# Google Cloud CertifiedGenerative AI Leader Exam

## Google Generative-AI-Leader

**Version Demo**

**Total Demo Questions: 20**

**Total Premium Questions: 200**

**Buy Premium PDF**

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

# Topic Break Down

| Topic | No. of Questions |
|---|---|
| Topic 1, Fundamentals of generative AI | 37 |
| Topic 2, Google Cloud's generative AI offerings | 85 |
| Topic 3, Techniques to improve generative AI model output | 38 |
| Topic 4, Business strategies for successful generative AI solutions | 40 |
| Total | 200 |

QUESTION NO: 1

A large company is creating their generative AI (gen AI) solution by using Google Cloud's offerings. They want to ensure that their mid-level managers contribute to a successful gen AI rollout by following Google-recommended practices. What should the mid-level managers do?

**A.** Perform continuous testing, measurement, and refinement based on user feedback and real-world performance data.

**B.** Create a robust data strategy to ensure teams can access high-quality, relevant data that is appropriate for training and fine-tuning gen AI models.

**C.** Drive gen AI adoption by identifying high-impact, feasible solutions that address specific challenges within their workflows.

**D.** Secure funding and resources for AI initiatives by demonstrating the potential return on investment to the chief financial officer (CFO).

**ANSWER: C**

**Explanation:**

Drive gen AI adoption by identifying high-impact, feasible solutions that address specific challenges within their workflows. Mid-level managers have essential operational knowledge: they understand their teams' processes, recurring bottlenecks, customer needs, and the practical constraints that determine whether a proposed AI solution can be adopted successfully. This positions them to translate an organization's broader AI strategy into focused use cases with clear business value.

Google's guidance for generative AI transformation emphasizes beginning with valuable, achievable use cases and iterating toward broader adoption. Managers can partner with business users to identify opportunities where generative AI improves a concrete workflow, such as accelerating knowledge retrieval, assisting with document creation, or improving service interactions. They can then help prioritize opportunities according to expected impact, feasibility, available data, responsible-AI considerations, and alignment with organizational objectives. This creates useful early results, encourages employee engagement, and supplies evidence for scaling successful solutions across the company.

For more guidance on organizing and advancing a generative AI transformation, see Google Cloud's Generative AI transformation resources and Google Cloud's generative AI architecture guidance.

**QUESTION NO: 2**

A company's development team is eager to start building generative AI solutions with Google Cloud, but has limited experience in AI development. They need to launch their gen AI solution quickly.

What Google Cloud benefit would help the company achieve their goal?

**A.** Google Cloud's pre-trained models and low- and no-code AI tools and services.

**B.** Google Cloud's collaborative AI community and support forums connect developers with AI experts.

**C.** Google Cloud's comprehensive training materials and tutorials to help developers.

**D.** Google Cloud's focus on continuous improvement provides access to the latest AI tools, features, and best practices.

**ANSWER: A**

**Explanation:**

Google Cloud's pre-trained models and low- and no-code AI tools and services directly support a rapid generative AI launch for a team with limited AI experience. Rather than collecting large training datasets, selecting a model architecture, provisioning training infrastructure, and training a model from scratch, developers can access Google's managed foundation models through Vertex AI. These models can perform common generative tasks such as text generation, summarization, question answering, and multimodal content generation through APIs.

Vertex AI also provides managed development capabilities that reduce implementation and operational effort. Teams can prototype with Vertex AI Studio, use prompt design and evaluation workflows, ground model responses in enterprise data, and deploy managed applications without building the underlying machine-learning platform themselves. This enables

developers to focus on the business use case, user experience, prompts, and responsible-AI controls instead of foundational model development and MLOps. Google describes Vertex AI as a unified platform for building with generative AI and provides access to Gemini and other foundation models. These managed, ready-to-use capabilities are therefore the most direct benefit for an organization seeking speed to market while developing its AI skills. See the Vertex AI generative AI overview and Google Cloud generative AI learning overview.

## QUESTION NO: 3

A marketing team wants to use a generative AI model to create product descriptions for their new line of eco-friendly water bottles. They provide a brief prompt stating, "Write a product description for our new water bottle." The model generates a generic, lackluster description that is factually accurate but lacks engaging language and doesn't highlight the environmental benefits that are key to their brand. What should the marketing team do to overcome this limitation of the generated product description?

**A.** Train the model on a dataset of marketing materials from other eco-friendly brands.

**B.** Add details to the prompt about the audience, tone, and keywords.

**C.** Increase the token count for the model to allow for longer descriptions.

**D.** Lower the temperature setting of the model to produce more consistent results.

## ANSWER: B

**Explanation:**

Add details to the prompt about the audience, tone, and keywords. The output is generic because the instruction provides almost no context about the desired marketing message, brand voice, intended customer, or product attributes that must be emphasized. Generative models use the information in their prompt to infer both content and style; therefore, a stronger prompt should specify the target audience, an engaging tone, and concrete sustainability claims such as reusable design, material composition, durability, or reducing single-use plastic. The team can also state a desired format and length, such as a concise e-commerce description with a headline and a call to action. Providing this direction makes the request measurable and gives the model the context needed to produce copy aligned with the brand's environmental positioning. This is a prompt-engineering use case: improve the instructions before considering more resource-intensive customization approaches. Google Cloud's prompt-design guidance recommends clear instructions, relevant context, and explicit output requirements to improve the relevance and quality of model responses. Google also notes that prompts can include examples where needed to demonstrate the desired style or structure. See Google Cloud: Introduction to prompting and Google Cloud: Prompt design strategies.

## QUESTION NO: 4

What does Vertex AI Search enable companies to do?

**A.** To index and retrieve information from the entire public web, providing a comprehensive view of publicly available data.

**B.** To surface the most popular and frequently accessed content based on global user search patterns and trends.

**C.** To compare products from numerous online retailers, allowing users to find the best deals and product options across the internet.

**D.** To ground LLM responses with first-party data, third-party data, and Google's knowledge graph.

## ANSWER: D

**Explanation:**

Vertex AI Search enables organizations to ground large language model responses in enterprise-relevant sources, including their own first-party data, connected third-party data, and Google's knowledge graph. Grounding links a model's response to retrieved, authoritative information rather than relying only on the model's pretrained knowledge. This allows a company to

create search and conversational experiences over sources such as internal documents, websites, knowledge bases, and product catalogs.

In practice, Vertex AI Search retrieves information relevant to a user query and uses it as context for an AI-generated answer. This supports more useful, domain-aware responses and can provide citations or links back to the supporting source material. Because the answer is informed by selected data sources that the organization manages or connects, it is better aligned with current business information, policies, and product content. Google describes Vertex AI Search as a platform for building AI-powered search and recommendation applications over enterprise data, with grounding capabilities for generative answers. See the Vertex AI Search enterprise search overview and the Vertex AI grounding overview.

## QUESTION NO: 5

A company wants employees to ask questions about thousands of internal documents. The documents contain tables, charts, and diagrams in addition to text. Which capability enables a Gemini application to interpret these different content types together?

**A.** Multimodal understanding

**B.** Batch data warehousing

**C.** Relational database normalization

**D.** Network packet inspection

## ANSWER: A

**Explanation:**

Multimodal understanding enables an application to process and reason over more than one type of input, such as text, images, documents, charts, and tables. This is useful when the meaning of a document depends on visual layout or information that is represented outside plain text.

Gemini supports multimodal inputs, allowing developers to build experiences that use information from different modalities in a single request. This can help employees obtain answers that account for both written content and visual elements in internal documents. See the Vertex AI multimodal overview.

## QUESTION NO: 6

A company is evaluating different generative AI (gen AI) platforms and wants to understand the role of the infrastructure layer in supporting the development and deployment of gen AI models. What is the function of the infrastructure layer in the gen AI landscape?

**A.** To provide access to pre-trained gen AI models.

**B.** To provide a user-friendly interface for interacting with gen AI models.

**C.** To provide the compute resources needed to run and train AI models and store training data.

**D.** To provide tools and resources for building and deploying gen AI applications.

## ANSWER: C

**Explanation:**

To provide the compute resources needed to run and train AI models and store training data is correct because the infrastructure layer is the foundational technical layer on which generative AI workloads operate. Generative AI training requires large-scale, highly parallel computation, commonly using accelerators such as GPUs or TPUs, along with high-throughput networking to coordinate distributed workloads. Inference, which generates responses from deployed models, also depends on scalable compute capacity that can meet latency and throughput requirements. This layer additionally provides durable, scalable storage for the large datasets, checkpoints, model weights, and artifacts used throughout training and serving. Google Cloud describes AI infrastructure as including purpose-built accelerators and the supporting compute,

networking, and storage capabilities needed for demanding AI workloads. These underlying resources enable organizations to develop and operate models reliably at scale, while higher layers provide managed model, development, and end-user capabilities. See Google Cloud: What is AI infrastructure? and Google Cloud Compute.

## QUESTION NO: 7

An organization is deploying generative AI applications on Google Cloud. They want to understand the fundamental security measures integrated into Google ' s global network and hardware. What term specifically describes these inherent protective measures?

**A.** Network protection

**B.** Operations management

**C.** Security oversight

**D.** Secure-by-design infrastructure

## ANSWER: D

**Explanation:**

Secure-by-design infrastructure is the term for security capabilities that Google builds directly into the foundational infrastructure used to deliver Google Cloud services. Rather than treating protection as an add-on applied after systems are deployed, this approach incorporates security throughout the design and operation of Google's global network, data centers, servers, and custom hardware. For an organization deploying generative AI workloads, these embedded protections provide a foundational layer beneath its own workload configuration, identity controls, data governance, and application security practices. Google describes its infrastructure as designed with defense in depth, using multiple layers of physical, hardware, software, and operational safeguards. This includes purpose-built servers and hardware components, a globally distributed private network, encryption protections, and continuous infrastructure security processes. The shared-responsibility model means that these Google-managed foundational measures work alongside the customer's responsibility to securely configure and govern its cloud resources. Therefore, when the question asks for the broad concept describing inherent protections integrated into Google's network and hardware from the outset, secure-by-design infrastructure is the precise term. See the Google infrastructure security design overview and Google Cloud shared responsibility model.

## QUESTION NO: 8

A manager wants to ensure that only quality data is used in their AI model. Which scenario is most likely to lead to an unfair and biased outcome?

**A.** A financial institution builds a fraud detection model using transaction data that is incomplete, lacking details about the merchants involved in a significant percentage of the transactions.

**B.** A company ' s customer service natural language processing model is trained on a dataset containing incorrectly encoded text from older website versions.

**C.** A marketing team creates a personalized recommendation engine based on website clickstream data that includes bot traffic.

**D.** A technology firm trains a facial recognition system primarily using images of individuals from one demographic group, with limited representation from other groups.

## ANSWER: D

**Explanation:**

A technology firm trains a facial recognition system primarily using images of individuals from one demographic group, with limited representation from other groups. This is most likely to produce an unfair and biased outcome because the training data does not adequately represent the population on which the system may be used. Machine learning models learn patterns from their examples; when some demographic groups have far fewer examples, the model can have less reliable

feature learning and higher error rates for those groups. In facial recognition, this can result in systematically different identification or verification performance across demographic populations, creating an equity and fairness concern rather than merely a general reduction in data quality. Google's responsible AI guidance emphasizes evaluating datasets for representativeness and measuring model behavior across relevant groups to identify and mitigate unfair outcomes before deployment. Data should be sufficiently diverse, evaluated for coverage gaps, and monitored after release so that performance disparities can be detected and addressed. This scenario directly illustrates representation bias caused by an imbalanced dataset. See Google Cloud's Responsible AI guidance and the Google AI Principles.

## QUESTION NO: 9

A marketing team wants to use a generative AI model to create product descriptions for their new line of eco-friendly water bottles. They provide a brief prompt stating, " Write a product description for our new water bottle. " The model generates a generic, lackluster description that is factually accurate but lacks engaging language and doesn ' t highlight the environmental benefits that are key to their brand. What should the marketing team do to overcome this limitation of the generated product description?

**A.** Train the model on a dataset of marketing materials from other eco-friendly brands.

**B.** Add details to the prompt about the audience, tone, and keywords.

**C.** Increase the token count for the model to allow for longer descriptions.

**D.** Lower the temperature setting of the model to produce more consistent results.

## ANSWER: B

**Explanation:**

Add details to the prompt about the audience, tone, and keywords. The model's response reflects the limited direction in the original request: it identifies the product type but provides no information about the intended customers, brand voice, differentiating product attributes, or required sustainability claims. Prompt engineering improves output quality by supplying this context explicitly. For example, the team could identify eco-conscious consumers as the audience, request an upbeat and persuasive tone, state factual bottle features such as reusable materials or BPA-free construction, and require messaging about reducing single-use plastic. They can also provide a desired structure, length, and example phrasing, while ensuring that environmental statements are accurate and substantiated. These instructions give the model clear constraints and goals, allowing it to generate copy that is both engaging and aligned with the brand's environmental positioning. Google's guidance for generative AI emphasizes being specific, providing relevant context, and clearly defining the desired format and style when prompting models. Iterating on the prompt and evaluating the resulting descriptions is a practical way to refine marketing content without changing the underlying model. See Google Cloud prompt design guidance and Google AI prompting strategies.

## QUESTION NO: 10

A company ' s sales team spends a significant amount of time researching potential leads and manually entering data into their customer relationship management (CRM) tool. They want to improve the team ' s efficiency and enable them to focus on building relationships and closing deals. What should the organization do?

**A.** Integrate the CRM with a popular sales intelligence platform to automatically enrich lead profiles with valuable data.

**B.** Develop a custom AI solution using Google Cloud's AutoML Natural Language to analyze lead communications and automatically update the CRM.

**C. Implement Gemini Enterprise " unified enterprise search " including a CRM agent to automate lead research and data entry.**

**D.** Implement Google Cloud ' s Contact Center AI to qualify leads and route them to the appropriate sales representatives.

## ANSWER: C

**Explanation:**

**Implement Gemini Enterprise "unified enterprise search" including a CRM agent to automate lead research and data entry.** is the best fit because it addresses both time-consuming activities as an integrated generative AI workflow. Gemini Enterprise can provide grounded access to enterprise information through search and data connectors, allowing sales users or agents to locate and synthesize relevant information about a prospective customer from approved organizational sources. This reduces the effort required to gather account context before outreach.

An agent can then apply that retrieved context to a sales workflow, such as preparing lead summaries, extracting relevant details, and invoking authorized actions against connected business systems. With appropriate identity, permissions, governance, and CRM integration in place, this lets the organization automate routine enrichment and record-update tasks while keeping sales representatives responsible for relationship building and deal execution. The result is a practical combination of enterprise search, generative reasoning, and agentic action rather than a standalone model used for only one isolated task.

Google describes Gemini Enterprise as a platform for connecting enterprise data, searching it, and building AI agents in its Gemini Enterprise overview. Its agent documentation further explains how organizations can create agents that use enterprise data and tools to complete business tasks.

## QUESTION NO: 11

A logistics company wants to use a generative AI (gen AI) agent to automatically check real-time inventory levels across its warehouses and adjust delivery schedules. The gen AI agent needs access to internal inventory data. They want the most cost-effective solution. What should the organization do?

**A.** Build a custom API instead of using the gen AI agent.

**B.** Use pre-built gen AI chatbots for inventory questions.

**C.** Use Vertex AI Studio to fine-tune a model with sample inventory data.

**D.** Use Google Cloud databases and Vertex AI for the agent to get live data.

## ANSWER: D

**Explanation:**

Use Google Cloud databases and Vertex AI for the agent to get live data. This approach directly supports the operational requirement: inventory and delivery decisions must be based on the current state of warehouse systems, not on information embedded in a model at an earlier point in time. The inventory system remains the authoritative source of truth in an appropriate managed data store, while the agent retrieves the required current values when it performs a task.

Vertex AI supports building generative AI applications and agents that connect models to enterprise systems through tools, functions, and APIs. The agent can invoke a controlled service that queries inventory data, apply the returned results in its reasoning, and then invoke an approved scheduling service to adjust deliveries. This architecture also supports least-privilege access, auditing, validation of requested actions, and separation between the model and sensitive operational data.

It is cost-effective because the organization can reuse its existing inventory data platform and integrate only the live-data capabilities required by the agent, rather than training and continually retraining a model as stock levels change. Google documents agent development capabilities in Vertex AI Agent Engine overview and enterprise data grounding approaches in Vertex AI grounding overview.

## QUESTION NO: 12

A company wants to adopt generative AI and is concerned about vendor lock-in. They want to maintain flexibility in their technology stack. What Google Cloud strength would ease their concerns?

**A.** Google Cloud’s AI solutions have an open approach that supports customer choice across offerings.

**B.** Google Cloud ' s AI solutions are pre-packaged for easy deployment, eliminating the need for customization and integration efforts.

**C.** Google Cloud ' s strict adherence to proprietary technologies ensures the highest level of security and performance.

**D.** Google Cloud ' s focus on automation aims to replace human jobs with AI systems, potentially leading to significant workforce reductions.

**ANSWER: A**

**Explanation:**

Google Cloud's AI solutions have an open approach that supports customer choice across offerings. This directly addresses vendor lock-in concerns because Google Cloud supports open technologies, interoperability, and portability across the AI lifecycle. Customers can use managed Google Cloud services where they add value while continuing to use familiar open-source frameworks and industry-standard tools in their existing environments. Google has made significant contributions to open-source technologies used broadly across modern application and AI platforms, including Kubernetes, TensorFlow, and the Gemma family of open models. For generative AI, this approach enables organizations to select models and deployment patterns that meet their technical, governance, and operational needs rather than requiring a single proprietary stack. Vertex AI also supports a range of foundation models and development tools, helping teams build, tune, evaluate, and deploy AI solutions with flexibility. Maintaining choice among models, frameworks, infrastructure, and integration approaches lets a company evolve its architecture over time and reduce dependency on any one implementation. This combination of open ecosystems and managed cloud capabilities is therefore a core Google Cloud strength for an organization seeking innovation without sacrificing technology-stack flexibility. See Vertex AI documentation and Google AI for Developers: Gemma.

**QUESTION NO: 13**

A marketing team wants to use a foundation model to create social media and advertising campaigns. They want to create written articles and images from text. They lack deep AI expertise and need a versatile solution. Which Google foundation model should they use?

**A.** Gemma

**B.** Imagen

**C.** Gemini

**D.** Veo

**ANSWER: C**

**Explanation:**

**Gemini** is the appropriate choice because it is Google's multimodal foundation-model family, designed to work across multiple content types and tasks through a unified model interface. For a marketing team, this supports a broad campaign workflow: developing article copy, headlines, calls to action, audience-oriented variations, creative briefs, and image-related content from natural-language prompts. This versatility is especially valuable when the team does not have deep AI expertise, because Gemini can be accessed through Google's managed generative AI products and APIs rather than requiring the team to train a specialized model.

Google describes Gemini as a multimodal model capable of understanding and combining information across modalities, including text and images. In Vertex AI, Gemini models support text generation and multimodal prompting, and image-generation capabilities are available in the Gemini model family for applicable use cases and regions. This makes Gemini the best fit for a single, flexible foundation-model solution supporting both written marketing assets and text-prompted visual creative work. See the Vertex AI multimodal overview and Gemini models documentation.

**QUESTION NO: 14**

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:**

Use role prompting. This technique gives the model a defined persona, perspective, or professional role and provides context about the intended audience. For example, the organization can instruct the model to act as an experienced marketing copywriter creating an approachable campaign for a particular customer segment, such as young technology enthusiasts, small-business owners, or enterprise decision-makers. The prompt can further identify the desired brand voice, reading level, cultural context, communication goals, and tone. These instructions help the model select language, style, examples, and messaging that fit the campaign's audience while remaining aligned with the organization's marketing intent.

Google's prompt-design guidance describes persona as a way to tell a model who it should act as, which steers the form and style of its response. Persona instructions are especially useful when a task requires outputs tailored to a particular user group or communication setting. In a marketing workflow, role prompting can be combined with clear audience and brand requirements to produce more consistently suitable draft copy. See Google Cloud's persona prompting guidance and its prompt design introduction.

**QUESTION NO: 15**

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 the appropriate model family because Google developed it as a family of lightweight, openly available generative AI models designed to make capable language-model development more accessible. Gemma models are based on the same research and technology used to create Gemini models, but are released with open weights and offered in sizes that can be run efficiently on developer hardware. This makes them well suited to local experimentation, rapid chatbot prototyping, and fine-tuning for a specific domain or task. Google provides tooling and integration paths for adapting Gemma, including parameter-efficient tuning methods, so developers can customize behavior without needing to train a model from scratch. The family includes compact variants intended for constrained environments, while still supporting high-quality natural-language understanding and text generation. These characteristics directly meet the need for an efficient local model with strong conversational language capabilities. Google describes Gemma as a collection of lightweight, state-of-the-art open models, and its documentation provides resources for running and tuning the models locally. See the Gemma documentation and Google Developers Blog overview of the Gemma family.

**QUESTION NO: 16**

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:**

A customized learning agent is the appropriate solution because it can coordinate multiple personalized learning tasks as part of a goal-oriented educational experience. The agent can use student inputs, assessment results, prior performance, and curriculum objectives to determine a learner's current level of knowledge. It can then select or recommend suitable resources, generate exercises targeted to identified gaps, and adapt subsequent activities based on progress and responses. This combines generative capabilities with decision-making and workflow orchestration, rather than treating each task as an isolated feature.

For this use case, the agent can be customized with the organization's pedagogical approach, lesson structure, approved content sources, assessment criteria, and rules for recording learner progress. It may also use tools or connected systems to retrieve learning materials and update student records. Google Cloud describes AI agents as systems that reason, plan, and take actions to achieve goals, including by using tools and enterprise data. This makes a customized learning agent well suited to delivering an ongoing, adaptive learning journey. See Vertex AI Agent Engine overview and Google Cloud: What are AI agents?.

## QUESTION NO: 17

The office of the CISO wants to use generative AI (gen AI) to help automate tasks like summarizing case information, researching threats, and taking actions like creating detection rules. What agent should they use?

**A.** Security agent

**B.** Data agent

**C.** Code agent

**D.** Customer service agent

**ANSWER: A**

**Explanation:**

**Security agent** is the appropriate choice because the requested work is a security-operations workflow: understanding and summarizing case or incident context, enriching an investigation with threat research, and assisting with response-oriented tasks such as creating detection rules. A security-focused agent is designed to reason over security telemetry and case context, support analysts during investigation, and help operationalize findings in security tools and processes. This directly supports a CISO organization's objectives of improving analyst productivity, accelerating detection and response, and making investigations more consistent.

Google Cloud's security offerings apply generative AI to security operations by helping teams investigate alerts and cases, obtain contextual explanations, conduct threat intelligence research, and streamline response activities. In Google Security Operations, Gemini capabilities are intended to assist security teams with investigating threats and operating their security workflows. Google also positions Security Command Center as a platform for identifying, prioritizing, and helping remediate cloud security risks. These capabilities align with an agent specialized for security rather than a general business workflow. See Google Security Operations and Security Command Center overview.

## QUESTION NO: 18

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:**

Reinforcement learning is the appropriate approach because the game character is an agent that must repeatedly choose actions while interacting with an environment. Rather than learning from a fixed set of labeled examples, it receives feedback after its actions in the form of positive scores for beneficial outcomes and negative scores for undesirable outcomes. This feedback is a reward signal. The learning objective is to discover a policy—a strategy for selecting actions based on the current game state—that maximizes cumulative reward over time.

This is particularly well suited to game navigation and objective completion because an action may have consequences beyond the immediate step. For example, a move that does not yield an instant reward can still be valuable if it enables the character to reach a goal later. By exploring possible actions, observing resulting state changes and rewards, and updating its policy through repeated episodes, the agent can improve its game-playing behavior through trial and error. Google Cloud describes reinforcement learning as training an agent to make decisions by interacting with an environment and receiving rewards or penalties. See Google Cloud: What is reinforcement learning? and Google Machine Learning Crash Course: Reinforcement learning.

**QUESTION NO: 19**

A global news agency is developing a generative AI tool to quickly summarize breaking news articles as they emerge online. The goal is to provide their audience with rapid updates on fast-developing stories from various global sources. What Google Cloud solution should they use?

**A.** Document AI

**B.** BigQuery

**C.** Vertex AI Natural Language API

**D.** Grounding with Google Search

**ANSWER: D**

**Explanation:**

Grounding with Google Search is the appropriate solution because a breaking-news summarization workflow depends on timely, externally current information rather than only the model's pre-existing training knowledge. When a supported Gemini model is grounded with Google Search in Vertex AI, it can use Google Search to retrieve relevant, up-to-date public web information while generating a response. This is particularly useful for rapidly changing global events, where new reporting and updates may appear after the model's knowledge cutoff. The service can also return grounding metadata, including search queries and supporting source links, which helps a news organization make generated summaries more transparent and enables editorial teams or audiences to inspect the information used. The agency should still apply editorial review, source-quality policies, and verification procedures before publishing, but Google Search grounding provides the needed current-information retrieval capability for this use case. Google documents the feature and its supported models in the Grounding with Google Search guide. Additional details about grounding responses and source metadata are available in the Grounding metadata documentation.

## QUESTION NO: 20

A software development team wants to use generative AI (gen AI) to code faster so they can launch their software prototype quicker. What should the team do?

**A.** Use gen AI to refactor and optimize existing code.

**B.** Use gen AI to suggest code snippets and complete functions.

**C.** Use gen AI to automatically generate comprehensive documentation for their code.

**D.** Use gen AI to identify potential bugs and security vulnerabilities in their code.

## ANSWER: B

**Explanation:**

Use gen AI to suggest code snippets and complete functions. This directly accelerates the core work required to build a software prototype: translating requirements into working code. AI-assisted coding tools can generate boilerplate, propose implementation patterns, autocomplete code, and create functions from natural-language prompts or surrounding context. This lets developers spend more of their limited prototype time integrating features, validating user flows, and testing whether the product idea solves the intended problem.

Google Cloud's Gemini Code Assist is designed for this development workflow. It provides code completion and code generation in supported IDEs, helping developers write and modify code with less manual effort. For an early prototype, these capabilities support rapid iteration: a developer can describe a desired behavior, receive an initial implementation, review it, and adapt it to the application's needs. Human review and testing remain essential, but AI suggestions reduce time spent on repetitive coding tasks and help teams reach a demonstrable version faster.

See the Gemini Code Assist overview and Google Cloud's documentation on writing code with Gemini Code Assist.