Exam Questions Professional-Cloud-Architect

Google Certified Professional - Cloud Architect (GCP)

https://www.2passeasy.com/dumps/Professional-Cloud-Architect/

NEW QUESTION 1
- (Exam Topic 1)
For this question, refer to the Mountkirk Games case study
Mountkirk Games needs to create a repeatable and configurable mechanism for deploying isolated application environments. Developers and testers can access each other's environments and resources, but they cannot access staging or production resources. The staging environment needs access to some services from production.
What should you do to isolate development environments from staging and production?

A. Create a project for development and test and another for staging and production.
B. Create a network for development and test and another for staging and production.
C. Create one subnetwork for development and another for staging and production.
D. Create one project for development, a second for staging and a third for production.

Answer: D

NEW QUESTION 2
- (Exam Topic 1)
For this question, refer to the Mountkirk Games case study.
Mountkirk Games wants you to design their new testing strategy. How should the test coverage differ from their existing backends on the other platforms?

A. Tests should scale well beyond the prior approaches.
B. Unit tests are no longer required, only end-to-end tests.
C. Tests should be applied after the release is in the production environment.
D. Tests should include directly testing the Google Cloud Platform (GCP) infrastructure.

Answer: A

Explanation:
 From Scenario:
A few of their games were more popular than expected, and they had problems scaling their application servers, MySQL databases, and analytics tools.
Requirements for Game Analytics Platform include: Dynamically scale up or down based on game activity

NEW QUESTION 3
- (Exam Topic 1)
For this question, refer to the Mountkirk Games case study.
Mountkirk Games wants to set up a real-time analytics platform for their new game. The new platform must meet their technical requirements. Which combination of Google technologies will meet all of their requirements?

A. Container Engine, Cloud Pub/Sub, and Cloud SQL
B. Cloud Dataflow, Cloud Storage, Cloud Pub/Sub, and BigQuery
C. Cloud SQL, Cloud Storage, Cloud Pub/Sub, and Cloud Dataflow
D. Cloud Dataproc, Cloud Pub/Sub, Cloud SQL, and Cloud Dataflow
E. Cloud Pub/Sub, Compute Engine, Cloud Storage, and Cloud Dataproc

Answer: B

Explanation:
A real time requires Stream / Messaging so Pub/Sub, Analytics by Big Query.
Ingest millions of streaming events per second from anywhere in the world with Cloud Pub/Sub, powered by Google's unique, high-speed private network. Process the streams with Cloud Dataflow to ensure reliable, exactly-once, low-latency data transformation. Stream the transformed data into BigQuery, the cloud-native data warehousing service, for immediate analysis via SQL or popular visualization tools.
From scenario: They plan to deploy the game's backend on Google Compute Engine so they can capture streaming metrics, run intensive analytics.
Requirements for Game Analytics Platform
≫ Dynamically scale up or down based on game activity
≫ Process incoming data on the fly directly from the game servers
≫ Process data that arrives late because of slow mobile networks
≫ Allow SQL queries to access at least 10 TB of historical data
≫ Process files that are regularly uploaded by users' mobile devices
≫ Use only fully managed services
References: https://cloud.google.com/solutions/big-data/stream-analytics/

NEW QUESTION 4
- (Exam Topic 2)
For this question, refer to the TerramEarth case study.
TerramEarth plans to connect all 20 million vehicles in the field to the cloud. This increases the volume to 20 million 600 byte records a second for 40 TB an hour. How should you design the data ingestion?

A. Vehicles write data directly to GCS.
B. Vehicles write data directly to Google Cloud Pub/Sub.
C. Vehicles stream data directly to Google BigQuery.
D. Vehicles continue to write data using the existing system (FTP).

Answer: B

Explanation:
https://cloud.google.com/solutions/data-lifecycle-cloud-platform

https://cloud.google.com/solutions/designing-connected-vehicle-platform

NEW QUESTION 5
- (Exam Topic 2)
For this question, refer to the TerramEarth case study.
TerramEarth has equipped unconnected trucks with servers and sensors to collet telemetry data. Next year they want to use the data to train machine learning models. They want to store this data in the cloud while reducing costs. What should they do?

A. Have the vehicle' computer compress the data in hourly snapshots, and store it in a Google Cloud storage (GCS) Nearline bucket.
B. Push the telemetry data in Real-time to a streaming dataflow job that compresses the data, and store it in Google BigQuery.
C. Push the telemetry data in real-time to a streaming dataflow job that compresses the data, and store it in Cloud Bigtable.
D. Have the vehicle's computer compress the data in hourly snapshots, a Store it in a GCS Coldline bucket.

Answer: D

Explanation:
Coldline Storage is the best choice for data that you plan to access at most once a year, due to its slightly lower availability, 90-day minimum storage duration, costs for data access, and higher per-operation costs. For example:
Cold Data Storage - Infrequently accessed data, such as data stored for legal or regulatory reasons, can be stored at low cost as Coldline Storage, and be available when you need it.
Disaster recovery - In the event of a disaster recovery event, recovery time is key. Cloud Storage provides low latency access to data stored as Coldline Storage.
References: https://cloud.google.com/storage/docs/storage-classes

NEW QUESTION 6
- (Exam Topic 2)
Your agricultural division is experimenting with fully autonomous vehicles. You want your architecture to promote strong security during vehicle operation. Which two architecture should you consider?
Choose 2 answers:

A. Treat every micro service call between modules on the vehicle as untrusted.
B. Require IPv6 for connectivity to ensure a secure address space.
C. Use a trusted platform module (TPM) and verify firmware and binaries on boot.
D. Use a functional programming language to isolate code execution cycles.
E. Use multiple connectivity subsystems for redundancy.
F. Enclose the vehicle's drive electronics in a Faraday cage to isolate chips.

Answer: AC

NEW QUESTION 7
- (Exam Topic 2)
For this question, refer to the TerramEarth case study
Your development team has created a structured API to retrieve vehicle data. They want to allow third parties to develop tools for dealerships that use this vehicle event data. You want to support delegated authorization against this data. What should you do?

A. Build or leverage an OAuth-compatible access control system.
B. Build SAML 2.0 SSO compatibility into your authentication system.
C. Restrict data access based on the source IP address of the partner systems.
D. Create secondary credentials for each dealer that can be given to the trusted third party.

Answer: A

Explanation:
https://cloud.google.com/appengine/docs/flexible/go/authorizing-apps https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#delegate_application_autho Delegate application authorization with OAuth2
Cloud Platform APIs support OAuth 2.0, and scopes provide granular authorization over the methods that are supported. Cloud Platform supports both service-account and user-account OAuth, also called three-legged OAuth.
References:
https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#delegate_application_autho
https://cloud.google.com/appengine/docs/flexible/go/authorizing-apps

NEW QUESTION 8
- (Exam Topic 2)
For this question, refer to the TerramEarth case study.
TerramEarth's 20 million vehicles are scattered around the world. Based on the vehicle's location its telemetry data is stored in a Google Cloud Storage (GCS) regional bucket (US. Europe, or Asia). The CTO has asked you to run a report on the raw telemetry data to determine why vehicles are breaking down after 100 K miles. You want to run this job on all the data. What is the most cost-effective way to run this job?

A. Move all the data into 1 zone, then launch a Cloud Dataproc cluster to run the job.
B. Move all the data into 1 region, then launch a Google Cloud Dataproc cluster to run the job.
C. Launch a cluster in each region to preprocess and compress the raw data, then move the data into a multi region bucket and use a Dataproc cluster to finish the job.
D. Launch a cluster in each region to preprocess and compress the raw data, then move the data into a region bucket and use a Cloud Dataproc cluster to finish the jo

Answer: D

Explanation:
Storageguarantees 2 replicates which are geo diverse (100 miles apart) which can get better remote latency and availability.
More importantly, is that multiregional heavily leverages Edge caching and CDNs to provide the content to the end users.

All this redundancy and caching means that Multiregional comes with overhead to sync and ensure consistency between geo-diverse areas. As such, it's much better for write-once-read-many scenarios. This means frequently accessed (e.g. "hot" objects) around the world, such as website content, streaming videos, gaming or mobile applications.
References:
https://medium.com/google-cloud/google-cloud-storage-what-bucket-class-for-the-best-performance-5c847ac8f

NEW QUESTION 9
- (Exam Topic 2)
For this question refer to the TerramEarth case study
Operational parameters such as oil pressure are adjustable on each of TerramEarth's vehicles to increase their efficiency, depending on their environmental conditions. Your primary goal is to increase the operating efficiency of all 20 million cellular and unconnected vehicles in the field How can you accomplish this goal?

A. Have your engineers inspect the data for patterns, and then create an algorithm with rules that makeoperational adjustments automatically.
B. Capture all operating data, train machine learning models that identify ideal operations, and run locally to make operational adjustments automatically.
C. Implement a Google Cloud Dataflow streaming job with a sliding window, and use Google Cloud Messaging (GCM) to make operational adjustments automatically.
D. Capture all operating data, train machine learning models that identify ideal operations, and host in Google Cloud Machine Learning (ML) Platform to make operational adjustments automatically.

Answer: B

NEW QUESTION 10
- (Exam Topic 3)
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?

A. Cloud Spanner
B. Google BigQuery
C. Google Cloud SQL
D. Google Cloud Datastore

Answer: D

Explanation:
https://cloud.google.com/datastore/docs/concepts/overview Common workloads for Google Cloud Datastore:
≫ User profiles
≫ Product catalogs
≫ Game state
References: https://cloud.google.com/storage-options/ https://cloud.google.com/datastore/docs/concepts/overview

NEW QUESTION 10
- (Exam Topic 3)
For this question, refer to the JencoMart case study.
The JencoMart security team requires that all Google Cloud Platform infrastructure is deployed using a least privilege model with separation of duties for administration between production and development resources. What Google domain and project structure should you recommend?

A. Create two G Suite accounts to manage users: one for development/test/staging and one for production.Each account should contain one project for every application.
B. Create two G Suite accounts to manage users: one with a single project for all development applications and one with a single project for all production applications.
C. Create a single G Suite account to manage users with each stage of each application in its own project.
D. Create a single G Suite account to manage users with one project for the development/test/staging environment and one project for the production environment.

Answer: D

Explanation:
Note: The principle of least privilege and separation of duties are concepts that, although semantically different, are intrinsically related from the standpoint of security. The intent behind both is to prevent people from having higher privilege levels than they actually need
≫ Principle of Least Privilege: Users should only have the least amount of privileges required to perform their job and no more. This reduces authorization exploitation by limiting access to resources such as targets, jobs, or monitoring templates for which they are not authorized.
≫ Separation of Duties: Beyond limiting user privilege level, you also limit user duties, or the specific jobs they can perform. No user should be given responsibility for more than one related function. This limits the ability of a user to perform a malicious action and then cover up that action.
References: https://cloud.google.com/kms/docs/separation-of-duties

NEW QUESTION 11
- (Exam Topic 3)
For this question, refer to the JencoMart case study.
JencoMart has decided to migrate user profile storage to Google Cloud Datastore and the application servers to Google Compute Engine (GCE). During the migration, the existing infrastructure will need access to Datastore to upload the data. What service account key-management strategy should you recommend?

A. Provision service account keys for the on-premises infrastructure and for the GCE virtual machines (VMs).
B. Authenticate the on-premises infrastructure with a user account and provision service account keys for the VMs.
C. Provision service account keys for the on-premises infrastructure and use Google Cloud Platform (GCP) managed keys for the VMs
D. Deploy a custom authentication service on GCE/Google Container Engine (GKE) for the on-premises infrastructure and use GCP managed keys for the VMs.

Answer: A

Explanation:
https://cloud.google.com/iam/docs/understanding-service-accounts
Migrating data to Google Cloud Platform
Let's say that you have some data processing that happens on another cloud provider and you want to transfer the processed data to Google Cloud Platform. You can use a service account from the virtual machines on the external cloud to push the data to Google Cloud Platform. To do this, you must create and download a service account key when you create the service account and then use that key from the external process to call the Cloud Platform APIs.
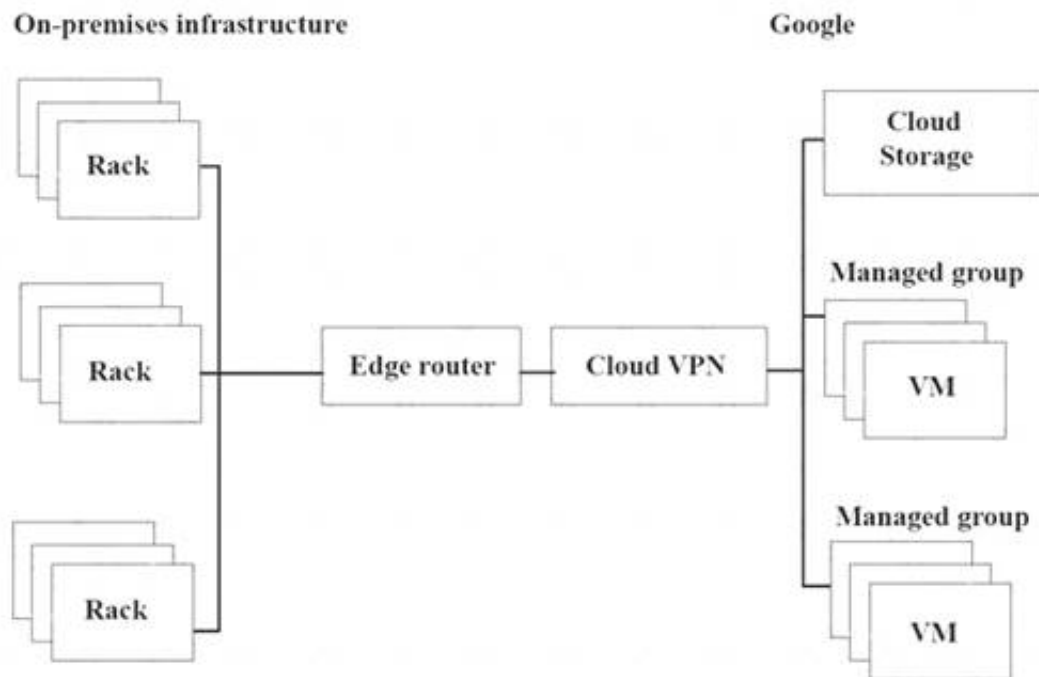References:
https://cloud.google.com/iam/docs/understanding-service-accounts#migrating_data_to_google_cloud_platform

NEW QUESTION 14
- (Exam Topic 3)
For this question, refer to the JencoMart case study.
The migration of JencoMart's application to Google Cloud Platform (GCP) is progressing too slowly. The infrastructure is shown in the diagram. You want to maximize throughput. What are three potential bottlenecks? (Choose 3 answers.)

A. A single VPN tunnel, which limits throughput
B. A tier of Google Cloud Storage that is not suited for this task
C. A copy command that is not suited to operate over long distances
D. Fewer virtual machines (VMs) in GCP than on-premises machines
E. A separate storage layer outside the VMs, which is not suited for this task
F. Complicated internet connectivity between the on-premises infrastructure and GCP

Answer: ADF

NEW QUESTION 18
- (Exam Topic 4)
For this question, refer to the Dress4Win case study.
Dress4Win has asked you for advice on how to migrate their on-premises MySQL deployment to the cloud. They want to minimize downtime and performance impact to their on-premises solution during the migration. Which approach should you recommend?

A. Create a dump of the on-premises MySQL master server, and then shut it down, upload it to the cloud environment, and load into a new MySQL cluster.
B. Setup a MySQL replica server/slave in the cloud environment, and configure it for asynchronous replication from the MySQL master server on-premises until cutover.
C. Create a new MySQL cluster in the cloud, configure applications to begin writing to both on-premises and cloud MySQL masters, and destroy the original cluster at cutover.
D. Create a dump of the MySQL replica server into the cloud environment, load it into: Google Cloud Datastore, and configure applications to read/write to Cloud Datastore at cutover.

Answer: B

NEW QUESTION 20
- (Exam Topic 4)
For this question, refer to the Dress4Win case study.
You want to ensure Dress4Win's sales and tax records remain available for infrequent viewing by auditors for at least 10 years. Cost optimization is your top priority. Which cloud services should you choose?

A. Google Cloud Storage Coldline to store the data, and gsutil to access the data.
B. Google Cloud Storage Nearline to store the data, and gsutil to access the data.
C. Google Bigtable with US or EU as location to store the data, and gcloud to access the data.
D. BigQuery to store the data, and a web server cluster in a managed instance group to access the data.Google Cloud SQL mirrored across two distinct regions to store the data, and a Redis cluster in a managed instance group to access the data.

Answer: A

Explanation:
References: https://cloud.google.com/storage/docs/storage-classes

NEW QUESTION 22
- (Exam Topic 4)
Dress4win has end to end tests covering 100% of their endpoints.
They want to ensure that the move of cloud does not introduce any new bugs.
Which additional testing methods should the developers employ to prevent an outage?

A. They should run the end to end tests in the cloud staging environment to determine if the code is working as intended.
B. They should enable google stack driver debugger on the application code to show errors in the code
C. They should add additional unit tests and production scale load tests on their cloud staging environment.
D. They should add canary tests so developers can measure how much of an impact the new release causes to latency

Answer: B

NEW QUESTION 24
- (Exam Topic 4)
For this question, refer to the Dress4Win case study.
Dress4Win has end-to-end tests covering 100% of their endpoints. They want to ensure that the move to the cloud does not introduce any new bugs. Which additional testing methods should the developers employ to prevent an outage?

A. They should enable Google Stackdriver Debugger on the application code to show errors in the code.
B. They should add additional unit tests and production scale load tests on their cloud staging environment.
C. They should run the end-to-end tests in the cloud staging environment to determine if the code is working as intended.
D. They should add canary tests so developers can measure how much of an impact the new release causes to latency.

Answer: B

NEW QUESTION 29
- (Exam Topic 4)
For this question, refer to the Dress4Win case study.
As part of their new application experience, Dress4Wm allows customers to upload images of themselves. The customer has exclusive control over who may view these images. Customers should be able to upload images with minimal latency and also be shown their images quickly on the main application page when they log in. Which configuration should Dress4Win use?

A. Store image files in a Google Cloud Storage bucke
B. Use Google Cloud Datastore to maintain metadata that maps each customer's ID and their image files.
C. Store image files in a Google Cloud Storage bucke
D. Add custom metadata to the uploaded images in Cloud Storage that contains the customer's unique ID.
E. Use a distributed file system to store customers' image
F. As storage needs increase, add more persistent disks and/or node
G. Assign each customer a unique ID, which sets each file's owner attribute, ensuring privacy of images.
H. Use a distributed file system to store customers' image
I. As storage needs increase, add more persistent disks and/or node
J. Use a Google Cloud SQL database to maintain metadata that maps each customer's ID to their image files.

Answer: A

NEW QUESTION 32
- (Exam Topic 5)
As part of implementing their disaster recovery plan, your company is trying to replicate their production MySQL database from their private data center to their GCP project using a Google Cloud VPN connection. They are experiencing latency issues and a small amount of packet loss that is disrupting the replication.
What
should they do?

A. Configure their replication to use UDP.
B. Configure a Google Cloud Dedicated Interconnect.
C. Restore their database daily using Google Cloud SQL.
D. Add additional VPN connections and load balance them.
E. Send the replicated transaction to Google Cloud Pub/Sub.

Answer: B

NEW QUESTION 34
- (Exam Topic 5)
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
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

Answer: B

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

NEW QUESTION 38
- (Exam Topic 5)

You deploy your custom Java application to Google App Engine. It fails to deploy and gives you the following stack trace.

```
java.lang.SecurityException: SHA1 digest error for
com/Altostrat/CloakedServlet.class
        at com.google.appengine.runtime.Request.process
-d36f818a24b8cf1d (Request.java)
        at
sun.security.util.ManifestEntryVerifier.verify
(ManifestEntryVerifier.java:210)
        at java.util.jar.JarVerifier.processEntry
(JarVerifier.java:218)
        at java.util.jar.JarVerifier.update
(JarVerifier.java:205)
        at
java.util.jar.JarVerifiersVerifierStream.read
(JarVerifier.java:428)
        at sun.misc.Resource.getBytes
(Resource.java:124)
        at java.net.URL.ClassLoader.defineClass
(URLClassLoader.java:273)
        at sun.reflect.GeneratedMethodAccessor5.invoke
(Unknown Source)
        at
sun.reflect.DelegatingMethodAccessorImpl.invoke
(DelegatingMethodAccessorImpl.java:43)
        at java.lang.reflect.Method.invoke
(Method.java:616)
        at java.lang.ClassLoader.loadClass
(ClassLoader.java:266)
```

What should you do?

A. Upload missing JAR files and redeploy your application.
B. Digitally sign all of your JAR files and redeploy your application
C. Recompile the CLoakedServlet class using and MD5 hash instead of SHA1

Answer: B

NEW QUESTION 42
- (Exam Topic 5)
Your company has an enterprise application running on Compute Engine that requires high availability and high performance. The application has been deployed on two instances in two zones in the same region m active passive mode. The application writes data to a persistent disk in the case of a single zone outage that data should be immediately made available to the other instance in the other zone. You want to maximize performance while minimizing downtime and data loss. What should you do?

A. * 1. Attach a persistent SSD disk to the first instance* 2. Create a snapshot every hour* 3. In case of a zone outage, recreate a persistent SSD disk in the second instance where data is coming from the created snapshot
B. * 1 Create a Cloud Storage bucket* 2. Mount the bucket into the first instance with gcs-fuse* 3. In case of a zone outage, mount the Cloud Storage bucket to the second instance with gcs-fuse
C. * 1 Attach a local SSD lo the first instance disk* 2. Execute an rsync command every hour where the target is a persistent SSD disk attached to the second instance* 3. In case of a zone outage, use the second instance
D. * 1. Attach a regional SSD persistent Ask to the first instance* 2. In case of a zone outage, force-attach the disk to the other instance

Answer: D

NEW QUESTION 43
- (Exam Topic 5)
You want to enable your running Google Kubernetes Engine cluster to scale as demand for your application changes.
What should you do?

A. Add additional nodes to your Kubernetes Engine cluster using the following command:gcloud container clusters resizeCLUSTER_Name – -size 10
B. Add a tag to the instances in the cluster with the following command:gcloud compute instances add-tagsINSTANCE - -tags enable-autoscaling max-nodes-10
C. Update the existing Kubernetes Engine cluster with the following command:gcloud alpha container clustersupdate mycluster - -enable-autoscaling - -min-nodes=1 - -max-nodes=10
D. Create a new Kubernetes Engine cluster with the following command:gcloud alpha container clusterscreate mycluster - -enable-autoscaling - -min-nodes=1 - -max-nodes=10and redeploy your application

Answer: C

Explanation:
https://cloud.google.com/kubernetes-engine/docs/concepts/cluster-autoscaler To enable autoscaling for an existing node pool, run the following command:
gcloud container clusters update [CLUSTER_NAME] --enable-autoscaling \--min-nodes 1 --max-nodes 10
--zone [COMPUTE_ZONE] --node-pool default-pool

NEW QUESTION 47
- (Exam Topic 5)
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.

Answer: A

NEW QUESTION 48
- (Exam Topic 5)
Your company has decided to make a major revision of their API in order to create better experiences for their developers. They need to keep the old version of the API available and deployable, while allowing new customers and testers to try out the new API. They want to keep the same SSL and DNS records in place to serve both APIs. What should they do?

A. Configure a new load balancer for the new version of the API.
B. Reconfigure old clients to use a new endpoint for the new API.
C. Have the old API forward traffic to the new API based on the path.
D. Use separate backend pools for each API path behind the load balancer.

Answer: D

Explanation:
https://cloud.google.com/endpoints/docs/openapi/lifecycle-management

NEW QUESTION 52
- (Exam Topic 5)
You have an App Engine application that needs to be updated. You want to test the update with production traffic before replacing the current application version. What should you do?

A. Deploy the update using the Instance Group Updater to create a partial rollout, which allows for canary testing.
B. Deploy the update as a new version in the App Engine application, and split traffic between the new and current versions.
C. Deploy the update in a new VPC, and use Google's global HTTP load balancing to split traffic between the update and current applications.
D. Deploy the update as a new App Engine application, and use Google's global HTTP load balancing to split traffic between the new and current applications.

Answer: B

Explanation:
https://cloud.google.com/appengine/docs/standard/python/splitting-traffic

NEW QUESTION 53
- (Exam Topic 5)
Your customer wants to capture multiple GBs of aggregate real-time key performance indicators (KPIs) from their game servers running on Google Cloud Platform and monitor the KPIs with low latency. How should they capture the KPIs?

A. Store time-series data from the game servers in Google Bigtable, and view it using Google Data Studio.
B. Output custom metrics to Stackdriver from the game servers, and create a Dashboard in Stackdriver Monitoring Console to view them.
C. Schedule BigQuery load jobs to ingest analytics files uploaded to Cloud Storage every ten minutes, and visualize the results in Google Data Studio.
D. Insert the KPIs into Cloud Datastore entities, and run ad hoc analysis and visualizations of them in Cloud Datalab.

Answer: A

Explanation:
https://cloud.google.com/monitoring/api/v3/metrics-details#metric-kinds

NEW QUESTION 56
- (Exam Topic 5)
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?

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

Explanation:
Reference: https://cloud.google.com/compute/docs/instance-groups

NEW QUESTION 60
- (Exam Topic 5)
You deploy your custom java application to google app engine. It fails to deploy and gives you the following stack trace:

```
Java.lang.securityException : SHA1 digest

At com.google.appengine.runtime.Request.pr

At

Sun.securityutil.manifestENtryVerifier.ver

At java . net . URLClassLoader . defineCla

At sun . reflect. GeneratedMethodAccessors

At

Sun.reflect . DelegatingMethodAccesorImpl.

At java . lang . reflect . MThod . invoke
```

A. Recompile the CLoakedServlet class using and MD5 hash instead of SHA1
B. Digitally sign all of your JAR files and redeploy your application.
C. Upload missing JAR files and redeploy your application

Answer: B

NEW QUESTION 64
- (Exam Topic 5)
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 BigQuer
B. Create a dimension for the users and a metric for the amount of queries per user.
C. In the BigQuery interface, execute a query on the JOBS table to get the required information.
D. Use 'bq show' to list all job
E. Per job, use 'bq Is' to list job information and get the required information.
F. Use Cloud Audit Logging to view Cloud Audit Logs, and create a filter on the query operation to get the required information.

Answer: C

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

NEW QUESTION 69
- (Exam Topic 5)
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?

A. * 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.
B. * 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.
C. * 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.
D. * 1. Create a Cloud VPN to the office network.* 2. Configure Private Google Access for on-premises hosts.

Answer: A

Explanation:
For all Google Cloud services secured with VPC Service Controls, you can ensure that: Resources within a perimeter are accessed only from clients within authorized VPC networks using Private Google Access with either Google Cloud or on-premises. https://cloud.google.com/vpc-service-controls/docs/overview https://cloud.google.com/vpc-service-controls/docs/overview. You create a service control across your VPC and any cloud bucket or any project resource to restrict access. Anything outside of it can't access the resources within service control perimeter

NEW QUESTION 74
- (Exam Topic 5)
Your company is designing its data lake on Google Cloud and wants to develop different ingestion pipelines to collect unstructured data from different sources. After the data is stored in Google Cloud, it will be processed in several data pipelines to build a recommendation engine for end users on the website. The structure of the data retrieved from the source systems can change at any time. The data must be stored exactly as it was retrieved for reprocessing purposes in case the data structure is incompatible with the current processing pipelines. You need to design an architecture to support the use case after you retrieve the data. What should you do?

A. Send the data through the processing pipeline, and then store the processed data in a BigQuery table for reprocessing.
B. Store the data in a BigQuery tabl
C. Design the processing pipelines to retrieve the data from the table.
D. Send the data through the processing pipeline, and then store the processed data in a Cloud Storage bucket for reprocessing.
E. Store the data in a Cloud Storage bucke
F. Design the processing pipelines to retrieve the data from the bucket

Answer: D

NEW QUESTION 76
- (Exam Topic 5)
You team needs to create a Google Kubernetes Engine (GKE) cluster to host a newly built application that requires access to third-party services on the internet. Your company does not allow any Compute Engine instance to have a public IP address on Google Cloud. You need to create a deployment strategy that adheres to these guidelines. What should you do?

A. Create a Compute Engine instance, and install a NAT Proxy on the instanc
B. Configure all workloads on GKE to pass through this proxy to access third-party services on the Internet
C. Configure the GKE cluster as a private cluster, and configure Cloud NAT Gateway for the cluster subnet
D. Configure the GKE cluster as a route-based cluste
E. Configure Private Google Access on the Virtual Private Cloud (VPC)
F. Configure the GKE cluster as a private cluste
G. Configure Private Google Access on the Virtual Private Cloud (VPC)

Answer: B

Explanation:
A Cloud NAT gateway can perform NAT for nodes and Pods in a private cluster, which is a type of
VPC-native cluster. The Cloud NAT gateway must be configured to apply to at least the following subnet IP address ranges for the subnet that your cluster uses:
Subnet primary IP address range (used by nodes)
Subnet secondary IP address range used for Pods in the cluster Subnet secondary IP address range used for Services in the cluster
The simplest way to provide NAT for an entire private cluster is to configure a Cloud NAT gateway to apply to all of the cluster's subnet's IP address ranges.
https://cloud.google.com/nat/docs/overview

NEW QUESTION 81
- (Exam Topic 5)
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

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.

NEW QUESTION 84
- (Exam Topic 5)
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.

Answer: C

NEW QUESTION 89
- (Exam Topic 5)
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?

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

NEW QUESTION 90
- (Exam Topic 5)
You have been asked to select the storage system for the click-data of your company's large portfolio of websites. This data is streamed in from a custom website analytics package at a typical rate of 6,000 clicks per minute, with bursts of up to 8,500 clicks per second. It must been stored for future analysis by your data science and user experience teams. Which storage infrastructure should you choose?

A. Google Cloud SQL
B. Google Cloud Bigtable
C. Google Cloud Storage
D. Google cloud Datastore

Answer: C

Explanation:
https://cloud.google.com/bigquery/docs/loading-data-cloud-storage

NEW QUESTION 95
- (Exam Topic 5)
You have been engaged by your client to lead the migration of their application infrastructure to GCP. One of their current problems is that the on-premises high performance SAN is requiring frequent and expensive upgrades to keep up with the variety of workloads that are identified as follows: 20TB of log archives retained for legal reasons; 500 GB of VM boot/data volumes and templates; 500 GB of image thumbnails; 200 GB of customer session state data that allows customers to restart sessions even if off-line for several days.
Which of the following best reflects your recommendations for a cost-effective storage allocation?

A. Local SSD for customer session state dat
B. Lifecycle-managed Cloud Storage for log archives, thumbnails, and VM boot/data volumes.
C. Memcache backed by Cloud Datastore for the customer session state dat
D. Lifecycle- managed Cloud Storage for log archives, thumbnails, and VM boot/data volumes.
E. Memcache backed by Cloud SQL for customer session state dat
F. Assorted local SSD-backed instances for VM boot/data volume
G. Cloud Storage for log archives and thumbnails.
H. Memcache backed by Persistent Disk SSD storage for customer session state dat
I. Assorted local SSDbacked instances for VM boot/data volume
J. Cloud Storage for log archives and thumbnails.

Answer: D

Explanation:
https://cloud.google.com/compute/docs/disks

NEW QUESTION 99
- (Exam Topic 5)
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.

Answer: A

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

NEW QUESTION 104
- (Exam Topic 5)
Auditors visit your teams every 12 months and ask to review all the Google Cloud Identity and Access Management (Cloud IAM) policy changes in the previous 12

months. You want to streamline and expedite the analysis and audit process. What should you do?

A. Create custom Google Stackdriver alerts and send them to the auditor.
B. Enable Logging export to Google BigQuery and use ACLs and views to scope the data shared with the auditor.
C. Use cloud functions to transfer log entries to Google Cloud SQL and use ACLS and views to limit an auditor's view.
D. Enable Google Cloud Storage (GCS) log export to audit logs Into a GCS bucket and delegate access to the bucket.

Answer: D

Explanation:
Export the logs to Google Cloud Storage bucket - Archive Storage, as it will not be used for 1 year, price for which is $0.004 per GB per Month. The price for long term storage in BigQuery is $0.01 per GB per Month (250% more). Also for analysis purpose, whenever Auditors are there(once per year), you can use BigQuery and use GCS bucket as external data source. BigQuery supports querying Cloud Storage data from these storage classes:
Standard Nearline Coldline Archive

NEW QUESTION 106
- (Exam Topic 5)
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?

A. Compress and upload both achieved files and files uploaded daily using the qsutil –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 Storag
C. Establish a connection with Google using a Dedicated Interconnect or Direct Peering connection and use it to upload files daily.
D. Lease a Transfer Appliance, upload archived files to it, and send it, and send it to Google to transferarchived data to Cloud Storag
E. Establish one Cloud VPN Tunnel to VPC networks over the public internet, and compares and upload files daily using the gsutil –m option.
F. Lease a Transfer Appliance, upload archived files to it, and send it to Google to transfer archived data to Cloud Storag
G. 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

NEW QUESTION 109
- (Exam Topic 5)
Your organization has stored sensitive data in a Cloud Storage bucket. For regulatory reasons, your company must be able to rotate the encryption key used to encrypt the data in the bucket. The data will be processed in Dataproc. You want to follow Google-recommended practices for security What should you do?

A. Create a key with Cloud Key Management Service (KMS) Encrypt the data using the encrypt method of Cloud KMS.
B. Create a key with Cloud Key Management Service (KMS). Set the encryption key on the bucket to the Cloud KMS key.
C. Generate a GPG key pai
D. Encrypt the data using the GPG ke
E. Upload the encrypted data to the bucket.
F. Generate an AES-256 encryption ke
G. Encrypt the data in the bucket using the customer-supplied encryption keys feature.

Answer: B

Explanation:
https://cloud.google.com/storage/docs/encryption/using-customer-managed-keys#add-object-key https://cloud.google.com/storage/docs/encryption/using-customer-managed-keys

NEW QUESTION 114
- (Exam Topic 5)
Your company sends all Google Cloud logs to Cloud Logging. Your security team wants to monitor the logs. You want to ensure that the security team can react quickly if an anomaly such as an unwanted firewall change or server breach is detected. You want to follow Google-recommended practices. What should you do?

A. Schedule a cron job with Cloud Schedule
B. The scheduled job queries the logs every minute for the relevant events.
C. Export logs to BigQuery, and trigger a query in BigQuery to process the log data for the relevant events.
D. Export logs to a Pub/Sub topic, and trigger Cloud Function with the relevant log events.
E. Export logs to a Cloud Storage bucket, and trigger Cloud Run with the relevant log events.

Answer: C

Explanation:
https://cloud.google.com/blog/products/management-tools/automate-your-response-to-a-cloud-logging-event

NEW QUESTION 115
- (Exam Topic 5)
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.

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.

NEW QUESTION 119
- (Exam Topic 5)
Your company is using Google Cloud. You have two folders under the Organization: Finance and Shopping. The members of the development team are in a Google Group. The development team group has been assigned the Project Owner role on the Organization. You want to prevent the development team from creating resources in projects in the Finance folder. What should you do?

A. Assign the development team group the Project Viewer role on the Finance folder, and assign the development team group the Project Owner role on the Shopping folder.
B. Assign the development team group only the Project Viewer role on the Finance folder.
C. Assign the development team group the Project Owner role on the Shopping folder, and remove the development team group Project Owner role from the Organization.
D. Assign the development team group only the Project Owner role on the Shopping folder.

Answer: C

Explanation:
https://cloud.google.com/resource-manager/docs/cloud-platform-resource-hierarchy
"Roles are always inherited, and there is no way to explicitly remove a permission for a lower-level resource that is granted at a higher level in the resource hierarchy. Given the above example, even if you were to remove the Project Editor role from Bob on the "Test GCP Project", he would still inherit that role from the "Dept Y" folder, so he would still have the permissions for that role on "Test GCP Project"."
Reference: https://cloud.google.com/resource-manager/docs/creating-managing-folders

NEW QUESTION 122
- (Exam Topic 5)
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 discount
B. Provision a staff position for service cost management.
C. Utilize free tier and sustained use discount
D. Provide training to the team about service cost management.
E. Utilize free tier and committed use discount
F. Provision a staff position for service cost management.
G. Utilize free tier and committed use discount
H. Provide training to the team about service cost management.

Answer: D

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

NEW QUESTION 124
- (Exam Topic 5)
Your company has a stateless web API that performs scientific calculations. The web API runs on a single Google Kubernetes Engine (GKE) cluster. The cluster is currently deployed in us-central1. Your company has expanded to offer your API to customers in Asia. You want to reduce the latency for the users in Asia. What should you do?

A. Use a global HTTP(s) load balancer with Cloud CDN enabled
B. Create a second GKE cluster in asia-southeast1, and expose both API's using a Service of type Load Balance
C. Add the public Ips to the Cloud DNS zone
D. Increase the memory and CPU allocated to the application in the cluster
E. Create a second GKE cluster in asia-southeast1, and use kubemci to create a global HTTP(s) load balancer

Answer: D

Explanation:
https://cloud.google.com/kubernetes-engine/docs/concepts/multi-cluster-ingress#how_works
https://github.com/GoogleCloudPlatform/k8s-multicluster-ingress https://cloud.google.com/blog/products/gcp/how-to-deploy-geographically-distributed-services-on-kubernetes-e

NEW QUESTION 125
- (Exam Topic 5)
Your web application must comply with the requirements of the European Union's General Data Protection Regulation (GDPR). You are responsible for the technical architecture of your web application. What should you do?

A. Ensure that your web application only uses native features and services of Google Cloud Platform,because Google already has various certifications and provides "pass-on" compliance when you use native features.
B. Enable the relevant GDPR compliance setting within the GCPConsole for each of the services in use within your application.
C. Ensure that Cloud Security Scanner is part of your test planning strategy in order to pick up any compliance gaps.
D. Define a design for the security of data in your web application that meets GDPR requirements.

Answer: D

Explanation:
https://cloud.google.com/security/gdpr/?tab=tab4
Reference: https://www.mobiloud.com/blog/gdpr-compliant-mobile-app/

NEW QUESTION 129
- (Exam Topic 5)
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?

A. Use the Admin Directory API to authenticate against the Active Directory domain controller.
B. Use Google Cloud Directory Sync to synchronize Active Directory usernames with cloud identities and configure SAML SSO.
C. Use Cloud Identity-Aware Proxy configured to use the on-premises Active Directory domain controller as an identity provider.
D. 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.

Answer: B

Explanation:
https://cloud.google.com/solutions/federating-gcp-with-active-directory-introduction#implementing_federation

NEW QUESTION 133
- (Exam Topic 5)
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?

A. 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 hour
B. Schedule the shell script on one of the production instances to automate the task.
C. 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.
D. Deploy the development and acceptance applications on a managed instance group and enable autoscaling.
E. Use regular Compute Engine instances for the production environment, and use preemptible VMs for the acceptance and development environments.

Answer: B

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

NEW QUESTION 134
- (Exam Topic 5)
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

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

Answer: AC

NEW QUESTION 137
- (Exam Topic 5)
Your company places a high value on being responsive and meeting customer needs quickly. Their primary business objectives are release speed and agility. You want to reduce the chance of security errors being accidentally introduced. Which two actions can you take? Choose 2 answers

A. Ensure every code check-in is peer reviewed by a security SME.
B. Use source code security analyzers as part of the CI/CD pipeline.
C. Ensure you have stubs to unit test all interfaces between components.
D. Enable code signing and a trusted binary repository integrated with your CI/CD pipeline.
E. Run a vulnerability security scanner as part of your continuous-integration /continuous-delivery (CI/CD) pipeline.

Answer: BE

Explanation:
https://docs.microsoft.com/en-us/vsts/articles/security-validation-cicd-pipeline?view=vsts

NEW QUESTION 139
- (Exam Topic 5)
You want to enable your running Google Container Engine cluster to scale as demand for your application changes.
What should you do?

A. Add additional nodes to your Container Engine cluster using the following command: gcloud container clusters resize CLUSTER_NAME --size 10
B. Add a tag to the instances in the cluster with the following command:gcloud compute instances add-tags INSTANCE --tags enable --autoscaling max-nodes-10
C. Update the existing Container Engine cluster with the following command:gcloud alpha container clusters update mycluster --enable-autoscaling --min-nodes=1 --max-nodes=10
D. Create a new Container Engine cluster with the following command:gcloud alpha container clusters create mycluster --enable-autocaling --min-nodes=1 --max-

nodes=10 and redeploy your application.

Answer: B

Explanation:
https://cloud.google.com/kubernetes-engine/docs/concepts/cluster-autoscaler Cluster autoscaling
--enable-autoscaling
Enables autoscaling for a node pool.
Enables autoscaling in the node pool specified by --node-pool or the default node pool if --node-pool is not provided.
Where:
--max-nodes=MAX_NODES
Maximum number of nodes in the node pool.
Maximum number of nodes to which the node pool specified by --node-pool (or default node pool if unspecified) can scale.

NEW QUESTION 143
- (Exam Topic 5)
You want to allow your operations learn to store togs from all the production protects in your Organization, without during logs from other projects All of the production projects are contained in a folder. You want to ensure that all logs for existing and new production projects are captured automatically. What should you do?

A. Create an aggregated export on the Production folde
B. Set the log sink to be a Cloud Storage bucket in an operations project
C. Create an aggregated export on the Organization resourc
D. Set the tog sink to be a Cloud Storage bucket in an operations project.
E. Create log exports in the production project
F. Set the log sinks to be a Cloud Storage bucket in an operations project.
G. Create tog exports in the production project
H. Set the tog sinks to be BigQuery datasets in the production projects and grant IAM access to the operations team to run queries on the datasets

Answer: B

NEW QUESTION 147
- (Exam Topic 5)
Your customer runs a web service used by e-commerce sites to offer product recommendations to users. The company has begun experimenting with a machine learning model on Google Cloud Platform to improve the quality of results.
What should the customer do to improve their model's results over time?

A. Export Cloud Machine Learning Engine performance metrics from Stackdriver to BigQuery, to be used to analyze the efficiency of the model.
B. Build a roadmap to move the machine learning model training from Cloud GPUs to Cloud TPUs, which offer better results.
C. Monitor Compute Engine announcements for availability of newer CPU architectures, and deploy the model to them as soon as they are available for additional performance.
D. Save a history of recommendations and results of the recommendations in BigQuery, to be used as training data.

Answer: D

Explanation:
https://cloud.google.com/solutions/building-a-serverless-ml-model

NEW QUESTION 148
- (Exam Topic 5)
You write a Python script to connect to Google BigQuery from a Google Compute Engine virtual machine. The script is printing errors that it cannot connect to BigQuery. What should you do to fix the script?

A. Install the latest BigQuery API client library for Python
B. Run your script on a new virtual machine with the BigQuery access scope enabled
C. Create a new service account with BigQuery access and execute your script with that user
D. Install the bq component for gccloud with the command gcloud components install bq.

Answer: B

Explanation:
The error is most like caused by the access scope issue. When create new instance, you have the default Compute engine default service account but most serves access including BigQuery is not enable. Create an instance Most access are not enabled by default You have default service account but don't have the permission (scope) you can stop the instance, edit, change scope and restart it to enable the scope access. Of course, if you Run your script on a new virtual machine with the BigQuery access scope enabled, it also works
https://cloud.google.com/compute/docs/access/service-accounts

NEW QUESTION 151
- (Exam Topic 5)
Your company pushes batches of sensitive transaction data from its application server VMs to Cloud Pub/Sub for processing and storage. What is the Google-recommended way for your application to authenticate to the required Google Cloud services?

A. Ensure that VM service accounts are granted the appropriate Cloud Pub/Sub IAM roles.
B. Ensure that VM service accounts do not have access to Cloud Pub/Sub, and use VM access scopes to grant the appropriate Cloud Pub/Sub IAM roles.
C. Generate an OAuth2 access token for accessing Cloud Pub/Sub, encrypt it, and store it in Cloud Storage for access from each VM.
D. Create a gateway to Cloud Pub/Sub using a Cloud Function, and grant the Cloud Function service account the appropriate Cloud Pub/Sub IAM roles.

Answer: A

NEW QUESTION 155
- (Exam Topic 5)
You are tasked with building an online analytical processing (OLAP) marketing analytics and reporting tool. This requires a relational database that can operate on hundreds of terabytes of data. What is the Google
recommended tool for such applications?

A. Cloud Spanner, because it is globally distributed
B. Cloud SQL, because it is a fully managed relational database
C. Cloud Firestore, because it offers real-time synchronization across devices
D. BigQuery, because it is designed for large-scale processing of tabular data

Answer: A

Explanation:
Reference: https://cloud.google.com/files/BigQueryTechnicalWP.pdf

NEW QUESTION 157
- (Exam Topic 5)
You are migrating third-party applications from optimized on-premises virtual machines to Google Cloud. You are unsure about the optimum CPU and memory options. The application have a consistent usage patterns across multiple weeks. You want to optimize resource usage for the lowest cost. What should you do?

A. Create a Compute engine instance with CPU and Memory options similar to your application's current on-premises virtual machin
B. Install the cloud monitoring agent, and deploy the third party applicatio
C. Run a load with normal traffic levels on third party application and follow the Rightsizing Recommendations in the Cloud Console
D. Create an App Engine flexible environment, and deploy the third party application using a Docker file and a custom runtim
E. Set CPU and memory options similar to your application's current on-premisesvirtual machine in the app.yaml file.
F. Create an instance template with the smallest available machine type, and use an image of the third party application taken from the current on-premises virtual machin
G. Create a managed instance group that uses average CPU to autoscale the number of instances in the grou
H. Modify the average CPU utilization threshold to optimize the number of instances running.
I. Create multiple Compute Engine instances with varying CPU and memory option
J. Install the cloud monitoring agent and deploy the third-party application on each of the
K. Run a load test with high traffic levels on the application and use the results to determine the optimal settings.

Answer: A

Explanation:
Create a Compute engine instance with CPU and Memory options similar to your application's current
on-premises virtual machine. Install the cloud monitoring agent, and deploy the third party application. Run a load with normal traffic levels on third party
application and follow the Rightsizing Recommendations in the Cloud Console
https://cloud.google.com/migrate/compute-engine/docs/4.9/concepts/planning-a-migration/cloud-instance-rights

NEW QUESTION 158
- (Exam Topic 5)
A development manager is building a new application He asks you to review his requirements and identify what cloud technologies he can use to meet them. The application must
* 1. Be based on open-source technology for cloud portability
* 2. Dynamically scale compute capacity based on demand
* 3. Support continuous software delivery
* 4. Run multiple segregated copies of the same application stack
* 5. Deploy application bundles using dynamic templates
* 6. Route network traffic to specific services based on URL
Which combination of technologies will meet all of his requirements?

A. Google Container Engine, Jenkins, and Helm
B. Google Container Engine and Cloud Load Balancing
C. Google Compute Engine and Cloud Deployment Manager
D. Google Compute Engine, Jenkins, and Cloud Load Balancing

Answer: A

Explanation:
Helm for managing Kubernetes
Kubernetes can base on the URL to route traffic to different location (path) https://cloud.google.com/kubernetes-engine/docs/tutorials/http-balancer eg.apiVersion:
networking.k8s.io/v1beta1
kind: Ingress metadata:
name: fanout-ingress spec:
rules:
- http: paths:
- path: /* backend: serviceName: web servicePort: 8080
- path: /v2/* backend: serviceName: web2 servicePort: 8080

NEW QUESTION 161
- (Exam Topic 5)
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?

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

NEW QUESTION 163
- (Exam Topic 5)
Your company is developing a web-based application. You need to make sure that production deployments are linked to source code commits and are fully auditable. What should you do?

A. Make sure a developer is tagging the code commit with the date and time of commit
B. Make sure a developer is adding a comment to the commit that links to the deployment.
C. Make the container tag match the source code commit hash.
D. Make sure the developer is tagging the commits with :latest

Answer: C

Explanation:
From: https://cloud.google.com/architecture/best-practices-for-building-containers Under: Tagging using the Git commit hash (bottom of page almost)
"In this case, a common way of handling version numbers is to use the Git commit SHA-1 hash (or a short version of it) as the version number. By design, the Git commit hash is immutable and references a specific version of your software.
You can use this commit hash as a version number for your software, but also as a tag for the Docker image built from this specific version of your software. Doing so makes Docker images traceable: because in this case the image tag is immutable, you instantly know which specific version of your software is running inside a given container."

NEW QUESTION 167
- (Exam Topic 5)
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 dis
B. Create an instance template from the snapsho
C. Create an autoscaled managed instance group from the instance template.
D. Create a snapshot of the existing dis
E. Create a custom image from the snapsho
F. Create an autoscaled managed instance group from the custom image.
G. Create a custom image from the existing dis
H. Create an instance template from the custom image.Create an autoscaled managed instance group from the instance template.
I. Create an instance template from the existing dis
J. Create a custom image from the instance template.Create an autoscaled managed instance group from the custom image.

Answer: B

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

NEW QUESTION 170
- (Exam Topic 5)
You are working with a data warehousing team that performs data analysis. The team needs to process data from external partners, but the data contains personally identifiable information (PII). You need to process and store the data without storing any of the PII data. What should you do?

A. Create a Dataflow pipeline to retrieve the data from the external source
B. As part of the pipeline use the Cloud Data Loss Prevention (Cloud DLP) API to remove any PII data Store the result in BigQuery
C. Create a Dataflow pipeline to retrieve the data from the external source
D. As part of the pipeline store all non-PII data in BigQuery and store all PII data in a Cloud Storage bucket that has a retention policy set.
E. Ask the external partners to upload an data on Cloud Storage Configure Bucket Lock for the bucket Create a Dataflow pipeline to read the data from the bucket As part of the pipeline, use the Cloud Data Loss Prevention (Cloud DIP) API to remove any PII data Store the result in BigQuery
F. Ask the external partners to import ail data in your BigQuery dataset Create a dataflow pipeline to copy the data into a new table As part of the Dataflow bucket skip all data in columns that have PII data

Answer: A

Explanation:
Create a Dataflow pipeline to retrieve the data from the external sources, he did not specify the way he is going to create it, it might be a pub/sub or external table or whatever.

NEW QUESTION 172
- (Exam Topic 5)
To reduce costs, the Director of Engineering has required all developers to move their development infrastructure resources from on-premises virtual machines (VMs) to Google Cloud Platform. These resources go through multiple start/stop events during the day and require state to persist. You have been asked to design the process of running a development environment in Google Cloud while providing cost visibility to the finance department. Which two steps should you take? Choose 2 answers

A. Use the --no-auto-delete flag on all persistent disks and stop the VM.
B. Use the -auto-delete flag on all persistent disks and terminate the VM.
C. Apply VM CPU utilization label and include it in the BigQuery billing export.

D. Use Google BigQuery billing export and labels to associate cost to groups.
E. Store all state into local SSD, snapshot the persistent disks, and terminate the VM.
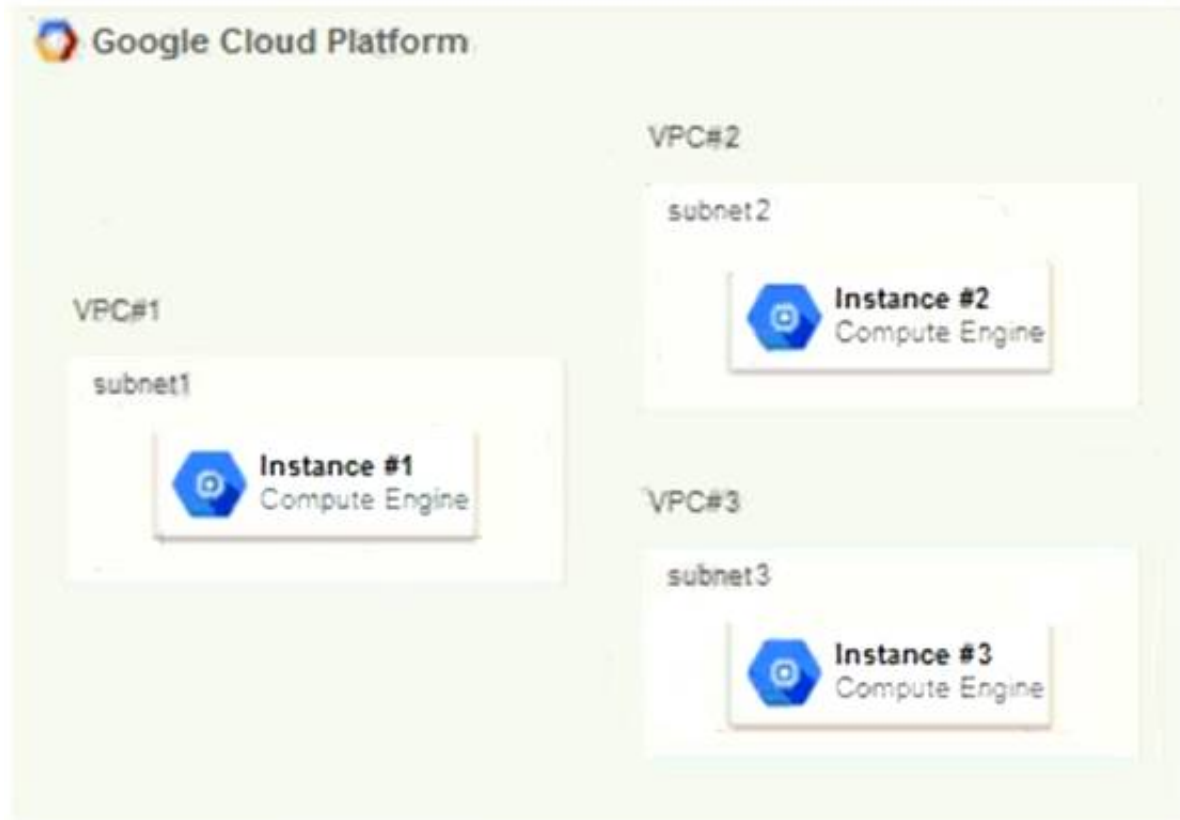F. Store all state in Google Cloud Storage, snapshot the persistent disks, and terminate the VM.

Answer: AD

Explanation:
https://cloud.google.com/billing/docs/how-to/export-data-bigquery

NEW QUESTION 174
- (Exam Topic 5)
Your company has a project in Google Cloud with three Virtual Private Clouds (VPCs). There is a Compute Engine instance on each VPC. Network subnets do not overlap and must remain separated. The network configuration is shown below.

Instance #1 is an exception and must communicate directly with both Instance #2 and Instance #3 via internal IPs. How should you accomplish this?

A. Create a cloud router to advertise subnet #2 and subnet #3 to subnet #1.
B. Add two additional NICs to Instance #1 with the following configuration:•NIC1VPC: VPC #2SUBNETWORK: subnet #2•NIC2VPC: VPC #3SUBNETWORK: subnet #3Update firewall rules to enable traffic between instances.
C. Create two VPN tunnels via CloudVPN:•1 between VPC #1 and VPC #2.•1 between VPC #2 and VPC #3.Update firewall rules to enable traffic between the instances.
D. Peer all three VPCs:•Peer VPC #1 with VPC #2.•Peer VPC #2 with VPC #3.Update firewall rules to enable traffic between the instances.

Answer: B

Explanation:
As per GCP documentation: "By default, every instance in a VPC network has a single network interface. Use these instructions to create additional network interfaces. Each interface is attached to a different VPC network, giving that instance access to different VPC networks in Google Cloud. You cannot attach multiple network interfaces to the same VPC network." Refer to:
https://cloud.google.com/vpc/docs/create-use-multiple-interfaces
https://cloud.google.com/vpc/docs/create-use-multiple-interfaces#i_am_not_able_to_connect_to_secondary_inte

NEW QUESTION 179
- (Exam Topic 5)
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 Clusterin
D. Place nodes in different subnets.
E. Set up SQL Server Always On Availability Groups using Windows Failover Clusterin
F. Place nodes in different zones.

Answer: D

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

NEW QUESTION 182
- (Exam Topic 5)
Your company has just recently activated Cloud Identity to manage users. The Google Cloud Organization has been configured as wed. The security learn needs to secure protects that will be part of the Organization. They want to prohibit IAM users outside the domain from gaining permissions from now on. What should they do?

A. Configure an organization policy to restrict identities by domain

B. Configure an organization policy to block creation of service accounts
C. Configure Cloud Scheduler to trigger a Cloud Function every hour that removes all users that don't belong to the Cloud identity domain from all projects.

Answer: A

NEW QUESTION 185
- (Exam Topic 5)
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?

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

NEW QUESTION 187
- (Exam Topic 5)
A development team at your company has created a dockerized HTTPS web application. You need to deploy the application on Google Kubernetes Engine (GKE) and make sure that the application scales automatically.
How should you deploy to GKE?

A. Use the Horizontal Pod Autoscaler and enable cluster autoscalin
B. Use an Ingress resource to loadbalance the HTTPS traffic.
C. Use the Horizontal Pod Autoscaler and enable cluster autoscaling on the Kubernetes cluste
D. Use a Service resource of type LoadBalancer to load-balance the HTTPS traffic.
E. Enable autoscaling on the Compute Engine instance grou
F. Use an Ingress resource to load balance the HTTPS traffic.
G. Enable autoscaling on the Compute Engine instance grou
H. Use a Service resource of type LoadBalancer to load-balance the HTTPS traffic.

Answer: B

Explanation:
https://cloud.google.com/kubernetes-engine/docs/tutorials/http-balancer https://cloud.google.com/kubernetes-engine/docs/concepts/network-overview#ext-lb

NEW QUESTION 191
- (Exam Topic 5)
Your solution is producing performance bugs in production that you did not see in staging and test environments. You want to adjust your test and deployment procedures to avoid this problem in the future. What should you do?

A. Deploy fewer changes to production.
B. Deploy smaller changes to production.
C. Increase the load on your test and staging environments.
D. Deploy changes to a small subset of users before rolling out to production.

Answer: C

NEW QUESTION 193
- (Exam Topic 5)
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 files3) 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 hashes3) 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 upload2) Use gsutil -m to upload all the files to the Cloud Storage3) Use gsutil cp to download the uploaded files4) 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 files3) Use gsutil ls -L gs://[YOUR_BUCKET_NAME] to collect CRC32C hashes of the uploaded files 4)Compare the hashes

Answer: D

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

NEW QUESTION 198
- (Exam Topic 5)
Your web application uses Google Kubernetes Engine to manage several workloads. One workload requires a consistent set of hostnames even after pod scaling and relaunches.

Which feature of Kubernetes should you use to accomplish this?

A. StatefulSets
B. Role-based access control
C. Container environment variables
D. Persistent Volumes

Answer: A

Explanation:
https://kubernetes.io/docs/tutorials/stateful-application/basic-stateful-set/

NEW QUESTION 203
- (Exam Topic 5)
You are working at a financial institution that stores mortgage loan approval documents on Cloud Storage.
Any change to these approval documents must be uploaded as a separate approval file, so you want to ensure that these documents cannot be deleted or overwritten for the next 5 years. What should you do?

A. Create a retention policy on the bucket for the duration of 5 year
B. Create a lock on the retention policy.
C. Create the bucket with uniform bucket-level access, and grant a service account the role of Object Write
D. Use the service account to upload new files.
E. Use a customer-managed key for the encryption of the bucke
F. Rotate the key after 5 years.
G. Create the bucket with fine-grained access control, and grant a service account the role of Object Writer.Use the service account to upload new files.

Answer: A

Explanation:
Reference: https://cloud.google.com/storage/docs/using-bucket-lock

NEW QUESTION 206
- (Exam Topic 5)
Your organization wants to control IAM policies for different departments independently, but centrally. Which approach should you take?

A. Multiple Organizations with multiple Folders
B. Multiple Organizations, one for each department
C. A single Organization with Folder for each department
D. A single Organization with multiple projects, each with a central owner

Answer: C

Explanation:
Folders are nodes in the Cloud Platform Resource Hierarchy. A folder can contain projects, other folders, or a combination of both. You can use folders to group projects under an organization in a hierarchy. For example, your organization might contain multiple departments, each with its own set of GCP resources. Folders allow you to group these resources on a per-department basis. Folders are used to group resources that share common IAM policies. While a folder can contain multiple folders or resources, a given folder or resource can have exactly one parent.
References: https://cloud.google.com/resource-manager/docs/creating-managing-folders

NEW QUESTION 208
- (Exam Topic 5)
A production database virtual machine on Google Compute Engine has an ext4-formatted persistent disk for data files The database is about to run out of storage space How can you remediate the problem with the least amount of downtime?

A. In the Cloud Platform Console, increase the size of the persistent disk and use the resize2fs command in Linux.
B. Shut down the virtual machine, use the Cloud Platform Console to increase the persistent disk size, then restart the virtual machine.
C. In the Cloud Platform Console, increase the size of the persistent disk and verify the new space is ready to use with the fdisk command in Linux.
D. In the Cloud Platform Console, create a new persistent disk attached to the virtual machine, format and mount it, and configure the database service to move the files to the new disk.
E. In the Cloud Platform Console, create a snapshot of the persistent disk, restore the snapshot to a new larger disk, unmount the old disk, mount the new disk, and restart the database service.

Answer: A

Explanation:
On Linux instances, connect to your instance and manually resize your partitions and file systems to use the additional disk space that you added.
Extend the file system on the disk or the partition to use the added space. If you grew a partition on your disk, specify the partition. If your disk does not have a partition table, specify only the disk ID.
sudo resize2fs /dev/[DISK_ID][PARTITION_NUMBER]
where [DISK_ID] is the device name and [PARTITION_NUMBER] is the partition number for the device where you are resizing the file system.
References: https://cloud.google.com/compute/docs/disks/add-persistent-disk

NEW QUESTION 210
- (Exam Topic 5)
You have an application that runs in Google Kubernetes Engine (GKE). Over the last 2 weeks, customers have reported that a specific part of the application returns errors very frequently. You currently have no logging or monitoring solution enabled on your GKE cluster. You want to diagnose the problem, but you have not been able to replicate the issue. You want to cause minimal disruption to the application. What should you do?

A. * 1. Update your GKE cluster to use Cloud Operations for GKE.* 2. Use the GKE Monitoring dashboard to investigate logs from affected Pods.

B. * 1. Create a new GKE cluster with Cloud Operations for GKE enabled.* 2. Migrate the affected Pods to the new cluster, and redirect traffic for those Pods to the new cluster.* 3. Use the GKE Monitoring dashboard to investigate logs from affected Pods.
C. * 1. Update your GKE cluster to use Cloud Operations for GKE, and deploy Prometheus.* 2. Set an alert to trigger whenever the application returns an error.
D. * 1. Create a new GKE cluster with Cloud Operations for GKE enabled, and deploy Prometheus.* 2. Migrate the affected Pods to the new cluster, and redirect traffic for those Pods to the new cluste
E. * 3. Set an alert to trigger whenever the application returns an error.

Answer: A

Explanation:
Reference: https://cloud.google.com/blog/products/management-tools/using-logging-your-apps-running- kubernetes-engine

NEW QUESTION 213
- (Exam Topic 5)
An application development team has come to you for advice.They are planning to write and deploy an HTTP(S) API using Go 1.12. The API will have a very unpredictable workload and must remain reliable during peaks in traffic. They want to minimize operational overhead for this application. What approach should you recommend?

A. Use a Managed Instance Group when deploying to Compute Engine
B. Develop an application with containers, and deploy to Google Kubernetes Engine (GKE)
C. Develop the application for App Engine standard environment
D. Develop the application for App Engine Flexible environment using a custom runtime

Answer: C

Explanation:
https://cloud.google.com/appengine/docs/the-appengine-environments

NEW QUESTION 214
- (Exam Topic 5)
You are creating a solution to remove backup files older than 90 days from your backup Cloud Storage bucket. You want to optimize ongoing Cloud Storage spend. What should you do?

A. Write a lifecycle management rule in XML and push it to the bucket with gsutil.
B. Write a lifecycle management rule in JSON and push it to the bucket with gsutil.
C. Schedule a cron script using gsutil is -lr gs://backups/** to find and remove items older than 90 days.
D. Schedule a cron script using gsutil ls -1 gs://backups/** to find and remove items older than 90 days and schedule it with cron.

Answer: B

Explanation:
https://cloud.google.com/storage/docs/gsutil/commands/lifecycle

NEW QUESTION 218
- (Exam Topic 5)
Your company has announced that they will be outsourcing operations functions. You want to allow developers to easily stage new versions of a cloud-based application in the production environment and allow the outsourced operations team to autonomously promote staged versions to production. You want to minimize the operational overhead of the solution. Which Google Cloud product should you migrate to?

A. App Engine
B. GKE On-Prem
C. Compute Engine
D. Google Kubernetes Engine

Answer: D

Explanation:
Reference: https://cloud.google.com/security/compliance/eba-outsourcing-mapping-gcp

NEW QUESTION 223
- (Exam Topic 5)
Your operations team has asked you to help diagnose a performance issue in a production application that runs on Compute Engine. The application is dropping requests that reach it when under heavy load. The process list for affected instances shows a single application process that is consuming all available CPU, and autoscaling has reached the upper limit of instances. There is no abnormal load on any other related systems, including the database. You want to allow production traffic to be served again as quickly as possible. Which action should you recommend?

A. Change the autoscaling metric to agent.googleapis.com/memory/percent_used.
B. Restart the affected instances on a staggered schedule.
C. SSH to each instance and restart the application process.
D. Increase the maximum number of instances in the autoscaling group.

Answer: D

Explanation:
Reference: https://cloud.google.com/blog/products/sap-google-cloud/best-practices-for-sap-app-server- autoscaling-on-google-cloud

NEW QUESTION 226

- (Exam Topic 5)
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.

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_orga
Reference: https://cloud.google.com/vpc/docs/shared-vpc

NEW QUESTION 231
- (Exam Topic 5)
You want to create a private connection between your instances on Compute Engine and your on-premises data center. You require a connection of at least 20 Gbps. You want to follow Google-recommended practices.
How should you set up the connection?

A. Create a VPC and connect it to your on-premises data center using Dedicated Interconnect.
B. Create a VPC and connect it to your on-premises data center using a single Cloud VPN.
C. Create a Cloud Content Delivery Network (Cloud CDN) and connect it to your on-premises data center using Dedicated Interconnect.
D. Create a Cloud Content Delivery Network (Cloud CDN) and connect it to your on-premises datacenter using a single Cloud VPN.

Answer: A

Explanation:
Reference: https://cloud.google.com/compute/docs/instances/connecting-advanced

NEW QUESTION 232
- (Exam Topic 5)
You are developing an application using different microservices that should remain internal to the cluster. You want to be able to configure each microservice with a specific number of replicas. You also want to be able to address a specific microservice from any other microservice in a uniform way, regardless of the number of replicas the microservice scales to. You need to implement this solution on Google Kubernetes Engine. What should you do?

A. Deploy each microservice as a Deploymen
B. Expose the Deployment in the cluster using a Service, and use the Service DNS name to address it from other microservices within the cluster.
C. Deploy each microservice as a Deploymen
D. Expose the Deployment in the cluster using an Ingress, and use the Ingress IP address to address the Deployment from other microservices within the cluster.
E. Deploy each microservice as a Po
F. Expose the Pod in the cluster using a Service, and use the Service DNS name to address the microservice from other microservices within the cluster.
G. Deploy each microservice as a Po
H. Expose the Pod in the cluster using an Ingress, and use the Ingress IP address name to address the Pod from other microservices within the cluster.

Answer: A

Explanation:
https://kubernetes.io/docs/concepts/services-networking/ingress/

NEW QUESTION 235
- (Exam Topic 5)
You company has a Kubernetes application that pulls messages from Pub/Sub and stores them in Firestore. Because the application is simple, it was deployed as a single pod. The infrastructure team has analyzed Pub/Sub metrics and discovered that the application cannot process the messages in real time. Most of them wait for minutes before being processed. You need to scale the elaboration process that is I/O-intensive. What should you do?

A. Configure a Kubernetes autoscaling based on the subscription/push_request metric.
B. Use the –enable- autoscaling flag when you create the Kubernetes cluster
C. Configure a Kubernetes autoscaling based on the subscription/num_undelivered message metric.
D. Use kubectl autoscale deployment APP_NAME –max 6 –min 2 –cpu- percent 50 to configure Kubernetes autoscaling deployment

Answer: A

Explanation:
https://cloud.google.com/kubernetes-engine/docs/concepts/custom-and-external-metrics#external_metrics

NEW QUESTION 236

- (Exam Topic 5)
Your customer wants to do resilience testing of their authentication layer. This consists of a regional managed instance group serving a public REST API that reads from and writes to a Cloud SQL instance.
What should you do?

A. Engage with a security company to run web scrapes that look your users' authentication data om malicious websites and notify you if any if found.
B. Deploy intrusion detection software to your virtual machines to detect and log unauthorized access.
C. Schedule a disaster simulation exercise during which you can shut off all VMs in a zone to see how your application behaves.
D. Configure a red replica for your Cloud SQL instance in a different zone than the master, and then manually trigger a failover while monitoring KPIs for our REST API.

Answer: C

NEW QUESTION 240
- (Exam Topic 5)
You want your Google Kubernetes Engine cluster to automatically add or remove nodes based on CPUload. What should you do?

A. Configure a HorizontalPodAutoscaler with a target CPU usag
B. Enable the Cluster Autoscaler from the GCP Console.
C. Configure a HorizontalPodAutoscaler with a target CPU usag
D. Enable autoscaling on the managed instance group for the cluster using the gcloud command.
E. Create a deployment and set the maxUnavailable and maxSurge propertie
F. Enable the Cluster Autoscaler using the gcloud command.
G. Create a deployment and set the maxUnavailable and maxSurge propertie
H. Enable autoscaling on the cluster managed instance group from the GCP Console.

Answer: B

NEW QUESTION 241
- (Exam Topic 5)
You are implementing a single Cloud SQL MySQL second-generation database that contains business-critical transaction data. You want to ensure that the minimum amount of data is lost in case of catastrophic failure. Which two features should you implement? (Choose two.)

A. Sharding
B. Read replicas
C. Binary logging
D. Automated backups
E. Semisynchronous replication

Answer: CD

Explanation:
Backups help you restore lost data to your Cloud SQL instance. Additionally, if an instance is having a problem, you can restore it to a previous state by using the backup to overwrite it. Enable automated backups for any instance that contains necessary data. Backups protect your data from loss or damage.
Enabling automated backups, along with binary logging, is also required for some operations, such as clone and replica creation.
Reference: https://cloud.google.com/sql/docs/mysql/backup-recovery/backups

NEW QUESTION 243
- (Exam Topic 5)
You are running a cluster on Kubernetes Engine to serve a web application. Users are reporting that a specific part of the application is not responding anymore. You notice that all pods of your deployment keep restarting after 2 seconds. The application writes logs to standard output. You want to inspect the logs to find the cause of the issue. Which approach can you take?

A. Review the Stackdriver logs for each Compute Engine instance that is serving as a node in the cluster.
B. Review the Stackdriver logs for the specific Kubernetes Engine container that is serving the unresponsive part of the application.
C. Connect to the cluster using gcloud credentials and connect to a container in one of the pods to read the logs.
D. Review the Serial Port logs for each Compute Engine instance that is serving as a node in the cluster.

Answer: B

NEW QUESTION 246
- (Exam Topic 5)
You have deployed an application to Kubernetes Engine, and are using the Cloud SQL proxy container to make the Cloud SQL database available to the services running on Kubernetes. You are notified that the application is reporting database connection issues. Your company policies require a post-mortem. What should you do?

A. Use gcloud sql instances restart.
B. Validate that the Service Account used by the Cloud SQL proxy container still has the Cloud Build Editor role.
C. In the GCP Console, navigate to Stackdriver Loggin
D. Consult logs for Kubernetes Engine and Cloud SQL.
E. In the GCP Console, navigate to Cloud SQ
F. Restore the latest backu
G. Use kubect1 to restart all pods.

Answer: C

NEW QUESTION 249
- (Exam Topic 5)

You want to establish a Compute Engine application in a single VPC across two regions. The application must communicate over VPN to an on-premises network. How should you deploy the VPN?

A. Use VPC Network Peering between the VPC and the on-premises network.
B. Expose the VPC to the on-premises network using IAM and VPC Sharing.
C. Create a global Cloud VPN Gateway with VPN tunnels from each region to the on-premises peer gateway.
D. Deploy Cloud VPN Gateway in each regio
E. Ensure that each region has at least one VPN tunnel to the on-premises peer gateway.

Answer: C

Explanation:
https://cloud.google.com/vpn/docs/how-to/creating-static-vpns

NEW QUESTION 253
- (Exam Topic 5)
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

Answer: C

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

NEW QUESTION 255
- (Exam Topic 5)
One of your primary business objectives is being able to trust the data stored in your application. You want to log all changes to the application data. How can you design your logging system to verify authenticity of your logs?

A. Write the log concurrently in the cloud and on premises.
B. Use a SQL database and limit who can modify the log table.
C. Digitally sign each timestamp and log entry and store the signature.
D. Create a JSON dump of each log entry and store it in Google Cloud Storage.

Answer: C

Explanation:
https://cloud.google.com/storage/docs/access-logs
References: https://cloud.google.com/logging/docs/reference/tools/gcloud-logging

NEW QUESTION 256
- (Exam Topic 5)
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

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

NEW QUESTION 260
- (Exam Topic 5)
You need to deploy an application on Google Cloud that must run on a Debian Linux environment. The application requires extensive configuration in order to operate correctly. You want to ensure that you can install Debian distribution updates with minimal manual intervention whenever they become available. What should you do?

A. Create a Compute Engine instance template using the most recent Debian imag
B. Create an instance from this template, and install and configure the application as part of the startup scrip
C. Repeat this process whenever a new Google-managed Debian image becomes available.
D. Create a Debian-based Compute Engine instance, install and configure the application, and use OS patch management to install available updates.
E. Create an instance with the latest available Debian imag
F. Connect to the instance via SSH, and install and configure the application on the instanc
G. Repeat this process whenever a new Google-managed Debian image becomes available.
H. Create a Docker container with Debian as the base imag
I. Install and configure the application as part of the Docker image creation proces
J. Host the container on Google Kubernetes Engine and restart the container whenever a new update is available.

Answer: B

Explanation:
Reference: https://cloud.google.com/compute/docs/os-patch-management

NEW QUESTION 261
- (Exam Topic 5)
You are migrating your on-premises solution to Google Cloud in several phases. You will use Cloud VPN to maintain a connection between your on-premises systems and Google Cloud until the migration is completed.
You want to make sure all your on-premises systems remain reachable during this period. How should you organize your networking in Google Cloud?

A. Use the same IP range on Google Cloud as you use on-premises
B. Use the same IP range on Google Cloud as you use on-premises for your primary IP range and use a secondary range that does not overlap with the range you use on-premises
C. Use an IP range on Google Cloud that does not overlap with the range you use on-premises
D. Use an IP range on Google Cloud that does not overlap with the range you use on-premises for your primary IP range and use a secondary range with the same IP range as you use on-premises

Answer: C

NEW QUESTION 266
- (Exam Topic 6)
For this question, refer to the Dress4Win case study. Considering the given business requirements, how would you automate the deployment of web and transactional data layers?

A. Deploy Nginx and Tomcat using Cloud Deployment Manager to Compute Engin
B. Deploy a Cloud SQL server to replace MySQ
C. Deploy Jenkins using Cloud Deployment Manager.
D. Deploy Nginx and Tomcat using Cloud Launche
E. Deploy a MySQL server using Cloud Launcher.Deploy Jenkins to Compute Engine using Cloud Deployment Manager scripts.
F. Migrate Nginx and Tomcat to App Engin
G. Deploy a Cloud Datastore server to replace the MySQL server in a high-availability configuratio
H. Deploy Jenkins to Compute Engine using Cloud Launcher.
I. Migrate Nginx and Tomcat to App Engin
J. Deploy a MySQL server using Cloud Launche
K. Deploy Jenkins to Compute Engine using Cloud Launcher.

Answer: A

NEW QUESTION 270
- (Exam Topic 6)
For this question, refer to the Dress4Win case study. Which of the compute services should be migrated as –is and would still be an optimized architecture for performance in the cloud?

A. Web applications deployed using App Engine standard environment
B. RabbitMQ deployed using an unmanaged instance group
C. Hadoop/Spark deployed using Cloud Dataproc Regional in High Availability mode
D. Jenkins, monitoring, bastion hosts, security scanners services deployed on custom machine types

Answer: C

NEW QUESTION 274
- (Exam Topic 7)
For this question, refer to the TerramEarth case study. Considering the technical requirements, how should you reduce the unplanned vehicle downtime in GCP?

A. Use BigQuery as the data warehous
B. Connect all vehicles to the network and stream data into BigQuery using Cloud Pub/Sub and Cloud Dataflo
C. Use Google Data Studio for analysis and reporting.
D. Use BigQuery as the data warehous
E. Connect all vehicles to the network and upload gzip files to a Multi-Regional Cloud Storage bucket using gclou
F. Use Google Data Studio for analysis and reporting.
G. Use Cloud Dataproc Hive as the data warehous
H. Upload gzip files to a MultiRegional Cloud Storage bucke
I. Upload this data into BigQuery using gclou
J. Use Google data Studio for analysis and reporting.
K. Use Cloud Dataproc Hive as the data warehous
L. Directly stream data into prtitioned Hive table
M. Use Pig scripts to analyze data.

Answer: A

NEW QUESTION 276
- (Exam Topic 7)
For this question, refer to the TerramEarth case study. You need to implement a reliable, scalable GCP solution for the data warehouse for your company, TerramEarth. Considering the TerramEarth business and technical requirements, what should you do?

A. Replace the existing data warehouse with BigQuer
B. Use table partitioning.

C. Replace the existing data warehouse with a Compute Engine instance with 96 CPUs.
D. Replace the existing data warehouse with BigQuer
E. Use federated data sources.
F. Replace the existing data warehouse with a Compute Engine instance with 96 CPU
G. Add an additional Compute Engine pre-emptible instance with 32 CPUs.

Answer: C

Explanation:
https://cloud.google.com/solutions/bigquery-data-warehouse#external_sources https://cloud.google.com/solutions/bigquery-data-warehouse

NEW QUESTION 278
- (Exam Topic 7)
For this question, refer to the TerramEarth case study. TerramEarth has decided to store data files in Cloud Storage. You need to configure Cloud Storage lifecycle rule to store 1 year of data and minimize file storage cost.
Which two actions should you take?

A. Create a Cloud Storage lifecycle rule with Age: "30", Storage Class: "Standard", and Action: "Set to Coldline", and create a second GCS life-cycle rule with Age: "365", Storage Class: "Coldline", and Action: "Delete".
B. Create a Cloud Storage lifecycle rule with Age: "30", Storage Class: "Coldline", and Action: "Set to Nearline", and create a second GCS life-cycle rule with Age: "91", Storage Class: "Coldline", and Action: "Set to Nearline".
C. Create a Cloud Storage lifecycle rule with Age: "90", Storage Class: "Standard", and Action: "Set to Nearline", and create a second GCS life-cycle rule with Age: "91", Storage Class: "Nearline", and Action: "Set to Coldline".
D. Create a Cloud Storage lifecycle rule with Age: "30", Storage Class: "Standard", and Action: "Set to Coldline", and create a second GCS life-cycle rule with Age: "365", Storage Class: "Nearline", and Action: "Delete".

Answer: A

NEW QUESTION 279
- (Exam Topic 8)
For this question, refer to the Mountkirk Games case study. You are in charge of the new Game Backend Platform architecture. The game communicates with the backend over a REST API.
You want to follow Google-recommended practices. How should you design the backend?

A. Create an instance template for the backen
B. For every region, deploy it on a multi-zone managed instance grou
C. Use an L4 load balancer.
D. Create an instance template for the backen
E. For every region, deploy it on a single-zone managed instance grou
F. Use an L4 load balancer.
G. Create an instance template for the backen
H. For every region, deploy it on a multi-zone managed instance grou
I. Use an L7 load balancer.
J. Create an instance template for the backen
K. For every region, deploy it on a single-zone managed instance grou
L. Use an L7 load balancer.

Answer: C

Explanation:
https://cloud.google.com/solutions/gaming/cloud-game-infrastructure#dedicated_game_server

NEW QUESTION 281
- (Exam Topic 8)
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?

A. Cloud Bigtable
B. Cloud Spanner
C. BigQuery
D. Cloud Datastore

Answer: A

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

NEW QUESTION 284
- (Exam Topic 8)
For this question, refer to the Mountkirk Games case study. Mountkirk Games wants you to design a way to test the analytics platform's resilience to changes in mobile network latency. What should you do?

A. Deploy failure injection software to the game analytics platform that can inject additional latency to mobile client analytics traffic.
B. Build a test client that can be run from a mobile phone emulator on a Compute Engine virtual machine, and run multiple copies in Google Cloud Platform regions all over the world to generate realistic traffic.
C. Add the ability to introduce a random amount of delay before beginning to process analytics files uploaded from mobile devices.
D. Create an opt-in beta of the game that runs on players' mobile devices and collects response times from analytics endpoints running in Google Cloud Platform regions all over the world.

Answer: D

NEW QUESTION 289
- (Exam Topic 8)
For this question, refer to the Mountkirk Games case study. You need to analyze and define the technical architecture for the database workloads for your company, Mountkirk Games. Considering the business and technical requirements, what should you do?

A. Use Cloud SQL for time series data, and use Cloud Bigtable for historical data queries.
B. Use Cloud SQL to replace MySQL, and use Cloud Spanner for historical data queries.
C. Use Cloud Bigtable to replace MySQL, and use BigQuery for historical data queries.
D. Use Cloud Bigtable for time series data, use Cloud Spanner for transactional data, and use BigQuery for historical data queries.

Answer: D

Explanation:
https://cloud.google.com/bigtable/docs/schema-design-time-series

NEW QUESTION 292
- (Exam Topic 8)
Your development team has created a mobile game app. You want to test the new mobile app on Android and iOS devices with a variety of configurations. You need to ensure that testing is efficient and cost-effective.
What should you do?

A. Upload your mobile app to the Firebase Test Lab, and test the mobile app on Android and iOS devices.
B. Create Android and iOS VMs on Google Cloud, install the mobile app on the VMs, and test the mobile app.
C. Create Android and iOS containers on Google Kubernetes Engine (GKE), install the mobile app on the containers, and test the mobile app.
D. Upload your mobile app with different configurations to Firebase Hosting and test each configuration.

Answer: C

NEW QUESTION 297
- (Exam Topic 8)
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?

A. 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
B. 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
C. 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.
D. Create a service account (SA) in the lgacy game's Google Cloud project, give the SA the Organization Admin rule and then give it the Firebase Admin role in both projects

Answer: B

NEW QUESTION 301
- (Exam Topic 9)
For this question, refer to the Helicopter Racing League (HRL) case study. HRL is looking for a cost-effective approach for storing their race data such as telemetry. They want to keep all historical records, train models using only the previous season's data, and plan for data growth in terms of volume and information collected. You need to propose a data solution. Considering HRL business requirements and the goals expressed by CEO S. Hawke, what should you do?

A. Use Firestore for its scalable and flexible document-based databas
B. Use collections to aggregate race data by season and event.
C. Use Cloud Spanner for its scalability and ability to version schemas with zero downtim
D. Split race data using season as a primary key.
E. Use BigQuery for its scalability and ability to add columns to a schem
F. Partition race data based on season.
G. Use Cloud SQL for its ability to automatically manage storage increases and compatibility with MySQ
H. Use separate database instances for each season.

Answer: C

Explanation:
Reference: https://cloud.google.com/bigquery/public-data

NEW QUESTION 303
- (Exam Topic 9)
For this question, refer to the Helicopter Racing League (HRL) case study. Recently HRL started a new regional racing league in Cape Town, South Africa. In an effort to give customers in Cape Town a better user experience, HRL has partnered with the Content Delivery Network provider, Fastly. HRL needs to allow traffic coming from all of the Fastly IP address ranges into their Virtual Private Cloud network (VPC network). You are a member of the HRL security team and you need to configure the update that will allow only the Fastly IP address ranges through the External HTTP(S) load balancer. Which command should you use?

A. glouc compute firewall rules update hlr-policy \--priority 1000 \target tags-sourceiplist fastly \--allow tcp:443
B. gcloud compute security policies rules update 1000 \--security-policy hlr-policy \--expression "evaluatePreconfiguredExpr('sourceiplist-fastly')" \--action " allow"
C. gcloud compute firewall rules update sourceiplist-fastly \priority 1000 \allow tcp: 443
D. gcloud compute priority-policies rules update 1000 \security policy from fastly--src- ip-ranges"-- action " allow"

Answer: B

Explanation:
Reference: https://cloud.google.com/load-balancing/docs/https D18912E1457D5D1DDCBD40AB3BF70D5D

NEW QUESTION 304
- (Exam Topic 9)
For this question, refer to the Helicopter Racing League (HRL) case study. HRL wants better prediction
accuracy from their ML prediction models. They want you to use Google's AI Platform so HRL can understand and interpret the predictions. What should you do?

A. Use Explainable AI.
B. Use Vision AI.
C. Use Google Cloud's operations suite.
D. Use Jupyter Notebooks.

Answer: A

Explanation:
Reference: https://cloud.google.com/ai-platform/prediction/docs/aiExplanation:s/preparing-metadata

NEW QUESTION 306
- (Exam Topic 9)
For this question, refer to the Helicopter Racing League (HRL) case study. A recent finance audit of cloud infrastructure noted an exceptionally high number of Compute Engine instances are allocated to do video encoding and transcoding. You suspect that these Virtual Machines are zombie machines that were not deleted after their workloads completed. You need to quickly get a list of which VM instances are idle. What should you do?

A. Log into each Compute Engine instance and collect disk, CPU, memory, and network usage statistics for analysis.
B. Use the gcloud compute instances list to list the virtual machine instances that have the idle: true label set.
C. Use the gcloud recommender command to list the idle virtual machine instances.
D. From the Google Console, identify which Compute Engine instances in the managed instance groups are no longer responding to health check probes.

Answer: C

Explanation:
Reference: https://cloud.google.com/compute/docs/instances/viewing-and-applying-idle-vm-recommendations

NEW QUESTION 307
- (Exam Topic 10)
For this question, refer to the EHR Healthcare case study. EHR has single Dedicated Interconnect
connection between their primary data center and Googles network. This connection satisfies EHR's network and security policies:
• On-premises servers without public IP addresses need to connect to cloud resources without public IP addresses
• Traffic flows from production network mgmt. servers to Compute Engine virtual machines should never traverse the public internet.
You need to upgrade the EHR connection to comply with their requirements. The new connection design must support business critical needs and meet the same network and security policy requirements. What should you do?

A. Add a new Dedicated Interconnect connection
B. Upgrade the bandwidth on the Dedicated Interconnect connection to 100 G
C. Add three new Cloud VPN connections
D. Add a new Carrier Peering connection

Answer: A

Explanation:
The case does not call out the throughput being an issue. However, to achieve 99.99%, you need to have 4 connections as per Google recommendations.
However, in the options only A has the option to add an additional Interconnect connection.
https://cloud.google.com/network-connectivity/docs/interconnect/concepts/dedicated-overview#availability

NEW QUESTION 311
- (Exam Topic 10)
For this question, refer to the EHR Healthcare case study. You need to define the technical architecture for securely deploying workloads to Google Cloud. You also need to ensure that only verified containers are deployed using Google Cloud services. What should you do? (Choose two.)

A. Enable Binary Authorization on GKE, and sign containers as part of a CI/CD pipeline.
B. Configure Jenkins to utilize Kritis to cryptographically sign a container as part of a CI/CD pipeline.
C. Configure Container Registry to only allow trusted service accounts to create and deploy containers from the registry.
D. Configure Container Registry to use vulnerability scanning to confirm that there are no vulnerabilities before deploying the workload.

Answer: A

Explanation:
Binary Authorization to ensure only verified containers are deployed To ensure deployment are secure and and consistent, automatically scan images for vulnerabilities with container analysis (https://cloud.google.com/docs/ci-cd/overview?hl=en&skip_cache=true)

NEW QUESTION 314
- (Exam Topic 10)
For this question, refer to the EHR Healthcare case study. You need to define the technical architecture for hybrid connectivity between EHR's on-premises systems and Google Cloud. You want to follow Google's recommended practices for production-level applications. Considering the EHR Healthcare business and

technical requirements, what should you do?

A. Configure two Partner Interconnect connections in one metro (City), and make sure the Interconnect connections are placed in different metro zones.
B. Configure two VPN connections from on-premises to Google Cloud, and make sure the VPN devices on-premises are in separate racks.
C. Configure Direct Peering between EHR Healthcare and Google Cloud, and make sure you are peering at least two Google locations.
D. Configure two Dedicated Interconnect connections in one metro (City) and two connections in another metro, and make sure the Interconnect connections are placed in different metro zones.

Answer: D

Explanation:
based on the requirement of secure and high-performance connection between on-premises systems to Google Cloud
https://cloud.google.com/network-connectivity/docs/interconnect/tutorials/partner-creating-9999-availability

NEW QUESTION 319
- (Exam Topic 10)
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?

A. Use a private cluster with a private endpoint with master authorized networks configured.
B. Use a public cluster with firewall rules and Virtual Private Cloud (VPC) routes.
C. Use a private cluster with a public endpoint with master authorized networks configured.
D. Use a public cluster with master authorized networks enabled and firewall rules.

Answer: A

Explanation:
https://cloud.google.com/kubernetes-engine/docs/concepts/private-cluster-concept#overview

NEW QUESTION 323
......

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual Professional-Cloud-Architect Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the Professional-Cloud-Architect Product From:

https://www.2passeasy.com/dumps/Professional-Cloud-Architect/

Money Back Guarantee

Professional-Cloud-Architect Practice Exam Features:

* Professional-Cloud-Architect Questions and Answers Updated Frequently

* Professional-Cloud-Architect Practice Questions Verified by Expert Senior Certified Staff

* Professional-Cloud-Architect Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* Professional-Cloud-Architect Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year