

## 2022 Latest Professional-Cloud-Architect dumps Exam Material with 251 Questions [Q11-Q25]



### 2022 Latest Professional-Cloud-Architect dumps Exam Material with 251 Questions Google Professional-Cloud-Architect Questions and Answers Guarantee you Oass the Test Easily QUESTION 11

For this question, refer to the TerramEarth case study. You are asked to design a new architecture for the ingestion of the data of the 200,000 vehicles that are connected to a cellular network. You want to follow Google-recommended practices.

Considering the technical requirements, which components should you use for the ingestion of the data?

- \* Google Kubernetes Engine with an SSL Ingress
- \* Cloud IoT Core with public/private key pairs
- \* Compute Engine with project-wide SSH keys
- \* Compute Engine with specific SSH keys

Explanation

<https://cloud.google.com/solutions/iot-overview>

<https://cloud.google.com/iot/quotas>

## QUESTION 12

All compute Engine instances in your VPC should be able to connect to an Active Directory server on specific ports. Any other traffic emerging from your instances is not allowed. You want to enforce this using VPC firewall rules.

How should you configure the firewall rules?

\* Create an egress rule with priority 1000 to deny all traffic for all instances. Create another egress rule

with priority 100 to allow the Active Directory traffic for all instances.

\* Create an egress rule with priority 100 to deny all traffic for all instances. Create another egress rule

with priority 1000 to allow the Active Directory traffic for all instances.

\* Create an egress rule with priority 1000 to allow the Active Directory traffic. Rely on the implied deny

egress rule with priority 100 to block all traffic for all instances.

\* Create an egress rule with priority 100 to allow the Active Directory traffic. Rely on the implied deny

egress rule with priority 1000 to block all traffic for all instances.

## QUESTION 13

You want to optimize the performance of an accurate, real-time, weather-charting application. The data comes from 50,000 sensors sending 10 readings a second, in the format of a timestamp and sensor reading. Where should you store the data?

\* Google BigQuery

\* Google Cloud SQL

\* Google Cloud Bigtable

\* Google Cloud Storage

Google Cloud Bigtable is a scalable, fully-managed NoSQL wide-column database that is suitable for both real-time access and analytics workloads.

Good for:

Low-latency read/write access

High-throughput analytics

Native time series support

Common workloads:

IoT, finance, adtech

Personalization, recommendations

Monitoring

Geospatial datasets

Graphs

References: <https://cloud.google.com/storage-options/>

Reference:

It is time-series data, So Big Table.

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

#### QUESTION 14

Your company places a high value on being responsive and meeting customer needs quickly. Their primary business objectives are release speed and agility. You want to reduce the chance of security errors being accidentally introduced.

Which two actions can you take? Choose 2 answers.

- \* Enable code signing and a trusted binary repository integrated with your CI/CD pipeline
- \* Ensure you have stubs to unit test all interfaces between components
- \* Run a vulnerability security scanner as part of your continuous-integration /continuous-delivery (CI/CD) pipeline
- \* Use source code security analyzers as part of the CI/CD pipeline
- \* Ensure every code check-in is peer reviewed by a security SME

#### QUESTION 15

You need to deploy a stateful workload on Google Cloud. The workload can scale horizontally, but each instance needs to read and write to the same POSIX filesystem. At high load, the stateful workload needs to support up to 100 MB/s of writes. What should you do?

- \* Use a persistent disk for each instance.
- \* Use a regional persistent disk for each instance.
- \* Create a Cloud Filestore instance and mount it in each instance.
- \* Create a Cloud Storage bucket and mount it in each instance using gcsfuse.

<https://cloud.google.com/storage/docs/gcs-fuse#notes>

Cloud Filestore: Cloud Filestore is a scalable and highly available shared file service fully managed by Google. Cloud Filestore provides persistent storage ideal for shared workloads. It is best suited for enterprise applications requiring persistent, durable, shared storage which is accessed by NFS or requires a POSIX compliant file system.

#### QUESTION 16

For this question, refer to the JencoMart case study.

JencoMart has built a version of their application on Google Cloud Platform that serves traffic to Asia. You want to measure success against their business and technical goals.

Which metrics should you track?

- \* Error rates for requests from Asia
- \* Latency difference between US and Asia
- \* Total visits, error rates, and latency from Asia

- \* Total visits and average latency for users in Asia
- \* The number of character sets present in the database

#### Topic 4, Dress4Win case study Company Overview

Dress4win is a web-based company that helps their users organize and manage their personal wardrobe using a website and mobile application. The company also cultivates an active social network that connects their users with designers and retailers. They monetize their services through advertising, e-commerce, referrals, and a freemium app model.

#### Company Background

Dress4win's application has grown from a few servers in the founder's garage to several hundred servers and appliances in a collocated data center. However, the capacity of their infrastructure is now insufficient for the application's rapid growth. Because of this growth and the company's desire to innovate faster, Dress4win is committing to a full migration to a public cloud.

#### Solution Concept

For the first phase of their migration to the cloud, Dress4win is considering moving their development and test environments. They are also considering building a disaster recovery site, because their current infrastructure is at a single location. They are not sure which components of their architecture they can migrate as is and which components they need to change before migrating them.

#### Existing Technical Environment

The Dress4win application is served out of a single data center location.

- \* Databases:
- \* Application servers:
- \* Storage appliances:
- \* Apache Hadoop/Spark servers:
- \* MQ servers:
- \* Miscellaneous servers:

#### Business Requirements

- \* Build a reliable and reproducible environment with scaled parity of production.
- \* Improve security by defining and adhering to a set of security and Identity and Access Management (IAM) best practices for cloud.
- \* Improve business agility and speed of innovation through rapid provisioning of new resources.
- \* Analyze and optimize architecture for performance in the cloud.
- \* Migrate fully to the cloud if all other requirements are met.

## Technical Requirements

- \* Evaluate and choose an automation framework for provisioning resources in cloud.
- \* Support failover of the production environment to cloud during an emergency.
- \* Identify production services that can migrate to cloud to save capacity.
- \* Use managed services whenever possible.
- \* Encrypt data on the wire and at rest.
- \* Support multiple VPN connections between the production data center and cloud environment.

## CEO Statement

Our investors are concerned about our ability to scale and contain costs with our current infrastructure. They are also concerned that a new competitor could use a public cloud platform to offset their up-front investment and freeing them to focus on developing better features.

## CTO Statement

We have invested heavily in the current infrastructure, but much of the equipment is approaching the end of its useful life. We are consistently waiting weeks for new gear to be racked before we can start new projects. Our traffic patterns are highest in the mornings and weekend evenings; during other times, 80% of our capacity is sitting idle.

## CFO Statement

Our capital expenditure is now exceeding our quarterly projections. Migrating to the cloud will likely cause an initial increase in spending, but we expect to fully transition before our next hardware refresh cycle. Our total cost of ownership (TCO) analysis over the next 5 years puts a cloud strategy between 30 to 50% lower than our current model.

## QUESTION 17

You set up an autoscaling instance group to serve web traffic for an upcoming launch. After configuring the instance group as a backend service to an HTTP(S) load balancer, you notice that virtual machine (VM) instances are being terminated and re-launched every minute. The instances do not have a public IP address. You have verified the appropriate web response is coming from each instance using the curl command. You want to ensure the backend is configured correctly. What should you do?

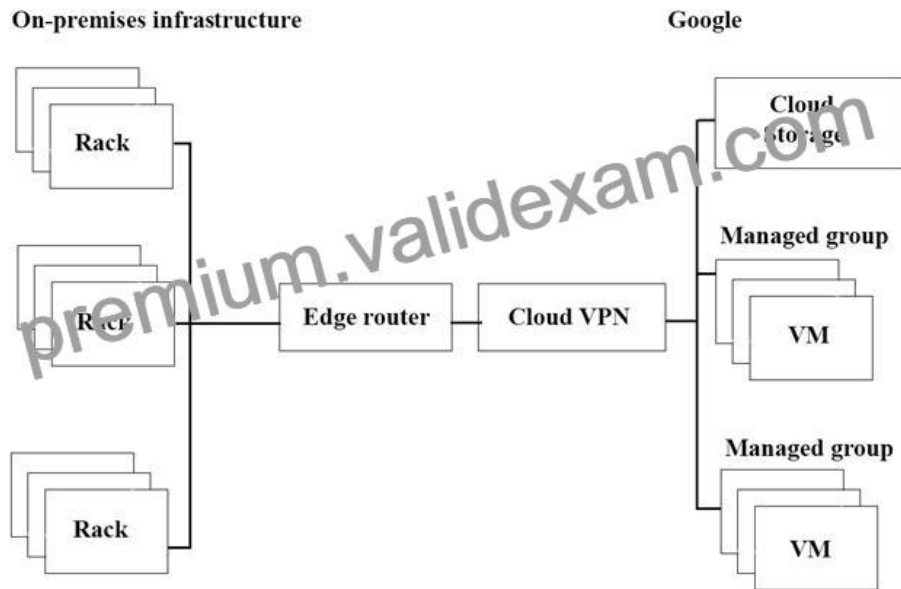
- \* Ensure that a firewall rule exists to allow source traffic on HTTP/HTTPS to reach the load balancer.
- \* Assign a public IP to each instance and configure a firewall rule to allow the load balancer to reach the instance public IP.
- \* Ensure that a firewall rule exists to allow load balancer health checks to reach the instances in the instance group.
- \* Create a tag on each instance with the name of the load balancer. Configure a firewall rule with the name of the load balancer as the source and the instance tag as the destination.

The best practice when configuration a health check is to check health and serve traffic on the same port. However, it is possible to perform health checks on one port, but serve traffic on another. If you do use two different ports, ensure that firewall rules and services running on instances are configured appropriately. If you run health checks and serve traffic on the same port, but decide to switch ports at some point, be sure to update both the backend service and the health check.

Backend services that do not have a valid global forwarding rule referencing it will not be health checked and will have no health status.

References: <https://cloud.google.com/compute/docs/load-balancing/http/backend-service>

### QUESTION 18



The migration of JencoMart's application to Google Cloud Platform (GCP) is progressing too slowly. The infrastructure is shown in the diagram. You want to maximize throughput.

What are three potential bottlenecks? Choose 3 answers.

- \* A single VPN tunnel, which limits throughput
- \* A tier of Google Cloud Storage that is not suited for this task
- \* A copy command that is not suited to operate over long distances
- \* Fewer virtual machines (VMs) in GCP than on-premises machines
- \* A separate storage layer outside the VMs, which is not suited for this task
- \* Complicated internet connectivity between the on-premises infrastructure and GCP

### QUESTION 19

Your marketing department wants to send out a promotional email campaign. The development team wants to minimize direct operation management. They project a wide range of possible customer responses, from 100 to 500,000 click-throughs per day. The link leads to a simple website that explains the promotion and collects user information and preferences. Which infrastructure should you recommend?

- \* Use Google App Engine to serve the website and Google Cloud Datastore to store user data.
- \* Use a Google Container Engine cluster to serve the website and store data to persistent disk.
- \* Use a managed instance group to serve the website and Google Cloud Bigtable to store user data.
- \* Use a single compute Engine virtual machine (VM) to host a web server, backed by Google Cloud SQL.

### QUESTION 20

You have a Python web application with many dependencies that requires 0.1 CPU cores and

128 MB of memory to operate in production. You want to monitor and maximize machine utilization. You also to reliably deploy new versions of the application. Which set of steps should you take?

\* Perform the following:

1. Create a managed instance group with f1-micro type machines.
2. Use a startup script to clone the repository, check out the production branch, install the dependencies, and start the Python app.
3. Restart the instances to automatically deploy new production releases.

\* Perform the following:

1. Create a managed instance group with n1-standard-1 type machines.
2. Build a Compute Engine image from the production branch that contains all of the dependencies and automatically starts the Python app.
3. Rebuild the Compute Engine image, and update the instance template to deploy new production releases.

\* Perform the following:

1. Create a Kubernetes Engine cluster with n1-standard-1 type machines.
2. Build a Docker image from the production branch with all of the dependencies, and tag it with the version number.
3. Create a Kubernetes Deployment with the `imagePullPolicy` set to `IfNotPresent`; in the staging namespace, and then promote it to the production namespace after testing.

\* Perform the following:

1. Create a Kubernetes Engine cluster with n1-standard-4 type machines.
2. Build a Docker image from the master branch with all of the dependencies, and tag it with `latest`.
3. Create a Kubernetes Deployment in the default namespace with the `imagePullPolicy` set to `Always`. Restart the pods to automatically deploy new production releases.

## QUESTION 21

For this question, refer to the JencoMart case study

A few days after JencoMart migrates the user credentials database to Google Cloud Platform and shuts down the old server, the new database server stops responding to SSH connections. It is still serving database requests to the application servers correctly. What three steps should you take to diagnose the problem? Choose 3 answers

- \* Delete the virtual machine (VM) and disks and create a new one.
- \* Delete the instance, attach the disk to a new VM, and investigate.
- \* Take a snapshot of the disk and connect to a new machine to investigate.
- \* Check inbound firewall rules for the network the machine is connected to.

- \* Connect the machine to another network with very simple firewall rules and investigate.
  - \* Print the Serial Console output for the instance for troubleshooting, activate the interactive console, and investigate.
- <https://cloud.google.com/compute/docs/troubleshooting/troubleshooting-ssh>

## QUESTION 22

Your company has decided to build a backup replica of their on-premises user authentication PostgreSQL database on Google Cloud Platform. The database is 4 TB, and large updates are frequent.

Replication requires private address space communication. Which networking approach should you use?

- \* Google Cloud Dedicated Interconnect
  - \* Google Cloud VPN connected to the data center network
  - \* A NAT and TLS translation gateway installed on-premises
  - \* A Google Compute Engine instance with a VPN server installed connected to the data center network
- <https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations>

## QUESTION 23

JencoMart has decided to migrate user profile storage to Google Cloud Datastore and the application servers to Google Compute Engine (GCE). During the migration, the existing infrastructure will need access to Datastore to upload the data.

What service account key-management strategy should you recommend?

- \* Provision service account keys for the on-premises infrastructure and for the GCE virtual machines (VMs)
- \* Authenticate the on-premises infrastructure with a user account and provision service account keys for the VMs
- \* Provision service account keys for the on-premises infrastructure and use Google Cloud Platform (GCP) managed keys for the VMs
- \* Deploy a custom authentication service on GCE/Google Kubernetes Engine (GKE) for the on-premises infrastructure and use GCP managed keys for the VMs

Migrating data to Google Cloud Platform

Let's say that you have some data processing that happens on another cloud provider and you want to transfer the processed data to Google Cloud Platform. You can use a service account from the virtual machines on the external cloud to push the data to Google Cloud Platform. To do this, you must create and download a service account key when you create the service account and then use that key from the external process to call the Cloud Platform APIs.

References:

[https://cloud.google.com/iam/docs/understanding-service-accounts#migrating\\_data\\_to\\_google\\_cloud\\_platform](https://cloud.google.com/iam/docs/understanding-service-accounts#migrating_data_to_google_cloud_platform)

## QUESTION 24

You want to enable your running Google Container Engine cluster to scale as demand for your application changes.

What should you do?

- \* Add additional nodes to your Container Engine cluster using the following command:

```
gcloud container clusters resize CLUSTER_NAME --size 10
```

- \* Add a tag to the instances in the cluster with the following command:

```
gcloud compute instances add-tags INSTANCE --tags enable --autoscaling max-nodes=10
```



\* Update the existing Container Engine cluster with the following command:

```
gcloud alpha container clusters update mycluster --enable-autoscaling --min-nodes=1 --max-nodes=10
```

\* Create a new Container Engine cluster with the following command:

```
gcloud alpha container clusters create mycluster --enable-autoscaling --min-nodes=1 --max-nodes=10
```

and redeploy your application.

Explanation

<https://cloud.google.com/kubernetes-engine/docs/concepts/cluster-autoscaler> Cluster autoscaling

```
--enable-autoscaling
```

Enables autoscaling for a node pool.

Enables autoscaling in the node pool specified by --node-pool or the default node pool if --node-pool is not provided.

Where:

```
--max-nodes=MAX_NODES
```

Maximum number of nodes in the node pool.

Maximum number of nodes to which the node pool specified by --node-pool (or default node pool if unspecified) can scale.

## QUESTION 25

You are developing an application using different microservices that should remain internal to the cluster. You want to be able to configure each microservice with a specific number of replicas. You also want to be able to address a specific microservice from any other microservice in a uniform way, regardless of the number of replicas the microservice scales to. You need to implement this solution on Google Kubernetes Engine. What should you do?

- \* Deploy each microservice as a Deployment. Expose the Deployment in the cluster using a Service, and use the Service DNS name to address it from other microservices within the cluster.
- \* Deploy each microservice as a Deployment. Expose the Deployment in the cluster using an Ingress, and use the Ingress IP address to address the Deployment from other microservices within the cluster.
- \* Deploy each microservice as a Pod. Expose the Pod in the cluster using a Service, and use the Service DNS name to address the microservice from other microservices within the cluster.
- \* Deploy each microservice as a Pod. Expose the Pod in the cluster using an Ingress, and use the Ingress IP address name to address the Pod from other microservices within the cluster.

**Share Latest Professional-Cloud-Architect DUMP Questions and Answers:**

<https://www.validexam.com/Professional-Cloud-Architect-latest-dumps.html>