Application Development with Cloud Run

Course 4336 – 24 Hours

Overview

This course introduces you to fundamentals, practices, capabilities and tools applicable to modern cloud-native application development using Google Cloud Run. Through a combination of lectures, hands-on labs, and supplemental materials, you will learn how to design, implement, deploy, secure, manage, and scale applications on Google Cloud using Cloud Run.

On Completion, Delegates will be able to

- Gain detailed understanding of Cloud Run, Google Cloud's fully managed compute platform for deploying and scaling containerized applications quickly and securely
- Write and migrate code your way using your favorite languages (Go, Python, Java, Ruby, Node.js, and more).
- Secure service to service communication based on service identities and grant applications only the permissions they need.
- Learn how to build highly available applications with low end-user latency, globally.
- Learn how to connect to, and persist data in the managed database offerings on Google Cloud.
- Understand how abstracting away all infrastructure management creates a simple developer experience.

Who Should Attend

- Cloud developers, API developers, customers and partners

Prerequisites

- Familiarity with Linux commands and command line interface.
- Basic understanding of Google Cloud.
- Basic understanding of networking.
- Basic understanding of one or more programming languages like Go, Python, Java, Ruby, or Node.js.
- Basic understanding of shell scripts, YAML, JSON, HTTP, and TLS.

Course Contents

- Introducing Application Development with Cloud Run
- Understanding Cloud Run
  - Use a docker repository on Artifact Registry to store your images: Cloud Run only deploys from there.
  - Cloud Run uses autoscaling to handle all incoming requests
  - Pay for use pricing model
  - No background tasks: Container lifetime is only guaranteed while handling requests
  - There is no persistent storage: Store data downstream
  - Cloud Run is portable (containers and Knative)
- Building Container Images
  - The contents of a container image (deep dive)
  - There are two ways to build container images
    - Buildpacks (hands-off)
    - Docker (you're in control)
  - Cloud Run supports both source-based and a container image based workflow
  - The most important considerations of building a secure container image
- Building Container Images
  - Container lifecycle
    - Idle vs serving
    - Shutdown lifecycle hook
  - Cold starts
    - Min instances
  - Container readiness
  - The service resource and what it describes
  - Configuring memory limits and CPU allocation
  - Deploying a new revision
  - Traffic steering (tagging, gradual rollouts)
- Configuring Service Identity and Authorization
  - Cloud IAM - Service account, policy binding, roles, types of members, resource hierarchy (in practice)
    - Service accounts
    - Cloud Run IAM roles
  - Cloud Run
    - Default service account
    - Risks of using the default service account
- Serving Requests
  - Custom Domains
  - Global Load Balancer
    - URL Map
    - Frontend
    - Backend services
  - Benefits and drawbacks of GLB over custom domain
  - Types of GLB Backends
  - Multi-region load balancing
  - Multi-regional applications challenges

    - Cloud CDN
- Using Inbound and Outbound Access Control
    - Ingress settings
    - Cloud Armor
    - Using Cloud IAM to protect services
      - Understand how authenticated requests (IAM + OIDC tokens) work (builds on Module 5)
    - VPC, VPC Access Connector
    - Egress settings
- Persisting Data
    - Understanding why you need to store data externally when running a workload on Cloud Run
    - Connect with Cloud SQL from Cloud Run
      - Understand how it works (managed Cloud SQL Proxy)
    - Managing concurrency as a way to safeguard performance (understand why and when)
    - Connecting with Memorystore
    - VPC Connector
      - Challenges with scaling Memorystore (throughput)
    - Briefly introduce Cloud Storage, Firestore and Cloud Spanner, while reinforcing how the client libraries use the built-in service account to connect (Module 5 is prerequisite knowledge).
    - Multi-region data storage (and what Spanner and Firestore can do for you)
- Implementing Service-to-Service Communication
    - Understanding Cloud Pub/Sub
      - Understanding topics, push subscriptions
      - Idempotency (Handling retries and at-least-once invocation)
      - Event ID, design for resume, or use a lease
      - Handling undeliverable messages
    - How to asynchronously schedule a background task on a different service
    - Cloud Tasks, and when to choose it over Cloud Pub/Sub
    - Benefits of using Pub/Sub to pass messages over making sync RPC requests
    - Learn about services in Google Cloud with a built-in integration to push events to Pub/Sub (Cloud Build, Artifact Registry, Cloud Storage, IOT Core, BigQuery)
    - Cloud Scheduler to invoke services on a schedule.
    - CloudEvents
    - EventArc, and how to consume Audit logs
      - What to expect now, and how EventArc will develop over time
- Orchestrating and Automating Serverless Workflows
    - Conceptual overview of Cloud Workflows
    - Invoking and passing parameters
    - Understand steps and jumps
    - Defining, using and passing values with variables
    - Using the switch statement to add logic
    - Workflow visualization
    - Calling HTTPS endpoints
    - Calling an authenticated Cloud Run service
    - Example: polling API for completion