

CERT-PASS

Google Generative AI Leader

Free Practice Questions Preview

Here are 35 sample questions to help you get started. Unlock the full exam to access all 1100+ questions with detailed explanations.

Question 1 : Section 1: Fundamentals of gen AI

A media company is evaluating four approaches for an executive decision-support pilot. Which choice most directly enables the organization to use a managed Google Cloud platform to develop, deploy, and manage machine learning models across the ML lifecycle?

- A. Use only Cloud Storage with no model-management capability
- B. Use Agent Platform for managed AI development and model operations**
- C. Use Imagen as the ML lifecycle platform
- D. Use the Gemini app as the model-operations platform

Answer: B

Agent Platform is the managed Google Cloud platform for building and operating AI solutions across lifecycle stages. The most tempting incorrect option fails because the Gemini app is a prebuilt assistant experience, not the managed platform for model development and operations.

Question 2 : Section 1: Fundamentals of gen AI

A research institute is planning an agent-enabled business process. The team needs to improve responses by revising instructions, context, examples, and output constraints without retraining the model. Which recommendation is best?

- A. Train a new model from scratch for every request
- B. Increase the number of IAM roles
- C. Replace the model with a database index
- D. Use prompt engineering**

Answer: D

Prompt engineering improves model behavior by designing effective prompts. The most tempting incorrect option fails because training a model from scratch is unnecessarily expensive when the issue is unclear instructions.

Question 3 : Section 1: Fundamentals of gen AI

Which option most directly addresses this requirement: improve AI results by addressing incomplete, inconsistent, irrelevant, unavailable, or incorrectly formatted source information?

- A. Increase temperature as the only action
- B. Ignore source quality because a foundation model corrects every issue
- C. Improve data quality and accessibility**
- D. Add more user interface animations

Answer: C

AI outcomes depend on relevant, complete, consistent, available, and usable data. The most tempting incorrect option fails because increasing temperature changes randomness and can worsen reliability when source quality is poor.

Question 4 : Section 2: Google Cloud's gen AI offerings

A telecommunications provider is improving a document-processing workflow. Which option contains the best primary choice and supporting practice to select the Google Cloud platform for building and operating custom AI agents with enterprise capabilities?

- A. Use only a spreadsheet macro; then Rely on a one-time demo and do not monitor the production workflow.
- B. Use Agent Platform; then Confirm tool permissions, integration behavior, and governance before rollout.**
- C. Use only Gemini for Google Workspace; then Skip validation because the first response looks reasonable.
- D. Use only Text-to-Speech API; then Assume a larger model automatically fixes the underlying issue.

Answer: B

Agent Platform provides capabilities for developing custom agents and related AI experiences. The most tempting incorrect option fails because Gemini for Google Workspace is a prebuilt productivity integration, not the platform for building custom enterprise agents.

Question 5 : Section 2: Google Cloud's gen AI offerings

A global manufacturer is designing a productivity-improvement program. A stakeholder suggests "Use only a spreadsheet macro." The actual requirement is to select the Google Cloud platform for building and operating custom AI agents with enterprise capabilities. Which alternative should the leader recommend?

- A. Use only Text-to-Speech API
- B. Use only Gemini for Google Workspace
- C. Use only a spreadsheet macro
- D. Use Agent Platform**

Answer: D

Agent Platform provides capabilities for developing custom agents and related AI experiences. The most tempting incorrect option fails because Gemini for Google Workspace is a prebuilt productivity integration, not the platform for building custom enterprise agents.

Question 6 : Section 2: Google Cloud's gen AI offerings

Which option best matches this requirement: identify Google's custom-designed accelerators used for AI workloads?

A. Tensor Processing Units (TPUs)

B. IAM conditions

C. Workspace templates

D. Cloud Storage buckets

Answer: A

TPUs are Google-designed accelerators optimized for machine-learning workloads. The most tempting incorrect option fails because Cloud Storage stores data but is not an AI accelerator.

Question 7 : Section 4: Business strategies for a successful gen AI solution

A public-sector agency has launched a pilot for a responsible AI adoption roadmap. The next design decision is how to reduce privacy risk when using personal data in an AI workflow. Which choice is most appropriate?

A. Use appropriate privacy controls, including anonymization or pseudonymization where suitable

B. Assume consent removes every privacy risk

C. Disable access controls

D. Publish raw identifiers in prompts

Answer: A

Anonymization and pseudonymization can reduce exposure of personal identifiers when applied appropriately. The most tempting incorrect option fails because publishing raw identifiers increases privacy risk.

Question 8 : Section 1: Fundamentals of gen AI

A legal-services firm is evaluating four approaches for an accessible digital experience. Which choice most directly enables the organization to choose Google's foundation-model family for multimodal understanding, conversational AI, content creation, and question answering?

A. Veo

B. Imagen

C. Gemini

D. Gemma

Answer: C

Gemini supports multimodal understanding and general-purpose generative AI experiences. The most tempting incorrect option fails because Imagen is focused on generating images from text, not broad multimodal conversational use.

Question 9 : Section 2: Google Cloud's gen AI offerings

A research institute is conducting a governance review of an accessible digital experience. The review confirms that the solution must provide an enterprise employee experience with internal information access, multimodal search, and custom agent capabilities, while supporting a controlled enterprise rollout. What is the best recommendation?

- A. Use only the public Gemini app with no enterprise context
- B. Use only Translation API
- C. Use only a relational database console
- D. Use Gemini Enterprise**

Answer: D

Gemini Enterprise supports enterprise AI experiences, including multimodal search and custom agent capabilities. The most tempting incorrect option fails because the public Gemini app alone does not provide the same enterprise-oriented internal search and agent capabilities.

Question 10 : Section 4: Business strategies for a successful gen AI solution

A professional-services firm needs a production-ready decision for a developer enablement initiative. The primary goal is to build adoption by preparing employees for new workflows and clarifying how their roles will change, while maintaining a practical path to measurable value. Which choice should be approved?

- A. Measure only model size
- B. Avoid employee communication
- C. Treat adoption as automatic after launch
- D. Include change management in the implementation plan**

Answer: D

Change management supports successful organizational adoption of gen AI solutions. The most tempting incorrect option fails because launching without employee preparation can reduce adoption even when the technology works.

Question 11 : Section 3: Techniques to improve gen AI model output

A hospitality group is improving a research knowledge portal. Which option contains the best primary choice and supporting practice to improve reliability for rare, atypical scenarios that were missed during initial testing?

- A. Evaluate only the most common happy path; then Assume a larger model automatically fixes the underlying issue.
- B. Assume production usage will match the test set exactly; then Rely on a one-time demo and do not monitor the production workflow.
- C. Disable feedback channels; then Skip validation because the first response looks reasonable.

D. Add edge-case evaluation, monitoring, and human oversight where risk warrants it; then Define measurable KPIs, track versions, and review changes over time.

Answer: D

Rare scenarios require deliberate evaluation and monitoring rather than happy-path testing alone. The most tempting incorrect option fails because testing only common cases leaves the known weakness unaddressed.

Question 12 : Section 4: Business strategies for a successful gen AI solution

Case study: A legal-services firm needs a two-part recommendation for a claims-support process. Which paired option best helps the team integrate gen AI into an organization in a way that moves from vision to prioritized use cases and measurable adoption?

- A. Select Start with an ungoverned company-wide rollout; then Rely on a one-time demo and do not monitor the production workflow.
- B. Select Focus only on a model demo; then Choose the most complex option even when it does not address the requirement.
- C. Establish a vision, prioritize use cases, build capabilities, manage change, measure value, and champion responsible AI; then Validate the choice against representative scenarios and measure the intended outcome.**
- D. Select Skip change management; then Remove governance controls to speed up the pilot.

Answer: C

Successful transformation combines strategic vision, prioritized delivery, capability building, change management, measurement, and responsibility. The most tempting incorrect option fails because an ungoverned rollout increases risk and makes value difficult to prove.

Question 13 : Section 1: Fundamentals of gen AI

An ecommerce marketplace is evaluating four approaches for a research knowledge portal. Which choice most directly enables the organization to choose Google's open model family for developers seeking customizable models for local deployments or specialized applications?

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

Answer: C

Gemma provides open models that developers can customize and use in specialized or local scenarios. The most tempting incorrect option fails because the Gemini app is a prebuilt assistant experience, not an open model family for local deployment.

Question 14 : Section 2: Google Cloud's gen AI offerings

For a document-processing workflow, A marketing agency is considering "Use only Text-to-Speech API." Which option is a better fit when the goal is to select the Google Cloud platform for building and operating custom AI agents with enterprise capabilities?

- A. Use only Text-to-Speech API
- B. Use Agent Platform**
- C. Use only a spreadsheet macro
- D. Use only Gemini for Google Workspace

Answer: B

Agent Platform provides capabilities for developing custom agents and related AI experiences. The most tempting incorrect option fails because Gemini for Google Workspace is a prebuilt productivity integration, not the platform for building custom enterprise agents.

Question 15 : Section 4: Business strategies for a successful gen AI solution

A financial-services firm is evaluating four approaches for a developer enablement initiative. Which choice most directly enables the organization to improve visibility into security posture and monitor workloads supporting AI systems?

- A. Use only a larger context window
- B. Use only a few-shot prompt
- C. Use Security Command Center and workload monitoring tools**
- D. Use only the Gemini app

Answer: C

Security Command Center and monitoring tools improve security visibility and operational oversight. The most tempting incorrect option fails because a larger context window does not provide security-posture visibility.

Question 16 : Section 1: Fundamentals of gen AI

Case study: A consumer-goods company needs a two-part recommendation for a marketing-content workflow. Which paired option best helps the team improve AI results by addressing incomplete, inconsistent, irrelevant, unavailable, or incorrectly formatted source information?

- A. Add more user interface animations; then Assume a larger model automatically fixes the underlying issue.
- B. Increase temperature as the only action; then Rely on a one-time demo and do not monitor the production workflow.
- C. Improve data quality and accessibility; then Validate the choice against representative scenarios and measure the intended outcome.**
- D. Ignore source quality because a foundation model corrects every issue; then Skip validation because the first response looks reasonable.

Answer: C

AI outcomes depend on relevant, complete, consistent, available, and usable data. The most tempting incorrect option fails because increasing temperature changes randomness and can worsen reliability when source quality is poor.

Question 17 : Section 2: Google Cloud's gen AI offerings

A global manufacturer is planning a developer enablement initiative. The team needs to modernize the contact-center foundation with an enterprise-grade cloud-native service. Which recommendation is best?

- A. Use only Cloud Vision API
- B. Use only the Gemini app

C. Use only a TPU reservation

D. Use Google Cloud Contact Center as a Service (CCaaS)

Answer: D

CCaaS provides an enterprise-grade cloud-native contact-center foundation. The most tempting incorrect option fails because the Gemini app is a general assistant and does not replace a contact-center platform.

Question 18 : Section 2: Google Cloud's gen AI offerings

For a multilingual support experience, A consumer-goods company is considering "Gems." Which option is a better fit when the goal is to identify the compute accelerators that can be selected for many model training and inference workloads?

A. Gems

B. Conversational Insights

C. Graphics Processing Units (GPUs)

D. Document AI processors

Answer: C

GPUs provide accelerated computation for many AI workloads. The most tempting incorrect option fails because gems are personalized assistants within Gemini, not hardware accelerators.

Question 19 : Section 2: Google Cloud's gen AI offerings

A travel company has launched a pilot for a contact-center transformation. The next design decision is how to select a prebuilt API instead of building a custom model for speech, translation, document, image, video, or text analysis tasks. Which choice is most appropriate?

A. Use the relevant prebuilt Google Cloud AI API

B. Use only a GPU cluster

C. Train a new foundation model for every narrow task

D. Use only IAM

Answer: A

Prebuilt APIs accelerate common AI capabilities such as speech, translation, document processing, vision, video, and language analysis. The most tempting incorrect option fails because training a new foundation model for a common narrow task is usually unnecessary and inefficient.

Question 20 : Section 1: Fundamentals of gen AI

A consumer-goods company needs a production-ready decision for an employee knowledge assistant. The primary goal is to explain why one large pretrained model can be adapted to many downstream business tasks, without adding unnecessary implementation complexity. Which choice should be approved?

A. A foundation model is only a hardware accelerator

B. A foundation model contains only labeled rows from a relational table

C. A foundation model is trained broadly and can be adapted to many tasks

D. A foundation model is a fixed workflow that cannot be customized

Answer: C

Foundation models are trained on large, diverse datasets and can be adapted to multiple use cases. The most tempting incorrect option fails because a fixed workflow lacks the broad learned capabilities and adaptability of a foundation model.

Question 21 : Section 3: Techniques to improve gen AI model output

A leader needs to manage reusable features consistently for AI workflows using the product name in the current exam guide. Which choice is most appropriate?

A. Use Model Garden as the exam-guide product name

B. Use Agent Platform Feature Store

C. Use Gemini Enterprise as a feature registry

D. Use the Gemini app as a feature registry

Answer: B

The current exam guide names Agent Platform Feature Store among continuous monitoring and evaluation considerations. The most tempting incorrect option fails because Model Garden catalogs models and is not a feature store.

Question 22 : Section 2: Google Cloud's gen AI offerings

An energy provider selected "Use Cloud Storage only," for a research knowledge portal, but the pilot still cannot provide customers with an AI-powered conversational self-service experience. What should the team use instead?

A. Use Agent Assist only

B. Use Conversational Insights only

C. Use Conversational Agents

D. Use Cloud Storage only

Answer: C

Conversational Agents can provide chatbot or virtual-agent experiences for customers. The most tempting incorrect option fails because Agent Assist supports human agents during interactions rather than replacing self-service with a conversational agent.

Question 23 : Section 1: Fundamentals of gen AI

A regional bank needs a production-ready decision for an agent-enabled business process. The primary goal is to identify the data type represented by emails, PDFs, audio recordings, and free-form documents, while supporting a controlled enterprise rollout.

Which choice should be approved?

- A. Unstructured data
- B. An IAM binding
- C. A reward signal
- D. Structured data only

Answer: A

Documents, recordings, and free-form text are typical examples of unstructured data. The most tempting incorrect option fails because structured data is organized into predefined fields such as rows and columns.

Question 24 : Section 3: Techniques to improve gen AI model output

A travel company selected "Use only a deterministic spreadsheet formula," for an employee knowledge assistant, but the pilot still cannot retrieve relevant enterprise documents at request time, add them to the model context, and generate a grounded answer. What should the team use instead?

- A. Use retrieval-augmented generation (RAG)
- B. Use fine-tuning as the only way to retrieve changing documents
- C. Use only a longer output limit
- D. Use only a deterministic spreadsheet formula

Answer: A

RAG retrieves relevant information, augments the prompt, and generates an answer using that context. The most tempting incorrect option fails because fine-tuning changes model behavior but does not perform request-time document retrieval.

Question 25 : Section 1: Fundamentals of gen AI

A telecommunications provider is evaluating four approaches for a model-governance program. Which choice most directly enables the organization to name the area of AI most directly concerned with understanding and working with human language?

- A. Reinforcement learning
- B. Natural language processing (NLP)
- C. Diffusion modeling
- D. Cloud infrastructure orchestration

Answer: B

NLP focuses on understanding, processing, and generating human language. The most tempting incorrect option fails because reinforcement learning is a learning approach based on rewards and penalties, not a language-specific field.

Question 26 : Section 3: Techniques to improve gen AI model output

A business team wants to reduce the risk that outputs systematically disadvantage particular groups. What should it select?

- A. Use a larger model and skip evaluation
- B. Remove human review
- C. Increase temperature
- D. Evaluate for bias and fairness, improve data and controls, and include appropriate oversight**

Answer: D

Responsible deployment requires bias and fairness evaluation plus appropriate mitigation and oversight. The most tempting incorrect option fails because a larger model does not automatically eliminate bias or fairness risks.

Question 27 : Section 3: Techniques to improve gen AI model output

A public-sector agency is evaluating four approaches for an employee knowledge assistant. Which choice most directly enables the organization to retrieve relevant enterprise documents at request time, add them to the model context, and generate a grounded answer?

- A. Use retrieval-augmented generation (RAG)**
- B. Use fine-tuning as the only way to retrieve changing documents
- C. Use only a deterministic spreadsheet formula
- D. Use only a longer output limit

Answer: A

RAG retrieves relevant information, augments the prompt, and generates an answer using that context. The most tempting incorrect option fails because fine-tuning changes model behavior but does not perform request-time document retrieval.

Question 28 : Section 2: Google Cloud's gen AI offerings

A retail company is evaluating four approaches for a model-governance program. Which choice most directly enables the organization to recommend a prebuilt Google assistant for writing, summarizing, translating, and other personal productivity tasks?

- A. Use the Gemini app**
- B. Use a TPU directly
- C. Use Security Command Center
- D. Use Cloud Video Intelligence API

Answer: A

The Gemini app provides a prebuilt generative AI assistant experience for common productivity tasks. The most tempting incorrect option fails because Security Command Center is for security posture management, not personal productivity.

Question 29 : Section 2: Google Cloud's gen AI offerings

A software vendor is improving a multilingual support experience. Which option contains the best primary choice and supporting practice to translate formatted business documents while preserving document structure where supported?

- A. Use the Cloud Vision API; then Remove governance controls to speed up the pilot.
- B. Use Gemini for Google Workspace as the translation backend; then Rely on a one-time demo and do not monitor the production workflow.
- C. Use the Document Translation API; then Confirm that the selected offering matches the user experience, data, and integration requirements.**
- D. Use only the Translation API for plain text strings; then Choose the most complex option even when it does not address the requirement.

Answer: C

Document Translation is designed for translating documents while preserving supported formatting and structure. The most tempting incorrect option fails because the basic Translation API is suitable for text but does not specifically address formatted-document translation.

Question 30 : Section 1: Fundamentals of gen AI

A financial-services firm is working on a customer-service modernization program and must choose between customization effort and speed to value when selecting a foundation model. Which option provides the strongest recommendation and follow-up practice?

- A. Ignore the task-specific behavior needed by users; then Remove governance controls to speed up the pilot.
- B. Choose solely by the largest parameter count; then Choose the most complex option even when it does not address the requirement.
- C. Evaluate fine-tuning and customization requirements; then Validate the choice against representative scenarios and measure the intended outcome.**
- D. Select a model only by its marketing name; then Rely on a one-time demo and do not monitor the production workflow.

Answer: C

Foundation-model selection should account for the level of customization the business use case requires. The most tempting incorrect option fails because the largest model is not automatically the best choice when customization, latency, and cost differ.

Question 31 : Section 4: Business strategies for a successful gen AI solution

A pharmaceutical company selected "Measure only the number of model parameters," for a productivity-improvement program, but the pilot still cannot evaluate people, money, and time before committing to an AI transformation plan. What should the team use instead?

- A. Ignore organizational readiness
- B. Assess available expertise, budget, and timeline**
- C. Measure only the number of model parameters

D. Assume unlimited resources

Answer: B

People, budget, and timeline shape what the organization can implement successfully. The most tempting incorrect option fails because model size does not replace a realistic resource assessment.

Question 32 : Section 3: Techniques to improve gen AI model output

A transportation company is designing a contact-center transformation. A stakeholder suggests "Use fine-tuning as the only way to retrieve changing documents." The actual requirement is to retrieve relevant enterprise documents at request time, add them to the model context, and generate a grounded answer. Which alternative should the leader recommend?

- A. Use only a deterministic spreadsheet formula
- B. Use retrieval-augmented generation (RAG)**
- C. Use fine-tuning as the only way to retrieve changing documents
- D. Use only a longer output limit

Answer: B

RAG retrieves relevant information, augments the prompt, and generates an answer using that context. The most tempting incorrect option fails because fine-tuning changes model behavior but does not perform request-time document retrieval.

Question 33 : Section 2: Google Cloud's gen AI offerings

A global manufacturer is designing an internal search experience. A stakeholder suggests "Use only a data warehouse dashboard." The actual requirement is to define a specific action that an agent can invoke, such as checking inventory or creating a ticket. Which alternative should the leader recommend?

- A. Use only a data warehouse dashboard
- B. Use top-p sampling
- C. Use only output-length limits
- D. Expose the action as a function**

Answer: D

Functions let an agent invoke well-defined actions. The most tempting incorrect option fails because top-p sampling influences token selection and does not implement a business action.

Question 34 : Section 2: Google Cloud's gen AI offerings

A university is planning an internal search experience. The team needs to define a specific action that an agent can invoke, such as checking inventory or creating a ticket. Which recommendation is best?

- A. Use only a data warehouse dashboard

- B. Use only output-length limits
- C. Use top-p sampling
- D. Expose the action as a function**

Answer: D

Functions let an agent invoke well-defined actions. The most tempting incorrect option fails because top-p sampling influences token selection and does not implement a business action.

Question 35 : Section 2: Google Cloud's gen AI offerings

Case study: A healthcare provider needs a two-part recommendation for a supply-chain operations assistant. Which paired option best helps the team allow a team to create a customized model with less manual model-development effort?

- A. Use only Cloud Storage; then Assume a larger model automatically fixes the underlying issue.
- B. Use Agent Platform AutoML; then Confirm tool permissions, integration behavior, and governance before rollout.**
- C. Use only IAM roles; then Rely on a one-time demo and do not monitor the production workflow.
- D. Use Gemini for Google Workspace; then Skip validation because the first response looks reasonable.

Answer: B

AutoML reduces the effort required to build customized ML solutions. The most tempting incorrect option fails because Cloud Storage is useful for data but does not create a customized model by itself.

Unlock All 1100+ Questions

Get the complete Q&A package with detailed explanations, topic analytics, and exam-accurate practice.

From €29.00

Visit: <https://cert-pass.com/exams/google-generative-ai-leader>

CERT-PASS

cert-pass.com

© 2026 Cert-Pass. This material is for personal use only. Do not distribute.