

[Sep-2024 Verified Google Professional-Cloud-Architect Bundle Real Exam Dumps PDF [Q138-Q156]



[Sep-2024] Verified Google Professional-Cloud-Architect Bundle Real Exam Dumps PDF
Professional-Cloud-Architect Dumps PDF New [2024] Ultimate Study Guide

Google Professional-Cloud-Architect certification is a valuable certification for professionals who work with GCP. Google Certified Professional - Cloud Architect (GCP) certification is designed to test the candidate's knowledge and skills in designing, developing, and managing secure, scalable, and reliable solutions on the GCP. Google Certified Professional - Cloud Architect (GCP) certification is highly regarded in the industry and is recognized by many organizations. Passing the certification exam is a testament to the candidate's expertise in GCP and their ability to design and implement solutions that meet the needs of organizations.

QUESTION 138

Your company has sensitive data in Cloud Storage buckets. Data analysts have Identity Access Management (IAM) permissions to read the buckets. You want to prevent data analysts from retrieving the data in the buckets from outside the office network. What should you do?

- * 1. Create a VPC Service Controls perimeter that includes the projects with the buckets.

2. Create an access level with the CIDR of the office network.

* 1. Create a firewall rule for all instances in the Virtual Private Cloud (VPC) network for source range.

2. Use the Classless Inter-domain Routing (CIDR) of the office network.

* 1. Create a Cloud Function to remove IAM permissions from the buckets, and another Cloud Function to add IAM permissions to the buckets.

2. Schedule the Cloud Functions with Cloud Scheduler to add permissions at the start of business and remove permissions at the end of business.

* 1. Create a Cloud VPN to the office network.

2. Configure Private Google Access for on-premises hosts.

QUESTION 139

You are deploying an application on App Engine that needs to integrate with an on-premises database. For security purposes, your on-premises database must not be accessible through the public Internet. What should you do?

* Deploy your application on App Engine standard environment and use App Engine firewall rules to limit access to the open on-premises database.

* Deploy your application on App Engine standard environment and use Cloud VPN to limit access to the on-premises database.

* Deploy your application on App Engine flexible environment and use App Engine firewall rules to limit access to the on-premises database.

* Deploy your application on App Engine flexible environment and use Cloud VPN to limit access to the on-premises database.

QUESTION 140

You are designing a large distributed application with 30 microservices. Each of your distributed microservices needs to connect to a database back-end. You want to store the credentials securely. Where should you store the credentials?

* In the source code

* In an environment variable

* In a secret management system

* In a config file that has restricted access through ACLs

QUESTION 141

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?

* Cloud Bigtable

* Cloud Spanner

* BigQuery

* Cloud Datastore

Explanation

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

QUESTION 142

Your company wants to start using Google Cloud resources but wants to retain their on-premises Active Directory domain controller for identity management. What should you do?

* Use the Admin Directory API to authenticate against the Active Directory domain controller.

- * Use Google Cloud Directory Sync to synchronize Active Directory usernames with cloud identities and configure SAML SSO.
- * Use Cloud Identity-Aware Proxy configured to use the on-premises Active Directory domain controller as an identity provider.
- * Use Compute Engine to create an Active Directory (AD) domain controller that is a replica of the onpremises AD domain controller using Google Cloud Directory Sync.

Reference:

<https://cloud.google.com/blog/products/identity-security/using-your-existing-identity-managementsystem-with-google-cloud-platform>

QUESTION 143

For this question, refer to the TerramEarth case study.

The TerramEarth development team wants to create an API to meet the company's business requirements. You want the development team to focus their development effort on business value versus creating a custom framework. Which method should they use?

- * Use Google App Engine with Google Cloud Endpoints. Focus on an API for dealers and partners.
- * Use Google App Engine with a JAX-RS Jersey Java-based framework. Focus on an API for the public.
- * Use Google App Engine with the Swagger (open API Specification) framework. Focus on an API for the public.
- * Use Google Container Engine with a Django Python container. Focus on an API for the public.
- * Use Google Container Engine with a Tomcat container with the Swagger (Open API Specification) framework. Focus on an API for dealers and partners.

Explanation

https://cloud.google.com/endpoints/docs/openapi/about-cloud-endpoints?hl=en_US&_ga=2.21787131.-1712523

<https://cloud.google.com/endpoints/docs/openapi/architecture-overview>

<https://cloud.google.com/storage/docs/gsutil/commands/test>

Develop, deploy, protect and monitor your APIs with Google Cloud Endpoints. Using an Open API Specification or one of our API frameworks, Cloud Endpoints gives you the tools you need for every phase of API development.

From scenario:

Business Requirements

Decrease unplanned vehicle downtime to less than 1 week, without increasing the cost of carrying surplus inventory Support the dealer network with more data on how their customers use their equipment to better position new products and services Have the ability to partner with different companies, especially with seed and fertilizer suppliers in the fast-growing agricultural business; to create compelling joint offerings for their customers.

Reference: <https://cloud.google.com/certification/guides/cloud-architect/casestudy-terramearth>

Topic 2, Mountkirk Games Case Study

Company Overview

Mountkirk Games makes online, session-based, multiplayer games for the most popular mobile platforms.

Company Background

Mountkirk Games builds all of their games with some server-side integration and has historically used cloud providers to lease physical servers. A few of their games were more popular than expected, and they had problems scaling their application servers, MySQL databases, and analytics tools.

Mountkirk's current model is to write game statistics to files and send them through an ETL tool that loads them into a centralized MySQL database for reporting.

Solution Concept

Mountkirk Games is building a new game, which they expect to be very popular. They plan to deploy the game's backend on Google Compute Engine so they can capture streaming metrics, run intensive analytics and take advantage of its autoscaling server environment and integrate with a managed NoSQL database.

Technical Requirements

Requirements for Game Backend Platform

1. Dynamically scale up or down based on game activity.
2. Connect to a managed NoSQL database service.
3. Run customized Linux distro.

Requirements for Game Analytics Platform

1. Dynamically scale up or down based on game activity.
2. Process incoming data on the fly directly from the game servers.
3. Process data that arrives late because of slow mobile networks.
4. Allow SQL queries to access at least 10 TB of historical data.
5. Process files that are regularly uploaded by users' mobile devices.
6. Use only fully managed services

CEO Statement

Our last successful game did not scale well with our previous cloud provider, resulting in lower user adoption and affecting the game's reputation. Our investors want more key performance indicators (KPIs) to evaluate the speed and stability of the game, as well as other metrics that provide deeper insight into usage patterns so we can adapt the games to target users.

CTO Statement

Our current technology stack cannot provide the scale we need, so we want to replace MySQL and move to an environment that provides autoscaling, low latency load balancing, and frees us up from managing physical servers.

CFO Statement

We are not capturing enough user demographic data usage metrics, and other KPIs. As a result, we do not engage the right users. We are not confident that our marketing is targeting the right users, and we are not selling enough premium Blast-Ups inside the games, which dramatically impacts our revenue.

QUESTION 144

For this question, refer to the EHR Healthcare case study. You are responsible for ensuring that EHR's use of Google Cloud will pass an upcoming privacy compliance audit. What should you do? (Choose two.)

- * Verify EHR's product usage against the list of compliant products on the Google Cloud compliance page.
- * Advise EHR to execute a Business Associate Agreement (BAA) with Google Cloud.
- * Use Firebase Authentication for EHR's user facing applications.
- * Implement Prometheus to detect and prevent security breaches on EHR's web-based applications.
- * Use GKE private clusters for all Kubernetes workloads.

<https://cloud.google.com/security/compliance/hipaa>

QUESTION 145

You are implementing Firestore for Mountkirk Games. Mountkirk Games wants to give a new game programmatic access to a legacy game's Firestore database. Access should be as restricted as possible. What should you do?

- * Create a service account (SA) in the legacy game's Google Cloud project, add this SA in the new game's IAM page, and then give it the Firebase Admin role in both projects
- * Create a service account (SA) in the legacy game's Google Cloud project, add a second SA in the new game's IAM page, and then give the Organization Admin role to both SAs
- * Create a service account (SA) in the legacy game's Google Cloud project, give it the Firebase Admin role, and then migrate the new game to the legacy game's project.
- * Create a service account (SA) in the legacy game's Google Cloud project, give the SA the Organization Admin rule and then give it the Firebase Admin role in both projects

QUESTION 146

You have an application that makes HTTP requests to Cloud Storage. Occasionally the requests fail with HTTP status codes of 5xx and 429.

How should you handle these types of errors?

- * Use gRPC instead of HTTP for better performance.
- * Implement retry logic using a truncated exponential backoff strategy.
- * Make sure the Cloud Storage bucket is multi-regional for geo-redundancy.
- * Monitor <https://status.cloud.google.com/feed.atom> and only make requests if Cloud Storage is not reporting an incident.

Explanation

Reference https://cloud.google.com/storage/docs/json_api/v1/status-codes

QUESTION 147

Your company is running its application workloads on Compute Engine. The applications have been deployed in production, acceptance, and development environments. The production environment is business-critical and is used 24/7, while the acceptance and development environments are only critical during office hours.

Your CFO has asked you to optimize these environments to achieve cost savings during idle times. What should you do?

- * Create a shell script that uses the `gcloud` command to change the machine type of the development and acceptance instances to a smaller machine type outside of office hours. Schedule the shell script on one of the production instances to automate the task.
- * Use Cloud Scheduler to trigger a Cloud Function that will stop the development and acceptance environments after office hours and start them just before office hours.
- * Deploy the development and acceptance applications on a managed instance group and enable autoscaling.
- * Use regular Compute Engine instances for the production environment, and use preemptible VMs for the acceptance and development environments.

Reference: <https://cloud.google.com/blog/products/it-ops/best-practices-for-optimizing-your-cloud-costs>

QUESTION 148

For this question, refer to the JencoMart case study.

JencoMart wants to move their User Profiles database to Google Cloud Platform. Which Google Database should they use?

- * Cloud Spanner
- * Google BigQuery
- * Google Cloud SQL
- * Google Cloud Datastore

<https://cloud.google.com/datastore/docs/concepts/overview>

Common workloads for Google Cloud Datastore:

User profiles

Product catalogs

Game state

Reference:

<https://cloud.google.com/datastore/docs/concepts/overview>

QUESTION 149

You need to reduce the number of unplanned rollbacks of erroneous production deployments in your company's web hosting platform. Improvement to the QA/Test processes accomplished an 80% reduction. Which additional two approaches can you take to further reduce the rollbacks? Choose 2 answers

- * Introduce a green-blue deployment model.
- * Replace the QA environment with canary releases.
- * Fragment the monolithic platform into microservices.
- * Reduce the platform's dependency on relational database systems.
- * Replace the platform's relational database systems with a NoSQL database.

QUESTION 150

You are creating an App Engine application that uses Cloud Datastore as its persistence layer. You need to retrieve several root entities for which you have the identifiers. You want to minimize the overhead in operations performed by Cloud Datastore. What should you do?

- * Create the Key object for each Entity and run a batch get operation

- * Create the Key object for each Entity and run multiple get operations, one operation for each entity
- * Use the identifiers to create a query filter and run a batch query operation
- * Use the identifiers to create a query filter and run multiple query operations, one operation for each entity

Explanation

<https://cloud.google.com/datastore/docs/concepts/entities#datastore-datastore-batch-upsert-nodejs>

QUESTION 151

You want your Google Kubernetes Engine cluster to automatically add or remove nodes based on CPUload.

What should you do?

- * Configure a HorizontalPodAutoscaler with a target CPU usage. Enable the Cluster Autoscaler from the GCP Console.
- * Configure a HorizontalPodAutoscaler with a target CPU usage. Enable autoscaling on the managed instance group for the cluster using the gcloud command.
- * Create a deployment and set the maxUnavailable and maxSurge properties. Enable the Cluster Autoscaler using the gcloud command.
- * Create a deployment and set the maxUnavailable and maxSurge properties. Enable autoscaling on the cluster managed instance group from the GCP Console.

QUESTION 152

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

What should you do?

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

```
gcloud container clusters resize
```

```
CLUSTER_Name &#8211; -size 10
```

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

```
gcloud compute instances add-tags
```

```
INSTANCE &#8211; -tags enable-
```

```
autoscaling max-nodes=10
```

- * Update the existing Kubernetes Engine cluster with the following command:

```
gcloud alpha container clusters
```

```
update mycluster &#8211; -enable-
```

```
autoscaling &#8211; -min-nodes=1 &#8211; -max-nodes=10
```

- * Create a new Kubernetes Engine cluster with the following command:

```
gcloud alpha container clusters
```

```
create mycluster &#8211; -enable-
```

autoscaling --min-nodes=1 --max-nodes=10

and redeploy your application

QUESTION 153

You have created several preemptible Linux virtual machine instances using Google Compute Engine. You want to properly shut down your application before the virtual machines are preempted. What should you do?

- * Create a shutdown script named `k99.shutdown` in the `/etc/rc.6.d/` directory.
- * Create a shutdown script registered as a `xinetd` service in Linux and configure a `Stackdriver` endpoint check to call the service.
- * Create a shutdown script and use it as the value for a new metadata entry with the key `shutdown-script` in the Cloud Platform Console when you create the new virtual machine instance.
- * Create a shutdown script, registered as a `xinetd` service in Linux, and use the `gcloud compute instances add-metadata` command to specify the service URL as the value for a new metadata entry with the key `shutdown-script-url`

QUESTION 154

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.

<https://cloud.google.com/vpc/docs/using-firewalls>

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.

QUESTION 155

For this question, refer to the EHR Healthcare case study. You are responsible for designing the Google Cloud network architecture for Google Kubernetes Engine. You want to follow Google best practices. Considering the EHR Healthcare business and technical requirements, what should you do to reduce the attack surface?

- * Use a private cluster with a private endpoint with master authorized networks configured.
- * Use a public cluster with firewall rules and Virtual Private Cloud (VPC) routes.
- * Use a private cluster with a public endpoint with master authorized networks configured.
- * Use a public cluster with master authorized networks enabled and firewall rules.

QUESTION 156

You are managing an application deployed on Cloud Run for Anthos, and you need to define a strategy for deploying new versions of the application. You want to evaluate the new code with a subset of production traffic to decide whether to proceed with the

rollout. What should you do?

- * Deploy a new revision to Cloud Run with the new version. Configure traffic percentage between revisions.
- * Deploy a new service to Cloud Run with the new version. Add a Cloud Load Balancing instance in front of both services.
- * In the Google Cloud Console page for Cloud Run, set up continuous deployment using Cloud Build for the development branch. As part of the Cloud Build trigger, configure the substitution variable TRAFFIC_PERCENTAGE with the percentage of traffic you want directed to a new version.
- * In the Google Cloud Console, configure Traffic Director with a new Service that points to the new version of the application on Cloud Run. Configure Traffic Director to send a small percentage of traffic to the new version of the application.

Pass Your Google Exam with Professional-Cloud-Architect Exam Dumps:

https://www.test4engine.com/Professional-Cloud-Architect_exam-latest-braindumps.html