

Associate Cloud Engineer Renewal

This is the new version of the Associate Cloud Engineer renewal exam guide. This renewal exam guide will be live on July 31. If you plan to take the Associate Cloud Engineer renewal exam in English on or after July 31, review this exam guide. **If you plan to take the Associate Cloud Engineer renewal exam before July 31, review the [current renewal exam guide](#).**

This renewal exam guide is a streamlined version of the standard Associate Cloud Engineer exam guide, focusing on key areas critical to the job role. The renewal exam is available to certified individuals who want to renew their active Associate Cloud Engineer certification and are within their renewal eligibility period.

Section and objective numbers in this renewal exam guide are aligned to the standard Associate Cloud Engineer exam guide. Omitted objectives are not assessed in the renewal exam.

Certification exam guide

An Associate Cloud Engineer deploys and secures applications, services, and infrastructure, monitors operations of multiple projects, and maintains enterprise solutions to ensure that they 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 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 (not included in renewal exam)

Section 2: Planning and implementing a cloud solution (~40% 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, Knative serving)
- 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)

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) 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

2.3 Planning and implementing networking resources. Considerations include:

- Creating a VPC with subnets (e.g., custom mode VPC, Shared VPC)
- Creating and applying 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 through infrastructure as code. Considerations include:

- Infrastructure as code tooling (e.g., Fabric FAST, Config Connector, Terraform, Helm)
- Planning and executing infrastructure as code deployments, including versioning, state management, and updates

Section 3: Ensuring successful operation of a cloud solution (~40% 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

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

3.3 Managing networking resources. Considerations include:

- Adding a subnet to an existing VPC
- Expanding a subnet to have more IP addresses
- Reserving static external or internal IP addresses
- Adding custom static routes in a VPC
- Working with Cloud DNS and Cloud NAT

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)
- 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 diagnostics to research an application issue (e.g., Cloud Trace, Cloud Profiler, Query Insights, index advisor)
- Viewing the Personalized Service Health dashboard
- Configuring and deploying Ops Agent
- Deploying Google Cloud Managed Service for Prometheus
- Configuring audit logs
- Using Gemini Cloud Assist for Cloud Monitoring
- Using Active Assist to optimize resource utilization

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

4.1 Managing Identity and Access Management (IAM). Considerations include:

- Viewing and creating IAM policies
- Managing the various role types and defining custom IAM roles (e.g., basic, predefined, custom)

4.2 Managing service accounts. Considerations include:

- Creating 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