



**Free Questions for Professional-Cloud-Architect by
actualtestdumps**

Shared by Sweeney on 12-12-2023

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

You are managing several internal applications that are deployed on Compute Engine. Business users inform you that an application has become very slow over the past few days. You want to find the underlying cause in order to solve the problem. What should you do first?

Options:

- A-** Inspect the logs and metrics from the instances in Cloud Logging and Cloud Monitoring.
- B-** Restore a backup of the application database from a time before the application became slow.
- C-** Deploy the applications on a managed instance group with autoscaling enabled. Add a load balancer in front of the managed instance group, and have the users connect to the IP of the load balancer.
- D-** Change the Compute Engine Instances behind the application to a machine type with more CPU and memory.

Answer:

A

Explanation:

When an application becomes slow, the first step you should take is to gather information about the underlying cause of the problem. One way to do this is by inspecting the logs and metrics from the instances where the application is deployed. Google Cloud Platform (GCP) provides tools such as Cloud Logging and Cloud Monitoring that can help you to collect and analyze this information. By reviewing the logs and metrics from the instances, you may be able to identify issues such as resource shortages (e.g. CPU, memory, or disk), network problems, or application errors that are causing the performance issues. Once you have identified the underlying cause of the problem, you can take steps to resolve it.

Question 2

Question Type: MultipleChoice

Your company has a Google Cloud project that uses BigQuery for data warehousing. The VPN tunnel between the on-premises environment and Google Cloud is configured with Cloud VPN. Your security team wants to avoid data exfiltration by malicious insiders, compromised code, and accidental oversharing. What should you do?

Options:

- A-** Configure VPC Service Controls and configure Private Google Access for on-premises hosts.
- B-** Create a service account, grant the BigQuery JobUser role and Storage Object Viewer role to the service account, and remove all other Identity and Access Management (IAM) access from the project.

C- Configure Private Google Access.

D- Configure Private Service Connect.

Answer:

A

Question 3

Question Type: MultipleChoice

Your company uses Google Kubernetes Engine (GKE) as a platform for all workloads. Your company has a single large GKE cluster that contains batch, stateful, and stateless workloads. The GKE cluster is configured with a single node pool with 200 nodes. Your company needs to reduce the cost of this cluster but does not want to compromise availability. What should you do?

Options:

A- Create a second GKE cluster for the batch workloads only. Allocate the 200 original nodes across both clusters.

B- Configure a HorizontalPodAutoscaler for all stateless workloads and for all compatible stateful workloads. Configure the cluster to use node auto scaling.

- C-** Configure CPU and memory limits on the namespaces in the cluster. Configure all Pods to have a CPU and memory limits.
- D-** Change the node pool to use spot VMs.

Answer:

C

Explanation:

One way to reduce the cost of a Google Kubernetes Engine (GKE) cluster without compromising availability is to use horizontal pod autoscalers (HPA) and node auto scaling. HPA allows you to automatically scale the number of Pods in a deployment based on the resource usage of the Pods. By configuring HPA for stateless workloads and for compatible stateful workloads, you can ensure that the number of Pods is automatically adjusted based on the actual resource usage, which can help to reduce costs. Node auto scaling allows you to automatically add or remove nodes from the node pool based on the resource usage of the cluster. By configuring node auto scaling, you can ensure that the cluster has the minimum number of nodes needed to meet the resource requirements of the workloads, which can also help to reduce costs.

Question 4

Question Type: MultipleChoice

You are deploying an application to Google Cloud. The application is part of a system. The application in Google Cloud must communicate over a private network with applications in a non-Google Cloud environment. The expected average throughput is 200 kbps. The business requires:

- * 99.99% system availability
- * cost optimization

You need to design the connectivity between the locations to meet the business requirements. What should you provision?

Options:

- A-** A Classic Cloud VPN gateway connected with one tunnel to an on-premises VPN gateway.
- B-** A Classic Cloud VPN gateway connected with two tunnels to an on-premises VPN gateway.
- C-** An HA Cloud VPN gateway connected with two tunnels to an on-premises VPN gateway.
- D-** Two HA Cloud VPN gateways connected to two on-premises VPN gateways. Configure each HA Cloud VPN gateway to have two tunnels, each connected to different on-premises VPN gateways.

Answer:

C

Explanation:

Question 5

Question Type: MultipleChoice

Your company has an application running on App Engine that allows users to upload music files and share them with other people. You want to allow users to upload files directly into Cloud Storage from their browser session. The payload should not be passed through the backend. What should you do?

Options:

- A-** 1. Set a CORS configuration in the target Cloud Storage bucket where the base URL of the App Engine application is an allowed origin.
2. Use the Cloud Storage Signed URL feature to generate a POST URL.
- B-** 1. Set a CORS configuration in the target Cloud Storage bucket where the base URL of the App Engine application is an allowed origin.
2. Assign the Cloud Storage WRITER role to users who upload files.
- C-** 1. Use the Cloud Storage Signed URL feature to generate a POST URL.
2. Use App Engine default credentials to sign requests against Cloud Storage.

- D-** 1. Assign the Cloud Storage WRITER role to users who upload files.
- 2. Use App Engine default credentials to sign requests against Cloud Storage.

Answer:

B

Question 6

Question Type: MultipleChoice

Your company has an application that is running on multiple instances of Compute Engine. It generates 1 TB per day of logs. For compliance reasons, the logs need to be kept for at least two years. The logs need to be available for active query for 30 days. After that, they just need to be retained for audit purposes. You want to implement a storage solution that is compliant, minimizes costs, and follows Google-recommended practices. What should you do?

Options:

- A-** 1. Install the Cloud Ops agent on all instances.
- 2. Create a sink to export logs into a partitioned BigQuery table.
- 3. Set a time_partitioning_expiration of 30 days.
- B-** 1. Install the Cloud Ops agent on all instances.

2. Create a sink to export logs into a regional Cloud Storage bucket.
 3. Create an Object Lifecycle rule to move files into a Coldline Cloud Storage bucket after one month.
 4. Configure a retention policy at the bucket level to create a lock.
- C-** 1. Create a daily cron job, running on all instances, that uploads logs into a partitioned BigQuery table. 2. Set a `time_partitioning_expiration` of 30 days.
- D-** 1. Write a daily cron job, running on all instances, that uploads logs into a Cloud Storage bucket.
2. Create a sink to export logs into a regional Cloud Storage bucket.
 3. Create an Object Lifecycle rule to move files into a Coldline Cloud Storage bucket after one month.

Answer:

B

Explanation:

The practice for managing logs generated on Compute Engine on Google Cloud is to install the Cloud Logging agent and send them to Cloud Logging.

The sent logs will be aggregated into a Cloud Logging sink and exported to Cloud Storage.

The reason for using Cloud Storage as the destination for the logs is that the requirement in question requires setting up a lifecycle based on the storage period.

In this case, the log will be used for active queries for 30 days after it is saved, but after that, it needs to be stored for a longer period of time for auditing purposes.

If the data is to be used for active queries, we can use BigQuery's Cloud Storage data query feature and move the data past 30 days to Coldline to build a cost-optimal solution.

Therefore, the correct answer is as follows

1. Install the Cloud Logging agent on all instances.

Create a sync that exports the logs to the region's Cloud Storage bucket.

3. Create an Object Lifecycle rule to move the files to the Coldline Cloud Storage bucket after one month. 4.

4. set up a bucket-level retention policy using bucket locking.'

Question 7

Question Type: MultipleChoice

Your company has a Google Cloud project that uses BigQuery for data warehousing on a pay-per-use basis. You want to monitor queries in real time to discover the most costly queries and which users spend the most. What should you do?

Options:

A- 1. Create a Cloud Logging sink to export BigQuery data access logs to Cloud Storage. 2. Develop a Dataflow pipeline to compute the cost of queries split by users.

B- 1. Create a Cloud Logging sink to export BigQuery data access logs to BigQuery.
2. Perform a BigQuery query on the generated table to extract the information you need.

C- 1. Activate billing export into BigQuery.
2. Perform a BigQuery query on the billing table to extract the information you need.

D- 1. In the BigQuery dataset that contains all the tables to be queried, add a label for each user that can launch a query.
2. Open the Billing page of the project.
3. Select Reports.
4. Select BigQuery as the product and filter by the user you want to check.

Answer:

C

Explanation:

<https://cloud.google.com/blog/products/data-analytics/taking-a-practical-approach-to-bigquery-cost-monitoring>

Question 8

Question Type: MultipleChoice

For this question, refer to the TerramEarth case study. You are building a microservice-based application for TerramEarth. The application is based on Docker containers. You want to follow Google-recommended practices to build the application continuously and store the build artifacts. What should you do?

Options:

- A-** 1. Configure a trigger in Cloud Build for new source changes.
2. Invoke Cloud Build to build one container image, and tag the image with the label 'latest.'
3. Push the image to the Artifact Registry.
- B-** 1. Configure a trigger in Cloud Build for new source changes.
2. Invoke Cloud Build to build container images for each microservice, and tag them using the code commit hash.
3. Push the images to the Artifact Registry.
- C-** 1 Create a Scheduler job to check the repo every minute.
2. For any new change, invoke Cloud Build to build container images for the microservices.
3. Tag the images using the current timestamp, and push them to the Artifact Registry.
- D-** 1. Configure a trigger in Cloud Build for new source changes.
2. The trigger invokes build jobs and build container images for the microservices.
3. Tag the images with a version number, and push them to Cloud Storage.

Answer:

C

To Get Premium Files for Professional-Cloud-Architect Visit

<https://www.p2pexams.com/products/professional-cloud-architect>

For More Free Questions Visit

<https://www.p2pexams.com/google/pdf/professional-cloud-architect>

