GOOGLE Generative-AI-Leader Exam Questions

Total Questions: 70+
Demo Questions: 10
Version: Updated for 2025

Prepared and Verified by Cert Empire – Your Trusted IT Certification Partner

For Access to the full set of Updated Questions – Visit:
Google Generative-AI-Leader Exam Questions by Cert Empire

Question: 1

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, in the context of generative AI, is an autonomous system that leverages a large language model (LLM) as its core reasoning engine. Its primary function is to perceive a situation or user request, break it down into actionable steps, and then select and utilize a variety of "tools" (such as APIs, search engines, or other models) to accomplish a complex goal. This process involves analysis, planning, and decision-making to act on behalf of a user with minimal human intervention, distinguishing it from a simple conversational interface or underlying infrastructure.

Why Incorrect Options are Wrong:

A. This describes the function of cloud infrastructure components like GPUs and TPUs, which an agent uses but does not provide.
B. This describes the function of data storage and management services (e.g., Google Cloud Storage, BigQuery), not the agent's role.
C. This describes a user interface (UI). While an agent may have a UI, its core function is the autonomous reasoning and task execution that happens behind the interface.

References:

1. Google Cloud Blog, "What are AI agents?" (May 2024): This article defines an AI agent as "a software program that can perform tasks on behalf of a user with some level of autonomy." It explicitly states that agents use LLMs for reasoning and can "access a set of tools to perform actions in the world," which directly supports the concept of analyzing situations and using multiple tools to make decisions.
2. Google Cloud Documentation, "Overview of Agent Builder": The documentation for Vertex AI Agent Builder describes how it enables the creation of agents that "can hold conversations" and are "equipped with tools" that allow them to "interact with external systems" to fulfill complex user

requests. This highlights the agent's function in orchestrating tools to perform actions, aligning with the correct answer.

3. Stanford University, CS221: Artificial Intelligence: Principles and Techniques, Autumn 2021, Lecture 1: Course materials define an agent as an entity that "perceives its environment through sensors and acts upon that environment through actuators." In the context of Gen AI, the LLM acts as the reasoning component, and "tools" are the actuators, allowing the agent to make informed decisions and act autonomously. This foundational definition supports the core concepts in option D.

Question: 2

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:

Google Cloud's Identity and Access Management (IAM) is the primary service for implementing fine-grained access control. It allows organizations to manage authorization by defining "who" (identity) has "what access" (role) for "which resource." For generative AI resources, such as models in Vertex AI, IAM enables administrators to grant specific permissions like viewing, using, or managing models to specific users, groups, or service accounts, thereby providing the required granular control.

Why Incorrect Options are Wrong:

B. Secure-by-design infrastructure: This is a foundational principle of how Google Cloud is built, encompassing physical security and data encryption, not a specific tool for managing user access to AI models.
C. Security Command Center: This is a centralized security management and risk platform for threat detection and posture assessment, not for setting granular user permissions on resources.
D. Workload monitoring tools: These tools, like Cloud Monitoring, are for observing performance and operational health. They do not manage or enforce access control policies.

References:

1. Google Cloud Documentation, "Overview of IAM": "Cloud Identity and Access Management (IAM) lets you grant granular access to specific Google Cloud resources and helps prevent access to other resources. IAM lets you adopt the security principle of least privilege, so you grant only the necessary access to your resources." (cloud.google.com/iam/docs/overview)
2. Google Cloud Documentation, "Vertex AI access control with IAM": "You can control access to Vertex AI resources by using Identity and Access Management (IAM). IAM lets you grant access to specific Google Cloud resources and prevents unwanted access to other resources... For example, you can grant a user the AI Platform Viewer (roles/aiplatform.viewer) role so the user

can view Vertex AI resources in a project."
(cloud.google.com/vertex-ai/docs/general/access-control)
3. Google Cloud Security Whitepaper, "Google Cloud security foundations guide": This guide describes the "secure-by-design" approach as a layered security strategy built into the infrastructure, from the physical data centers to the purpose-built security hardware. This confirms it is a design philosophy, not an access control service.
(cloud.google.com/resources/google-cloud-security-foundations-guide, Section: "Secure-by-design infrastructure")

Question: 3

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:

The research team's goal is to discover inherent structures, such as natural groupings and anomalies, within a dataset that lacks predefined labels. This is the primary objective of unsupervised learning. This machine learning paradigm is specifically designed to analyze unlabeled data to identify hidden patterns or intrinsic structures. Techniques like clustering would be used to find the "natural groupings of machine behavior," and anomaly detection would identify "potential anomalies, malfunctions." The problem description directly aligns with the definition and application of unsupervised learning.

Why Incorrect Options are Wrong:

A. Reinforcement learning involves an agent learning to make decisions by interacting with an environment to maximize a reward signal, which is not the scenario described.
C. Deep learning is a subfield of machine learning that uses multi-layered neural networks; it is a set of techniques, not a fundamental learning type defined by data labeling.
D. Supervised learning requires a dataset with labeled outcomes (e.g., data points explicitly marked as 'normal' or 'faulty') to train a model, which the team does not have.

References:

1. Google Cloud Documentation: In "What is machine learning?", Google Cloud defines the three types of machine learning. It states, "Unsupervised learning helps you find structure in your data where you don't have any labels... A common use of unsupervised learning is to cluster data into groups with similar characteristics." This directly supports using unsupervised learning for finding "natural groupings" in the unlabeled sensor data.

Source: Google Cloud, "AI & Machine Learning Products," What is machine learning? section.
2. Stanford University Courseware (CS229): The course notes for CS229 Machine Learning define unsupervised learning as the task where "we are given data that does not have any labels associated with it." It lists clustering as a canonical example, where the goal is to group the data into coherent "clusters." This aligns perfectly with the team's objective.
Source: Ng, A. (2023). CS229 Machine Learning Course Notes, Stanford University, Part I: Supervised Learning, Section 1: Introduction, page 3.
3. MIT OpenCourseWare: In the "Introduction to Machine Learning" course, unsupervised learning is described as the process of learning from "raw, unlabeled data." The goal is to "discover the underlying structure or distribution of the data," including tasks like clustering and dimensionality reduction, which are relevant to the described problem.
Source: MIT OpenCourseWare, 6.036 Introduction to Machine Learning, Fall 2020, Lecture 1: Introduction, Section on "Paradigms of Machine Learning."

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 a family of lightweight, open-source models developed by Google, built from the same research and technology used for the Gemini models. They are specifically designed to be accessible to developers and researchers, enabling them to run and fine-tune the models on their own local machines, such as laptops or workstations. This makes Gemma ideal for rapid prototyping and experimentation with limited hardware, directly addressing the developer's need for an efficient, open-source model for a chatbot application. The models are available in various sizes (e.g., 2B and 7B parameters) to balance performance and computational requirements.

Why Incorrect Options are Wrong:

A. Veo: Veo is a text-to-video generation model, not a large language model designed for chatbot applications or natural language understanding tasks.
B. Gemini: Gemini is Google's flagship family of powerful, multimodal models (e.g., Gemini Pro, Ultra). They are primarily accessed via APIs and are not open-source or designed for fine-tuning on local developer hardware.
D. Imagen: Imagen is a family of text-to-image diffusion models. Its purpose is to generate photorealistic images from text descriptions, not to power conversational AI.

References:

1. Google AI Blog. (2024, February 21). Gemma: Introducing new state-of-the-art open models. Google. Retrieved from https://blog.google/technology/developers/gemma-open-models/. Reference Point: The article states, "Today, we're releasing Gemma, a family of lightweight, state-of-the-art open models... They can run on a developer laptop or desktop computer... Gemma models are also available across your favorite tools: Colab and Kaggle notebooks, plus

frameworks like JAX, PyTorch, and TensorFlow through native Keras 3.0." This directly supports Gemma's suitability for local development and fine-tuning.

2. Gemma Team, Google. (2024). Gemma: Open Models Based on Gemini Research and Technology. arXiv:2403.08295 cs.CL. https://doi.org/10.48550/arXiv.2403.08295
Reference Point: Section 1 (Introduction) states, "We are releasing two sizes of Gemma models: Gemma 2B and Gemma 7B... Both models are released with weights and are suitable for fine-tuning on a variety of tasks... they can be deployed on a variety of hardware types, including laptop, desktop, and cloud." This academic paper confirms the model's design for accessibility and fine-tuning on developer-grade hardware.
3. Google for Developers. (n.d.). Gemma. Retrieved from https://ai.google.dev/gemma
Reference Point: The official product page highlights, "A family of lightweight, state-of-the-art open models from Google... Responsible by design, Gemma models are easy to deploy and develop with." This reinforces its status as an open, developer-focused model family.

Question: 5

What is an example of unsupervised machine learning?

A. Analyzing customer purchase patterns to identify natural groupings.

B. Training a system to recognize product images using labeled categories.

C. Predicting subscription renewal based on past renewal status data.

D. Forecasting sales figures using historical sales and marketing spend.

Answer:

A

Explanation:

Unsupervised machine learning involves training a model on data that is not labeled or categorized. The primary goal is to identify hidden patterns, structures, or groupings within the data itself. Analyzing customer purchase patterns to find "natural groupings" is a classic example of a clustering task. In this scenario, the algorithm identifies segments of customers with similar buying habits without any predefined labels for what those segments should be, which is the core principle of unsupervised learning.

Why Incorrect Options are Wrong:

B. Training a system to recognize product images using labeled categories. This is supervised learning (classification), as the model learns from data that has explicit "labeled categories" to map inputs (images) to outputs (labels).
C. Predicting subscription renewal based on past renewal status data. This is supervised learning (classification), where "past renewal status" (e.g., renewed/not renewed) serves as the label to predict a future binary outcome.
D. Forecasting sales figures using historical sales and marketing spend. This is supervised learning (regression), as it uses labeled historical data (past sales figures) to predict a continuous numerical value for the future.

References:

1. Google Cloud Documentation: In the "What is unsupervised learning?" article, it states, "A common unsupervised learning technique is clustering, which groups data points with similar characteristics. For example, an unsupervised learning algorithm can group together customers with similar purchasing behaviors..." This directly corresponds to the correct answer.
Source: Google Cloud, "What is unsupervised learning?", Section: "Unsupervised learning techniques and algorithms".
2. Stanford University Courseware (CS229): Andrew Ng's course notes on machine learning define unsupervised learning and provide clustering as a canonical example. "In the clustering

problem, we are given an unlabeled dataset and we would like to have an algorithm automatically group the data into coherent subsets or clusters." This aligns with identifying natural groupings in customer data.
Source: Ng, A. (2023). CS229 Machine Learning Course Notes, Part V Unsupervised Learning, Page 1. Stanford University.
3. MIT OpenCourseWare: The "Introduction to Machine Learning" course distinguishes learning problems. Unsupervised learning is described as finding patterns in data without explicit labels, such as clustering data points. In contrast, supervised learning, which covers options B, C, and D, is defined by learning a function from labeled input-output pairs.
Source: MIT OpenCourseWare, 6.036 Introduction to Machine Learning, Fall 2020, Lecture 1: Introduction, Section: "Types of Learning Problems".

Question: 6

An organization wants to use generative AI to create a marketing campaign. They need to ensure that the AI model generates text that is appropriate for the target audience. What should the organization do?

A. Use role prompting.

B. Use prompt chaining.

C. Use few-shot prompting.

D. Adjust the temperature parameter.

Answer:

A

Explanation:

Role prompting is a technique used to instruct a large language model (LLM) to adopt a specific persona, character, or role. By defining the role of an expert marketer targeting a specific demographic (e.g., "You are a marketing expert creating a campaign for environmentally-conscious millennials"), the organization can directly guide the model's tone, style, vocabulary, and overall approach. This ensures the generated text aligns with the intended audience's preferences and is more likely to be effective, making it the most direct method to solve the stated problem.

Why Incorrect Options are Wrong:

B. Use prompt chaining: This technique breaks a complex task into multiple, sequential prompts. It is for managing workflow, not for defining the stylistic tone for a specific audience.
C. Use few-shot prompting: This provides the model with a few examples to guide its output format. While it can influence style, it is less direct than explicitly assigning a role.
D. Adjust the temperature parameter: Temperature controls the randomness or creativity of the model's output. It does not specifically control the persona or ensure appropriateness for a target audience.

References:

1. Google Cloud Documentation, "Introduction to prompt design": In the section "Prompt components," the documentation explains the use of a "Persona." It states, "The persona is the role you want the model to take on when it generates a response... This technique is also known as role prompting." This directly supports using a persona to tailor output.
2. Google Cloud Skills Boost, "Introduction to Generative AI": In the module "Prompt Design," the courseware explains that providing a persona in a prompt is a key strategy to "guide the model to produce output from a certain point of view or in a certain style," which is precisely what is needed

for a targeted marketing campaign.
3. Vanderbilt University, "Prompt Engineering for ChatGPT" Courseware: The guide on "Prompting Best Practices" lists "Assign a Role" as a primary technique. It notes, "Asking the model to adopt a persona or role can result in better and more targeted responses." This academic source validates role prompting for creating audience-specific content.

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 in a high-stakes process like job application screening, transparency is paramount. Candidates need to trust that the process is equitable, and the HR team must be able to verify that the AI is not introducing or amplifying bias. By prioritizing transparency in how the AI evaluates applications and uses candidate data, the team addresses both goals directly. This approach aligns with core principles of responsible AI, such as accountability and fairness, by making the system's logic understandable and auditable.

Why Incorrect Options are Wrong:

A. Integrating with job boards is an operational goal focused on increasing the applicant pool, not on ensuring the fair and trustworthy evaluation of those applicants.
B. Minimizing processing time is an efficiency metric. Prioritizing speed over careful, transparent, and fair evaluation can undermine trust and lead to poor, biased outcomes.
D. Removing human review (human-in-the-loop) is contrary to responsible AI practices for critical decisions. It increases the risk of unchecked bias and errors, thereby eroding trust and fairness.

References:

1. Google AI Principles: The official Google AI principles state, "We will design AI systems to be appropriately cautious and to seek human direction as appropriate." and "AI algorithms and datasets can reflect, reinforce, or reduce unfair biases... We will seek to avoid unjust impacts on people, particularly those related to sensitive characteristics..." Transparency is a key mechanism for ensuring these principles are met. (Source: Google AI, "Our Principles", Section: "AI Principles", Points 2 and 5).
2. Google Cloud, "AI adoption framework: Part 3-Governing AI": This official documentation

emphasizes the need for AI governance, which includes "transparency and explainability." It states, "For AI systems to be trustworthy, it's important that stakeholders understand their output. Explainable AI (XAI) can help you understand and interpret predictions from your models." This directly supports the need for transparency in evaluation. (Source: Google Cloud Architecture Center, Document ID: ai-adoption-framework-governance, Section: "Establish AI principles").
3. Stanford University, Human-Centered AI (HAI), "An AI Bill of Rights for A.I. that Serves the American People": The report highlights that individuals should "know that an automated system is being used and understand how and why it contributes to outcomes that impact you." This principle of "Notice and Explanation" is fundamental to building trust in AI systems used for critical decisions like hiring. (Source: Stanford HAI, "The Blueprint for an AI Bill of Rights," Section: "Notice and Explanation").

Question: 8

A company has a machine learning project that involves diverse data types like streaming data and structured databases. How does Google Cloud support data gathering for this project?

A. Google Cloud provides tools such as Pub/Sub, Cloud Storage, and Cloud SQL.

B. The Gemini app is the primary Google Cloud tool for directly collecting data.

C. Google Cloud's strengths are in the data analysis tools such as BigQuery.

D. Google Cloud relies on Vertex AI to connect to external data.

Answer:

A

Explanation:

Google Cloud provides a comprehensive suite of services specifically designed for ingesting and storing diverse data types, which is fundamental for any machine learning project. For the scenario described, Pub/Sub is the ideal service for capturing real-time streaming data. Cloud SQL offers a fully managed relational database service for structured data. Cloud Storage serves as a highly scalable and durable object store for a wide variety of data, including unstructured files, backups, and staging data for ML pipelines. Together, these tools provide a robust foundation for data gathering.

Why Incorrect Options are Wrong:

B. The Gemini app is a user-facing conversational AI application. It is not an infrastructure tool designed for data collection or ingestion in an ML pipeline.
C. BigQuery is a serverless data warehouse primarily used for data analysis, querying, and processing at scale, which occurs after the data gathering stage.
D. Vertex AI is a unified MLOps platform for building, training, and deploying models. It consumes data from sources like Cloud Storage and BigQuery but is not the primary tool for data gathering.

References:

1. Google Cloud Documentation, "Data lifecycle on Google Cloud": This document outlines the stages of the data lifecycle, identifying "Ingest" as the first step. It lists services like Pub/Sub for streaming data and discusses storage options like Cloud SQL and Cloud Storage. (Reference: Google Cloud, "Data lifecycle on Google Cloud", Ingest section).
2. Google Cloud Documentation, "Pub/Sub": "Pub/Sub is an asynchronous and scalable messaging service that decouples services producing events from services processing events... It is a common data ingestion pipeline component for data that is created in real time." (Reference: Google Cloud, "What is Pub/Sub?", Product overview).
3. Google Cloud Documentation, "Cloud SQL": "Cloud SQL is a fully managed relational database

service for MySQL, PostgreSQL, and SQL Server." This directly addresses the need to handle structured databases. (Reference: Google Cloud, "Cloud SQL", Product overview).
4. Google Cloud Documentation, "Cloud Storage": "Cloud Storage is a service for storing your objects in Google Cloud. An object is an immutable piece of data consisting of a file of any format." This highlights its role in storing diverse data types for ML projects. (Reference: Google Cloud, "Cloud Storage", Key concepts).
5. Patterson, D. A., & Hennessy, J. L. (2021). Computer Organization and Design RISC-V Edition: The Hardware Software Interface (2nd ed.). Morgan Kaufmann. While a broad reference, university-level texts on cloud architecture consistently categorize services by function. In chapters covering cloud services, ingestion tools (like Pub/Sub), structured storage (like Cloud SQL), and object storage (like Cloud Storage) are presented as distinct foundational components for data-intensive applications like machine learning. (Reference: Chapter 6, "Parallel Processors from Client to Cloud").

Question: 9

A company is exploring Google Agentspace to improve how its employees search for information on their enterprise systems and automate certain tasks. What is the key business advantage of using Agentspace?

A. Enhanced real-time communication and collaboration among team members.

B. Greater interoperability with legacy software systems and databases.

C. Improved productivity and data interaction using AI assistants and advanced document analysis.

D. More granular control over support team access and permissions for sensitive data.

Answer:

C

Explanation:

The term "Agentspace" conceptually refers to an ecosystem of AI agents designed for enterprise use. The primary business advantage of deploying such agents, built on platforms like Google's Vertex AI Agent Builder, is to enhance employee productivity. These AI assistants achieve this by streamlining how employees interact with complex enterprise data. They leverage advanced document analysis and natural language understanding to provide quick, accurate answers from internal knowledge bases and can automate repetitive tasks, freeing up employees to focus on higher-value work. This directly translates to improved operational efficiency and productivity.

Why Incorrect Options are Wrong:

A. Enhanced real-time communication and collaboration are primary benefits of tools like Google Workspace (Chat, Meet), not the core function of AI agents for enterprise search and automation.
B. Interoperability with legacy systems is a crucial technical capability that enables the AI agent to function, but it is not the ultimate business advantage itself.
D. Granular access control is a critical security and governance feature, but it is a supporting requirement rather than the primary business driver for adopting AI agents.

References:

1. Google Cloud Blog (May 14, 2024). In the announcement for Vertex AI Agent Builder, it is stated that "Agent Builder helps developers easily build and deploy enterprise-ready generative AI experiences... These agents can help employees and customers get the answers they need and take action, improving customer and employee satisfaction and productivity." This directly links the use of AI agents to improved productivity and data interaction.
Source: Google Cloud Blog, "Introducing Vertex AI Agent Builder: build and deploy enterprise-ready generative AI experiences".

2. Google Cloud Next '24 Keynote. Google CEO Sundar Pichai and Google Cloud CEO Thomas Kurian detailed the vision for AI agents across the enterprise. Kurian stated, "We are introducing a new set of agents for common business functions... to help employees with a variety of tasks, from creative and development to security operations." The entire presentation framed the value of these agents in terms of transforming business processes and radically improving productivity.
Source: Google Cloud Next '24 Opening Keynote, "The New Way to Cloud," April 9, 2024.
3. Vertex AI Agent Builder Documentation. The official documentation describes the service as a tool to "build and deploy enterprise-grade generative AI agents that have access to your enterprise data." The use cases, such as enterprise search and conversational agents, are fundamentally aimed at making information more accessible and actions more automated, which are direct drivers of productivity.
Source: Google Cloud, Vertex AI Agent Builder Documentation, "Overview of Vertex AI Agent Builder".

Question: 10

A company wants to choose a generative AI (gen AI) use case that will be successful and have the most impact. What key factor should they determine first according to Google Cloud-recommended practices?

A. The availability of pre-trained models that are offered on various cloud computing platforms.

B. The frequency of updates to the underlying foundation models used by different gen AI platforms.

C. The specific business problems the company aims to solve and the desired outcomes.

D. The number of employees who will be trained to use the new gen AI tools.

Answer:

C

Explanation:

According to Google Cloud's recommended practices, the most critical first step in selecting a generative AI use case is to clearly define the business problem you are trying to solve and the desired outcomes. This problem-first approach ensures that the technology is applied to create tangible business value and that the project is aligned with strategic objectives. Identifying a specific, high-impact business need provides the necessary foundation and justification for the project, guiding all subsequent technical and implementation decisions. Without a clear business driver, a gen AI project risks becoming a "solution in search of a problem," leading to wasted resources and minimal impact.

Why Incorrect Options are Wrong:

A. The availability of pre-trained models is a technical feasibility assessment that should occur after a specific business problem has been identified.
B. The frequency of model updates is a detailed operational consideration, not the primary strategic factor for initially choosing a use case.
D. The number of employees to be trained is part of the implementation and change management plan, which is developed after the use case is selected.

References:

1. Google Cloud AI Adoption Framework: The framework's first phase, "Identify," emphasizes starting with business objectives. "The first step is to identify potential use cases by understanding your business objectives and challenges to pinpoint areas where generative AI can deliver the most value." This directly supports prioritizing business problems first. (Source: Google Cloud, "The AI Adoption Framework: A guide to scaling AI successfully," Page 5, "Phase 1: Identify").

2. Google Cloud Skills Boost Courseware: In the "Introduction to Generative AI" course, the module on "Generative AI Use Cases" consistently frames the adoption process around solving specific business challenges, such as improving customer service or accelerating content creation. The principle is to match the technology's capability to a pre-defined business need. (Source: Google Cloud Skills Boost, "Introduction to Generative AI," Module: "Generative AI Use Cases").
3. Google Cloud for Leaders Guide: "To get started with generative AI, leaders should first identify the key business challenges and opportunities where this technology can have the most significant impact." This guidance reinforces that the starting point is always the business context, not the technology itself. (Source: Google Cloud, "A C-suite guide to generative AI," Section: "How to get started with generative AI").