CodeCruise

SAIL SMOOTH IN TECH OCEAN

A74, TechnoPark, Andheri, Mumbai. Phone: +91 80809 75897 | +91 70345 62050 Email: ask@codecruise.in

Google Cloud Engineer

The course equips participants with the skills required to effectively deploy, manage, and secure cloud solutions using Google Cloud Platform. Covering setting up environments, planning and configuring solutions, deployment, operation, and security measures, ensuring proficiency in managing enterprise cloud solutions.

An Associate Cloud Engineer deploys and secures applications 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 and on-premises solutions. They are able to use the Google Cloud console and the command-line interface 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.

Recommended Candidate:

- Has understanding fundamentals of Google Cloud
- Has experience deploying cloud applications and monitoring operations
- Has experience managing cloud enterprise solutions

Course Plan

Section 1: Setting up a cloud solution environment
Section 2: Planning and configuring a cloud solution
Section 3: Deploying and implementing a cloud solution
Section 4: Ensuring successful operation of a cloud solution
Section 5: Configuring access and security

Tech Stack To Be Covered

Section 1: Setting up a cloud solution environment

1.1 Setting up cloud projects and accounts

- Creating a resource hierarchy:
- Organizing resources (folders, projects) to reflect the structure of the organization.
- Setting up an organization node.
- Applying organizational policies to the resource hierarchy:
- Implementing policies for security, compliance, and operational governance.
- Granting members IAM roles within a project:
- Assigning predefined and custom roles to users, groups, and service accounts.
- Managing users and groups in Cloud Identity:
- Manually adding, modifying, and removing users and groups.
- Automating user and group management using scripts or tools.
- Enabling APIs within projects:
- Activating necessary Google Cloud APIs for project functionality.
- Provisioning and setting up products in Google Cloud's operations suite:
- Configuring monitoring, logging, and alerting for cloud resources.

1.2 Managing billing configuration

- Creating one or more billing accounts:
- Setting up billing accounts to manage costs and track usage.
- Linking projects to a billing account:
- Associating projects with billing accounts for cost management.
- Establishing billing budgets and alerts:

- Defining budgets and setting alerts to monitor spending.
- Setting up billing exports:
- Exporting billing data to BigQuery, Pub/Sub, or Cloud Storage for analysis.

1.3 Installing and configuring the command line interface (CLI)

- Specifically the Cloud SDK (e.g., setting the default project):
- Installing and initializing the Cloud SDK.
- Configuring project settings and preferences.

Section 2: Planning and configuring a cloud solution

2.1 Planning and estimating Google Cloud product use using the Pricing Calculator

- Using the Pricing Calculator to estimate costs for various Google Cloud services.

2.2 Planning and configuring compute resources

- Selecting appropriate compute choices for a given workload:
- Deciding between Compute Engine, Google Kubernetes Engine (GKE), Cloud Run, Cloud Functions.
- Using preemptible VMs and custom machine types as appropriate:
- Utilizing cost saving options like preemptible VMs.

2.3 Planning and configuring data storage options

- Product choice:
- Selecting between Cloud SQL, BigQuery, Firestore, Cloud Spanner, Cloud Bigtable based on use case.
- Choosing storage options:
- Deciding between Zonal, Regional disks, Standard, Nearline, Coldline, Archive storage based on access needs and cost considerations.

2.4 Planning and configuring network resources

- Differentiating load balancing options:
- Understanding global vs regional load balancers.
- Identifying resource locations in a network for availability:
- Choosing appropriate regions and zones for deploying resources.
- Configuring Cloud DNS:

- Setting up domain name system configurations for applications.

Section 3: Deploying and implementing a cloud solution

3.1 Deploying and implementing Compute Engine resources

- Launching a compute instance:
- Using Google Cloud console and Cloud SDK (gcloud) for instance management.
- Creating an autoscaled managed instance group:
- Setting up instance templates and autoscaling policies.
- Generating/uploading a custom SSH key for instances:
- Managing SSH keys for secure access.
- Installing and configuring the Cloud Monitoring and Logging Agent:
- Ensuring proper monitoring and logging setup.
- Assessing compute quotas and requesting increases:
- Managing and requesting resource quotas.

3.2 Deploying and implementing Google Kubernetes Engine resources

- Installing and configuring the CLI for Kubernetes (kubectl):
- Setting up kubectl for cluster management.
- Deploying a Google Kubernetes Engine cluster:
- Configuring clusters with different setups (e.g., AutoPilot, regional, private clusters).
- Deploying a containerized application to GKE:
- Managing application deployments and updates.
- Configuring GKE monitoring and logging:
- Setting up observability for Kubernetes resources.

3.3 Deploying and implementing Cloud Run and Cloud Functions resources

- Deploying an application and updating scaling configuration, versions, and traffic splitting:
- Managing deployments and scaling configurations for serverless applications.
- Deploying an application that receives Google Cloud events:
- Integrating applications with Pub/Sub and Cloud Storage events.

3.4 Deploying and implementing data solutions

- Initializing data systems with products:
- Setting up Cloud SQL, Firestore, BigQuery, Cloud Spanner, Pub/Sub, Cloud Bigtable, Dataproc, Dataflow, Cloud Storage.
- Loading data: Using command line upload, API transfer, import/export, loading data from Cloud Storage, streaming data to Pub/Sub.

3.5 Deploying and implementing networking resources

- Creating a VPC with subnets:
- Configuring custom mode VPCs and shared VPCs.
- Launching a Compute Engine instance with custom network configuration:
- Setting up internal only IP addresses, Google private access, static external and private IP addresses, network tags.
- Creating ingress and egress firewall rules for a VPC:
- Managing firewall rules based on IP subnets, network tags, service accounts.
- Creating a VPN between a Google VPC and an external network using Cloud VPN:
- Setting up secure VPN connections.
- Creating a load balancer to distribute application network traffic:
- Implementing various load balancing options (e.g., Global HTTP(S), SSL Proxy, TCP Proxy, regional network load balancer, regional internal load balancer).

3.6 Deploying a solution using Cloud Marketplace

- Browsing the Cloud Marketplace catalog and viewing solution details:
- Exploring and selecting pre built solutions.
- Deploying a Cloud Marketplace solution:
- Implementing solutions from the Cloud Marketplace.

3.6 Deploying a solution using Cloud Marketplace

- Building infrastructure via Cloud Foundation Toolkit templates and implementing best practices:
- Using templates for consistent and repeatable deployments.
- Installing and configuring Config Connector in Google Kubernetes Engine:
- Managing resources using Config Connector for GKE.

Section 4: Ensuring successful operation of a cloud solution

3.6 Deploying a solution using Cloud Marketplace

- Managing a single VM instance:
- Starting, stopping, editing configuration, or deleting an instance.
- Remotely connecting to the instance:
- Using SSH for remote management.
- Attaching a GPU to a new instance and installing necessary dependencies:
- Configuring instances with GPU for high performance computing tasks.
- Viewing current running VM inventory:
- Monitoring instance IDs and details.
- Working with snapshots:
- Creating, viewing, and deleting snapshots.
- Working with images:
- Creating, viewing, and deleting images from VMs or snapshots.
- Working with instance groups:
- Setting autoscaling parameters, managing instance templates and instance groups.
- Working with management interfaces:
- Using Google Cloud console, Cloud Shell, and Cloud SDK for management.

4.2 Managing Google Kubernetes Engine resources

- Viewing current running cluster inventory:
- Monitoring nodes, pods, and services.
- Browsing Docker images and viewing their details in the Artifact Registry:
- Managing container images.
- Working with node pools:
- Adding, editing, or removing node pools.
- Working with pods:
- Adding, editing, or removing pods.
- Working with services:
- Managing Kubernetes services.
- Working with stateful applications:
- Configuring persistent volumes and stateful sets.
- Managing Horizontal and Vertical autoscaling configurations:
- Adjusting autoscaling settings.
- Working with management interfaces:
- Using Google Cloud console, Cloud Shell, Cloud SDK, kubectl for cluster management.

4.3 Managing Cloud Run resources

- Adjusting application traffic splitting parameters:
- Managing traffic distribution for applications.
- Setting scaling parameters for autoscaling instances:
- Configuring autoscaling settings.
- Determining whether to run Cloud Run (fully managed) or Cloud Run for Anthos:
- Choosing deployment options based on requirements.

4.4 Managing storage and database solutions

- Managing and securing objects in and between Cloud Storage buckets:
- Ensuring data security and integrity.
- Setting object life cycle management policies for Cloud Storage buckets:
- Automating data lifecycle management.
- Executing queries to retrieve data from data instances:
- Using SQL, BigQuery, Spanner, Datastore, Cloud Bigtable for data retrieval.
- Estimating costs of data storage resources:
- Calculating storage costs.
- Backing up and restoring database instances:
- Implementing backup and restore procedures for Cloud SQL, Datastore.
- Reviewing job status in Dataproc, Dataflow, or BigQuery:
- Monitoring data processing jobs.

4.5 Managing networking resources

- Adding a subnet to an existing VPC:
- Expanding network configurations.
- Expanding a subnet to have more IP addresses:
- Adjusting subnet configurations.
- Reserving static external or internal IP addresses:
- Managing IP address allocations.
- Working with CloudDNS, CloudNAT, Load Balancers, and firewall rules:
- Configuring and managing network services.

4.6 Monitoring and logging

- Creating Cloud Monitoring alerts based on resource metrics:
- Setting up alerts for resource monitoring.
- Creating and ingesting Cloud Monitoring custom metrics:
- Integrating custom metrics for detailed monitoring.
- Configuring log sinks to export logs to external systems:
- Exporting logs to on premises systems or BigQuery.
- Configuring log routers:
- Managing log routing for analysis.
- Viewing and filtering logs in Cloud Logging:
- Analyzing logs for troubleshooting.
- Viewing specific log message details in Cloud Logging:
- Investigating detailed log entries.
- Using cloud diagnostics to research an application issue:
- Leveraging Cloud Trace, Cloud Debug for issue resolution.
- Viewing Google Cloud status: Monitoring the overall health of Google Cloud services.

Section 5: Configuring access and security

5.1 Managing Identity and Access Management (IAM)

- Viewing IAM policies:
- Reviewing existing policies.
- Creating IAM policies:
- Defining new access policies.
- Managing the various role types and defining custom IAM roles:
- Handling primitive, predefined, and custom roles.

5.2 Managing service accounts

- Creating service accounts:
- Setting up service accounts for applications and services.
- Using service accounts in IAM policies with minimum permissions:
- Applying the principle of least privilege.
- Assigning service accounts to resources:
- Associating service accounts with specific resources.
- Managing IAM of a service account:

- Adjusting permissions for service accounts.
- Managing service account impersonation:
- Configuring impersonation settings.
- Creating and managing short lived service account credentials:
- Generating temporary credentials for secure access.

5.3 Viewing audit logs

- Monitoring and reviewing audit logs for security and compliance:
- Ensuring visibility into resource access and changes.