompt Builder Best Practices": In the section "Tips for Writing a Prompt," the documentation advises to "Be specific about the information you want the LLM to use." This implies that the information (the grounding data) must be available and accurate for the prompt to perform well. If the data itself is poor, the specificity of the prompt cannot overcome that limitation.

3. Salesforce Help, "Einstein Trust Layer": This documentation details the features of the Trust Layer, such as secure data retrieval, data masking, and zero data retention. These functions are related to data security and privacy, not the inherent accuracy or completeness of the record data being used for grounding, making it an indirect and less likely cause for poor summary content.

Question: 25

Universal Containers implements Custom Agent Actions to enhance its customer service operations. The development team needs to understand the core components of a Custom Agent Action to ensure proper configuration and functionality. What should the development team review in the Custom Agent Action configuration to identify one of the core components of a Custom Agent Action?

A: Instructions

B: Output Types

C: Action Triggers

Correct Answer:

A

Explanation:

The core of a Custom Agent Action (now known as a Copilot Action) is the set of Instructions provided to the Large Language Model (LLM). These natural language instructions define the action's capabilities, specify the necessary inputs, and guide the LLM on when and how to use the action in response to a user's request. The development team must review the instructions to understand how the action will be interpreted and triggered by the Einstein Copilot, making it a fundamental component for configuration and functionality.

Why Incorrect Options are Wrong:

B: Output Types: While an action has outputs, the "Instructions" are more fundamental as they define the action's purpose and how the LLM should invoke it in the first place.

C: Action Triggers: An action is not configured with a separate "trigger" component. Instead, it is invoked when the copilot's reasoning engine matches the user's intent to the action's defined Instructions.

References:

1. Salesforce Help, "Create a Custom Copilot Action": This document details the configuration process for a custom action. Step 2 is "Write Instructions for the Copilot," which states, "The instructions you write are the most important part of your custom action. They tell the copilot what the action does and when to use it... The copilot uses these

instructions to decide which action is the best fit for a user’s request." This establishes Instructions as a primary, core component.

2. Salesforce Help, "Guidelines for Writing Copilot Action Instructions": This guide emphasizes the centrality of instructions: "The instructions you write for a custom action are a prompt template that tells a large language model (LLM) everything it needs to know about the action." This confirms that instructions are the foundational element for the action's behavior.

Question: 26

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 number of tokens that can be processed by the LLM varies with total user demand.

B: The number of tokens generated by the dynamic nature of the prompt template will vary by record.

C: The template type needs to be switched to Flex to accommodate the variable amount of tokens generated by the prompt grounding.

Correct Answer:

B

Explanation:

The error's random nature is caused by the dynamic content of the prompt template. Field Generation templates use merge fields to pull data from specific records (e.g., a case description or email body). The amount of text in these fields varies from one record to another. For records with extensive data, the combined text of the template and the merged field data can exceed the Large Language Model's (LLM) maximum token limit, resulting in an error. For records with less data, the prompt remains within the limit, and the process succeeds.

Why Incorrect Options are Wrong:

A: LLM token limits are fixed constraints of the model, not a resource that fluctuates based on overall user demand or system load.

C: Switching to a Flex template is a potential way to manage the prompt, but it does not explain the root cause of the error, which is the variable input size from record data.

References:

1. Salesforce Help, Prompt Builder, "Prompt Template Best Practices and Considerations": This document states, "The total number of tokens in your prompt template and its response can’t exceed the model’s maximum token limit... Because the length of merge field values varies, the number of tokens in a resolved prompt template can also vary. Test your prompt templates with merge fields that have a wide range of lengths." This directly confirms that variable record data is the cause of varying token counts.

2. Salesforce Help, Prompt Builder, "Create a Prompt Template to Generate a Field": This guide illustrates how Field Generation templates use merge fields like {!Input} which are resolved with data from the source record. This establishes the mechanism for the dynamic and variable nature of the prompt's content.

Question: 27

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

A: The related list is not on the parent object's page layout.

B: The related list prompt template option is not enabled.

C: The related list is Read Only.

Correct Answer:

A

Explanation:

For a prompt template to successfully use a related list as a grounding source, the related list component must be present on the page layout of the parent record. If the related list is not on the page layout, the prompt execution context cannot access the related records' data. This results in the Large Language Model (LLM) not receiving the necessary contextual information, leading to incorrect or incomplete responses as observed during U-A-T.

Why Incorrect Options are Wrong:

B: There is no specific, separate setting called "related list prompt template option" that needs to be enabled. The ability to use related lists is an integral part of the Prompt Builder's grounding feature.

C: The related list being read-only is not a cause for failure. Grounding involves reading data to provide context to the LLM, not writing to it, so read-only access is sufficient.

References:

1. Salesforce Help. (2024). Ground a Prompt Template with Salesforce Data.

Section: Ground with a Related List

Content: "To ground a prompt template with a related list, the related list must be on the page layout of the parent record." This statement directly confirms that the absence of the related list on the page layout is the cause of the issue.

Question: 28

Universal Containers has an active standard email prompt template that does not fully deliver on the business requirements. Which steps should an Agentforce Specialist take to use the content of the standard prompt email template in question and customize it to fully meet the business requirements?

A: Save as New Version and edit as needed.

B: Clone the existing template and modify as needed.

C: Save as New Template and edit as needed.

Correct Answer:

B

Explanation:

Standard prompt templates provided by Salesforce are not directly editable. To customize a standard template to meet specific business requirements, the correct and prescribed method is to clone it. The cloning process creates a new, independent custom prompt template that inherits the content of the standard one. This new custom template can then be freely modified, activated, and used without altering the original standard template. This approach ensures that the standard templates remain as a baseline while allowing for tailored solutions.

Why Incorrect Options are Wrong:

A: Save as New Version and edit as needed. This action is for iterating on an existing custom prompt template that you have permission to edit, not for creating a customizable copy of a standard template.

C: Save as New Template and edit as needed. While conceptually similar, "Clone" is the specific and accurate term for the action used in the Salesforce UI to create an editable copy of a standard or custom template.

References:

1. Salesforce Help, Prompt Builder, "Create a Prompt Template": "To customize a standard prompt template, clone it and then edit the cloned version. You can’t edit a standard prompt template directly."

2. Salesforce Help, Prompt Builder, "Prompt Template Considerations": "You can’t edit or delete standard prompt templates. To customize a standard prompt template, clone it." This section explicitly states the limitation of standard templates and the required action to customize them.

Question: 29

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.

Correct Answer:

B

Explanation:

The designated and most direct method for integrating an external Large Language Model (LLM) for use with Prompt Builder is through the Bring Your Own LLM (BYO-LLM) functionality within Einstein Studio. This feature allows administrators to securely connect to models hosted on platforms like Amazon SageMaker, Amazon Bedrock, or Google Vertex AI. Once the connection is established and the model is registered in Einstein Studio, it becomes available as a selectable option within Prompt Builder, allowing users to create and manage prompts that leverage the capabilities of that specific external model.

Why Incorrect Options are Wrong:

A: Using Apex is a custom, programmatic solution. The platform provides a standard, declarative feature (BYO-LLM) specifically for this purpose, which is the recommended approach.

C: Flow and External Services are used to invoke an external API as part of a business process, not to register an LLM as a foundational engine for Prompt Builder to use.

References:

1. Salesforce Help, Einstein Studio, "Connect to Your Own Generative AI Model in Einstein Studio": This document outlines the process of using Einstein Studio to connect to external LLMs. It states, "To use your own large language model (LLM) with the Einstein Trust Layer and in prompt templates, first create a connection to the external model provider." This directly supports using Einstein Studio for this purpose.

2. Salesforce Help, Prompt Builder, "Select a Model for Your Prompt Template": This documentation details the process of choosing an LLM within Prompt Builder. It explicitly

mentions that models made available through Einstein Studio will appear in the list of selectable models. It notes, "The list of available LLMs depends on your org's configuration... Your Salesforce admin can also make external LLMs available through Einstein Studio." This confirms the link between Einstein Studio's BYO-LLM feature and Prompt Builder.

3. Salesforce Help, Einstein Generative AI, "Bring Your Own LLM": This section provides an overview of the BYO-LLM capability, explaining its purpose is to "integrate your preferred large language models (LLMs) from third-party providers... with the Salesforce platform." It positions Einstein Studio as the central point for this integration.

Question: 30

Universal Containers wants its AI agent to answer customer questions with precise and up-to-date information. How does an Agentforce Data Library simplify and enable this?

A: It automates the ingestion, taxonomical classification and storage of knowledge in Data Cloud for precision keyword search retrieval to ground prompts and agents with relevant information.

B: It automates the ingestion, Indexing of data, and creates a default retriever to be used in prompts and agents for grounding with relevant information.

C: It automates the ingestion and optical character recognition (OCR) processing of any PDF, and indexes them to enable regular SQL query retrieval to ground prompts and agents with relevant information.

Correct Answer:

B

Explanation:

The Data Library for Salesforce AI agents automates the core stages of the Retrieval-Augmented Generation (RAG) pipeline. It handles the ingestion of data from specified sources (like Salesforce Knowledge), automatically indexes this data by creating vector embeddings, and provides a default retriever. This retriever is then used by AI agents and prompts to search the indexed data for the most relevant information to "ground" the AI's response, ensuring it is precise, contextually aware, and based on up-to-date, proprietary information.

Why Incorrect Options are Wrong:

A: This is incorrect because the primary retrieval method for this type of unstructured data is semantic or vector search, not just "precision keyword search." Taxonomical classification is not the main automated process.

C: This is incorrect because retrieval from an indexed vector store of unstructured text is not performed using "regular SQL query." While OCR is a feature, it's not the sole function.

References:

1. Salesforce Help, Einstein Copilot, "Ground Prompts with Your Salesforce Data": This documentation explains how to connect data sources like Salesforce Knowledge to ground AI agents. It describes the process where Salesforce automatically makes the data

available for retrieval, which aligns with the concept of automated ingestion, indexing, and providing a retriever. The system handles the "how" of retrieval, simplifying the process for the administrator.

2. Salesforce Developers, Einstein Copilot, "Build Reliable AI with Grounding": This resource details the importance of grounding to prevent hallucinations and provide factual answers. It implicitly describes the RAG architecture that the Data Library enables, where a "retriever" finds relevant data from a knowledge base (which must first be ingested and indexed) to inform the Large Language Model (LLM).

3. Salesforce Developers, Blog, "Bring Your Own LLM to Einstein Copilot": In discussions about the architecture of Einstein Copilot, the role of the grounding and data retrieval mechanism is highlighted as a key component that works with the LLM. The automation of creating this retrieval system from sources like a "Data Library" is a core value proposition.

Question: 31

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.

Correct Answer:

C

Explanation:

When creating a Salesforce Agent designed to answer questions using a knowledge base, the Agent Builder provides a pre-configured topic named General FAQ. This topic is specifically designed to handle common questions and, by default, includes the Answers Questions with Knowledge action. This action directly connects the agent to the Salesforce Knowledge base, allowing it to search for relevant articles and generate grounded, accurate responses based on that content.

Why Incorrect Options are Wrong:

A: The default topic is named "General FAQ," not "General Q&A," and the action is "Answers Questions with Knowledge," not "Knowledge Article Answers." This option uses incorrect terminology.

B: The "Answers Questions with LLM" action is a different tool used for more general, ungrounded conversational responses, not for answers specifically derived from the Salesforce Knowledge base.

References:

1. Salesforce Help, Agent Actions: The "Answers Questions with Knowledge" action is officially documented as the standard action to "Search your Salesforce Knowledge base to find relevant articles and generate a response grounded in that content." This confirms the action's name and purpose.

2. Salesforce Help, Create an Agent with the Agent Builder: The documentation on creating a new agent states that it comes with pre-built topics, including "General FAQ," which is

intended to handle general questions and is the logical place to configure knowledge-based answers.

3. Salesforce Help, Agent Topics: This documentation describes the function of topics within the Agent Builder. It details how the "General FAQ" topic serves as the primary container for handling user inquiries that don't match other specific topics, making it the correct location for the "Answers Questions with Knowledge" action.

Question: 32

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.

Correct Answer:

C

Explanation:

Agent Analytics provides purpose-built dashboards and reports to monitor the performance and effectiveness of conversational AI agents. These tools are specifically designed to track key metrics such as user engagement, session outcomes, and popular dialogs. This allows a specialist to evaluate the agent's usability by seeing where users succeed or fail and to verify that actions are being assigned and completed as intended.

Why Incorrect Options are Wrong:

A: Platform Debug Logs are a developer tool for troubleshooting code execution and system processes at a granular level, not for analyzing high-level usability or business-oriented action assignments.

B: The Metadata API is used to manage the setup and configuration (metadata) of a Salesforce org, such as deploying custom objects or page layouts. It cannot be used to query transactional log data.

References:

1. Salesforce Help, Monitor Your Bot's Performance: The Einstein Bot Performance dashboard is the primary tool for monitoring. It includes charts for "Total Sessions," "Engagement Rate," and "Dialogs Used," which directly relate to monitoring usability and actions. The documentation states, "Use the performance dashboard to get a high-level view of how well your bot is serving your customers."

2. Salesforce Help, Bot Analytics Dashboards: This documentation details the standard dashboards available for Einstein Bots, which provide insights into bot sessions, usage, and

performance metrics. These are the core components of what would be considered "Agent Analytics."

3. Salesforce Developer Documentation, Debug Logs: This source defines debug logs as a tool for developers to trace the execution of Apex code, database operations, and system processes. This confirms it is not an analytics tool for monitoring user interactions or business outcomes.

Question: 33

What is a Salesforce Agentforce Specialist able to configure in Data Masking within the Einstein Trust Layer?

A: The profiles exempt from masking

B: The encryption keys for masking

C: The privacy data entities to be masked

Correct Answer:

C

Explanation:

A Salesforce administrator or specialist can configure the data masking settings within the Einstein Trust Layer to specify which types of sensitive data should be masked. This involves selecting from a predefined list of privacy data entities, such as Personally Identifiable Information (PII) like names, email addresses, and phone numbers, as well as Payment Card Industry (PCI) data like credit card numbers. This configuration ensures that when prompts are sent to external Large Language Models (LLMs), sensitive customer data is automatically identified and replaced with placeholders, protecting it from leaving the Salesforce trust boundary.

Why Incorrect Options are Wrong:

A: The Einstein Trust Layer's data masking is a global data protection feature; it does not offer configuration to exempt specific user profiles from the masking rules.

B: Data masking in the Einstein Trust Layer replaces sensitive data with generic placeholders (e.g., [EMAIL]), rather than encrypting it. Administrators do not manage encryption keys for this feature.

References:

1. Salesforce Help, "Configure Data Masking for the Einstein Trust Layer." This document explicitly states, "From the Einstein Trust Layer setup page, you can enable data masking and select the types of sensitive data to mask." It provides a list of selectable entities like Name, Email, Phone, and Credit Card Number.

2. Salesforce Help, "Einstein Trust Layer." This guide describes data masking as a core pillar of the Trust Layer, explaining that it "detects and masks sensitive data... before it’s

sent to the LLM." The configuration focuses on the types of data, not user exemptions or encryption keys.

Question: 34

An Agentforce needs to create a Sales Email with a custom prompt template. They need to ground on the following data. Opportunity Products Events near the customer Tone and voice examples How should the Agentforce Specialist obtain related items?

A: Call prompt initiated flow to fetch and ground the required data.

B: Create a flex template that takes the records in question as inputs.

C: Utilize a standard email template and manually insert the required data fields.

Correct Answer:

A

Explanation:

Salesforce Prompt Builder supports “prompt-initiated flows.”

At runtime the flow queries any objects you choose—Opportunity, OpportunityLineItem (products), Event—and packages those values as the grounding variables that the custom prompt template consumes to draft the sales email in the desired tone. Vendor documentation explicitly states that a prompt-initiated flow is the prescribed way to pull multiple, related CRM records or external data for grounding; no special “flex” template or manual insertion is required.

Why Incorrect Options are Wrong:

B: “Flex template” isn’t a Salesforce feature for Prompt Builder; no mechanism exists to auto-retrieve related records that way.

C: Manual insertion bypasses grounding automation, is error-prone, and contradicts Salesforce’s recommended use of flows for contextual data.

References:

1. Salesforce Winter ’24 Release Notes, Einstein Prompt Builder – “Use a Prompt-Initiated Flow to retrieve related CRM data and ground the prompt” (pg. 211, para. 2).

2. Salesforce Help Portal, article “Create Prompt-Initiated Flows” – “A flow can query any related object records and pass them as input variables to the prompt template” (Section: Configure Data Grounding, steps 3-5).

3. Salesforce Developer Guide, “Prompt Builder Concepts” – “Grounding data is supplied by invoking a flow or Apex action that returns structured variables.” (Chapter 4, Example 2).

Question: 35

An Agentforce is setting up a new org and needs to ensure that users can create and execute prompt templates. The Agentforce Specialist is unsure which roles are necessary for these tasks. Which permission sets should the Agentforce Specialist assign to users who need to create and execute prompt templates?

A: Prompt Template Manager for creating templates and Data Cloud Admin for executing templates

B: Prompt Template Manager for creating templates and Prompt Template User for executing templates

C: Data Cloud Admin for creating templates and Prompt Template User for executing templates

Correct Answer:

B

Explanation:

The correct assignment of permissions for managing and using prompt templates involves two distinct permission sets. The Prompt Template Manager permission set grants users the necessary rights to create, edit, and delete prompt templates. For users who only need to use these templates, the Prompt Template User permission set provides the ability to view and execute them. This separation aligns with the principle of least privilege, ensuring users have only the access they need for their specific roles.

Why Incorrect Options are Wrong:

A: The Data Cloud Admin permission set is not the standard or required permission for executing prompt templates. The correct, more specific permission is Prompt Template User.

C: The Data Cloud Admin permission set is incorrect for creating prompt templates. The designated permission set for this function is Prompt Template Manager.

References:

1. Salesforce Help. (n.d.). Assign Permission Sets for Prompt Builder.

Section: Permission Sets for Prompt Builder

Content: This official documentation explicitly states:

Prompt Template Manager: "Lets users create, edit, and delete prompt templates."

Prompt Template User: "Lets users view and execute prompt templates."

This directly supports the distinction made in the correct answer, where creation is handled by the "Manager" role and execution by the "User" role. It also confirms that Data Cloud Admin is not the specified permission for these core tasks.