

CERT-PASS

Google Google Cloud Associate Cloud Engineer

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: Setting up a cloud solution environment

A site reliability engineer is setting up project payments-core. A deployment command fails because a newly created project has not activated the service needed to create Cloud Run resources. Which option best meets the requirement?

- A. Enable Cloud Billing export to BigQuery for detailed analysis.
- B. Set the default project in the active gcloud configuration.
- C. Verify that the selected service and feature are available in the target region.
- D. Enable the required service API in the project before retrying the deployment.**

Answer: D

Google Cloud services must have their APIs enabled in the target project before their resources can be provisioned. Why the best distractor fails: Billing configuration controls cost tracking and payment relationships; it does not implement the requested runtime or security behavior.

Question 2 : Section 2: Planning and implementing a cloud solution

During an initial deployment in europe-west4, a production VM must survive host maintenance when possible and use the correct SSH access approach defined by the organization. What is the most appropriate action?

- A. Configure OS Login and grant the appropriate IAM roles.
- B. Create a custom machine type when predefined CPU-to-memory ratios are inefficient.
- C. Launch the Compute Engine instance with the required availability policy and approved SSH or OS Login configuration.**
- D. Use a zonal Persistent Disk appropriate for the workload.

Answer: C

VM creation includes availability and access settings that should be selected deliberately. Why the best distractor fails: The option is related, but it does not match the scenario's required service type, control boundary, or operational goal.

Question 3 : Section 3: Ensuring the successful operation of a cloud solution

A new workload is being prepared for a logistics provider. An administrator needs to connect securely to a Compute Engine VM for troubleshooting. Which solution should the team implement?

A. Use an approved Compute Engine remote-access method, such as OS Login-backed SSH, consistent with the organization's policy.

B. Use a regional Persistent Disk when cross-zone block-storage replication is required.

C. Use Spot VMs only for restart-tolerant interruptible workloads.

D. Use Filestore when the requirement is a managed shared NFS-style file system.

Answer: A

Remote access should use the organization's approved Compute Engine login approach, commonly integrated with IAM through OS Login. Why the best distractor fails: The option is related, but it does not match the scenario's required service type, control boundary, or operational goal.

Question 4 : Section 4: Configuring access and security

The team responsible for analytics-prod has a straightforward requirement. A service account can currently administer an entire project but only needs to publish messages to one Pub/Sub topic. Which option is the best fit?

A. Use Bigtable for high-throughput, low-latency wide-column access patterns.

B. Use Google Cloud Managed Lustre for throughput-intensive parallel POSIX file access.

C. Use Storage Transfer Service for a managed large-scale dataset transfer into Cloud Storage.

D. Replace the broad role with the narrowest suitable Pub/Sub role at the appropriate scope.

Answer: D

Least privilege requires reducing both the permission set and resource scope to what the workload needs. Why the best distractor fails: The option is related, but it does not match the scenario's required service type, control boundary, or operational goal.

Question 5 : Section 1: Setting up a cloud solution environment

Before the first production release, finance wants notifications when a project's forecasted spend crosses a threshold, but does not want workloads automatically stopped. What should be done first?

A. Create and ingest a custom metric for the application-specific signal.

B. Review Active Assist recommendations and validate them before applying changes.

C. Review service quotas and request an increase for the target region.

D. Create a billing budget with appropriate alert thresholds and notification channels.

Answer: D

Billing budgets and alerts notify stakeholders about spend; they do not automatically cap or stop consumption by default. Why the best distractor fails: The option is related, but it does not match the scenario's required service type, control boundary, or operational goal.

Question 6 : Section 2: Planning and implementing a cloud solution

A manufacturing group is preparing project security-ops. Administrators need centrally managed SSH access to Compute Engine instances using IAM rather than manually distributing SSH keys. What should the engineer do?

- A. Use Local SSD only when ephemeral high-performance local storage is acceptable.
- B. Create an instance template and deploy an autoscaled managed instance group.
- C. Configure OS Login and grant the appropriate IAM roles.**
- D. Evaluate an appropriate Hyperdisk type for the required block-storage performance profile.

Answer: C

OS Login integrates VM SSH access with IAM and reduces unmanaged key distribution. Why the best distractor fails: The option is related, but it does not match the scenario's required service type, control boundary, or operational goal.

Question 7 : Section 3: Ensuring the successful operation of a cloud solution

A network engineer is setting up project shared-platform. Operations needs to view currently running Compute Engine instances in a project before a maintenance window. Which option best meets the requirement?

- A. Set the default project in the active gcloud configuration.
- B. Create separate billing accounts and link each project's billing account appropriately.
- C. Deploy Google Cloud Managed Service for Prometheus for Prometheus-compatible monitoring.
- D. List or view the Compute Engine instances in the target project and region or zone context.**

Answer: D

Compute Engine inventory should be checked directly through the console, gcloud, or API. Why the best distractor fails: The option is related, but it does not match the scenario's required service type, control boundary, or operational goal.

Question 8 : Section 4: Configuring access and security

During an initial deployment in europe-west4, a developer needs read access to objects in one bucket and no permissions elsewhere. What is the most appropriate action?

- A. Apply an organization policy constraint at the narrowest hierarchy node that covers the target resources.
- B. Grant an appropriate Cloud Storage read role on the specific bucket.**
- C. Configure Workload Identity Federation for an external non-human workload.
- D. Use a Cloud Identity group for workforce access that changes frequently.

Answer: B

A bucket-scoped predefined read role follows least privilege. Why the best distractor fails: The option is related, but it does not match the scenario's required service type, control boundary, or operational goal.

Question 9 : Section 1: Setting up a cloud solution environment

A new workload is being prepared for an e-commerce marketplace. Analysts need detailed billing data in a queryable dataset for cost reporting and chargeback. Which solution should the team implement?

- A. Link the project to the correct billing account.
- B. Review Active Assist recommendations and validate them before applying changes.
- C. Enable Cloud Billing export to BigQuery.**
- D. Create and ingest a custom metric for the application-specific signal.

Answer: C

Billing export to BigQuery provides detailed cost and usage data for analysis. Why the best distractor fails: Billing configuration controls cost tracking and payment relationships; it does not implement the requested runtime or security behavior.

Question 10 : Section 2: Planning and implementing a cloud solution

The team responsible for customer-api-prod has a straightforward requirement. A fault-tolerant batch workload can restart when capacity is reclaimed and the team wants lower compute cost. Which option is the best fit?

- A. Use a Standard GKE cluster when the team needs direct node-pool control.
- B. Run the batch workers on Spot VMs.**
- C. Use a zonal Persistent Disk appropriate for the workload.
- D. Configure VM Manager for centralized VM operating-system management.

Answer: B

Spot VMs are appropriate for interruptible workloads that can tolerate VM termination. Why the best distractor fails: GKE is appropriate when Kubernetes behavior is required, but it adds or assumes an orchestration model that does not fit this scenario.

Question 11 : Section 3: Ensuring the successful operation of a cloud solution

Before the first production release, a VM data disk needs regular point-in-time backups that can be used to create replacement disks. What should be done first?

- A. Evaluate an appropriate Hyperdisk type for the required block-storage performance profile.
- B. Evaluate Google Cloud NetApp Volumes for enterprise NFS or SMB file workloads.
- C. Use a zonal Persistent Disk appropriate for the workload.
- D. Create and schedule Persistent Disk snapshots.**

Answer: D

Persistent Disk snapshots provide point-in-time disk backups and can be scheduled. Why the best distractor fails: The option is related, but it does not match the scenario's required service type, control boundary, or operational goal.

Question 12 : Section 4: Configuring access and security

A regional retailer is preparing project inventory-prod. A folder-level IAM role is granted to an operations group and a project is moved beneath that folder. What should the engineer do?

- A. Create a dedicated service account, grant minimum permissions, and attach it to the runtime.
- B. Expect the project to inherit the applicable folder-level IAM binding unless a policy mechanism changes the effective access.**
- C. Grant an appropriate predefined role at the narrowest suitable resource scope.
- D. Use service account impersonation for controlled short-lived access.

Answer: B

IAM allow policies inherit down the resource hierarchy from organization to folder to project and eligible child resources. Why the best distractor fails: The option is related, but it does not match the scenario's required service type, control boundary, or operational goal.

Question 13 : Section 1: Setting up a cloud solution environment

A platform engineer is setting up project developer-platform. A newly created project cannot use billable services because it is not associated with a payment account. Which option best meets the requirement?

- A. Link the project to the correct billing account.**
- B. Deploy the Ops Agent to collect supported VM logs and metrics.
- C. Create a Cloud Monitoring alerting policy based on the relevant metric and threshold.
- D. Use Cloud Hub to review active events and application health data.

Answer: A

A project must be linked to a billing account to use billable Google Cloud services. Why the best distractor fails: A VM-based approach adds direct infrastructure responsibility or does not provide the requested managed capability.

Question 14 : Section 2: Planning and implementing a cloud solution

During an initial deployment in europe-west4, a VM workload needs an unusual CPU-to-memory ratio that predefined machine types do not fit efficiently. What is the most appropriate action?

- A. Use a regional GKE cluster when the control plane must be resilient across zones.
- B. Configure kubectl credentials for the target cluster before administering Kubernetes resources.
- C. Create a custom machine type sized for the workload.**
- D. Use a Standard GKE cluster when the team needs direct node-pool control.

Answer: C

Custom machine types let you match VM CPU and memory more closely to workload requirements. Why the best distractor fails: GKE is appropriate when Kubernetes behavior is required, but it adds or assumes an orchestration model that does not fit this scenario.

Question 15 : Section 3: Ensuring the successful operation of a cloud solution

A new workload is being prepared for a logistics provider. A Kubernetes workload needs stable network identities and ordered handling for replicated stateful application instances. Which solution should the team implement?

- A. Create a GKE Autopilot cluster to reduce node-management responsibility.
- B. Create a Compute Engine managed instance group only when the workload is VM-based.
- C. Use a StatefulSet for replicated stateful Pods that need stable identities.
- D. Use a Kubernetes StatefulSet.**

Answer: D

StatefulSets are Kubernetes controllers designed for stateful applications needing stable identities and ordered behavior. Why the best distractor fails: GKE is appropriate when Kubernetes behavior is required, but it adds or assumes an orchestration model that does not fit this scenario.

Question 16 : Section 4: Configuring access and security

The team responsible for security-ops has a straightforward requirement. A team asks for Editor because it is easy, but the workload needs only a narrow set of actions. Which option is the best fit?

- A. Use a Cloud Identity group for workforce access that changes frequently.
- B. Prefer an appropriate predefined or custom least-privilege role instead of a broad basic role.**
- C. Grant an appropriate predefined role at the narrowest suitable resource scope.
- D. Apply an organization policy constraint at the narrowest hierarchy node that covers the target resources.

Answer: B

Basic roles are broad; predefined or custom roles are usually better for least privilege. Why the best distractor fails: The option is related, but it does not match the scenario's required service type, control boundary, or operational goal.

Question 17 : Section 1: Setting up a cloud solution environment

Before the first production release, a new production project must support cost governance and a service deployment that currently fails because the relevant API is disabled. Which solution should be selected?. What should be done first?

- A. Deploy Google Cloud Managed Service for Prometheus for Prometheus-compatible monitoring.
- B. Enable the required service API and configure a billing budget with alerts.**
- C. Create a folder under the organization and place the relevant projects in that folder.
- D. Configure Cloud Asset Inventory to analyze resources across the hierarchy.

Answer: B

API enablement fixes provisioning, while a budget and alerts establish cost governance. Why the best distractor fails: Observability can expose a condition, but it does not directly implement the requested configuration change.

Question 18 : Section 2: Planning and implementing a cloud solution

A manufacturing group is preparing project edge-services. Analysts need a serverless data warehouse for large-scale SQL analytics over historical data. What should the engineer do?

- A. Use Dataflow for managed batch and streaming data-processing pipelines.
- B. Use Cloud Storage Standard for frequently accessed object data.
- C. Use BigQuery.**
- D. Use Cloud SQL for a managed traditional relational database workload.

Answer: C

BigQuery is designed for large-scale analytical queries and serverless data warehousing. Why the best distractor fails: The option is related, but it does not match the scenario's required service type, control boundary, or operational goal.

Question 19 : Section 3: Ensuring the successful operation of a cloud solution

A application team is setting up project ml-training. A stateless GKE Deployment should increase or decrease its number of Pods based on workload metrics. Which option best meets the requirement?

- A. Configure traffic splitting so that a controlled percentage reaches the new revision.
- B. Use a private GKE cluster when nodes must avoid public IP addresses.
- C. Set Cloud Run autoscaling limits, including a maximum instance value where downstream protection is required.
- D. Configure a Horizontal Pod Autoscaler.**

Answer: D

Horizontal Pod Autoscaling changes the number of Pod replicas based on metrics. Why the best distractor fails: The option is related, but it does not match the scenario's required service type, control boundary, or operational goal.

Question 20 : Section 4: Configuring access and security

During an initial deployment in europe-west4, no predefined IAM role matches a tightly scoped internal support function. What is the most appropriate action?

- A. Create a custom IAM role containing only the required supported permissions.**
- B. Grant an appropriate predefined role at the narrowest suitable resource scope.
- C. Create separate billing accounts and link each project's billing account appropriately.
- D. Review inherited IAM bindings from folders and the organization as well as direct project bindings.

Answer: A

Custom roles are appropriate when predefined roles do not fit and the required permissions are supported. Why the best distractor fails: The option is related, but it does not match the scenario's required service type, control boundary, or operational goal.

Question 21 : Section 1: Setting up a cloud solution environment

A new workload is being prepared for an e-commerce marketplace. A growing company wants one policy boundary for all projects owned by its finance department while keeping finance projects separate from engineering projects. Which solution should the team implement?

- A. Create a Finance folder under the organization and place the finance projects in that folder.
- B. Configure Cloud Asset Inventory to analyze resources across the hierarchy.
- C. Set the default project in the active gcloud configuration.
- D. Provision the required Google Cloud Observability products before production deployment.

Answer: A

Folders provide a hierarchy level where organization policies and IAM bindings can be applied consistently to related projects. Why the best distractor fails: The option is related, but it does not match the scenario's required service type, control boundary, or operational goal.

Question 22 : Section 2: Planning and implementing a cloud solution

The team responsible for inventory-prod has a straightforward requirement. A mobile application needs a scalable document database with flexible document structures. Which option is the best fit?

- A. Use Bigtable for high-throughput, low-latency wide-column access patterns.
- B. Use the Google Cloud CLI to upload local objects to Cloud Storage.
- C. Use Firestore.
- D. Use Pub/Sub for asynchronous event delivery between decoupled producers and consumers.

Answer: C

Firestore is a managed document database suitable for application data with document-oriented access patterns. Why the best distractor fails: The option is related, but it does not match the scenario's required service type, control boundary, or operational goal.

Question 23 : Section 3: Ensuring the successful operation of a cloud solution

Before the first production release, a data scientist needs a managed notebook environment associated with the renamed Gemini Enterprise Agent Platform Workbench offering. What should be done first?

- A. Create a BigQuery load job to import files already stored in Cloud Storage.
- B. Use Cloud SQL for a managed traditional relational database workload.
- C. Evaluate Google Cloud NetApp Volumes for enterprise NFS or SMB file workloads.
- D. Use Gemini Enterprise Agent Platform Workbench, formerly Vertex AI Workbench, or an appropriate BigQuery notebook option for the task.

Answer: D

The updated guide includes managing notebooks in Gemini Enterprise Agent Platform Workbench and BigQuery. Why the best distractor fails: BigQuery is an analytical warehouse, so it should not replace an operational database or

unrelated runtime service.

Question 24 : Section 4: Configuring access and security

A regional retailer is preparing project analytics-prod. A Cloud Run service needs its own workload identity with only the permissions required to call a backend API. What should the engineer do?

- A. Create a dedicated service account, grant the minimum required role, and attach it to the Cloud Run service.
- B. Grant the role on the department folder when it should apply to that department's projects only.
- C. Use the platform's attached service account identity instead of distributing a long-lived key.
- D. Configure Workload Identity Federation for an external non-human workload.

Answer: A

A dedicated least-privilege service account is the correct workload identity pattern. Why the best distractor fails: The option is related, but it does not match the scenario's required service type, control boundary, or operational goal.

Question 25 : Section 1: Setting up a cloud solution environment

A site reliability engineer is setting up project data-foundation. Security requires that new workloads cannot receive public external IP addresses unless an approved exception exists. Which option best meets the requirement?

- A. Apply the appropriate organization policy constraint at the highest suitable hierarchy level.
- B. Set the default project in the active gcloud configuration.
- C. Review service quotas and request an increase for the target region.
- D. Create a billing budget with alert thresholds and notification channels.

Answer: A

Organization policies enforce guardrails across resources below the selected organization, folder, or project node. Why the best distractor fails: The option is related, but it does not match the scenario's required service type, control boundary, or operational goal.

Question 26 : Section 2: Planning and implementing a cloud solution

During an initial deployment in europe-west4, multiple services need asynchronous message delivery and event ingestion without tightly coupling producers and consumers. What is the most appropriate action?

- A. Use BigQuery for serverless analytical warehousing.
- B. Use Pub/Sub.
- C. Use Spanner for horizontally scalable relational data with strong consistency.
- D. Use Persistent Disk when a Compute Engine VM needs block storage.

Answer: B

Pub/Sub is the managed messaging service for asynchronous event ingestion and delivery. Why the best distractor fails:

BigQuery is an analytical warehouse, so it should not replace an operational database or unrelated runtime service.

Question 27 : Section 3: Ensuring the successful operation of a cloud solution

A new workload is being prepared for a logistics provider. A team must allow a workload to read objects from one bucket without granting broad project-wide permissions. Which solution should the team implement?

- A. Grant a narrowly scoped Cloud Storage IAM role on the bucket to the workload identity.**
- B. Create a dedicated service account, grant minimum permissions, and attach it to the runtime.
- C. Configure Cloud Asset Inventory to analyze resources across the hierarchy.
- D. Create separate billing accounts and link each project's billing account appropriately.

Answer: A

Bucket-scoped IAM follows least privilege for object access. Why the best distractor fails: The option is related, but it does not match the scenario's required service type, control boundary, or operational goal.

Question 28 : Section 4: Configuring access and security

The team responsible for edge-services has a straightforward requirement. A Compute Engine application needs to authenticate to Google Cloud APIs using a specific workload identity. Which option is the best fit?

- A. Grant the role on the department folder when it should apply to that department's projects only.
- B. Configure Workload Identity Federation for an external non-human workload.
- C. Use the platform's attached service account identity instead of distributing a long-lived key.
- D. Attach the appropriate service account to the VM and grant only the permissions it needs.**

Answer: D

The VM uses the attached service account as its workload identity for Google Cloud API access. Why the best distractor fails: The option is related, but it does not match the scenario's required service type, control boundary, or operational goal.

Question 29 : Section 1: Setting up a cloud solution environment

Before the first production release, a project team must deploy Compute Engine instances but must not modify organization-wide policy settings. What should be done first?

- A. Grant the role on the department folder when it should apply to that department's projects only.
- B. Create separate billing accounts and link each project's billing account appropriately.
- C. Review inherited IAM bindings from folders and the organization as well as direct project bindings.
- D. Grant the team an appropriate predefined Compute Engine role on the project.**

Answer: D

A project-level predefined role follows least privilege and limits the scope to the required project. Why the best distractor

fails: The option is related, but it does not match the scenario's required service type, control boundary, or operational goal.

Question 30 : Section 2: Planning and implementing a cloud solution

A manufacturing group is preparing project customer-api-prod. Compliance archives must be retained for years and are rarely read, so retrieval latency and access charges are acceptable tradeoffs for lower storage cost. What should the engineer do?

- A. Use the Google Cloud CLI to upload local objects to Cloud Storage.
- B. Use Spanner for horizontally scalable relational data with strong consistency.
- C. Store the objects in the Archive storage class.**
- D. Use Cloud Storage Archive for rarely accessed long-retention object data.

Answer: C

Archive storage is intended for rarely accessed objects retained for long periods. Why the best distractor fails: The option is related, but it does not match the scenario's required service type, control boundary, or operational goal.

Question 31 : Section 3: Ensuring the successful operation of a cloud solution

A network engineer is setting up project payments-core. A long-running analytical operation may have failed and the engineer needs to inspect its execution state. Which option best meets the requirement?

- A. Create a BigQuery load job to import files already stored in Cloud Storage.
- B. Use AlloyDB for a demanding PostgreSQL-compatible relational workload.
- C. Use BigQuery for serverless analytical warehousing.
- D. Review the BigQuery job status and relevant error details.**

Answer: D

BigQuery exposes job status and errors for analytical operations. Why the best distractor fails: BigQuery is an analytical warehouse, so it should not replace an operational database or unrelated runtime service.

Question 32 : Section 4: Configuring access and security

During an initial deployment in europe-west4, an administrator must allow a user to operate as a service account without granting unrelated project administration rights. What is the most appropriate action?

- A. Grant the appropriate service-account impersonation permission or role on the service account.**
- B. Review inherited IAM bindings from folders and the organization as well as direct project bindings.
- C. Use Workload Identity Federation for GKE for keyless Pod access to Google Cloud APIs.
- D. Grant an appropriate predefined role at the narrowest suitable resource scope.

Answer: A

Service account impersonation grants controlled use of the identity without broad project administration. Why the best

distractor fails: The option is related, but it does not match the scenario's required service type, control boundary, or operational goal.

Question 33 : Section 1: Setting up a cloud solution environment

A new workload is being prepared for an e-commerce marketplace. The operations department changes frequently and access must be managed without editing dozens of individual IAM bindings. Which solution should the team implement?

- A. Set the default project in the active gcloud configuration.
- B. Configure Cloud Asset Inventory to analyze resources across the hierarchy.
- C. Use a Cloud Identity group for workforce access that changes frequently.
- D. Grant IAM roles to a Cloud Identity group and manage membership in the group.**

Answer: D

Group-based access centralizes membership management while keeping IAM bindings stable. Why the best distractor fails: The option is related, but it does not match the scenario's required service type, control boundary, or operational goal.

Question 34 : Section 2: Planning and implementing a cloud solution

The team responsible for analytics-prod has a straightforward requirement. Several Compute Engine VMs need a managed shared file system with familiar NFS-style access. Which option is the best fit?

- A. Use Filestore.**
- B. Use Pub/Sub when the requirement is asynchronous message delivery rather than bulk file loading.
- C. Use Spanner for horizontally scalable relational data with strong consistency.
- D. Use Cloud SQL for a managed traditional relational database workload.

Answer: A

Filestore provides managed file storage suitable for shared file access patterns. Why the best distractor fails: The option is related, but it does not match the scenario's required service type, control boundary, or operational goal.

Question 35 : Section 3: Ensuring the successful operation of a cloud solution

Before the first production release, an application needs authoritative DNS records for a domain hosted using Google Cloud. What should be done first?

- A. Configure Shared VPC with a host project and attach service projects.
- B. Create a narrowly scoped ingress firewall rule for the required source, target, protocol, and port.
- C. Configure Cloud NAT for outbound connectivity from private resources without external IP addresses.
- D. Create or update a Cloud DNS managed zone and the required DNS records.**

Answer: D

Cloud DNS managed zones host DNS records. Why the best distractor fails: The option is related, but it does not match the scenario's required service type, control boundary, or operational goal.

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-google-cloud-associate-cloud-engineer>

CERT-PASS

cert-pass.com

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