

DUMPSBOSS.

Salesforce Certified AI Specialist Exam

Salesforce Salesforce-AI-Specialist

Version Demo

Total Demo Questions: 17

Total Premium Questions: 174

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, AI Fundamentals	27
Topic 2, Data for AI	5
Topic 3, Einstein Copilot	63
Topic 4, Prompt Engineering	65
Topic 5, Model Builder	14
Total	174

QUESTION NO: 1

When configuring a prompt template, an AI Specialist previews the results of the prompt template they've written. They see two distinct text outputs: Resolution and Response.

Which information does the Resolution text provide?

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

ANSWER: A

Explanation:

It shows the full text that is sent to the Trust Layer. In Prompt Builder preview, Salesforce first resolves the prompt template by replacing merge fields and grounding references with values from the selected sample record. The resulting resolved prompt is the complete set of instructions and contextual data that will proceed through the Einstein Trust Layer on its way to the large language model. Reviewing this Resolution is important because it lets an AI Specialist verify that record data, related-list data, and template instructions have been assembled as intended before generation occurs. It is therefore the best place to validate the actual prompt content created from the template and the sample data. Salesforce's prompt-template preview workflow distinguishes this resolved prompt from the generated model output, enabling authors to inspect the input assembled for the AI request and refine the template when necessary. See Salesforce's [Prompt Builder documentation](#) and the [Einstein Trust Layer documentation](#) for details about prompt processing and trusted generation.

QUESTION NO: 2

Universal Containers has a prompt template that produces a concise account update for sales leaders. In Model Playground, the AI Specialist observes that the generated response repeats the same customer concern several times even when the prompt and grounding data contain the concern only once.

Which hyperparameter should the AI Specialist adjust to discourage this repetition?

- A. Temperature
- B. Frequency Penalty
- C. Presence Penalty

ANSWER: B

Explanation:

Frequency Penalty is the most appropriate hyperparameter because it discourages the model from repeatedly using words or phrases that have already appeared in the generated response. Increasing it can help produce a less repetitive account update while the specialist continues to test against representative records.

Temperature primarily controls the degree of variation or randomness in output and is not the direct repetition control. Presence Penalty encourages the model to introduce new concepts, but Frequency Penalty specifically addresses repeated tokens and phrases.

QUESTION NO: 3

Universal Containers implements three custom actions to get three distinct types of sales summaries for its users. Users are complaining that they are not getting the right summary based on their utterances. What should the AI Specialist investigate as the root cause?

- A. Review that the custom action is assigned to an Agent.

B. Review the action Instructions to ensure they are unique.

C. Ensure the input and output types are correctly chosen.

ANSWER: B

Explanation:

Review the action Instructions to ensure they are unique. In Agentforce, action instructions tell the large language model the purpose of an action, when it should select that action, and the information it needs to carry it out. When several custom actions produce related sales summaries, vague or substantially overlapping instructions can make it difficult for the agent to reliably determine which summary a user is requesting. The result can be selection of a valid sales-summary action that is not the one intended by the user's wording.

The AI Specialist should compare the instructions for all three actions and make each one explicitly distinguishable. Instructions should identify the specific summary type, relevant dimensions, and applicable time frame or channel, such as regional quarterly performance, product-level annual performance, or channel comparison. Including clear selection criteria and representative user phrasing gives the model stronger signals for choosing the intended action. The instructions should focus on the business outcome and conditions that differentiate each action, rather than using nearly identical language such as "provide a sales summary." Salesforce's Agentforce guidance emphasizes writing clear, specific action instructions so the agent can choose and invoke the appropriate action. See [Salesforce Help: Agentforce Custom Actions](#) and [Trailhead: Agentforce Basics](#).

QUESTION NO: 4

Universal Containers needs to determine which of three retention offers should be presented to each customer. The organization must also show service managers the factors that influenced the selected offer, such as customer segment, contract value, and recent support history.

Which AI approach is most appropriate?

A. Use a generative model to write a new retention offer for every customer.

B. Use a predictive model to classify the most appropriate retention offer.

C. Use a Record Summary prompt template to select the retention offer.

ANSWER: B

Explanation:

Use a predictive model to classify the most appropriate retention offer. The requirement is to choose among a defined set of outcomes using customer attributes and to explain the factors that influenced the selection. Predictive AI is suited to classification and recommendation use cases and supports examination of the factors that drive a prediction.

A generative model could draft the wording of an offer, but it is not the primary solution for selecting and explaining one outcome from a controlled set. A Record Summary prompt template summarizes record information and does not create an explainable prediction for offer selection.

QUESTION NO: 5

What is the role of the large language model (LLM) in understanding intent and executing an Agent Action?

A. Identify the best matching topic and actions and correct order of execution

B. Determine a user's topic access and sort actions by priority to be executed

C. Find similar requested topics and provide the actions that need to be executed

ANSWER: A

Explanation:

The correct answer is “Identify the best matching topic and actions and correct order of execution.” In Agentforce, the LLM acts as the reasoning engine that interprets the user’s natural-language request in context and maps it to the most appropriate configured topic. Topics establish the scope of work an agent can handle, while actions are the operations available to fulfill work within that scope, such as retrieving CRM data, invoking a flow, or calling an API. After selecting the applicable topic, the LLM uses the topic instructions, action descriptions, and conversation context to decide which action or actions are needed and to orchestrate them in an appropriate sequence. This enables an agent to handle varied user phrasing without relying on a rigid, keyword-only routing model. The LLM’s planning is grounded in the agent’s configured topics, instructions, and available actions, helping the agent move from intent recognition to an executable plan. Salesforce describes Agentforce as using reasoning to understand requests and take actions, with topics and actions defining the agent’s capabilities. See [Salesforce Trailhead: Learn About Agentforce](#) and [Salesforce Developer Documentation: Einstein Generative AI Overview](#).

QUESTION NO: 6

An AI Specialist is testing a prompt template that creates executive account briefs. The output includes plausible recommendations that are not supported by the account record or the approved sales methodology.

Which refinement should the AI Specialist make first?

- A. Increase the model temperature to provide a wider range of recommendations.
- B. Ground the template with required account data and approved methodology content, and instruct it to use only those sources.
- C. Add an executive-sales persona without changing the grounding resources.

ANSWER: B

Explanation:

Ground the template with the required account data and approved sales methodology content, and instruct it to use only those sources. Grounding supplies the factual context the model needs to create an account-specific brief. Explicit instructions to use only supplied information further reduce unsupported recommendations and make the expected response boundary clear.

Increasing temperature would generally increase variation and is unsuitable when factual consistency is required. A more detailed persona can improve tone or perspective, but it does not establish the factual source for recommendations. The specialist should test the refined template with representative accounts before activating a new version.

QUESTION NO: 7

A sales representative can run an activated prompt template, but several customer attributes referenced by the template appear as unresolved placeholders in the generated response. An administrator confirms that the template is active and that the representative has the Prompt Template User permission set.

What should the AI Specialist investigate next?

- A. The model temperature configured for the prompt template.
- B. Field-level security and record access for the customer attributes used as grounding data.
- C. The prompt template version number that is currently active.

ANSWER: B

Explanation:

Field-level security and record access for the customer attributes used as grounding data. Secure Data Retrieval respects the permissions of the user who executes the prompt. If that user cannot read a referenced field or underlying record, Salesforce cannot retrieve the value for prompt grounding and does not expose it through the generated response.

The Prompt Template User permission set enables a user to execute templates, but it does not grant access to every object, record, or field referenced by a template. Changing the model temperature or creating a new template version would not resolve missing data caused by the user's Salesforce access controls.

QUESTION NO: 8

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:

To streamline the sales process and increase conversion rates is correct because Agentforce for Sales is designed to help sales organizations act on CRM data and AI-generated guidance in the flow of work. By surfacing relevant context, prioritizing sales activity, and supporting sellers with recommended next steps, it reduces the administrative effort that can slow down customer engagement. This enables representatives to focus more of their time on meaningful conversations and timely follow-up.

For example, sales agents can assist with work such as researching accounts, preparing for meetings, managing leads, and progressing opportunities. These capabilities help teams use current customer and pipeline information more consistently when deciding whom to contact, what action to take, and how to move a deal forward. Better prioritization and more relevant interactions support a more efficient sales process and can improve the likelihood that qualified opportunities advance toward conversion.

Salesforce describes Agentforce for Sales as providing autonomous and assistive capabilities across the sales lifecycle, while its Sales Cloud AI materials explain how AI helps sellers focus on high-value work and improve productivity. See [Salesforce Agentforce for Sales](#) and [Salesforce AI for Sales](#).

QUESTION NO: 9

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:

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. This reflects Agentforce's agentic architecture: the model first understands the user's natural-language goal and uses the available topic and action definitions as grounded capabilities. It then determines the appropriate sequence of steps, such as identifying that the request concerns opportunity records, applying the meaning of "my" and "open" in the current user context, invoking the relevant Salesforce action, and presenting the returned records. Agentforce topics define the areas an agent can handle, while actions provide the concrete operations the agent can perform. The reasoning and planning process enables the agent to select and orchestrate those capabilities dynamically instead of requiring the user to follow a fixed command format. Actions also execute within Salesforce's established permissions and data-access controls, so the resulting opportunity list is scoped to data the sales manager is authorized to access. Salesforce describes this model of grounding agents with topics and actions in its [Agentforce developer documentation](#) and explains agent configuration concepts in the [Salesforce Help documentation for Agentforce topics](#).

QUESTION NO: 10

Universal Containers (UC) is tracking web activities in Data Cloud for a unified contact, and wants to use that in a prompt template to help extract insights from the data.

Assuming that the Contact object is one of the objects associated with the prompt template, what is a valid way for DC to do this?

- A. Call the prompt directly from Data Cloud with a web tracing activity included in the prompt definition.
- B. Add the activity records as an enrichment related list to the Contact then pass the Contact into a prompt template workspace using related list grounding.
- C. Create a prompt template that takes a list of all Data Cloud activity records as input to pass to the large language model (LLM).

ANSWER: B

Explanation:

Add the activity records as an enrichment related list to the Contact then pass the Contact into a prompt template workspace using related list grounding. is valid because it makes the unified contact's relevant Data Cloud web-activity data available in the record context used by the prompt template. Data Cloud can surface data associated with a unified individual through related-list enrichment, allowing the Contact record to act as the prompt's primary business context. Related-list grounding then retrieves the specified related activity data when the prompt runs and supplies it as grounded context for the generated response.

This approach lets the template use meaningful behavioral signals, such as recent browsing or engagement activity, while retaining the Contact as the record being processed. The prompt author can select the related list and its fields deliberately, so the model receives a focused, structured representation of activity rather than an uncontrolled payload. That supports insight-generation use cases while aligning prompt execution with Salesforce's record-based grounding model. Salesforce's Prompt Builder learning materials describe grounding prompts with CRM data and related records, and Data Cloud documentation explains how unified customer data can be activated for Salesforce use cases. See [Salesforce Trailhead: Prompt Builder Basics](#) and [Salesforce Help: Data Cloud](#).

QUESTION NO: 11

Universal Containers wants to be able to detect with a high level confidence if content generated by a large language model (LLM) contains toxic language.

Which action should an AI Specialist take in the Trust Layer to confirm toxicity is being appropriately managed?

- A. Access the Toxicity Detection log in Setup and export all entries where isToxicityDetected is true.
- B. Create a flow that sends an email to a specified address each time the toxicity score from the response exceeds a predefined threshold.

C. Create a Trust Layer audit report within Data Cloud that uses a toxicity detector type filter to display toxic responses and their respective scores.

ANSWER: C

Explanation:

Create a Trust Layer audit report within Data Cloud that uses a toxicity detector type filter to display toxic responses and their respective scores. The Einstein Trust Layer records generative AI interactions in its audit trail, including the safety-detector outcomes applied to generated content. A report filtered for the toxicity detector gives the AI Specialist a centralized, reviewable view of the responses identified as toxic and the associated toxicity scores. This directly supports ongoing validation that the configured toxicity guardrail is detecting and handling unsafe output as intended.

Using the audit data also makes the review operationally useful: the specialist can inspect patterns over time, identify prompts or use cases that produce higher-risk output, and use the recorded scores to assess whether the organization's safety posture is appropriate. The Trust Layer is designed to apply safety protections such as toxicity detection and to retain audit information so organizations can monitor AI activity while maintaining governance controls. Salesforce describes these Trust Layer protections and their role in securing and monitoring generative AI interactions in its [Einstein Trust Layer Trailhead unit](#) and [Einstein Generative AI Trust Layer documentation](#).

QUESTION NO: 12

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:

UC wants a model fine-tuned using company data. This requirement calls for a model that can be tailored to the organization's domain, terminology, patterns, and business context rather than relying solely on the behavior of a general-purpose foundation model. Einstein Studio Model Builder supports a bring-your-own-model approach: an organization can connect Salesforce to supported external model providers and make an appropriate externally hosted model available for use in Salesforce generative AI workflows. This gives UC flexibility to select a model whose capabilities, hosting arrangement, and customization approach fit its specific use case.

Fine-tuning is especially relevant when prompt instructions and grounding alone do not sufficiently align a model with specialized company language or recurring task behavior. The fine-tuning process itself is performed according to the external provider's supported capabilities and data-handling process; Model Builder provides the Salesforce connection and governance path for using the resulting model. This approach should be evaluated alongside Salesforce trust, security, data-access, and responsible-AI requirements before production use. Salesforce describes Model Builder as the capability for connecting and managing models from external providers in Einstein Studio. See [Salesforce Einstein](#) and [Salesforce Help: Model Builder](#).

QUESTION NO: 13

Universal Containers is evaluating a custom external model for use in Prompt Builder. The model has already been trained and is hosted by an approved provider. The organization wants the model to be centrally configured and governed for use with Salesforce generative AI capabilities.

Which solution should the AI Specialist recommend?

- A. Embed the provider's model application in Salesforce using an iFrame.
- B. Configure the external model in Einstein Studio Model Builder.
- C. Create direct callouts from each prompt-enabled business process to the model endpoint.

ANSWER: B

Explanation:

Configure the external model in Einstein Studio Model Builder. Model Builder provides the supported Salesforce approach for bringing an externally hosted model into the platform so it can be managed and used in generative AI workflows. This supports centralized model configuration rather than requiring each business solution to implement its own direct integration.

Embedding a provider application in an iFrame does not make the model available as a governed generative AI model in Salesforce. Direct custom callouts can be useful for bespoke integrations, but they do not meet the requirement for centralized model management through Salesforce AI capabilities.

QUESTION NO: 14

Universal Containers (UC) is looking to enhance its operational efficiency. UC has recently adopted Salesforce and is considering implementing Einstein Copilot to improve its processes.

What is a key reason for implementing Einstein Copilot?

- A. Improving data entry and data cleansing
- B. Allowing AI to perform tasks without user interaction
- C. Streamlining workflows and automating repetitive tasks

ANSWER: C

Explanation:

Streamlining workflows and automating repetitive tasks is a key reason to implement Einstein Copilot. Einstein Copilot is designed to assist users conversationally within Salesforce, helping them complete work more efficiently by using trusted CRM data and invoking configured actions. For example, it can help users summarize records, draft or refine content, retrieve relevant information, and initiate business processes through copilot actions. This reduces the manual effort required for common, repeatable activities and lets employees spend more time on customer-facing or higher-value work.

Operational efficiency comes not merely from generating text, but from connecting generative AI assistance to an organization's Salesforce workflows, data, and actions. Administrators can configure capabilities that align with business processes and apply Salesforce's security, permission, and trust controls. This makes the assistant useful for accelerating everyday work while keeping people appropriately involved in reviewing and acting on business outcomes. Salesforce describes how Einstein Copilot uses actions to help users complete tasks in context, and its AI guidance emphasizes grounding AI experiences in enterprise data and workflows. See [Salesforce Trailhead: Einstein Copilot Basics](#) and [Salesforce Agentforce](#).

QUESTION NO: 15

An AI Specialist is tasked to optimize a business process flow by assigning actions to agents within the Salesforce Agentforce Platform.

What is the correct method for the AI Specialist to assign actions to an Agent?

- A. Assign the action to a topic in Agentforce Builder.
- B. Assign the action to a Topic first on the Agent Actions detail page.

C. Assign the action to a Topic first on Action Builder.

ANSWER: A

Explanation:

Assign the action to a topic in Agentforce Builder. In Agentforce, a topic defines the business area or job an agent can handle, including its classification guidance, instructions, and the actions that are available while the agent is handling that topic. The AI Specialist opens the target agent in Agentforce Builder, selects the appropriate topic, and adds the action from that topic's Actions area. This association gives the agent the action in the relevant conversational and business context, allowing the reasoning engine to determine when the action should be invoked for a request that matches the topic. Actions can include standard Agentforce actions, flows, Apex-based actions, and other supported integrations, but they must be made available through a topic before the agent can use them for that topic's work. After configuring the topic and action, the specialist should save and test the agent to confirm that the action is selected and executed as intended. Salesforce's Agentforce configuration guidance describes topics as the place where agents are equipped with actions, while the builder provides the interface for managing the agent's topics and their action assignments. See [Salesforce Help: Agentforce Topics](#) and [Trailhead: Meet Agentforce](#).

QUESTION NO: 16

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:

Create the serving endpoint in Databricks, then configure the model using Model Builder. is correct because the trained XGBoost model remains hosted in Databricks and must first be made available for inference through a Databricks Model Serving endpoint. A serving endpoint provides the callable, managed deployment interface that receives prediction requests and returns model outputs. Databricks documents how Model Serving deploys registered machine learning models as REST API endpoints: [Databricks Model Serving documentation](#).

After the endpoint exists, the AI Specialist uses Einstein Studio Model Builder to register and configure the external model connection for Data Cloud. Model Builder is the Data Cloud capability used to define how an externally hosted model is invoked, identify its inputs and outputs, and operationalize its prediction results for use in Data Cloud. In this scenario, that configuration enables product-recommendation inferences to be represented and stored as a stand-alone Data Model Object, making the results available for downstream segmentation, activation, and analytics. Salesforce describes Einstein Studio as the environment for connecting and managing models in Data Cloud: [Salesforce Einstein Studio documentation](#).

QUESTION NO: 17

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

Which solution should UC consider?

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

C. Enable Agentforce and deploy to sales users.

ANSWER: C

Explanation:

Enable Agentforce and deploy to sales users is the appropriate solution because Agentforce provides Salesforce-native AI agents that can use the organization's trusted business context, including CRM records, Salesforce knowledge, and other configured data sources, when responding to users. This addresses the key limitation of a general public AI assistant: it has no inherent awareness of UC's accounts, opportunities, products, customer interactions, or internal business processes. A sales user can therefore receive responses and assistance that are relevant to the customer and work already managed in Salesforce.

Agentforce is designed to connect AI reasoning and actions to Salesforce data and business workflows while operating within Salesforce's trust and security framework. UC can configure the agent's topics, actions, and grounding sources so that the agent supports practical sales work, such as preparing for customer conversations, summarizing account context, and guiding next steps. Access to data continues to respect the permissions and controls configured in Salesforce, helping UC make generative AI useful without relying on an uncontextualized public assistant.

Salesforce describes Agentforce and its use of trusted enterprise data at [Salesforce Agentforce](#). For implementation concepts and trusted grounding, see [Salesforce Help: Agentforce](#).