

#### QUESTION 7

Your development team has created a mobile game app. You want to test the new mobile app on Android and iOS devices with a variety of configurations. You need to ensure that testing is efficient and cost-effective. What should you do?

- A. Upload your mobile app to the Firebase Test Lab, and test the mobile app on Android and iOS devices.
- B. Create Android and iOS VMs on Google Cloud, install the mobile app on the VMs, and test the mobile app.
- C. Create Android and iOS containers on Google Kubernetes Engine (GKE), install the mobile app on the containers, and test the mobile app.
- D. Upload your mobile app with different configurations to Firebase Hosting and test each configuration.

**Correct Answer: C**

#### QUESTION 8

For this question, refer to the Mountkirk Games case study. You are in charge of the new Game Backend Platform architecture. The game communicates with the backend over a REST API. You want to follow Google-recommended practices. How should you design the backend?

- A. Create an instance template for the backend. For every region, deploy it on a multi-zone managed instance group. Use an L4 load balancer.
- B. Create an instance template for the backend. For every region, deploy it on a single-zone managed instance group. Use an L4 load balancer.
- C. Create an instance template for the backend. For every region, deploy it on a multi-zone managed instance group. Use an L7 load balancer.
- D. Create an instance template for the backend. For every region, deploy it on a single-zone managed instance group. Use an L7 load balancer.

**Correct Answer: C**

**Explanation:**

[https://cloud.google.com/solutions/gaming/cloud-game-infrastructure#dedicated\\_game\\_server](https://cloud.google.com/solutions/gaming/cloud-game-infrastructure#dedicated_game_server)

#### QUESTION 9

Your development teams release new versions of games running on Google Kubernetes Engine (GKE) daily. You want to create service level indicators (SLIs) to evaluate the quality of the new versions from the user's perspective. What should you do?

- A. Create CPU Utilization and Request Latency as service level indicators.
- B. Create GKE CPU Utilization and Memory Utilization as service level indicators.
- C. Create Request Latency and Error Rate as service level indicators.
- D. Create Server Uptime and Error Rate as service level indicators.

**Correct Answer: C**

#### QUESTION 10

For this question, refer to the Mountkirk Games case study. You need to analyze and define the technical architecture for the compute workloads for your company, Mountkirk Games.

[Professional-Cloud-Architect Exam Dumps](#)   [Professional-Cloud-Architect PDF Dumps](#)

[Professional-Cloud-Architect VCE Dumps](#)   [Professional-Cloud-Architect Q&As](#)

<https://www.ensurepass.com/PCA.html>

Considering the Mountkirk Games business and technical requirements, what should you do?

- A. Create network load balancers. Use preemptible Compute Engine instances.
- B. Create network load balancers. Use non-preemptible Compute Engine instances.
- C. Create a global load balancer with managed instance groups and autoscaling policies. Use preemptible Compute Engine instances.
- D. Create a global load balancer with managed instance groups and autoscaling policies. Use non-preemptible Compute Engine instances.

**Correct Answer:** D

#### **QUESTION 11**

For this question, refer to the Mountkirk Games case study. Mountkirk Games wants to migrate from their current analytics and statistics reporting model to one that meets their technical requirements on Google Cloud Platform. Which two steps should be part of their migration plan? (Choose two.)

- A. Evaluate the impact of migrating their current batch ETL code to Cloud Dataflow.
- B. Write a schema migration plan to denormalize data for better performance in BigQuery.
- C. Draw an architecture diagram that shows how to move from a single MySQL database to a MySQL cluster.
- D. Load 10 TB of analytics data from a previous game into a Cloud SQL instance, and run test queries against the full dataset to confirm that they complete successfully.
- E. Integrate Cloud Armor to defend against possible SQL injection attacks in analytics files uploaded to Cloud Storage.

**Correct Answer:** AB

**Explanation:**

[https://cloud.google.com/bigquery/docs/loading-data#loading\\_denormalized\\_nested\\_and\\_repeated\\_data](https://cloud.google.com/bigquery/docs/loading-data#loading_denormalized_nested_and_repeated_data)

#### **QUESTION 12**

For this question, refer to the Mountkirk Games case study. You need to analyze and define the technical architecture for the database workloads for your company, Mountkirk Games. Considering the business and technical requirements, what should you do?

- A. Use Cloud SQL for time series data, and use Cloud Bigtable for historical data queries.
- B. Use Cloud SQL to replace MySQL, and use Cloud Spanner for historical data queries.
- C. Use Cloud Bigtable to replace MySQL, and use BigQuery for historical data queries.
- D. Use Cloud Bigtable for time series data, use Cloud Spanner for transactional data, and use BigQuery for historical data queries.

**Correct Answer:** D

**Explanation:**

<https://cloud.google.com/bigtable/docs/schema-design-time-series>

#### **QUESTION 13**

Mountkirk Games wants you to secure the connectivity from the new gaming application platform to Google Cloud. You want to streamline the process and follow Google-recommended practices. What should you do?

- A. Configure Workload Identity and service accounts to be used by the application platform.

[Professional-Cloud-Architect Exam Dumps](#)   [Professional-Cloud-Architect PDF Dumps](#)

[Professional-Cloud-Architect VCE Dumps](#)   [Professional-Cloud-Architect Q&As](#)

<https://www.ensurepass.com/PCA.html>

- B. Use Kubernetes Secrets, which are obfuscated by default. Configure these Secrets to be used by the application platform.
- C. Configure Kubernetes Secrets to store the secret, enable Application-Layer Secrets Encryption, and use Cloud Key Management Service (Cloud KMS) to manage the encryption keys. Configure these Secrets to be used by the application platform.
- D. Configure HashiCorp Vault on Compute Engine, and use customer managed encryption keys and Cloud Key Management Service (Cloud KMS) to manage the encryption keys. Configure these Secrets to be used by the application platform.

**Correct Answer:** A

#### **QUESTION 14**

For this question, refer to the Mountkirk Games case study. Which managed storage option meets Mountkirk's technical requirement for storing game activity in a time series database service?

- A. Cloud Bigtable
- B. Cloud Spanner
- C. BigQuery
- D. Cloud Datastore

**Correct Answer:** A

#### **Explanation:**

<https://cloud.google.com/blog/products/databases/getting-started-with-time-series-trend-predictions-using-gcp>

## **Topic 8, Helicopter Racing League Case**

### **Company overview**

Helicopter Racing League (HRL) is a global sports league for competitive helicopter racing. Each year HRL holds the world championship and several regional league competitions where teams compete to earn a spot in the world championship. HRL offers a paid service to stream the races all over the world with live telemetry and predictions throughout each race.

### **Solution concept**

HRL wants to migrate their existing service to a new platform to expand their use of managed AI and ML services to facilitate race predictions. Additionally, as new fans engage with the sport, particularly in emerging regions, they want to move the serving of their content, both real-time and recorded, closer to their users.

### **Existing technical environment**

HRL is a public cloud-first company; the core of their mission-critical applications runs on their current public cloud provider. Video recording and editing is performed at the race tracks, and the content is encoded and transcoded, where needed, in the cloud. Enterprise-grade connectivity and local compute is provided by truck-mounted mobile data centers. Their race prediction services are hosted exclusively on their existing public cloud provider. Their existing technical environment is as follows:

- Existing content is stored in an object storage service on their existing public cloud provider.
- Video encoding and transcoding is performed on VMs created for each job.
- Race predictions are performed using TensorFlow running on VMs in the current public cloud provider.

### **Business requirements**

HRL's owners want to expand their predictive capabilities and reduce latency for their viewers in emerging markets. Their requirements are:

- Support ability to expose the predictive models to partners.
- Increase predictive capabilities during and before races:

Race results

Mechanical failures

### **Crowd sentiment**

- Increase telemetry and create additional insights.
- Measure fan engagement with new predictions.
- Enhance global availability and quality of the broadcasts.
- Increase the number of concurrent viewers.
- Minimize operational complexity.
- Ensure compliance with regulations.
- Create a merchandising revenue stream.

### **Technical requirements**

- Maintain or increase prediction throughput and accuracy.
- Reduce viewer latency.
- Increase transcoding performance.
- Create real-time analytics of viewer consumption patterns and engagement.
- Create a data mart to enable processing of large volumes of race data.

### **Executive statement**

Our CEO, S. Hawke, wants to bring high-adrenaline racing to fans all around the world. We listen to our fans, and they want enhanced video streams that include predictions of events within the race (e.g., overtaking). Our current platform allows us to predict race outcomes but lacks the facility to support real-time predictions during races and the capacity to process season-long results.

### **QUESTION 1**

For this question, refer to the Helicopter Racing League (HRL) case study. Recently HRL started a new regional racing league in Cape Town, South Africa. In an effort to give customers in Cape Town a better user experience, HRL has partnered with the Content Delivery Network provider, Fastly. HRL needs to allow traffic coming from all of the Fastly IP address ranges into their Virtual Private Cloud network (VPC network). You are a member of the HRL security team and you need to configure the update that will allow only the Fastly IP address ranges through the External HTTP(S) load balancer. Which command should you use?

- A. `gcloud compute firewall rules update hlr-policy \ -priority 1000 \ target tags-sourceiplist fastly \ -allow tcp:443`
- B. `gcloud compute security policies rules update 1000 \ -security-policy hlr-policy \ -expression "evaluatePreconfiguredExpr('sourceiplist-fastly')\" \ -action " allow"`
- C. `gcloud compute firewall rules update sourceiplist-fastly \ priority 1000 \ allow tcp: 443`
- D. `gcloud compute priority-policies rules update 1000 \`

security policy from fastly -src- ip-ranges" - action " allow"

**Correct Answer:** B

**Explanation:**

<https://cloud.google.com/load-balancing/docs/https>

## QUESTION 2

For this question, refer to the Helicopter Racing League (HRL) case study. A recent finance audit of cloud infrastructure noted an exceptionally high number of Compute Engine instances are allocated to do video encoding and transcoding. You suspect that these Virtual Machines are zombie machines that were not deleted after their workloads completed. You need to quickly get a list of which VM instances are idle. What should you do?

- A. Log into each Compute Engine instance and collect disk, CPU, memory, and network usage statistics for analysis.
- B. Use the gcloud compute instances list to list the virtual machine instances that have the idle: true label set.
- C. Use the gcloud recommender command to list the idle virtual machine instances.
- D. From the Google Console, identify which Compute Engine instances in the managed instance groups are no longer responding to health check probes.

**Correct Answer:** C

**Explanation:**

<https://cloud.google.com/compute/docs/instances/viewing-and-applying-idle-vm-recommendations>

## QUESTION 3

For this question, refer to the Helicopter Racing League (HRL) case study. The HRL development team releases a new version of their predictive capability application every Tuesday evening at 3 a.m. UTC to a repository. The security team at HRL has developed an in-house penetration test Cloud Function called Airwolf. The security team wants to run Airwolf against the predictive capability application as soon as it is released every Tuesday. You need to set up Airwolf to run at the recurring weekly cadence. What should you do?

- A. Set up Cloud Tasks and a Cloud Storage bucket that triggers a Cloud Function.
- B. Set up a Cloud Logging sink and a Cloud Storage bucket that triggers a Cloud Function.
- C. Configure the deployment job to notify a Pub/Sub queue that triggers a Cloud Function.
- D. Set up Identity and Access Management (IAM) and Confidential Computing to trigger a Cloud Function.

**Correct Answer:** A

## QUESTION 4

For this question, refer to the Helicopter Racing League (HRL) case study. HRL wants better prediction accuracy from their ML prediction models. They want you to use Google's AI Platform so HRL can understand and interpret the predictions. What should you do?

- A. Use Explainable AI.

[Professional-Cloud-Architect Exam Dumps](#)   [Professional-Cloud-Architect PDF Dumps](#)

[Professional-Cloud-Architect VCE Dumps](#)   [Professional-Cloud-Architect Q&As](#)

<https://www.ensurepass.com/PCA.html>