



Vendor: Google

Exam Code: Generative-AI-Leader

Exam Name: Generative AI Leader

Version: DEMO

QUESTION 1

A company wants to create an AI-powered educational solution that provides personalized learning experiences for students. This platform will assess a student's knowledge, recommend relevant learning materials, and generate personalized exercises. The application would provide the structure for lessons and track progress. What type of AI solution should they use?

- A. An AI-powered recommendation system for learning resources
- B. A large language model fine-tuned on educational content
- C. A learning management system (LMS)
- D. A customized learning agent

Answer: D

Explanation:

The request goes beyond just recommendations or content generation. It involves assessing knowledge, recommending materials, generating personalized exercises, providing lesson structure, and tracking progress. This implies a more comprehensive, intelligent system that acts as an assistant or tutor for the student, which is best described as a customized learning agent. This agent would likely leverage LLMs and recommendation systems as components, but the overall solution is an agent.

QUESTION 2

What is a primary benefit of using a multi-agent system?

- A. To simplify the most basic and repetitive rule-based tasks.
- B. To consolidate all unique AI functions into a single, undifferentiated model.
- C. To serve as a platform for hosting traditional, non-AI applications.
- D. To manage complex tasks that demand coordinated AI functions.

Answer: D

Explanation:

Multi-agent systems are designed to tackle complex problems by breaking them down into sub-tasks, where each agent specializes in a specific function. These agents then coordinate and collaborate to achieve a larger, more intricate goal that a single, monolithic AI model might struggle with.

QUESTION 3

An organization needs an AI tool to analyze and summarize lengthy customer feedback text transcripts. You need to choose a Google foundation model with a large context window. What foundation model should the organization choose?

- A. Gemini
- B. CodeGemma
- C. Imagen
- D. Chirp

Answer: A

Explanation:

Gemini models are known for their large context windows, making them highly suitable for processing and summarizing lengthy texts like customer feedback transcripts. CodeGemma is specialized for code, Imagen for image generation, and Chirp for speech.

QUESTION 4

A software developer needs a highly efficient, open-source large language model that can be fine-tuned on a local machine for rapid prototyping of a chatbot application. They require a model that offers strong performance in natural language understanding and generation, while being lightweight enough to run on limited hardware. Which Google-developed family of models should they use?

- A. Veo
- B. Gemini
- C. Gemma
- D. Imagen

Answer: C

Explanation:

Gemma is Google's family of lightweight, state-of-the-art open models, built from the same research and technology used to create the Gemini3 models. They are designed for developers to build innovative AI applications on their local machines or in the cloud, offering a balance of performance and efficiency suitable for limited hardware and rapid prototyping. Veo is for video generation, Gemini is typically larger and more general-purpose, and Imagen is for image generation.

QUESTION 5

An organization wants granular control over who can use and see their generative AI models and related resources on Google Cloud. Which Google Cloud security offering is specifically for this purpose?

- A. Identity and Access Management
- B. Secure-by-design infrastructure
- C. Security Command Center
- D. Workload monitoring tools

Answer: A

Explanation:

Identity and Access Management (IAM) is the fundamental Google Cloud service that allows you to define who has what access to which resources. It provides granular control over permissions for users, groups, and service accounts, including access to generative AI models and related data.

QUESTION 6

A company is developing a conversational AI chatbot. They need to ensure the chatbot can engage in human-like conversations and provide accurate information. What should they do to enhance the chatbot's ability to understand and respond effectively to user prompts?

- A. Use prompt engineering techniques, like few-shot prompting, to provide the chatbot with examples of successful interactions.
- B. Limit the chatbot's training data to prevent it from learning irrelevant information.
- C. Use strict keyword matching to ensure that the chatbot only responds to specific commands.
- D. Lower model temperature setting to produce more consistent and predictable responses.

Answer: A

Explanation:

Prompt engineering, especially techniques like few-shot prompting (providing examples of desired input-output pairs), is crucial for guiding a generative AI model to understand context and generate relevant, human-like responses. Limiting data or using strict keyword matching would severely restrict the chatbot's conversational ability, and lowering temperature makes responses less creative, not necessarily more understanding.

QUESTION 7

A human resources team is implementing a new generative AI application to assist the department in screening a large volume of job applications. They want to ensure fairness and build trust with potential candidates. What should the team prioritize?

- A. Integrating the AI application with various job boards to maximize candidate reach.
- B. Focusing on minimizing the processing time for each application to improve efficiency.
- C. Ensuring AI operates transparently, especially regarding application evaluation and data usage.
- D. Ensuring that the AI application can automatically rank all candidates without requiring human review.

Answer: C

Explanation:

To ensure fairness and build trust, especially in sensitive areas like job applications, transparency in how AI evaluates applications and uses data is paramount. This involves understanding potential biases, explaining decisions (where possible), and ensuring human oversight.

QUESTION 8

A company's large learning model (LLM) is producing hallucinations that are a result of the Knowledge cutoff. How does retrieval-augmented generation (RAG) overcome this limitation?

- A. RAG fine-tunes the LLM on specific customer query patterns to improve the speed and efficiency of response generation.
- B. RAG enhances the creative writing capabilities of the LLM to generate more engaging and informative responses.
- C. RAG enables the LLM to retrieve relevant and up-to-date information from knowledge sources.
- D. RAG uses human oversight to ensure accuracy before presenting information to the customer.

Answer: C

Explanation:

The primary purpose of RAG is to address the "knowledge cutoff" and hallucination issues of LLMs. It does this by retrieving relevant, up-to-date information from external knowledge sources (like databases or documents) at inference time and then using this retrieved information to ground the LLM's generation, ensuring factual accuracy and relevance to the specific query.

QUESTION 9

An organization with a team of live customer service agents wants to improve agent efficiency and customer satisfaction during support interactions. They are looking for a tool that can provide real-time guidance to agents, suggest helpful information, and streamline the support process without fully automating customer conversations. Which component of Google's Customer Engagement Suite should they use?

- A. Agent Assist
- B. Conversational Agents

- C. Conversational Insights
- D. Google Cloud Contact Center as a Service

Answer: A

Explanation:

As previously mentioned, Agent Assist is specifically designed for real-time support to human agents, providing them with suggestions and relevant information during live customer interactions. Conversational Agents (chatbots) automate interactions, Conversational Insights analyze conversations after they occur, and Contact Center as a Service is the broader infrastructure.

QUESTION 10

A financial institution uses generative AI (gen AI) to approve and reject loan applications, but gives no reasons for rejection. Customers are starting to file complaints. The company needs to implement a solution to reduce the complaints. What should the company do?

- A. Collect a larger and more diverse dataset for the gen AI model.
- B. Implement explainable gen AI policies.
- C. Fine-tune the gen AI model.
- D. Develop fairness assessments for the gen AI model.

Answer: B

Explanation:

The core problem is the lack of reasons for rejection, leading to customer complaints. This falls under the domain of explainable AI (XAI). Implementing explainable gen AI policies or mechanisms would allow the institution to provide transparency into how the AI made its decision, addressing the customer complaints directly. While other options might improve the model, they don't directly solve the transparency issue.

QUESTION 11

A large multinational corporation with geographically dispersed teams struggles with knowledge silos and inconsistent access to crucial internal information. What is a key business benefit of using Google Agentspace in this scenario?

- A. Improved IT infrastructure management across offices.
- B. Seamless knowledge sharing and collaboration across internal systems.
- C. Enhanced data encryption and compliance for internal communications.
- D. Automation of employee performance reviews using AI.

Answer: B

Explanation:

Google Agentspace (or similar agent-based frameworks) aims to connect and orchestrate various AI capabilities and data sources. In a scenario with knowledge silos, a key benefit would be to enable seamless knowledge sharing and collaboration by allowing agents to access, process, and disseminate information across different internal systems and teams.

QUESTION 12

A security team needs a centralized platform to gain a comprehensive overview of their organization's security health across their entire Google Cloud environment, including potential threats to their generative AI deployments. Which Google Cloud security offering is specifically for this purpose?

- A. Workload monitoring tools
- B. Security Command Center
- C. Identity and Access Management
- D. Secure-by-design infrastructure

Answer: B

Explanation:

Security Command Center is Google Cloud's comprehensive security management and data risk platform. It provides centralized visibility into security posture, identifies vulnerabilities, detects threats, and helps manage compliance across the entire Google Cloud environment, including services and deployments like generative AI.

QUESTION 13

A company wants to use an AI agent to automate some tasks. They want everyone to understand the different functions of an AI agent. What is the function of an AI agent in the context of gen AI?

- A. To provide the computational resources needed to train and run gen AI models.
- B. To store and manage large datasets used for training and running gen AI models.
- C. To provide a user-friendly interface for interacting with gen AI models.
- D. To analyze situations, use multiple tools, and make informed decisions without requiring constant human input.

Answer: D

Explanation:

An AI agent, especially in the context of generative AI, is designed to be more autonomous and capable than a simple model. Its function is to understand a goal, analyze a situation, leverage various tools (including other generative AI models or external APIs), and make decisions or take actions to achieve that goal, often with minimal human intervention.

QUESTION 14

A research team has collected a large dataset of sensor readings from various industrial machines. This dataset includes measurements like temperature, pressure, vibration levels, and electrical current, recorded at regular intervals. The team has not yet assigned any labels or categories to these readings and wants to identify potential anomalies, malfunctions, or natural groupings of machine behavior based on the sensor data alone. What type of machine learning should they use?

- A. Reinforcement learning
- B. Unsupervised learning
- C. Deep learning
- D. Supervised learning

Answer: B

Explanation:

Since the team has not yet assigned any labels or categories to the sensor readings and wants to identify "anomalies, malfunctions, or natural groupings" based on the data alone, this is a classic unsupervised learning problem. Unsupervised learning techniques like clustering or anomaly detection are used to find hidden patterns or structures in unlabeled data.

Thank You for Trying Our Product

Braindump2go Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.



View list of all certification exams: <http://www.braindump2go.com/all-products.html>



Microsoft



ORACLE



JUNIPER
NETWORKS



EMC²
where information lives

10% Discount Coupon Code: ASTR14

