

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).

Question: 103

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

Answer: C

Explanation:

Agentforce Sales Agent provides real-time insights and AI-powered recommendations, which are designed to streamline the sales process and help sales representatives focus on key tasks to increase conversion rates. It offers features like lead scoring, opportunity prioritization, and proactive recommendations, ensuring that sales teams can interact with customers efficiently and close deals faster.

Option A: While tracking customer interactions is beneficial, it is only part of the broader capabilities offered by Agentforce Sales Agent and is not the primary objective for improving real-time productivity.

Option B: Agentforce Sales Agent does not automate the entire sales process but provides actionable recommendations to assist the sales team.

Option C: This aligns with the tool's core purpose of enhancing productivity and driving sales success.

Reference:

"Einstein Next Best Action for Sales Teams | Salesforce Trailhead" .

Question: 104

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.

Answer: C

Explanation:

The Prompt Builder in Salesforce has some specific limitations, one of which is that custom objects are supported only for Flex template types. This means that users must rely on Flex templates to integrate custom objects into their prompts.

Option A: While rich text area fields have certain restrictions, this does not pertain to the core limitation of integrating custom objects.

Option B: Updates and creations for prompt templates are indeed recorded in the Setup Audit Trail, so this statement is incorrect.

Option C: This is the correct answer as it reflects a documented limitation of the Prompt Builder.

Reference:

"Prompt Builder Limitations | Salesforce Documentation" .

Question: 105

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.

Answer: B

Explanation:

Einstein Studio (Model Builder) allows organizations to connect and utilize external foundational models while fine-tuning them with company-specific data. This capability is particularly suited to businesses like Universal Containers (UC) that require customization of foundational models to better align with their unique data and use cases.

Option A: Adjusting model temperature is a parameter-level setting for controlling randomness in AI-generated responses but does not necessitate connecting to an external foundational model.

Option B: This is the correct answer because Einstein Studio supports fine-tuning external models with proprietary company data, enabling a tailored and more accurate AI solution for UC.

Option C: Changing frequency penalties is another parameter-level adjustment and does not require external foundational models or Einstein Studio.

Reference:

"Using Einstein Studio to Connect Foundational Models | Salesforce Trailhead" .

Question: 106

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.

Answer: A

Explanation:

To integrate inferences from an XGBoost model into Salesforce's Data Cloud as a stand-alone Data Model Object (DMO):

Create the Serving Endpoint in Databricks:

The serving endpoint is necessary to make the trained model available for real-time inference. Databricks provides tools to host and expose the model via an endpoint.

Configure the Model Using Model Builder:

After creating the endpoint, the AI Specialist should configure it within Einstein Studio's Model Builder, which integrates external endpoints with Salesforce Data Cloud for processing and storing inferences as DMOs.

Option B: Serving endpoints are not created in Einstein Studio; they are set up in external platforms like Databricks before integration.

Option C: A Python SDK connector is not used to bring model inferences into Salesforce Data Cloud; Model Builder is the correct tool.

Reference:

"Einstein Studio and Model Integration with External Endpoints | Salesforce Trailhead" .

Question: 107

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:

Einstein analyzes the details of the customer interaction, including chat or email threads, to extract relevant information for the summary.

Knowledge Articles:

It considers linked Knowledge Articles or articles referred to during the case resolution process, ensuring the summary incorporates accurate resolutions or additional resources provided to the customer.

Cases:

Einstein also examines historical cases and related case records to ground the summary in context from past resolutions or interactions.

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" .

Question: 108

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.

Answer: A

Explanation:

When a custom Agent action is not being prioritized correctly by the planner service, the root cause is often missing or improperly defined action dependencies. The planner service determines the execution order of actions based on dependencies defined in the action instructions. To resolve this, the AI Specialist must explicitly specify dependent actions using their API names in the custom action's configuration. This ensures the planner understands the sequence in which actions must be executed to meet business logic requirements.

Salesforce documentation highlights that dependencies are critical for orchestrating workflows in Einstein Bots and Agentforce. For example, if Action B requires data from Action A, Action A's API name must be listed as a dependency in Action B's instructions. The Einstein Bot Developer Guide states that failing to define dependencies can lead to race conditions or incorrect execution order.

In contrast:

Profiles or custom permissions (B) control access to the action but do not influence execution order. LLM model provider and version (C) determine the AI model used for processing but are unrelated to the planner's sequencing logic.

Reference:

Salesforce Help Article: [Configure Custom Actions for Einstein Bots](#) (Section: "Defining Action Dependencies").

Einstein Bot Developer Guide: "Orchestrating Workflows with the Planner Service" (Dependency

Management best practices).

Question: 109

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

Answer: B

Explanation:

To use the Record Snapshots grounding feature in a prompt template, you must create a field set that includes all fields required for grounding. Field sets define which fields from an object are accessible to the AI model, ensuring the prompt template has structured data to generate contextually accurate responses. Salesforce documentation emphasizes that grounding relies on explicitly defined field sets to avoid exposing unintended data and to comply with security policies.

Page Layout configuration (A) controls UI organization but does not directly enable grounding. Dynamic Forms (C) customize record pages dynamically but are unrelated to data grounding for prompts.

Reference:

Salesforce Help Article: [Grounding with Record Snapshots](#) ("Configuring Field Sets for Grounding" section).

Einstein GPT Implementation Guide: "Preparing Data for AI Prompts" (Field Set requirements).

Question: 110

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.

Answer: C

Explanation:

Action Instructions are critical for defining how a custom Agent action should be executed, ensuring alignment with the intended user experience. They provide step-by-step guidance to the bot or LLM on logic, data handling, and integration workflows, directly impacting how users interact with the action. For example, clear instructions prevent errors in API calls or data processing, ensuring seamless interactions.

Salesforce documentation states that poorly defined instructions lead to mismatched expectations, while well-structured instructions ensure the action behaves predictably. This aligns with delivering a

consistent user experience.

A refers to user invocation, which is handled by dialogue flows, not instructions.

B is incorrect because the LLM selects actions based on context/intent, not instructions.

Reference:

Salesforce Help Article: [Custom Action Configuration](#) ("Defining Action Instructions" section).

Einstein Bot Developer Guide: "Designing User-Centric Actions with Instructions" (User Experience alignment).

Question: 111

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

Answer: B

Explanation:

Knowledge Replies is the Einstein Generative AI capability that generates recommended responses for service agents by grounding responses in the organization's Knowledge articles. This reduces handling time by providing contextually accurate suggestions sourced directly from approved content. Salesforce documentation states that Knowledge Replies leverage natural language processing (NLP) to match customer inquiries with relevant articles and draft replies, ensuring consistency and compliance.

Case Replies (A) focus on generating responses based on case data (e.g., case fields, history) but do not explicitly ground responses in Knowledge articles.

Service Replies (C) is not a standard Einstein capability tied to Knowledge-driven responses.

Reference:

Salesforce Help Article: [Einstein for Service – Knowledge Replies](#) ("How Knowledge Replies Work" section).

Einstein for Service Implementation Guide: "Reducing Handle Time with Knowledge-Driven Responses."

Question: 112

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

Answer: C

Explanation:

Dynamic Grounding in the Einstein Trust Layer architecture ensures that large language model (LLM) prompts are enriched with organization-specific data (e.g., Salesforce records, Knowledge articles) to generate accurate and relevant responses. By dynamically injecting contextual data into prompts, it reduces hallucinations and aligns outputs with trusted business data.

Prompt Defense (A) focuses on blocking malicious inputs or prompt injections but does not enhance responses with organizational data.

Data Masking (B) redacts sensitive information but does not contribute to grounding responses in business context.

Reference:

Salesforce Help Article: [Einstein Trust Layer – Dynamic Grounding](#) ("How Dynamic Grounding Works" section).

Einstein Trust Layer Technical Overview: "Contextual Accuracy with Dynamic Grounding."

Question: 113

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 Salesforce 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).

Question: 114

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.

Answer: C

Explanation:

Einstein Trust Layer's Toxicity Scoring categorizes content on a scale of 0 to 1, where 1 indicates the highest level of toxicity (e.g., harmful, biased, or inappropriate language). This scoring helps organizations filter unsafe AI-generated content. A score of 1 triggers mitigation actions, such as blocking the response or alerting administrators.

A score of 0 would indicate no toxicity (B is incorrect).

The scoring system does not use "least toxic" as a category (A is misleading).

Reference:

Salesforce Help Article: [Einstein Trust Layer – Toxicity Scoring](#) ("Interpreting Toxicity Scores" section).

Einstein GPT Safety Overview: "Mitigating Harmful Content with Toxicity Detection."

Question: 115

An AI Specialist at Universal Containers is trying to set up a new Field Generation prompt template. They take the following steps.

1. Create a new Field Generation prompt template.
2. Choose Case as the object type.
3. Select the custom field AI_Analysis_c as the target field.

After creating the prompt template, the AI Specialist saves, tests, and activates it. However, 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 Einstein Service Replies should be used instead.
- C. They forgot to edit the Lightning page layout and associate the field to a prompt template

Answer: C

Explanation:

For Field Generation prompt templates to display the Sparkle icon (indicating AI-generated content), the target field must be explicitly associated with the prompt template on the Lightning page layout. Even if the prompt template is activated, failing to add the field to the page layout and link it to the

template will result in the field behaving as a standard field. Salesforce documentation emphasizes that page layout configuration is mandatory to enable AI-driven field interactions.

Reactivating the layout (A) is unnecessary unless the layout itself was modified after activation.

Case objects are supported for Field Generation (B is incorrect).

Reference:

Salesforce Help Article: [Configure Field Generation Prompt Templates](#) ("Associating Fields with Page Layouts" section).

Einstein GPT Implementation Guide: "Enabling AI-Generated Fields in Lightning Pages."

Question: 116

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 errors, incorrectly triggered actions, and incomplete plans?

- A. Event Logs
- B. Plan Canvas
- C. Agent Settings

Answer: B

Explanation:

Event Logs in Salesforce capture detailed interaction data, including agent errors, triggered actions, and incomplete plans. These logs provide visibility into user-Agent interactions for troubleshooting performance issues. The Einstein Bot Analytics documentation highlights Event Logs as the primary source for auditing bot behavior and diagnosing issues like misconfigured actions or plan execution failures.

Plan Canvas (B) is for designing workflows, not auditing.

Agent Settings (C) control configuration but do not store interaction history.

Reference:

Salesforce Help Article: [Einstein Bot Analytics and Logs](#) ("Accessing Event Logs" section).

Einstein Bot Developer Guide: "Monitoring and Debugging Agent Performance."

Question: 117

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

- A. Enable a sales 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 faster.
- 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.

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 AI Specialist Trailhead: "Sales Automation with Agentforce" (Natural Language Task Execution).

Question: 118

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.

Answer: A

Explanation:

For no-code implementations, Prompt Builder allows AI Specialists to create prompt templates that dynamically ground responses in Salesforce CRM data (e.g., past purchases) and Data Cloud insights (e.g., product interests) without custom coding. The standard foundation model (e.g., Einstein GPT) can synthesize this data into summaries, leveraging structured and unstructured sources (e.g., email/phone transcripts). Fine-tuning (B) or custom models (C) require code and are unnecessary here, as the use case does not involve unique data patterns requiring model retraining.

Reference:

Salesforce Help Article: [Prompt Builder for No-Code AI](#) ("Grounding in CRM and Data Cloud" section).

Einstein GPT Implementation Guide: "Generating Summaries with Pre-Built Models."

Question: 119

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 create AI-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."

Question: 120

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 appropriate 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 list 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 AI Specialist Trailhead: "How Agentforce Processes User Requests."

Question: 121

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 Resolution

Answer: C

Explanation:

Einstein Work Summaries automatically generate concise summaries of customer interactions (e.g., chat transcripts). Beyond the "summary" field, it extracts and populates Issue (key problem discussed) and Resolution (action taken to resolve the issue). These fields help agents and supervisors quickly grasp the conversation's context without reviewing the full transcript.

Sentiment Analysis and Emotion Detection (Option A): While Einstein Conversation Insights provides sentiment scores and emotion detection, these are separate from Work Summaries. Work Summaries focus on factual summaries, not sentiment.

Draft Survey Request Email (Option B): Not part of Work Summaries. This would require automation tools like Flow or Email Studio.

Issue and Resolution (Option C): Directly referenced in Salesforce documentation as fields populated by Einstein Work Summaries.

Reference:

Salesforce Help Article: [Einstein Work Summaries](#)

Einstein Work Summaries focus on "key details like Issue and Resolution" alongside summaries.

Contrast with [Einstein Conversation Insights](#) for sentiment/emotion analysis.

Question: 122

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 in the advanced settings for the permissions, field-level security, and sharing settings to be respected.
- B. The custom action adheres to the permissions, field-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.

Answer: B
