

MO53 - Final exam 2024P

Duration: 2h

Note 1: Documents and electronic devices are not authorized.

Note 2: Given points are for information only.

Exercise 1 (5 points): Choose one single answer per question.

A. You are migrating your workload from on-premises deployment to Google Kubernetes Engine (GKE). You want to minimize costs and stay within budget. What should you do?

1. Configure Autopilot in GKE to monitor node utilization and eliminate idle nodes.
2. Configure the needed capacity; the sustained use discount will make you stay within budget.
3. Scale individual nodes up and down with the Horizontal Pod Auto scaler.
4. Create several nodes using Compute Engine, add them to a managed instance group, and set the group to scale up and down depending on load.

B. You are creating a Cloud IOT application requiring data storage of up to 10 petabytes (PB). The application must support high-speed reads and writes of small pieces of data, but your data schema is simple. You want to use the most economical solution for data storage. What should you do?

1. Store the data in Cloud Spanner, and add an in-memory cache for speed.
2. Store the data in Cloud Storage, and distribute the data through Cloud CDN for speed.
3. Store the data in Cloud Bigtable, and implement the business logic in the programming language of your choice.
4. Use BigQuery, and implement the business logic in SQL.

C. Your team needs to directly connect your on-premises resources to several virtual machines inside a virtual private cloud (VPC). You want to provide your team with fast and secure access to the VMs with minimal maintenance and cost. What should you do?

1. Set up Cloud Interconnect.
2. Use Cloud VPN to create a bridge between the VPC and your network.
3. Assign a public IP address to each VM, and assign a strong password to each one.
4. Start a Compute Engine VM, install a software router, and create a direct tunnel to each VM.

D. You are responsible for monitoring all changes in your Cloud Storage and Firestore instances. For each change, you need to invoke an action that will verify the compliance of the change in near real time. You want to accomplish this with minimal setup. What should you do?

1. Use the trigger mechanism in each datastore to invoke the security script.
2. Use Cloud Function events, and call the security script from the Cloud Function triggers.
3. Use a Python script to get logs of the datastores, analyze them, and invoke the security script.
4. Redirect your data-changing queries to an App Engine application, and call the security script from the application.

E. Your team is building the development, test, and production environments for your project deployment in Google Cloud. You need to efficiently deploy and manage these environments and ensure that they are consistent. You want to follow Google-recommended practices. What should you do?

1. Create a Cloud Shell script that uses gcloud commands to deploy the environments.
2. Create one Terraform configuration for all environments. Parameterize the differences between environments.
3. For each environment, create a Terraform configuration. Use them for repeated deployment. Reconcile the templates periodically.
4. Use the Cloud Foundation Toolkit to create one deployment template that will work for all environments, and deploy with Terraform.

Exercise 2 (8 points): Choose between (True/False) for the following statements. Leave it empty if not sure because the wrong answers are noted negatively (you lose points for this exercise).

- A.** Elasticity is a technical feature of cloud computing allowing to assign the resources of the same physical machine to several applications.
- B.** Quality of Service (QoS) is a set of objective (technology-centric) quality metrics, while Quality of Experience (QoE) is a set of subjective (user-centric) quality metrics.
- C.** Grid target communities are mainly professionals while P2P target communities are mass-use, diverse and anonymous individuals.
- D.** In cloud computing, a Service Level Objective (SLO) contains many Service Level Agreements (SLAs).
- E.** Resource allocator is a component of elasticity manager that distributes the requests among the existing VMs.
- F.** Using a cross-platform API, the developer of an application calls a common and unified API and gets a standard based answer regardless of the implementations of different providers.
- G.** Resource provisioning is a mechanism where the Application Service Provider (ASP) assigns a job to running VM, while scheduling is a mechanism where the ASP buys more resources from the cloud or releases existing ones.
- H.** Infrastructure as a Service (IaaS) is targeted towards operational teams. It is about renting virtualized computing and storage resources (raw virtualized resources).

Exercise 3 (7 points): Answer the following questions.

- A.** Draw the grid architecture stating all layers and explain shortly the functionality of each layer.
- B.** Provide a state machine diagram about the VM lifecycle in Google cloud.
- C.** List the services of the Google cloud's operations suite for resource monitoring and provide the main functionality per each.

End of exam