DUMPSARENA

Salesforce Certified AI Associate

Salesforce Salesforce-AI-Associate

Version Demo

Total Demo Questions: 12

Total Premium Questions: 125

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, AI Fundamentals	9
Topic 2, AI Capabilities	21
Topic 3, Ethical Considerations of AI	49
Topic 4, Data for AI	44
Topic 5, Einstein Generative AI	2
Total	125

QUESTION NO: 1

Cloud Kicks wants to reduce the time service agents spend reading long case histories before responding to customers.

Which use of generative AI is most appropriate for this goal?

A. Generate a summary of the case history

B. Predict the next month's case volume

C. Assign conversion scores to customer leads

ANSWER: A

Explanation:

Generating a summary of the case history is the most appropriate use because generative AI can synthesize unstructured information into a concise, readable overview. An agent can use the summary to understand the customer's issue, prior interactions, and actions already taken before deciding how to respond.

Predicting case volume supports workforce planning rather than understanding an individual case. Lead scoring is a sales use case and does not address service-case review.

QUESTION NO: 2

A Salesforce administrator creates a new field to capture an order's destination country.

Which field type should they use to ensure data quality?

A. Text

B. Picklist

C. Number

ANSWER: B

Explanation:

Picklist is the appropriate field type because it presents users with a defined set of country values rather than allowing unrestricted entry. For an order's destination country, this promotes consistent, standardized records: users select the same approved value for a country each time, instead of introducing variations caused by spelling, abbreviations, capitalization, or alternate names. Standardized country data improves the reliability of reports, list views, automations, validation logic, integrations, and any AI or analytics processes that rely on the field. An administrator can maintain the available values centrally, make the field required where appropriate, and restrict users to the configured values so that only intended entries are saved. Salesforce describes picklists as fields that let users select a value from a predefined list, and administrators can control whether users may enter values outside that list. This makes a picklist well suited to categorical information such as a destination country. See Salesforce's Custom Field Types documentation and Picklist Fields documentation.

QUESTION NO: 3

A sales manager is looking to enhance the quality of lead data in their CRM system.

Which process will most likely help the team accomplish this goal?

A. Redesign the lead conversion process,

B. Review and update missing lead information.

C. Prioritize active leads quarterly.

ANSWER: B

Explanation:

Review and update missing lead information is the process most directly associated with improving lead-data quality. High-quality CRM data is complete, accurate, current, and usable for business processes. When lead records lack key details—such as contact information, company attributes, lead source, industry, or qualification data—sales teams have less context for outreach and may be unable to segment, route, score, or prioritize leads reliably. Establishing a regular process to identify incomplete lead records and enrich or correct them improves completeness and makes the information more useful throughout the lead lifecycle.

In Salesforce, complete lead data supports consistent reporting, effective automation, and better sales engagement. Teams can use required fields, validation rules, duplicate-management capabilities, list views, reports, and data-governance practices to detect and prevent gaps in the information captured for leads. Maintaining these records over time is important because data can become outdated as prospects change roles, companies, or contact details. Salesforce's data-quality learning materials emphasize managing data quality as an ongoing activity, including identifying missing or incomplete values and maintaining trusted records. See Salesforce Trailhead's Data Quality module and Salesforce Help on Leads for related guidance.

QUESTION NO: 4

What is a potential outcome of using poor-quality data in AI application?

A. AI model training becomes slower and less efficient

B. AI models may produce biased or erroneous results.

C. AI models become more interpretable

ANSWER: B

Explanation:

AI models may produce biased or erroneous results. is the correct outcome because AI systems learn patterns from the data used to train, ground, or evaluate them. When that data is inaccurate, incomplete, inconsistent, outdated, unrepresentative, or improperly labeled, the system can learn flawed relationships and carry those flaws into its predictions, generated content, classifications, or recommendations. For example, a dataset that fails to represent relevant customer groups can lead to outputs that perform unevenly across those groups. Similarly, incorrect or stale records can cause an AI application to produce unreliable responses even when the underlying model is technically functioning as designed.

Salesforce's trusted AI guidance emphasizes that data quality and governance are foundational to reliable AI outcomes. Organizations should establish data-quality controls, use relevant and representative data, and monitor AI outputs so that issues can be identified and addressed over time. This supports the accuracy, fairness, and trustworthiness expected from AI applications. See Salesforce's Trusted AI resources and the NIST AI Risk Management Framework for guidance on managing data-related AI risks.

QUESTION NO: 5

What are the potential consequences of an organization suffering from poor data quality?

A. Low employee morale, stock devaluation, and inability to attract top talent

B. Revenue loss, poor customer service, and reputational damage

C. Technical debt, monolithic system architecture, and slow ETL throughput

ANSWER: B

Explanation:

Revenue loss, poor customer service, and reputational damage is correct because reliable data is essential to every customer-facing and business decision. When customer, product, sales, or service data is incomplete, duplicated, outdated, or inaccurate, teams can target the wrong people, make incorrect recommendations, miss sales opportunities, and base forecasts on misleading information. These issues can directly reduce revenue through inefficient marketing, lost conversions, avoidable operational costs, and poor prioritization of sales and service activity.

Customer service also depends on having a complete and trustworthy view of the customer. High-quality data helps representatives understand prior interactions, preferences, entitlements, and current issues, allowing them to provide relevant, consistent support. Inaccurate records can instead create frustrating interactions, such as repeating questions, using incorrect contact details, or presenting irrelevant offers. Over time, these experiences erode confidence in the organization and can harm its reputation. This is especially important for AI-enabled experiences: Salesforce emphasizes that trusted, well-governed data is the foundation for useful, accurate, and responsible AI outputs. Salesforce's guidance on duplicate management and matching and its data-quality overview describe the importance of maintaining accurate, consistent, complete data for effective operations and customer relationships.

QUESTION NO: 6

Cloud Kicks learns of complaints from customers who are receiving too many sales calls and emails.

Which data quality dimension should be assessed to reduce these communication Inefficiencies?

A. Duplication

B. Usage

C. Consent

ANSWER: A

Explanation:

Duplication is the appropriate data quality dimension because repeated customer records can cause the same person to be included in outreach more than once. When a customer exists as multiple Lead, Contact, Person Account, or related marketing records, different sales representatives, campaigns, automations, or integrated systems may each treat those records as separate people. The result is unnecessary repeat calls and emails, fragmented interaction history, and an inconsistent customer experience. Assessing duplication means identifying records that represent the same customer and determining which fields can reliably be used for matching, such as email address, phone number, name, or account identifiers. Cloud Kicks can then apply matching rules and duplicate rules to flag likely duplicate records, merge appropriate records, and establish an ongoing process for preventing new duplicates from entering Salesforce. This creates a more complete customer profile and helps ensure communication activity is coordinated against one trusted record rather than several copies. Salesforce provides duplicate-management capabilities, including matching rules and duplicate rules, to support this type of data-quality control. See Salesforce Trailhead: Duplicate Management and Salesforce Help: Manage Duplicate Records.

QUESTION NO: 7

What role does data quality play in the ethical us of AI applications?

A. High-quality data is essential for ensuring unbiased and fair AI decisions, promoting ethical use, and preventing discrimination.

B. High-quality data ensures the process of demographic attributes requires for personalized campaigns.

C. Low-quality data reduces the risk of unintended bias as the data is not overfitted to demographic groups.

ANSWER: A

Explanation:

High-quality data is essential for ensuring unbiased and fair AI decisions, promoting ethical use, and preventing discrimination. AI systems learn patterns from the data used to train, ground, test, and evaluate them. Therefore, data must be accurate, sufficiently complete, relevant to the intended use case, consistently defined, timely, and representative of the population affected by the system. These qualities help teams produce more reliable outputs and identify gaps that could result in unfair treatment or disparate outcomes for particular groups.

Data quality is also an ongoing governance responsibility rather than a one-time preparation task. Organizations should assess source reliability, document data provenance, monitor for drift or changing conditions, and use appropriate controls for sensitive and personal information. Maintaining these practices supports transparency, accountability, and trustworthy AI deployment throughout the application lifecycle. Salesforce frames trusted AI around principles such as accuracy, safety, honesty, empowerment, and sustainability; dependable, well-governed data is foundational to meeting those principles in practice. See Salesforce's Trusted AI overview and its Ethical AI resources.

QUESTION NO: 8

To avoid introducing unintended bias to an AI model, which type of data should be omitted?

A. Transactional

B. Engagement

C. Demographic

ANSWER: C

Explanation:

Demographic data is the correct choice because it captures characteristics associated with individuals or population groups, such as age, gender, race, ethnicity, income, education, or similar attributes. When these attributes are not necessary for the stated business purpose, omitting them from model inputs helps reduce the risk that an AI system will make predictions or recommendations based directly on protected or sensitive characteristics. This supports responsible AI design by limiting opportunities for disparate treatment and by ensuring that data collection is purpose-driven and proportionate to the use case.

Salesforce's guidance for trusted AI emphasizes using high-quality, relevant data and considering fairness throughout the AI lifecycle. Assessing which fields are truly needed before training or grounding an AI solution is an important governance practice. Demographic information may also embody historical inequities, so excluding unnecessary demographic fields is a practical early safeguard against encoding those patterns into a model. This decision should be documented as part of data governance and paired with testing for fairness, since other fields can sometimes correlate with demographic traits. See Salesforce's Trusted AI resources and its ethical and humane use of AI guidance.

QUESTION NO: 9

A customer using Einstein Prediction Builder is confused about why a certain prediction was made.

Following Salesforce's Trusted AI Principle of Transparency, which customer information should be accessible on the Salesforce Platform?

A. An explanation of how Prediction Builder works and a link to Salesforce's Trusted AI Principles

B. An explanation of the prediction's rationale and a model card that describes how the model was created

C. A marketing article of the product that clearly outlines the oroduct's capabilities and features

ANSWER: B

Explanation:

An explanation of the prediction's rationale and a model card that describes how the model was created is the information that should be accessible. Salesforce's transparency principle requires meaningful clarity about how an AI system works and

why it produces its outputs. For a Prediction Builder result, the rationale helps a user understand the factors that most influenced an individual prediction, making the output more interpretable and actionable. A model card complements that result-level explanation by documenting the model's purpose, creation, performance context, intended use, and relevant limitations. Together, these resources enable users to assess whether a prediction is appropriate to rely on for the business decision at hand and support responsible human oversight. Salesforce describes transparency as making AI systems understandable and providing people with information about how AI is developed and used. This focus on accessible rationale and model documentation supports informed, accountable use of predictive AI on the platform. See Salesforce's Ethical and Humane Use of Technology page and the Prediction Builder overview in Salesforce Help.

QUESTION NO: 10

What is a key benefit of effective interaction between humans and AI systems?

A. Leads to more informed and balanced decision making

B. Alerts humans to the presence of biased data

C. Reduces the need for human involvement

ANSWER: A

Explanation:

Leads to more informed and balanced decision making is correct because effective human-AI interaction combines complementary capabilities. AI can rapidly analyze large volumes of data, identify patterns, and provide predictions or recommendations, while people contribute contextual understanding, domain expertise, judgment, and accountability. When users can understand, appropriately question, and incorporate AI outputs into their work, they can make decisions using more relevant evidence rather than relying solely on either human intuition or automated output.

This collaborative approach is a core principle of trusted AI use. Salesforce's guidance on responsible AI emphasizes that people should remain empowered to make informed decisions and should be able to understand how AI-supported outcomes affect their work. Clear interaction and human oversight help users apply AI recommendations thoughtfully, including when evaluating confidence, context, and potential impacts. This produces decisions that are both better informed by data and more balanced by human values and business knowledge. See Salesforce's Trusted AI Principles and the Trailhead module on responsible AI.

QUESTION NO: 11

What is a key challenge of human AI collaboration in decision-making?

A. Leads to move informed and balanced decision-making

B. Creates a reliance on AI, potentially leading to less critical thinking and oversight

C. Reduce the need for human involvement in decision-making processes

ANSWER: B

Explanation:

Creates a reliance on AI, potentially leading to less critical thinking and oversight is correct because effective human-AI collaboration requires people to remain accountable decision-makers, rather than treating an AI output as automatically authoritative. AI can rapidly synthesize information, identify patterns, and generate recommendations, but its results can be incomplete, inaccurate, biased, or unsuitable for the business context. When users become overly reliant on recommendations, they may stop challenging assumptions, validating source data, considering important contextual factors, or escalating high-impact decisions for appropriate review. This is commonly described as automation bias: the tendency to over-trust automated guidance. Maintaining meaningful human oversight helps ensure that AI-assisted decisions remain responsible, accurate, and aligned with organizational values and customer needs. In Salesforce environments, this means reviewing generated content and recommendations, applying human judgment to consequential use cases, and using governance processes that preserve accountability. Salesforce's responsible AI guidance emphasizes human involvement

and oversight as essential practices for trustworthy AI, while its Einstein Trust Layer materials describe controls intended to support secure, responsible use of generative AI. See Salesforce Responsible AI and Salesforce Einstein Trust Layer.

QUESTION NO: 12

Which AI tool is a web of connections, guided by weights and biases?

A. Neural networks

B. Predictive Analytics

C. Rules- based systems

D. Mark this item for later review,

ANSWER: A

Explanation:

Neural networks are a web of interconnected processing units, often called neurons or nodes, organized into layers. Each connection between nodes has a numerical weight that represents how strongly information from one node should influence the next. Nodes also use bias values, which help shift the activation threshold and let the model represent more complex relationships in data. During training, the network adjusts its weights and biases to reduce the difference between its predictions and known outcomes. This learning process enables neural networks to detect patterns that would be difficult to define with manually written logic, including relationships in images, language, audio, and customer data. The term “web of connections” therefore directly describes the network architecture, while “guided by weights and biases” identifies the learned parameters that drive its predictions. In practical business AI use cases, neural-network-based models can contribute to capabilities such as classification, recommendation, and natural-language processing. Salesforce describes AI as using data and algorithms to identify patterns and make predictions or decisions; neural networks are a foundational machine-learning approach for doing that. See Salesforce’s Artificial Intelligence for Customer Relationship Management Trailhead module and its overview of what artificial intelligence is.