

Correct Answer: A

Explanation:

<https://cloud.google.com/serverless-options>

QUESTION 13

Your application needs to process credit card transactions. You want the smallest scope of Payment Card Industry (PCI) compliance without compromising the ability to analyze transactional data and trends relating to which payment methods are used. How should you design your architecture?

- A. Create a tokenizer service and store only tokenized data.
- B. Create separate projects that only process credit card data.
- C. Create separate subnetworks and isolate the components that process credit card data.
- D. Streamline the audit discovery phase by labeling all of the virtual machines (VMs) that process PCI data.
- E. Enable Logging export to Google BigQuery and use ACLs and views to scope the data shared with the auditor.

Correct Answer: A

Explanation:

<https://cloud.google.com/solutions/pci-dss-compliance-in-gcp>

QUESTION 14

Your company wants you to build a highly reliable web application with a few public APIs as the backend. You don't expect a lot of user traffic, but traffic could spike occasionally. You want to leverage Cloud Load Balancing, and the solution must be cost-effective for users. What should you do?

- A. Store static content such as HTML and images in Cloud CDN. Host the APIs on App Engine and store the user data in Cloud SQL.
- B. Store static content such as HTML and images in a Cloud Storage bucket. Host the APIs on a zonal Google Kubernetes Engine cluster with worker nodes in multiple zones, and save the user data in Cloud Spanner.
- C. Store static content such as HTML and images in Cloud CDN. Use Cloud Run to host the APIs and save the user data in Cloud SQL.
- D. Store static content such as HTML and images in a Cloud Storage bucket. Use Cloud Functions to host the APIs and save the user data in Firestore.

Correct Answer: D

Explanation:

<https://cloud.google.com/load-balancing/docs/https/setting-up-https-serverless#gcloud:-cloud-functions>

<https://cloud.google.com/blog/products/networking/better-load-balancing-for-app-engine-cloud-run-and-functions>

QUESTION 15

Your applications will be writing their logs to BigQuery for analysis. Each application should have its own table. Any logs older than 45 days should be removed. You want to optimize storage and follow Google recommended practices. What should you do?

- A. Configure the expiration time for your tables at 45 days

[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. Make the tables time-partitioned, and configure the partition expiration at 45 days
- C. Rely on BigQuery's default behavior to prune application logs older than 45 days
- D. Create a script that uses the BigQuery command line tool (bq) to remove records older than 45 days

Correct Answer: B

Explanation:

<https://cloud.google.com/bigquery/docs/managing-partitioned-tables>

QUESTION 16

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?

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

Correct Answer: C

Explanation:

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

QUESTION 17

Your company is forecasting a sharp increase in the number and size of Apache Spark and Hadoop jobs being run on your local datacenter. You want to utilize the cloud to help you scale this upcoming demand with the least amount of operations work and code change. Which product should you use?

- A. Google Cloud Dataflow
- B. Google Cloud Dataproc
- C. Google Compute Engine
- D. Google Container Engine

Correct Answer: B

Explanation:

Google Cloud Dataproc is a fast, easy-to-use, low-cost and fully managed service that lets you run the Apache Spark and Apache Hadoop ecosystem on Google Cloud Platform. Cloud Dataproc provisions big or small clusters rapidly, supports many popular job types, and is integrated with other Google Cloud Platform services, such as Google Cloud Storage and Stackdriver Logging, thus helping you reduce TCO.

References:

<https://cloud.google.com/dataproc/docs/resources/faq>

QUESTION 18

Your company has successfully migrated to the cloud and wants to analyze their data stream to

[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>

optimize operations. They do not have any existing code for this analysis, so they are exploring all their options. These options include a mix of batch and stream processing, as they are running some hourly jobs and live-processing some data as it comes in. Which technology should they use for this?

- A. Google Cloud Dataproc
- B. Google Cloud Dataflow
- C. Google Container Engine with Bigtable
- D. Google Compute Engine with Google BigQuery

Correct Answer: B

Explanation:

Dataflow is for processing both the Batch and Stream.

Cloud Dataflow is a fully-managed service for transforming and enriching data in stream (real time) and batch (historical) modes with equal reliability and expressiveness -- no more complex workarounds or compromises needed.

References:

<https://cloud.google.com/dataflow/>

QUESTION 19

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?

- A. Ensure that a firewall rule exists to allow source traffic on HTTP/HTTPS to reach the load balancer.
- B. Assign a public IP to each instance and configure a firewall rule to allow the load balancer to reach the instance public IP.
- C. Ensure that a firewall rule exists to allow load balancer health checks to reach the instances in the instance group.
- D. 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.

Correct Answer: C

Explanation:

<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.

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

[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>

QUESTION 20

You need to set up Microsoft SQL Server on GCP. Management requires that there's no downtime in case of a data center outage in any of the zones within a GCP region. What should you do?

- A. Configure a Cloud SQL instance with high availability enabled.
- B. Configure a Cloud Spanner instance with a regional instance configuration.
- C. Set up SQL Server on Compute Engine, using Always On Availability Groups using Windows Failover Clustering. Place nodes in different subnets.
- D. Set up SQL Server Always On Availability Groups using Windows Failover Clustering. Place nodes in different zones.

Correct Answer: D

Explanation:

<https://cloud.google.com/sql/docs/sqlserver/configure-ha>

QUESTION 21

Your company has a networking team and a development team. The development team runs applications on Compute Engine instances that contain sensitive data. The development team requires administrative permissions for Compute Engine. Your company requires all network resources to be managed by the networking team. The development team does not want the networking team to have access to the sensitive data on the instances. What should you do?

- A.
 1. Create a project with a standalone VPC and assign the Network Admin role to the networking team.
 2. Create a second project with a standalone VPC and assign the Compute Admin role to the development team.
 3. Use Cloud VPN to join the two VPCs.
- B.
 1. Create a project with a standalone Virtual Private Cloud (VPC), assign the Network Admin role to the networking team, and assign the Compute Admin role to the development team.
- C.
 1. Create a project with a Shared VPC and assign the Network Admin role to the networking team.
 2. Create a second project without a VPC, configure it as a Shared VPC service project, and assign the Compute Admin role to the development team.
- D.
 1. Create a project with a standalone VPC and assign the Network Admin role to the networking team.
 2. Create a second project with a standalone VPC and assign the Compute Admin role to the development team.
 3. Use VPC Peering to join the two VPCs.

Correct Answer: C

Explanation:

In this scenario, a large organization has a central team that manages security and networking controls for the entire organization. Developers do not have permissions to make changes to any network or security settings defined by the security and networking team but they are granted permission to create resources such as virtual machines in shared subnets. To facilitate this the organization makes use of a shared VPC (Virtual Private Cloud). A shared VPC allows creation of a VPC network of RFC 1918 IP spaces that associated projects (service projects) can then use. Developers using the associated projects can create VM instances in the shared VPC network spaces. The organization's network and security admins can create subnets, VPNs, and firewall rules usable by all the projects in the VPC network.

https://cloud.google.com/iam/docs/job-functions/networking#single_team_manages_security_network_for_organization
<https://cloud.google.com/vpc/docs/shared-vpc>

QUESTION 22

You are using Cloud SQL as the database backend for a large CRM deployment. You want to scale as usage increases and ensure that you don't run out of storage, maintain 75% CPU usage cores, and keep replication lag below 60 seconds. What are the correct steps to meet your requirements?

- A.
 1. Enable automatic storage increase for the instance.
 2. Create a Stackdriver alert when CPU usage exceeds 75%, and change the instance type to reduce CPU usage.
 3. Create a Stackdriver alert for replication lag, and shard the database to reduce replication time.
- B.
 1. Enable automatic storage increase for the instance.
 2. Change the instance type to a 32-core machine type to keep CPU usage below 75%.
 3. Create a Stackdriver alert for replication lag, and shard the database to reduce replication time.
- C.
 1. Create a Stackdriver alert when storage exceeds 75%, and increase the available storage on the instance to create more space.
 2. Deploy memcached to reduce CPU load.
 3. Change the instance type to a 32-core machine type to reduce replication lag.
- D.
 1. Create a Stackdriver alert when storage exceeds 75%, and increase the available storage on the instance to create more space.
 2. Deploy memcached to reduce CPU load.
 3. Create a Stackdriver alert for replication lag, and change the instance type to a 32-core machine type to reduce replication lag.

Correct Answer: A

QUESTION 23

Your company is running a stateless application on a Compute Engine instance. The application is used heavily during regular business hours and lightly outside of business hours. Users are reporting that the application is slow during peak hours. You need to optimize the application's performance. What should you do?

- A. Create a snapshot of the existing disk. Create an instance template from the snapshot. Create an autoscaled managed instance group from the instance template.
- B. Create a snapshot of the existing disk. Create a custom image from the snapshot. Create an autoscaled managed instance group from the custom image.
- C. Create a custom image from the existing disk. Create an instance template from the custom image. Create an autoscaled managed instance group from the instance template.
- D. Create an instance template from the existing disk. Create a custom image from the instance template. Create an autoscaled managed instance group from the custom image.

Correct Answer: B

Explanation:

<https://cloud.google.com/compute/docs/instance-templates/create-instance-templates>

QUESTION 24

You have an application deployed on Kubernetes Engine using a Deployment named echo-deployment. The deployment is exposed using a Service called echo-service. You need to

[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>