

# Associate Cloud Engineer

## Certification exam guide

An Associate Cloud Engineer deploys and secures applications, services, and infrastructure, monitors the operations of multiple projects, and maintains enterprise solutions to meet target performance metrics. This individual has experience working with public clouds or on-premises solutions. They are able to perform common platform-based tasks, supported by AI tooling, to maintain and scale one or more deployed solutions that leverage Google-managed or self-managed services on Google Cloud.

### **Section 1: Setting up a cloud solution environment (~20% of the exam)**

1.1 Setting up cloud projects and accounts. Considerations include:

- Creating a resource hierarchy
- Applying organizational policies to the resource hierarchy
- Granting members Identity and Access Management (IAM) roles within a project
- Managing users and groups in Cloud Identity (manually and automated)
- Enabling APIs within projects
- Provisioning and setting up products in Google Cloud Observability
- Assessing quotas and requesting increases
- Setting up standalone organizations
- Setting up cloud networking
- Verifying product availability across geographical locations (e.g., regions, zones)
- Configuring Cloud Asset Inventory and using Gemini Cloud Assist to analyze resources
- Configuring Workforce Identity Federation

1.2 Managing billing configuration. Considerations include:

- Creating one or more billing accounts
- Linking projects to a billing account
- Establishing billing budgets and alerts
- Setting up billing exports

## Section 2: Planning and implementing a cloud solution (~30% of the exam)

2.1 Planning and implementing compute resources. Considerations include:

- Selecting appropriate compute choices for a given workload (e.g., Compute Engine, Google Kubernetes Engine [GKE], Cloud Run, Cloud Run functions, Agent Runtime on Gemini Enterprise Agent Platform [formerly Vertex AI Agent Engine])
- Launching a compute instance (e.g., availability policy, SSH keys)
- Choosing the appropriate storage for Compute Engine (e.g., zonal Persistent Disk, regional Persistent Disk, Google Cloud Hyperdisk)
- Creating an autoscaled managed instance group by using an instance template
- Configuring OS Login
- Configuring VM Manager
- Using Spot VM instances and custom machine types
- Installing and configuring the command-line interface (CLI) for Kubernetes (kubectl)
- Deploying a GKE cluster with different configurations (e.g., GKE Autopilot, regional clusters, private clusters)
- Deploying a containerized application to GKE
- Deploying an application to serverless compute platforms, including for the processing of Google Cloud events (e.g., Pub/Sub events, Cloud Storage object change notification events, Eventarc)
- Identifying whether to use GPUs or TPUs

2.2 Planning and implementing storage and data solutions. Considerations include:

- Choosing and deploying data products (e.g., Cloud SQL, BigQuery, Firestore, Spanner, Bigtable, AlloyDB, Dataflow, Pub/Sub, Google Cloud Managed Service for Apache Kafka, Memorystore)
- Choosing and deploying storage products (e.g., Cloud Storage, Filestore, Google Cloud NetApp Volumes, Google Cloud Managed Lustre) and Cloud Storage options (e.g., Standard, Nearline, Coldline, Archive)
- Loading data (e.g., command-line upload, load data from Cloud Storage, Storage Transfer Service)
- Maintaining multi-region redundancy across data solutions

# Google Cloud

2.3 Planning and implementing networking resources. Considerations include:

- Creating a VPC with subnets (e.g., custom mode VPC, Shared VPC, VPC Network Peering)
- Creating and applying VPC firewalls rules and Cloud Next Generation Firewall (Cloud NGFW) policies with ingress and egress rules and attributes (e.g., action, source, destination, targets, protocols, ports)
- Using Tags (e.g., secure Tags) and service accounts in Cloud NGFW policy rules
- Establishing network connectivity (e.g., Cloud VPN, VPC Network Peering, Cloud Interconnect)
- Choosing and deploying load balancers
- Differentiating Network Service Tiers

2.4 Planning and implementing resources using tooling. Considerations include:

- Infrastructure as Code tooling (e.g., Fabric FAST, Config Connector, Terraform, Helm)
- AI-assisted planning and implementation (e.g., Gemini CLI, Google Antigravity, Gemini Cloud Assist, Application Design Center)

## **Section 3: Ensuring the successful operation of a cloud solution (~30% of the exam)**

3.1 Managing compute resources. Considerations include:

- Remotely connecting to a Compute Engine instance
- Viewing current running Compute Engine instances
- Working with snapshots and images (e.g., create, view, and delete images or snapshots; schedule a snapshot)
- Viewing current running GKE cluster inventory (e.g., nodes, Pods, Services)
- Configuring GKE to access Artifact Registry
- Working with GKE node pools (e.g., add, edit, or remove a node pool; autoscaling node pool)
- Working with Kubernetes resources (e.g., Pods, Services, StatefulSets)
- Managing horizontal and vertical Pod autoscaling configurations
- Managing GKE Autopilot Pod resource requests
- Deploying new versions of a Cloud Run application
- Adjusting application traffic splitting parameters (e.g., Cloud Run, Cloud Run functions, GKE)
- Configuring autoscaling for a Cloud Run application

# Google Cloud

- Attaching GPUs and TPUs
- Deploying an agent to Agent Runtime on Gemini Enterprise Agent Platform (formerly Vertex AI Agent Engine)
- Managing notebooks in Gemini Enterprise Agent Platform Workbench (formerly Vertex AI Workbench) and BigQuery
- Managing developer environments (e.g., Cloud Workstations)

## 3.2 Managing storage and data solutions. Considerations include:

- Managing and securing objects in Cloud Storage buckets
- Setting object lifecycle management policies for Cloud Storage buckets
- Executing queries to retrieve data from data instances (e.g., Cloud SQL, BigQuery, Bigtable, Spanner, Firestore, AlloyDB)
- Estimating costs of data storage resources
- Backing up and restoring database instances (e.g., Cloud SQL, Firestore, Spanner, AlloyDB, Bigtable)
- Reviewing job status (e.g., Dataflow, BigQuery)
- Using Database Center to manage the Google Cloud database fleet
- Configuring customer-managed encryption keys (CMEK)

## 3.3 Managing networking resources. Considerations include:

- Resizing a subnet's IPv4 address range
- Reserving static external or internal IP addresses
- Adding custom static routes in a VPC
- Using Cloud DNS and Cloud NAT
- Managing VPC firewall rules and Cloud NGFW policies

## 3.4 Monitoring and logging. Considerations include:

- Creating Cloud Monitoring alerts based on resource metrics
- Creating and ingesting Cloud Monitoring custom metrics (e.g., from applications or logs)
- Configuring audit logs (e.g., VPC Flow Logs, audit logs, firewall logs)
- Exporting logs to external systems (e.g., on-premises, BigQuery)
- Configuring log buckets, log analytics, and log routers
- Viewing and filtering logs in Cloud Logging
- Viewing specific log message details in Cloud Logging
- Using cloud diagnostic tools (e.g., Cloud Trace, Cloud Profiler, Query Insights, index advisor) to investigate an application issue
- Viewing the Personalized Service Health dashboard

# Google Cloud

- Configuring and deploying Ops Agent
- Deploying Google Cloud Managed Service for Prometheus
- Using Gemini Cloud Assist for Cloud Monitoring
- Using Active Assist to optimize resource utilization
- Using Cloud Hub to monitor active events and application health data

## **Section 4: Configuring access and security (~20% of the exam)**

### 4.1 Managing IAM. Considerations include:

- Viewing and creating IAM policies
- Attaching roles and policy inheritance in the Organization hierarchy
- Managing the various role types and defining custom IAM roles

### 4.2 Managing service accounts. Considerations include:

- Creating service accounts, including Google-managed service accounts
- Using service accounts in IAM policies with minimum permissions
- Assigning service accounts to resources
- Managing IAM permissions of a service account
- Managing service account impersonation
- Creating and managing short-lived service account credentials
- Using a Google Cloud service account with a GKE application
- Provisioning Workload Identity Federation