DUMPSBOSS.

Google Cloud CertifiedGenerative AI Leader Exam

Google Generative-AI-Leader

Version Demo

Total Demo Questions: 10

Total Premium Questions: 84

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

QUESTION NO: 1

A learning and development team wants to quickly create a new hire training video with a custom avatar and voiceover that matches their company's branding and key messaging. They did not receive any money to spend on the production. What should they do?

A. Generate the video frames with Imagen.

B. Prompt the Gemini app to create a video.

C. Train a model with Vertex AI and produce a video.

D. Create a video with Google Vids.

ANSWER: D

Explanation:

They should use Google Vids. It's built for exactly this kind of "make me a training video fast" scenario, and it includes AI help for things like generating a first draft, adding narration/voiceover, and pulling everything into a shareable video without needing a full production budget.

The other choices don't really fit the goal. Imagen is mainly for generating images, not producing a complete narrated training video. The Gemini app can help you write a script or outline, but it's not a dedicated end-to-end video creation tool. And training a model in Vertex AI would be slow, technical, and definitely not free.

Reference: https://workspace.google.com/products/vids/

QUESTION NO: 2

A retail company with a large online catalog wants to improve customer experience and drive sales by implementing multimodal search capabilities (image, voice, and text). What is a primary business benefit of this capability?

A. Improved customer engagement and product discovery leading to increased satisfaction and potential sales.

B. Reduced dependency on keyword optimization for product listings and improved search engine rankings.

C. Lowered operational costs associated with managing and updating product information across different platforms and channels.

D. Streamlined inventory management processes and more accurate demand forecasting for popular items.

ANSWER: A

Explanation:

The big win with multimodal search is that shoppers can find what they want in the most natural way for them—typing, speaking, or even snapping a photo. That removes friction in the shopping journey (especially when people don't know the exact product name), which usually means better product discovery and more completed purchases.

Option A matches that direct business outcome: higher engagement, easier discovery, happier customers, and a lift in sales. The other options might happen indirectly in some companies, but they aren't the primary, most reliable benefit tied to multimodal search itself. For example, SEO rankings (B) are a different problem, and inventory forecasting (D) is more about analytics and supply chain than how customers search.

Google highlights multimodal capabilities (text, image, etc.) as a way to improve how users search and interact with content, which is exactly what's being asked here: https://cloud.google.com/vertex-ai/generative-ai/docs/multimodal/overview and https://cloud.google.com/vision/product-search/docs/overview

QUESTION NO: 3

What is a key advantage of using Google's custom-designed TPUs?

A. TPUs are lightweight processors intended for deployment on edge devices.

B. TPUs increase the storage capacity and data retrieval speeds within Google Cloud data centers.

C. TPUs are specialized AI processors that excel at parallel processing for machine learning workloads.

D. TPUs are primarily designed to improve the general processing speed of virtual machines in the cloud.

ANSWER: C

Explanation:

Google's TPUs (Tensor Processing Units) are built specifically for machine learning, so their big advantage is speed and efficiency on the kinds of math ML models do all day long—especially large matrix multiplications. That's why they're a strong fit for training and running deep learning models and LLMs at scale.

The key idea is parallelism: TPUs are designed to run lots of operations at the same time, which can make both training and inference much faster (and often more cost-efficient) than using general-purpose CPUs for the same workload. That's what option C is getting at.

The other choices miss the mark. Edge TPUs do exist, but "lightweight processors for edge devices" isn't the main advantage of Google's cloud TPUs (A). TPUs don't increase storage capacity or retrieval speed (B), and they're not meant as a general "make every VM faster" accelerator (D)—they're purpose-built for ML computations.

References: https://cloud.google.com/tpu/docs/tpus and https://cloud.google.com/blog/products/ai-machine-learning/introducing-cloud-tpu-v5e

QUESTION NO: 4

A development team is configuring a generative AI model for a customer-facing application and wants to ensure the generated content is appropriate and harmless. What is the primary function of the safety settings parameter in a generative AI model?

A. To limit the maximum text length that the model generates by ensuring concise responses.

B. To determine the number of tokens the model can process at once by influencing the complexity and length of inputs and outputs.

C. To filter out potentially harmful or inappropriate content from the model's output based on the desired level of filtering.

D. To control the creativity and randomness of the model's output by adjusting the diversity of word choices.

ANSWER: C

Explanation:

Safety settings are basically the “content guardrails” for a generative model. They’re there to reduce the chance the model outputs harmful, unsafe, or inappropriate text (for example hate/harassment, sexual content, violence, or self-harm), by applying thresholds that block or restrict those categories.

The other options describe different knobs: max output length (tokens) controls how long the answer can be, and temperature/top-p controls how random or creative the wording is. Those can change style and length, but they don’t directly stop unsafe content. Safety settings are the parameter meant specifically for that filtering job.

References: https://cloud.google.com/vertex-ai/generative-ai/docs/multimodal/configure-safety-attributes and https://ai.google.dev/gemini-api/docs/safety-settings

QUESTION NO: 5

A pharmaceutical company's research and development department spends significant time manually reviewing new scientific papers to identify potential drug targets. They need a solution that can answer questions about these documents and provide summarized insights to researchers without requiring extensive coding expertise. What should the organization do?

A. Use Gemini for Google Workspace to facilitate collaborative document review.

B. Use Vertex AI Search to index the papers and enable keyword-based searches.

C. Use Vertex AI AutoML to train a model that classifies papers into predefined research areas.

D. Use Vertex AI Agent Builder to create a custom AI agent.

ANSWER: D

Explanation:

They should use **Vertex AI Agent Builder** because it’s meant for building question-answering and summarization experiences over a set of documents with minimal coding. In this scenario, researchers don’t just want to “search” papers—they want to ask natural-language questions (like “What proteins are linked to pathway X?”) and get a useful summary back. Agent Builder is designed for exactly that kind of conversational, document-grounded agent.

The other choices don’t fit as well. Gemini for Google Workspace helps individuals write and summarize inside Docs/Gmail, but it’s not a purpose-built system for building a custom Q&A agent over a controlled corpus. Vertex AI Search is great for finding documents, but keyword-style search alone won’t reliably produce grounded, synthesized answers. AutoML is for training predictive models (like classification), which doesn’t solve the “ask questions and summarize papers” need.

References: https://cloud.google.com/vertex-ai/generative-ai/docs/agent-builder/overview and https://cloud.google.com/vertex-ai/generative-ai/docs/enterprise-search/overview

QUESTION NO: 6

A research company needs to analyze several lengthy PDF documents containing financial reports and identify key performance indicators (KPIs) and their trends over the past year. They want a Google Cloud prebuilt generative AI tool that can process these documents and provide summarized insights directly from the source material with citations. What should the analyst do?

A. Create a custom Gem in Gemini Advanced with predefined KPIs to look across different financial reports.

B. Use the Gemini app to ask general financial trend questions.

C. Use NotebookLM to upload and analyze the documents.

D. Use Gemini for Google Workspace within Google Docs to copy and paste sections of the reportsfor summary and analysis.

ANSWER: C

Explanation:

They should use **NotebookLM** and upload the PDFs there. NotebookLM is built for exactly this “research assistant” workflow: you give it long source documents, then ask questions like “What happened to revenue and margin quarter over quarter?” and it will summarize what it finds while showing **citations back to the original text**. That citation piece is the big deal here because it lets the analyst verify each KPI claim directly in the report.

The other options don’t fit as well. Gemini app / Gemini Advanced are great general chat tools, but they’re not as purpose-built for working from a curated set of uploaded sources with reliable inline citations. And copying/pasting chunks into Gemini for Google Workspace is clunky for “several lengthy PDFs” and increases the chance you miss context.

References: https://notebooklm.google/ and https://support.google.com/notebooklm/answer/14276468

QUESTION NO: 7

What will Google Cloud's Agent Assist help a company achieve?

A. The infrastructure to provide an enterprise-grade contact center solution with omnichannel support, routing, and integration with CRM systems.

B. The ability to analyze conversational data to identify customer sentiment, common topics of discussion, and insights into agent performance and customer experience.

C. The ability to provide real-time assistance and recommended responses to live customer service agents during their interactions.

D. The ability to build and deploy deterministic and generative chatbot agents for automated customer support.

ANSWER: C

Explanation:

Agent Assist is meant to support human contact-center agents while they're actively talking with customers. Instead of replacing the agent, it listens to the conversation and surfaces helpful things in the moment—like suggested replies, relevant knowledge-base articles, and next-best actions—so the agent can respond faster and more consistently.

The other choices describe different (but related) products or capabilities. Full contact-center infrastructure (routing, omnichannel, CRM integration) is more in the space of Contact Center AI platform components, not Agent Assist itself. Conversation analytics and sentiment/topic insights are typically handled by analytics tools (for example, CCAI Insights). And building fully automated chatbots is usually handled by Dialogflow/Vertex AI Agent Builder, not Agent Assist.

References: https://cloud.google.com/agent-assist and https://cloud.google.com/architecture/ccai-agent-assist

QUESTION NO: 8

A company is developing an AI character for a video game. The AI character needs to learn how to navigate a complex environment and make decisions to achieve certain objectives within the game. When the AI takes actions that lead to positive outcomes, like finding a reward or overcoming an obstacle, it receives a positive score. When it takes actions that lead to negative outcomes, like hitting a wall or losing progress, it receives a negative score. Through this process of trial and error, the AI gradually improves the character's ability to play the game effectively. What machine learning should the company use?

A. Reinforcement learning

B. Unsupervised learning

C. Supervised learning

D. Deep learning

ANSWER: A

Explanation:

This is a classic reinforcement learning setup: the game character (the "agent") interacts with the game world (the "environment"), takes actions, and then gets a reward signal—positive for good outcomes and negative for bad ones. Over many rounds of trial and error, it learns a policy (a strategy) that maximizes total reward, which is exactly what you want for navigation and decision-making in a game.

The other choices don't fit as well. Supervised learning needs labeled examples of the "right" action for each situation, which you typically don't have in an open-ended game. Unsupervised learning is more about finding patterns or clusters without rewards. Deep learning is a technique (neural nets) that can be used inside reinforcement learning, but it's not the learning setup described here.

References: https://cloud.google.com/learn/what-is-reinforcement-learning and https://www.deepmind.com/learning-resources/reinforcement-learning

QUESTION NO: 9

When building a generative AI solution, what is a recommended practice to mitigate hallucinations (i.e. incorrect or made-up outputs)?

A. Increase model temperature arbitrarily

B. Do not monitor outputs after deployment

C. Use retrieval-augmented generation (RAG) or grounding

D. Always rely purely on the model's internal knowledge

ANSWER: C

Explanation:

The most practical way to cut down hallucinations is to "ground" the model in real, trusted information at answer time. That's exactly what retrieval-augmented generation (RAG) does: before the model responds, your system fetches relevant passages from approved sources (like your docs, knowledge base, or database) and gives that context to the model. The model is then far more likely to stick to facts instead of guessing.

The other options go in the wrong direction. Turning up temperature makes outputs more random and creative, which usually increases the risk of made-up details. Not monitoring after deployment is risky because you won't catch failures, drift, or new hallucination patterns. And relying only on the model's internal knowledge is a common cause of hallucinations, especially for niche, changing, or company-specific info.

References: https://cloud.google.com/vertex-ai/generative-ai/docs/grounding/overview and https://cloud.google.com/vertex-ai/generative-ai/docs/rag-overview

QUESTION NO: 10

What does a diffusion model do?

A. Analyzes data and predicts future trends and patterns.

B. Optimizes business processes and resource allocation.

C. Facilitates the storage and management of structured data.

D. Generates high-quality content by refining noise into structured data.

ANSWER: D

Explanation:

A diffusion model is a generative model that learns to create new content (most famously images) by working "backwards" from noise. During training, it learns what it looks like to gradually add noise to real data until it becomes basically random. Then it learns the reverse step: how to remove that noise little by little.

When you generate something, you start with pure random noise and run that denoising process over many steps until a clear, structured result appears. If you give it a text prompt, that prompt helps steer the denoising so the final image matches what you asked for. That's why option D is the best fit.

The other choices describe different things: forecasting (A), operations optimization (B), or data storage (C). None of those capture the "noise → denoise → image/content" idea that diffusion models are known for.

References: https://en.wikipedia.org/wiki/Diffusion_model and https://developers.google.com/machine-learning/resources/intro-llms (for general foundation model context).