service to fail over to the standby instance in case of a disaster.

Correct Answer: C

QUESTION 51
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?

A. Deploy your application on App Engine standard environment and use App Engine firewall rules to limit access to the open on-premises database.
B. Deploy your application on App Engine standard environment and use Cloud VPN to limit access to the onpremises database.
C. Deploy your application on App Engine flexible environment and use App Engine firewall rules to limit access to the on-premises database.
D. Deploy your application on App Engine flexible environment and use Cloud VPN to limit access to the on-premises database.

Correct Answer: D
Explanation:
https://cloud.google.com/appengine/docs/flexible/python/using-third-party-databases

QUESTION 52
A lead software engineer tells you that his new application design uses websockets and HTTP sessions that are not distributed across the web servers. You want to help him ensure his application will run property on Google Cloud Platform. What should you do?

A. Help the engineer to convert his websocket code to use HTTP streaming.
B. Review the encryption requirements for websocket connections with the security team.
C. Meet with the cloud operations team and the engineer to discuss load balancer options.
D. Help the engineer redesign the application to use a distributed user session service that does not rely on websockets and HTTP sessions.

Correct Answer: C
Explanation:
Google Cloud Platform (GCP) HTTP(S) load balancing provides global load balancing for HTTP(S) requests destined for your instances.

The HTTP(S) load balancer has native support for the WebSocket protocol.

QUESTION 53
Your BigQuery project has several users. For audit purposes, you need to see how many queries each user ran in the last month.

A. Connect Google Data Studio to BigQuery. Create a dimension for the users and a metric for the amount of queries per user.
B. In the BigQuery interface, execute a query on the JOBS table to get the required information.
C. Use `bq show' to list all jobs. Per job, use `bq Is' to list job information and get the required information.
D. Use Cloud Audit Logging to view Cloud Audit Logs, and create a filter on the query operation to get the required information.

Correct Answer: C
Explanation:
https://cloud.google.com/bigquery/docs/managing-jobs

QUESTION 54
Your company is planning to upload several important files to Cloud Storage. After the upload is completed, they want to verify that the upload content is identical to what they have on- premises. You want to minimize the cost and effort of performing this check. What should you do?

A. 1. Use gsutil -m to upload all the files to Cloud Storage.
   2. Use gsutil cp to download the uploaded files
   3. Use Linux diff to compare the content of the files
B. 1. Use gsutil -m to upload all the files to Cloud Storage.
   2. Develop a custom Java application that computes CRC32C hashes
   3. Use gsutil ls -L gs://[YOUR_BUCKET_NAME] to collect CRC32C hashes of the uploaded files
   4. Compare the hashes
C. 1. Use Linux shasum to compute a digest of files you want to upload
   2. Use gsutil -m to upload all the files to the Cloud Storage
   3. Use gsutil cp to download the uploaded files
   4. Use Linux shasum to compute a digest of the downloaded files 5.Compre the hashes
D. 1. Use gsutil -m to upload all the files to Cloud Storage.
   2. Use gsutil hash -c FILE_NAME to generate CRC32C hashes of all on-premises files
   3. Use gsutil ls -L gs://[YOUR_BUCKET_NAME] to collect CRC32C hashes of the uploaded files
   4. Compare the hashes

Correct Answer: D
Explanation:
https://cloud.google.com/storage/docs/gsutil/commands/hash

QUESTION 55
You are working at an institution that processes medical data. You are migrating several workloads onto Google Cloud. Company policies require all workloads to run on physically separated hardware, and workloads from different clients must also be separated You created a sole-tenant node group and added a node for each client. You need to deploy the workloads on these dedicated hosts. What should you do?

A. Add the node group name as a network tag when creating Compute Engine instances in order to host each workload on the correct node group.
B. Add the node name as a network tag when creating Compute Engine instances in order to host each workload on the correct node.
C. Use node affinity labels based on the node group name when creating Compute Engine instances in order to host each workload on the correct node group
D. Use node affinity labels based on the node name when creating Compute Engine instances in order to host each workload on the correct node.

Correct Answer: C
Explanation:
https://cloud.google.com/compute/docs/nodes/provisioning-sole-tenant-vms#provision_a_sole-tenant_vm
https://cloud.google.com/compute/docs/nodes/provisioning-sole-tenant-vms#gcloud_2

When you create a VM, you request sole-tenancy by specifying node affinity or anti-affinity, referencing one or more node affinity labels. You specify custom node affinity labels when you create a node template, and Compute Engine automatically includes some default affinity labels on each node. By specifying affinity when you create a VM, you can schedule VMs together on a specific node or nodes in a node group. By specifying anti-affinity when you create a VM, you can ensure that certain VMs are not scheduled together on the same node or nodes in a node group.

QUESTION 56
You are moving an application that uses MySQL from on-premises to Google Cloud. The application will run on Compute Engine and will use Cloud SQL. You want to cut over to the Compute Engine deployment of the application with minimal downtime and no data loss to your customers. You want to migrate the application with minimal modification. You also need to determine the cutover strategy. What should you do?

A. 1. Set up Cloud VPN to provide private network connectivity between the Compute Engine application and the on-premises MySQL server.
2. Stop the on-premises application.
3. Create a mysqldump of the on-premises MySQL server.
4. Upload the dump to a Cloud Storage bucket.
5. Import the dump into Cloud SQL.
6. Modify the source code of the application to write queries to both databases and read from its local database.
7. Start the Compute Engine application.
8. Stop the on-premises application.

B. 1. Set up Cloud SQL proxy and MySQL proxy.
2. Create a mysqldump of the on-premises MySQL server.
3. Upload the dump to a Cloud Storage bucket.
4. Import the dump into Cloud SQL.
5. Stop the on-premises application.
6. Start the Compute Engine application.

C. 1. Set up Cloud VPN to provide private network connectivity between the Compute Engine application and the on-premises MySQL server.
2. Stop the on-premises application.
3. Start the Compute Engine application, configured to read and write to the on-premises MySQL server.
4. Create the replication configuration in Cloud SQL.
5. Configure the source database server to accept connections from the Cloud SQL replica.
6. Finalize the Cloud SQL replica configuration.
7. When replication has been completed, stop the Compute Engine application.
8. Promote the Cloud SQL replica to a standalone instance.
9. Restart the Compute Engine application, configured to read and write to the Cloud SQL standalone instance.

D. 1. Stop the on-premises application.
2. Create a mysqldump of the on-premises MySQL server.
3. Upload the dump to a Cloud Storage bucket.
4. Import the dump into Cloud SQL.
5. Start the application on Compute Engine.

Correct Answer: C
Explanation:
External replica promotion migration In the migration strategy of external replica promotion, you create an external database replica and synchronize the existing data to that replica. This can happen with minimal downtime to the existing database. When you have a replica database, the

two databases have different roles that are referred to in this document as primary and replica. After the data is synchronized, you promote the replica to be the primary in order to move the management layer with minimal impact to database uptime. In Cloud SQL, an easy way to accomplish the external replica promotion is to use the automated migration workflow. This process automates many of the steps that are needed for this type of migration.

https://cloud.google.com/architecture/migrating-mysql-to-cloudsql-concept

The best option for migrating your MySQL database is to use an external replica promotion. In this strategy, you create a replica database and set your existing database as the primary. You wait until the two databases are in sync, and you then promote your MySQL replica database to be the primary. This process minimizes database downtime related to the database migration. - https://cloud.google.com/architecture/migrating-mysql-to-cloudsql-concept#external_replica_promotion_migration

QUESTION 57
You are analyzing and defining business processes to support your startup's trial usage of GCP, and you don't yet know what consumer demand for your product will be. Your manager requires you to minimize GCP service costs and adhere to Google best practices. What should you do?

A. Utilize free tier and sustained use discounts. Provision a staff position for service cost management.
B. Utilize free tier and sustained use discounts. Provide training to the team about service cost management.
C. Utilize free tier and committed use discounts. Provision a staff position for service cost management.
D. Utilize free tier and committed use discounts. Provide training to the team about service cost management.

Correct Answer: D
Explanation:
https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#billing_and_management

QUESTION 58
You are designing a mobile chat application. You want to ensure people cannot spoof chat messages, by providing a message were sent by a specific user. What should you do

A. Tag messages client side with the originating user identifier and the destination user.
B. Encrypt the message client side using block-based encryption with a shared key.
C. Use public key infrastructure (PKI) to encrypt the message client side using the originating user's private key.
D. Use a trusted certificate authority to enable SSL connectivity between the client application and the server.

Correct Answer: C

QUESTION 59
Your company has just acquired another company, and you have been asked to integrate their existing Google Cloud environment into your company's data center. Upon investigation, you discover that some of the RFC 1918 IP ranges being used in the new company's Virtual Private

Cloud (VPC) overlap with your data center IP space. What should you do to enable connectivity and make sure that there are no routing conflicts when connectivity is established?

A. Create a Cloud VPN connection from the new VPC to the data center, create a Cloud Router, and apply new IP addresses so there is no overlapping IP space.
B. Create a Cloud VPN connection from the new VPC to the data center, and create a Cloud NAT instance to perform NAT on the overlapping IP space.
C. Create a Cloud VPN connection from the new VPC to the data center, create a Cloud Router, and apply a custom route advertisement to block the overlapping IP space.
D. Create a Cloud VPN connection from the new VPC to the data center, and apply a firewall rule that blocks the overlapping IP space.

Correct Answer: A
Explanation:
To connect two networks together we need (1) either VPN or interconnect and (2) peering. When there is peering, you cannot have conflicting IP addresses. You can use either Cloud VPN or Cloud Interconnect to securely connect your on-premises network to your VPC network.

https://cloud.google.com/vpc/docs/vpc-peering#transit-network

At the time of peering, Google Cloud checks to see if there are any subnet IP ranges that overlap subnet IP ranges in the other network. If there is any overlap, peering is not established.

https://cloud.google.com/vpc/docs/vpc-peering#considerations

NAT is used to translate private to public IP and vice versa, however because we are connecting 2 networks together, they become private IPs. So it is not applicable.

QUESTION 60
Your company's test suite is a custom C++ application that runs tests throughout each day on Linux virtual machines. The full test suite takes several hours to complete, running on a limited number of on premises servers reserved for testing. Your company wants to move the testing infrastructure to the cloud, to reduce the amount of time it takes to fully test a change to the system, while changing the tests as little as possible. Which cloud infrastructure should you recommend?

A. Google Compute Engine unmanaged instance groups and Network Load Balancer
B. Google Compute Engine managed instance groups with auto-scaling
C. Google Cloud Dataproc to run Apache Hadoop jobs to process each test
D. Google App Engine with Google Stackdriver for logging

Correct Answer: B
Explanation:
https://cloud.google.com/compute/docs/instance-groups/

Google Compute Engine enables users to launch virtual machines (VMs) on demand. VMs can be launched from the standard images or custom images created by users.

Managed instance groups offer autoscaling capabilities that allow you to automatically add or remove instances from a managed instance group based on increases or decreases in load. Autoscaling helps your applications gracefully handle increases in traffic and reduces cost when the need for resources is lower.