



Free Questions for Professional-Cloud-Architect by
vceexamstest

Shared by Watkins on 15-04-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Your company has an application running on multiple Compute Engine instances. You need to ensure that the application can communicate with an on-premises service that requires high throughput via internal IPs, while minimizing latency. What should you do?

Options:

- A- Use OpenVPN to configure a VPN tunnel between the on-premises environment and Google Cloud.
- B- Configure a direct peering connection between the on-premises environment and Google Cloud.
- C- Use Cloud VPN to configure a VPN tunnel between the on-premises environment and Google Cloud.
- D- Configure a Cloud Dedicated Interconnect connection between the on-premises environment and Google Cloud.

Answer:

C

Explanation:

Reference <https://cloud.google.com/architecture/setting-up-private-access-to-cloud-apis-through-vpn-tunnels>

Question 2

Question Type: MultipleChoice

You need to deploy an application to Google Cloud. The application receives traffic via TCP and reads and writes data to the filesystem. The application does not support horizontal scaling. The application process requires full control over the data on the file system because concurrent access causes corruption. The business is willing to accept a downtime when an incident occurs, but the application must be available 24/7 to support their business operations. You need to design the architecture of this application on Google Cloud.

What should you do?

Options:

- A-** Use a managed instance group with instances in multiple zones, use Cloud Filestore, and use an HTTP load balancer in front of the instances.
- B-** Use a managed instance group with instances in multiple zones, use Cloud Filestore, and use a network load balancer in front of the instances.
- C-** Use an unmanaged instance group with an active and standby instance in different zones, use a regional persistent disk, and use an HTTP load balancer in front of the instances.
- D-** Use an unmanaged instance group with an active and standby instance in different zones, use a regional persistent disk, and use a network load balancer in front of the instances.

Answer:

D

Question 3

Question Type: MultipleChoice

Your company is planning to perform a lift and shift migration of their Linux RHEL 6.5+ virtual machines. The virtual machines are running in an on-premises VMware environment. You want to migrate them to Compute Engine following Google-recommended practices. What should you do?

Options:

- A-** 1. Define a migration plan based on the list of the applications and their dependencies.
2. Migrate all virtual machines into Compute Engine individually with Migrate for Compute Engine.
- B-** 1. Perform an assessment of virtual machines running in the current VMware environment.
2. Create images of all disks. Import disks on Compute Engine.
3. Create standard virtual machines where the boot disks are the ones you have imported.
- C-** 1. Perform an assessment of virtual machines running in the current VMware environment.

2. Define a migration plan, prepare a Migrate for Compute Engine migration RunBook, and execute the migration.

D- 1. Perform an assessment of virtual machines running in the current VMware environment.

2. Install a third-party agent on all selected virtual machines.

3. Migrate all virtual machines into Compute Engine.

Answer:

C

Explanation:

The framework illustrated in the preceding diagram has four phases:

* **Assess.** In this phase, you assess your source environment, assess the workloads that you want to migrate to Google Cloud, and assess which VMs support each workload.

* **Plan.** In this phase, you create the basic infrastructure for Migrate for Compute Engine, such as provisioning the resource hierarchy and setting up network access.

* **Deploy.** In this phase, you migrate the VMs from the source environment to Compute Engine.

* **Optimize.** In this phase, you begin to take advantage of the cloud technologies and capabilities.

Question 4

Question Type: MultipleChoice

You are developing a globally scaled frontend for a legacy streaming backend data API. This API expects events in strict chronological order with no repeat data for proper processing.

Which products should you deploy to ensure guaranteed-once FIFO (first-in, first-out) delivery of data?

Options:

- A- Cloud Pub/Sub alone
- B- Cloud Pub/Sub to Cloud DataFlow
- C- Cloud Pub/Sub to Stackdriver
- D- Cloud Pub/Sub to Cloud SQL

Answer:

B

Explanation:

Reference <https://cloud.google.com/pubsub/docs/ordering>

Question 5

Question Type: MultipleChoice

Your company is building a new architecture to support its data-centric business focus. You are responsible for setting up the network. Your company's mobile and web-facing applications will be deployed on-premises, and all data analysis will be conducted in GCP. The plan is to process and load 7 years of archived .csv files totaling 900 TB of data and then continue loading 10 TB of data daily. You currently have an existing 100-MB internet connection.

What actions will meet your company's needs?

Options:

- A-** Compress and upload both archived files and files uploaded daily using the `gsutil --m` option.
- B-** Lease a Transfer Appliance, upload archived files to it, and send it, and send it to Google to transfer archived data to Cloud Storage. Establish a connection with Google using a Dedicated Interconnect or Direct Peering connection and use it to upload files daily.

- C-** Lease a Transfer Appliance, upload archived files to it, and send it, and send it to Google to transfer archived data to Cloud Storage. Establish one Cloud VPN Tunnel to VPC networks over the public internet, and compares and upload files daily using the gsutil --m option.
- D-** Lease a Transfer Appliance, upload archived files to it, and send it to Google to transfer archived data to Cloud Storage. Establish a Cloud VPN Tunnel to VPC networks over the public internet, and compress and upload files daily.

Answer:

B

Explanation:

<https://cloud.google.com/interconnect/docs/how-to/direct-peering>

Question 6

Question Type: MultipleChoice

You need to ensure reliability for your application and operations by supporting reliable task a scheduling for compute on GCP. Leveraging Google best practices, what should you do?

Options:

- A-** Using the Cron service provided by App Engine, publishing messages directly to a message-processing utility service running on Compute Engine instances.
- B-** Using the Cron service provided by App Engine, publish messages to a Cloud Pub/Sub topic. Subscribe to that topic using a message-processing utility service running on Compute Engine instances.
- C-** Using the Cron service provided by Google Kubernetes Engine (GKE), publish messages directly to a message-processing utility service running on Compute Engine instances.
- D-** Using the Cron service provided by GKE, publish messages to a Cloud Pub/Sub topic. Subscribe to that topic using a message-processing utility service running on Compute Engine instances.

Answer:

B

Explanation:

<https://cloud.google.com/solutions/reliable-task-scheduling-compute-engine>

Question 7

Question Type: MultipleChoice

You are deploying a PHP App Engine Standard service with SQL as the backend. You want to minimize the number of queries to the database.

What should you do?

Options:

- A-** Set the memcache service level to dedicated. Create a key from the hash of the query, and return database values from memcache before issuing a query to Cloud SQL.
- B-** Set the memcache service level to dedicated. Create a cron task that runs every minute to populate the cache with keys containing query results.
- C-** Set the memcache service level to shared. Create a cron task that runs every minute to save all expected queries to a key called "cached-queries".
- D-** Set the memcache service level to shared. Create a key called "cached-queries", and return database values from the key before using a query to Cloud SQL.

Answer:

A

Explanation:

<https://cloud.google.com/appengine/docs/standard/php/memcache/using>

Question 8

Question Type: MultipleChoice

Your company acquired a healthcare startup and must retain its customers' medical information for up to 4 more years, depending on when it was created. Your corporate policy is to securely retain this data, and then delete it as soon as regulations allow.

Which approach should you take?

Options:

- A-** Store the data in Google Drive and manually delete records as they expire.
- B-** Anonymize the data using the Cloud Data Loss Prevention API and store it indefinitely.
- C-** Store the data using the Cloud Storage and use lifecycle management to delete files when they expire.
- D-** Store the data in Cloud Storage and run a nightly batch script that deletes all expired data.

Answer:

C

Explanation:

<https://cloud.google.com/storage/docs/lifecycle>

Question 9

Question Type: MultipleChoice

Your company creates rendering software which users can download from the company website. Your company has customers all over the world. You want to minimize latency for all your customers. You want to follow Google-recommended practices.

How should you store the files?

Options:

- A-** Save the files in a Multi-Regional Cloud Storage bucket.
- B-** Save the files in a Regional Cloud Storage bucket, one bucket per zone of the region.
- C-** Save the files in multiple Regional Cloud Storage buckets, one bucket per zone per region.

D- Save the files in multiple Multi-Regional Cloud Storage buckets, one bucket per multi-region.

Answer:

A

Explanation:

<https://cloud.google.com/storage/docs/locations#location-mr>

Question 10

Question Type: MultipleChoice

You need to develop procedures to test a disaster plan for a mission-critical application. You want to use

Google-recommended practices and native capabilities within GCP.

What should you do?

Options:

- A- Use Deployment Manager to automate service provisioning. Use Activity Logs to monitor and debug your tests.
- B- Use Deployment Manager to automate provisioning. Use Stackdriver to monitor and debug your tests.
- C- Use gcloud scripts to automate service provisioning. Use Activity Logs monitor and debug your tests.
- D- Use automated scripts to automate service provisioning. Use Activity Logs monitor and debug your tests.

Answer:

B

Explanation:

<https://cloud.google.com/solutions/dr-scenarios-planning-guide>

Question 11

Question Type: MultipleChoice

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?

Options:

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

Answer:

A

Explanation:

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

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>

