Google

Exam Questions Professional-Cloud-Architect

Google Certified Professional - Cloud Architect (GCP)

NEW QUESTION 1
- (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 2
- (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 3
- (Topic 2)
For this question, refer to the TerramEarth case study.
To speed up data retrieval, more vehicles will be upgraded to cellular connections and be able to transmit data to the ETL process. The current FTP process is error-prone and restarts the data transfer from the start of the file when connections fail, which happens often. You want to improve the reliability of the solution and minimize data transfer time on the cellular connections. What should you do?

A. Use one Google Container Engine cluster of FTP server
B. Save the data to a Multi- Regional bucke
C. Run the ETL process using data in the bucket.
D. Use multiple Google Container Engine clusters running FTP servers located in different region
E. Save the data to Multi-Regional buckets in us, eu, and asi
F. Run the ETL process using the data in the bucket.
G. Directly transfer the files to different Google Cloud Multi-Regional Storage bucket locations in us, eu, and asia using Google APIs over HTTP(S). Run the ETL process using the data in the bucket.
H. Directly transfer the files to a different Google Cloud Regional Storage bucket location in us, eu, and asia using Google APIs over HTTP(S). Run the ETL process to retrieve the data from each Regional bucket.

Answer: D

Explanation:
https://cloud.google.com/storage/docs/locations

NEW QUESTION 4
- (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 5
- (Topic 2)
For this question, refer to the TerramEarth case study
You analyzed TerramEarth's business requirement to reduce downtime, and found that they can achieve a majority of time saving by reducing customers' wait time for parts You decided to focus on reduction of the 3 weeks aggregate reporting time Which modifications to the company's processes should you recommend?

A. Migrate from CSV to binary format, migrate from FTP to SFTP transport, and develop machine learning analysis of metrics.
B. Migrate from FTP to streaming transport, migrate from CSV to binary format, and develop machine learning analysis of metrics.
C. Increase fleet cellular connectivity to 80%, migrate from FTP to streaming transport, and develop machine learning analysis of metrics.
D. Migrate from FTP to SFTP transport, develop machine learning analysis of metrics, and increase dealer local inventory by a fixed factor.

Answer: C

Explanation:
The Avro binary format is the preferred format for loading compressed data. Avro data is faster to load because the data can be read in parallel, even when the data blocks are compressed.
Cloud Storage supports streaming transfers with the gsutil tool or boto library, based on HTTP chunked transfer encoding. Streaming data lets you stream data to and from your Cloud Storage account as soon as it becomes available without requiring that the data be first saved to a separate file. Streaming transfers are useful if you have a process that generates data and you do not want to buffer it locally before uploading it, or if you want to send the result from a computational pipeline directly into Cloud Storage.
References: https://cloud.google.com/storage/docs/streaming https://cloud.google.com/bigquery/docs/loading-data

NEW QUESTION 6
- (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 7
- (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 andone for productio
B. Each account should contain one project for every application.
C. 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.
D. Create a single G Suite account to manage users with each stage of each application in its own project.
E. 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 8
- (Topic 4)
For this question, refer to the Dress4Win case study.
As part of Dress4Win's plans to migrate to the cloud, they want to be able to set up a managed logging and monitoring system so they can handle spikes in their traffic load.
They want to ensure that:
• The infrastructure can be notified when it needs to scale up and down to handle the ebb and flow of usage throughout the day

• Their administrators are notified automatically when their application reports errors.
• They can filter their aggregated logs down in order to debug one piece of the application across many hosts
Which Google StackDriver features should they use?

A. Logging, Alerts, Insights, Debug
B. Monitoring, Trace, Debug, Logging
C. Monitoring, Logging, Alerts, Error Reporting
D. Monitoring, Logging, Debug, Error Report

Answer: D

NEW QUESTION 9
- (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 10
- (Topic 5)
You want to store critical business information in Cloud Storage buckets. The information is regularly changed but previous versions need to be referenced on a regular basis. You want to ensure that there is a record of all changes to any information in these buckets. You want to ensure that accidental edits or deletions can be easily roiled back. Which feature should you enable?

A. Bucket Lock
B. Object Versioning
C. Object change notification
D. Object Lifecycle Management

Answer: B

NEW QUESTION 10
- (Topic 5)
You have developed an application using Cloud ML Engine that recognizes famous paintings from uploaded images. You want to test the application and allow specific people to upload images for the next 24 hours. Not all users have a Google Account. How should you have users upload images?

A. Have users upload the images to Cloud Storag
B. Protect the bucket with a passwordthat expires after 24 hours.
C. Have users upload the images to Cloud Storage using a signed URL that expires after 24 hours.
D. Create an App Engine web application where users can upload image
E. Configure App Engine to disable the application after 24 hour
F. Authenticate users via Cloud Identity.
G. Create an App Engine web application where users can upload images for the next 24 hour
H. Authenticate users via Cloud Identity.

Answer: A

Explanation:
https://cloud.google.com/blog/products/storage-data-transfer/uploading-images-directly-to-cloud-storage-by-using-signed-url

NEW QUESTION 12
- (Topic 5)
Your company has a Google Workspace account and Google Cloud Organization Some developers in the company have created Google Cloud projects outside of the Google Cloud Organization
You want to create an Organization structure that allows developers to create projects, but prevents them from modifying production projects You want to manage policies for all projects centrally and be able to set more restrictive policies for production projects
You want to minimize disruption to users and developers when business needs change in the future You want to follow Google-recommended practices How should you design the Organization structure?

A. * 1 Create a second Google Workspace account and Organization* 2 Grant all developers the Project Creator IAM role on the new Organization * 3 Move the developer projects into the new Organization* 4 Set the policies for all projects on both Organizations.* 5 Additionally set the production policies on the original Organization
B. * 1 Create a folder under the Organization resource named "Production '* 2 Grant all developers the Project Creator IAM role on the Organization * 3. Move the developer projects into the Organization* 4 Set the policies for all projects on the Organization* 5 Additionally set the production policies on the 'Production" folder
C. * 1 Create folders under the Organization resource named "Development" and Production'* 2 Grant all developers the Project Creator IAM role on the ""Development1 folder * 3. Move the developer projects into the "Development" folder* 4 Set the policies for all projects on the Organization* 5 Additionally set the production policies on the "Production" folder
D. * 1 Designate the Organization for production projects only* 2 Ensure that developers do not have the Project Creator IAM role on the Organization * 3 Create development projects outside of the Organization using the developer Google Workspace accounts* 4 Set the policies for all projects on the Organization* 5 Additionally set the production policies on the individual production projects

Answer: C

Explanation:
This option can help create an organization structure that allows developers to create projects, but prevents them from modifying production projects. Folders are containers for projects and other folders within Google Cloud organizations. Folders allow resources to be structured hierarchically and inherit policies from their parent resources. By creating folders under the organization resource named “Development” and “Production”, you can organize your projects by environment and apply different policies to them. By granting all developers the Project Creator IAM role on the “Development” folder, you can allow them to create projects under that folder, but not under the “Production” folder. By moving the developer projects into the “Development” folder, you can ensure that they are subject to the policies set on that folder. By setting the policies for all projects on the organization, you can manage policies centrally and efficiently. By additionally setting the production policies on the “Production” folder, you can enforce more restrictive policies for production projects and prevent developers from modifying them. The other options are not optimal for this scenario, because they either create a second Google Workspace account and organization, which increases complexity and cost (A), or do not use folders to organize projects by environment, which makes it harder to manage policies and permissions (B, D). References:
? https://cloud.google.com/resource-manager/docs/creating-managing-folders
? https://cloud.google.com/architecture/framework/system-design

NEW QUESTION 16
- (Topic 5)
You have a Python web application with many dependencies that requires 0.1 CPU cores and 128 MB of memory to operate in production. You want to monitor and maximize machine utilization. You also to reliably deploy new versions of the application. Which set of steps should you take?

A. Perform the following:1) Create a managed instance group with f1-micro type machines.2) Use a startup script to clone the repository, check out the production branch, install the dependencies, and start the Python app.3) Restart the instances to automatically deploy new production releases.
B. Perform the following:1) Create a managed instance group with n1-standard-1 type machines.2) Build a Compute Engine image from the production branch that contains all of the dependencies andautomatically starts the Python app.3) Rebuild the Compute Engine image, and update the instance template to deploy new productionreleases.
C. Perform the following:1) Create a Kubernetes Engine cluster with n1-standard-1 type machines.2) Build a Docker image from the production branch with all of the dependencies, and tag it with theversion number.3) Create a Kubernetes Deployment with the imagePullPolicy set to “IfNotPresent” in the stagingnamespace, and then promote it to the production namespace after testing.
D. Perform the following:1) Create a Kubernetes Engine (GKE) cluster with n1-standard-4 type machines.2) Build a Docker image from the master branch will all of the dependencies, and tag it with “latest”.3) Create a Kubernetes Deployment in the default namespace with the imagePullPolicy set to “Always”.Restart the pods to automatically deploy new production releases.

Answer: D

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

NEW QUESTION 18
- (Topic 5)
You need to ensure reliability for your application and operations by supporting reliable task a scheduling for compute on GCP. Leveraging Google best practices, what should you do?

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

Answer: B

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

NEW QUESTION 23
- (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 28
- (Topic 5)
Your company has a Google Cloud project that uses BigQuery for data warehousing They have a VPN tunnel between the on-premises environment and Google Cloud that is configured with Cloud VPN. The security team wants to avoid data exfiltration by malicious insiders, compromised code, and accidental oversharing.
What should they do?

A. Configure Private Google Access for on-premises only.
B. Perform the following tasks:1) Create a service account.2) Give the BigQuery JobUser role and Storage Reader role to the service account.3) Remove all other IAM access from the project.
C. Configure VPC Service Controls and configure Private Google Access.
D. Configure Private Google Access.

Answer: C

Explanation:
https://cloud.google.com/vpc-service-controls/docs/overview
VPC Service Controls improves your ability to mitigate the risk of data exfiltration from Google Cloud services such as Cloud Storage and BigQuery.

NEW QUESTION 33
- (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 therequired information.

Answer: C

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

NEW QUESTION 37
- (Topic 5)
You are helping the QA team to roll out a new load-testing tool to test the scalability of your primary cloud services that run on Google Compute Engine with Cloud Bigtable. Which three requirements should they include? Choose 3 answers

A. Ensure that the load tests validate the performance of Cloud Bigtable.
B. Create a separate Google Cloud project to use for the load-testing environment.
C. Schedule the load-testing tool to regularly run against the production environment.
D. Ensure all third-party systems your services use are capable of handling high load.
E. Instrument the production services to record every transaction for replay by the load- testing tool.
F. Instrument the load-testing tool and the target services with detailed logging and metrics collection.

Answer: ABF

NEW QUESTION 40
- (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 44
- (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 returndatabase 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 databasevalues 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 48
- (Topic 5)
You want to optimize the performance of an accurate, real-time, weather-charting application. The data comes from 50,000 sensors sending 10 readings a second, in the format of a timestamp and sensor reading. Where should you store the data?

A. Google BigQuery
B. Google Cloud SQL
C. Google Cloud Bigtable
D. Google Cloud Storage

Answer: C

Explanation:
It is time-series data, So Big Table. https://cloud.google.com/bigtable/docs/schema-design-time-series
Google Cloud Bigtable is a scalable, fully-managed NoSQL wide-column database that is suitable for both real-time access and analytics workloads.
Good for:
? Low-latency read/write access
? High-throughput analytics
? Native time series support
? Common workloads:
? IoT, finance, adtech
? Personalization, recommendations
? Monitoring
? Geospatial datasets
? Graphs
References: https://cloud.google.com/storage-options/

NEW QUESTION 51
- (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 52
- (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: AD

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 57

- (Topic 5)
One of the developers on your team deployed their application in Google Container Engine with the Dockerfile below. They report that their application deployments are taking too long.

```
FROM ubuntu:16.04

COPY . /src

RUN apt-get update && apt-get install -y python python-pip

RUN pip install -r requirements.txt
```

You want to optimize this Dockerfile for faster deployment times without adversely affecting the app's functionality.
Which two actions should you take? Choose 2 answers.

A. Remove Python after running pip.
B. Remove dependencies from requirements.txt.
C. Use a slimmed-down base image like Alpine linux.
D. Use larger machine types for your Google Container Engine node pools.
E. Copy the source after the package dependencies (Python and pip) are installed.

Answer: CE

Explanation:
The speed of deployment can be changed by limiting the size of the uploaded app, limiting the complexity of the build necessary in the Dockerfile, if present, and by ensuring a fast and reliable internet connection.
Note: Alpine Linux is built around musl libc and busybox. This makes it smaller and more resource efficient than traditional GNU/Linux distributions. A container requires no more
than 8 MB and a minimal installation to disk requires around 130 MB of storage. Not only do you get a fully-fledged Linux environment but a large selection of packages from the repository.
References: https://groups.google.com/forum/#!topic/google-appengine/hZMEkmmObDU https://www.alpinelinux.org/about/

NEW QUESTION 61
- (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 62
- (Topic 5)
You need to upload files from your on-premises environment to Cloud Storage. You want the files to be encrypted on Cloud Storage using customer-supplied encryption keys. What should you do?

A. Supply the encryption key in a .boto configuration fil
B. Use gsutil to upload the files.
C. Supply the encryption key using gcloud confi
D. Use gsutil to upload the files to that bucket.
E. Use gsutil to upload the files, and use the flag --encryption-key to supply the encryption key.
F. Use gsutil to create a bucket, and use the flag --encryption-key to supply the encryption ke
G. Use gsutil to upload the files to that bucket.

Answer: A

Explanation:
https://cloud.google.com/storage/docs/encryption/customer-supplied- keys#gsutil

NEW QUESTION 63
- (Topic 5)
Your company creates rendering software which users can download from the company website. Your company has customers all over the world. You want to minimize latency for all your customers. You want to follow Google-recommended practices.
How should you store the files?

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

Answer: A

Explanation:

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

NEW QUESTION 67
- (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 cluster.* 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 68
- (Topic 5)
Your company is migrating its on-premises data center into the cloud. As part of the migration, you want to integrate Kubernetes Engine for workload orchestration. Parts of your architecture must also be PCI DSScompliant.
Which of the following is most accurate?

A. App Engine is the only compute platform on GCP that is certified for PCI DSS hosting.
B. Kubernetes Engine cannot be used under PCI DSS because it is considered shared hosting.
C. Kubernetes Engine and GCP provide the tools you need to build a PCI DSS-compliant environment.
D. All Google Cloud services are usable because Google Cloud Platform is certified PCI- compliant.

Answer: D

Explanation:
https://cloud.google.com/security/compliance/pci-dss

NEW QUESTION 72
- (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 . DelegatingMethodAccesorImp1.

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 75
- (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 Storagebucket 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: A

Explanation:
? An aggregated export is a type of sink that combines and routes log entries from the Google Cloud resources contained by an organization or folder1. By creating an aggregated export on the Production folder, you can capture all the logs from the existing and new production projects in that folder automatically1.
? A log sink is a destination for log entries that match a filter1. By setting the log sink to be a Cloud Storage bucket in an operations project, you can store the log entries in Cloud Storage and allow your operations team to access them1.

NEW QUESTION 80
- (Topic 5)
You are deploying an application to Google Cloud. The application is part of a system. The application in Google Cloud must communicate over a private network with applications in a non-Google Cloud environment. The expected average throughput is 200 kbps. The business requires:
• 99.99% system availability
• cost optimization
You need to design the connectivity between the locations to meet the business requirements. What should you provision?

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

Answer: C

Explanation:
https://cloud.google.com/network-connectivity/docs/vpn/concepts/topologies#configurations_that_support_9999_availability

NEW QUESTION 85
- (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 88
- (Topic 5)
Your development team has installed a new Linux kernel module on the batch servers in Google Compute Engine (GCE) virtual machines (VMs) to speed up the nightly batch process. Two days after the installation, 50% of web application deployed in the same
nightly batch run. You want to collect details on the failure to pass back to the development team. Which three actions should you take? Choose 3 answers

A. Use Stackdriver Logging to search for the module log entries.
B. Read the debug GCE Activity log using the API or Cloud Console.
C. Use gcloud or Cloud Console to connect to the serial console and observe the logs.
D. Identify whether a live migration event of the failed server occurred, using in the activity log.
E. Adjust the Google Stackdriver timeline to match the failure time, and observe the batch server metrics.
F. Export a debug VM into an image, and run the image on a local server where kernel log messages will be displayed on the native screen.

Answer: ACE

Explanation:
https://www.flexera.com/blog/cloud/2013/12/google-compute-engine-live-migration-passes- the-test/
"With live migration, the virtual machines are moved without any downtime or noticeable service degradation"

NEW QUESTION 93
- (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 95
- (Topic 5)
Your company has a Google Cloud project that uses BigQuery for data warehousing. The VPN tunnel between the on-premises environment and Google Cloud is configured with Cloud VPN. Your security team wants to avoid data exfiltration by malicious insiders, compromised code, and accidental oversharing. What should you do?

A. Configure VPC Service Controls and configure Private Google Access for on-premises hosts.
B. Create a service account, grant the BigQuery JobUser role and Storage Object Viewer role to the service account, and remove all other Identity and Access Management (1AM) access from the project.
C. Configure Private Google Access.
D. Configure Private Service Connect.

Answer: A

NEW QUESTION 100
- (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: A

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

NEW QUESTION 101
- (Topic 5)
Google Cloud Platform resources are managed hierarchically using organization, folders, and projects. When Cloud Identity and Access Management (IAM) policies exist at these different levels, what is the effective policy at a particular node of the hierarchy?

A. The effective policy is determined only by the policy set at the node
B. The effective policy is the policy set at the node and restricted by the policies of its ancestors
C. The effective policy is the union of the policy set at the node and policies inherited from its ancestors
D. The effective policy is the intersection of the policy set at the node and policies inherited from its ancestors

Answer: B

Explanation:
Reference: https://cloud.google.com/resource-manager/docs/cloud-platform-resource- hierarchy

NEW QUESTION 103
- (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 transferarchived data to Cloud Storag
C. Establish a connection with Google using a Dedicated Interconnect orDirect 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 108
- (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 snapshot.Create anautoscaled managed instance group from the instance template.
C. Create a snapshot of the existing dis
D. Create a custom image from the snapsho
E. Create an autoscaled managed instance group from the custom image.
F. Create a custom image from the existing dis
G. Create an instance template from the custom imag
H. 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 111
- (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 114
- (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 116
- (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 oftype 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-engine-with-kubemci

NEW QUESTION 117
- (Topic 5)
You need to migrate Hadoop jobs for your company's Data Science team without modifying the underlying infrastructure. You want to minimize costs and infrastructure management effort. What should you do?

A. Create a Dataproc cluster using standard worker instances.
B. Create a Dataproc cluster using preemptible worker instances.
C. Manually deploy a Hadoop cluster on Compute Engine using standard instances.
D. Manually deploy a Hadoop cluster on Compute Engine using preemptible instances.

Answer: B

Explanation:
Reference: https://cloud.google.com/architecture/hadoop/hadoop-gcp-migration-jobs

NEW QUESTION 120
- (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 canarytesting.
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 125
- (Topic 5)
A recent audit that a new network was created in Your GCP project. In this network, a GCE instance has an SSH port open the world. You want to discover this network's origin. What should you do?

A. Search for Create VM entry in the Stackdriver alerting console.
B. Navigate to the Activity page in the Home sectio
C. Set category to Data Access and search for Create VM entry.
D. In the logging section of the console, specify GCE Network as the logging sectio
E. Search for the Create Insert entry.
F. Connect to the GCE instance using project SSH Key
G. Identify previous logins in system logs, and match these with the project owners list.

Answer: C

NEW QUESTION 126
- (Topic 5)
You need to evaluate your team readiness for a new GCP project. You must perform the evaluation and create a skills gap plan incorporates the business goal of cost optimization. Your team has deployed two GCP projects successfully to date. What should you do?

A. Allocate budget for team trainin
B. Set a deadline for the new GCP project.
C. Allocate budget for team trainin
D. Create a roadmap for your team to achieve Google Cloud certification based on job role.
E. Allocate budget to hire skilled external consultant
F. Set a deadline for the new GCP project.
G. Allocate budget to hire skilled external consultant
H. Create a roadmap for your team to achieve Google Cloud certification based on job role.

Answer: B

Explanation:
https://services.google.com/fh/files/misc/cloud_center_of_excellence.pdf

NEW QUESTION 130
- (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 134
- (Topic 5)
You are designing a Data Warehouse on Google Cloud and want to store sensitive data in BigQuery. Your company requires you to generate encryption keys outside of Google Cloud. You need to implement a solution. What should you do?

A. Generate a new key in Cloud Key Management Service (Cloud KMS). Store all data in Cloud Storage using the customer-managed key option and select the created ke
B. Set up a Dataflow pipeline to decrypt the data and to store it in a BigQuery dataset.
C. Generate a new key in Cloud Key Management Service (Cloud KMS). Create a dataset in BigQuery using the customer-managed key option and select the created key
D. Import a key in Cloud KM
E. Store all data in Cloud Storage using the customer- managed key option and select the created ke
F. Set up a Dataflow pipeline to decrypt the data and to store it in a new BigQuery dataset.
G. Import a key in Cloud KM
H. Create a dataset in BigQuery using the customer-supplied key option and select the created key.

Answer: D

Explanation:
https://cloud.google.com/bigquery/docs/customer-managed-encryption

NEW QUESTION 138
- (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

A.

Answer: D

NEW QUESTION 139
- (Topic 5)
You are managing an application deployed on Cloud Run for Anthos, and you need to define a strategy for deploying new versions of the application. You want to evaluate the new code with a subset of production traffic to decide whether to proceed with the rollout. What should you do?

A. Deploy a new revision to Cloud Run with the new versio
B. Configure traffic percentage between revisions.
C. Deploy a new service to Cloud Run with the new versio
D. Add a Cloud Load Balancing instance in front of both services.
E. In the Google Cloud Console page for Cloud Run, set up continuous deployment using Cloud Build for the development branc
F. As part of the Cloud Build trigger, configure the substitution variable TRAFFIC_PERCENTAGE with the percentage of traffic you want directed to a new version.
G. In the Google Cloud Console, configure Traffic Director with a new Service that points to the new version of the application on Cloud Ru
H. Configure Traffic Director to send a small percentage of traffic to the new version of the application.

Answer: A

Explanation:
https://cloud.google.com/run/docs/rollouts-rollbacks-traffic-migration

NEW QUESTION 142
- (Topic 5)
Your company has an application running on App Engine that allows users to upload music files and share them with other people. You want to allow users to upload files directly into Cloud Storage from their browser session. The payload should not be passed through the backend. What should you do?
A.
* 1. Set a CORS configuration in the target Cloud Storage bucket where the base URL of the App
Engine application is an allowed origin.
* 2. Use the Cloud Storage Signed URL feature to generate a POST URL.
B.
* 1. Set a CORS configuration in the target Cloud Storage bucket where the base URL of the App
Engine application is an allowed origin.
* 2. Assign the Cloud Storage WRITER role to users who upload files.
C.
* 1. Use the Cloud Storage Signed URL feature to generate a POST URL.

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

A.

Answer: B

NEW QUESTION 146
- (Topic 5)
You are using a single Cloud SQL instance to serve your application from a specific zone.
You want to
introduce high availability. What should you do?

A. Create a read replica instance in a different region
B. Create a failover replica instance in a different region
C. Create a read replica instance in the same region, but in a different zone
D. Create a failover replica instance in the same region, but in a different zone

Answer: B

Explanation:
https://cloud.google.com/sql/docs/mysql/high-availability

NEW QUESTION 148
- (Topic 5)
Your company uses the Firewall Insights feature in the Google Network Intelligence Center. You have several firewall rules applied to Compute Engine instances. You need to evaluate the efficiency of the applied firewall ruleset. When you bring up the Firewall Insights page in the Google Cloud Console, you notice that there are no log rows to display. What should you do to troubleshoot the issue?

A. Enable Virtual Private Cloud (VPC) flow logging.
B. Enable Firewall Rules Logging for the firewall rules you want to monitor.
C. Verify that your user account is assigned the compute.networkAdmin Identity and Access Management (IAM) role.
D. Install the Google Cloud SDK, and verify that there are no Firewall logs in the command line output.

Answer: B

Explanation:

Reference: https://cloud.google.com/network-intelligence-center/docs/firewall-insights/how-to/using-firewall- insights

NEW QUESTION 153
- (Topic 5)
Your company has an application running on a deployment in a GKE cluster. You have a separate cluster for development, staging and production. You have discovered that the team is able to deploy a Docker image to the production cluster without first testing the deployment in development and then staging. You want to allow the team to have autonomy but want to prevent this from happening. You want a Google Cloud solution that can be implemented quickly with minimal effort. What should you do?

A. Create a Kubernetes admission controller to prevent the container from starting if it is not approved for usage in the given environment
B. Configure a Kubernetes lifecycle hook to prevent the container from starting if it is not approved for usage in the given environment
C. Implement a corporate policy to prevent teams from deploying Docker image to an environment unless the Docker image was tested in an earlier environment
D. Configure the binary authorization policies for the development, staging and production cluster
E. Create attestations as part of the continuous integration pipeline”

Answer: D

Explanation:
https://cloud.google.com/architecture/prep-kubernetes-engine-for- prod#binary-authorization
The most common Binary Authorization use cases involve attestations. An attestation certifies that a specific image has completed a previous stage, as described previously. You configure the Binary Authorization policy to verify the attestation before allowing the image to be deployed. At deploy time, instead of redoing activities that were completed in earlier stages, Binary Authorization only needs to verify the attestation. https://cloud.google.com/binary-authorization/docs/overview

NEW QUESTION 156
- (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 159
- (Topic 5)
The operations manager asks you for a list of recommended practices that she should consider when migrating a J2EE application to the cloud. Which three practices should you recommend? Choose 3 answers

A. Port the application code to run on Google App Engine.
B. Integrate Cloud Dataflow into the application to capture real-time metrics.
C. Instrument the application with a monitoring tool like Stackdriver Debugger.
D. Select an automation framework to reliably provision the cloud infrastructure.
E. Deploy a continuous integration tool with automated testing in a staging environment.
F. Migrate from MySQL to a managed NoSQL database like Google Cloud Datastore or Bigtable.

Answer: AEF

Explanation:
References: https://cloud.google.com/appengine/docs/standard/java/tools/uploadinganapp https://cloud.google.com/appengine/docs/standard/java/building-app/cloud-sql

NEW QUESTION 164
- (Topic 5)
Your company has decided to build a backup replica of their on-premises user authentication PostgreSQL database on Google Cloud Platform. The database is 4 TB, and large updates are frequent. Replication requires private address space communication. Which networking approach should you use?

A. Google Cloud Dedicated Interconnect
B. Google Cloud VPN connected to the data center network
C. A NAT and TLS translation gateway installed on-premises
D. A Google Compute Engine instance with a VPN server installed connected to the data center network

Answer: A

Explanation:
https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations
Google Cloud Dedicated Interconnect provides direct physical connections and RFC 1918 communication between your on-premises network and Google's network. Dedicated Interconnect enables you to transfer large amounts of data between networks, which can be more cost effective than purchasing additional bandwidth over the public Internet or using VPN tunnels.
Benefits:
? Traffic between your on-premises network and your VPC network doesn't traverse the public Internet. Traffic traverses a dedicated connection with fewer hops, meaning there are less points of failure where traffic might get dropped or disrupted.
? Your VPC network's internal (RFC 1918) IP addresses are directly accessible from your on-premises network. You don't need to use a NAT device or VPN tunnel to reach internal IP addresses. Currently, you can only reach internal IP addresses over a dedicated connection. To reach Google external IP addresses, you must use a separate connection.
? You can scale your connection to Google based on your needs. Connection capacity is delivered over one or more 10 Gbps Ethernet connections, with a maximum of eight connections (80 Gbps total per interconnect).
? The cost of egress traffic from your VPC network to your on-premises network is reduced. A dedicated connection is generally the least expensive method if you have a high-volume of traffic to and from Google's network.
References: https://cloud.google.com/interconnect/docs/details/dedicated

NEW QUESTION 167
- (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-premises virtual machine in the app.yaml file.
F. Create an instance template with the smallest available machine type, and use an imageof 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-rightsizing?hl=en

NEW QUESTION 169
- (Topic 5)
You have an application that will run on Compute Engine. You need to design an architecture that takes into account a disaster recovery plan that requires your application to fail over to another region in case of a regional outage. What should you do?

A. Deploy the application on two Compute Engine instances in the same project but in a different regio
B. Use the first instance to serve traffic, and use the HTTP load balancing service to fail over to the standby instance in case of a disaster.
C. Deploy the application on a Compute Engine instanc
D. Use the instance to serve traffic, and use the HTTP load balancing service to fail over to an instance on your premises in case of a disaster.
E. Deploy the application on two Compute Engine instance groups, each in the same project but in a different regio
F. Use the first instance group to serve traffic, and use the HTTP load balancing service to fail over to the standby instance group in case of a disaster.
G. Deploy the application on two Compute Engine instance groups, each in separate project and a different regio
H. Use the first instance group to server traffic, and use the HTTP load balancing service to fail over to the standby instance in case of a disaster.

Answer: C

NEW QUESTION 174
- (Topic 5)
Your company is developing a new application that will allow globally distributed users to upload pictures and share them with other selected users. The application will support millions of concurrent users. You want to allow developers to focus on just building code without having to create and maintain the underlying infrastructure. Which service should you use to deploy the application?

A. App Engine
B. Cloud Endpoints
C. Compute Engine
D. Google Kubernetes Engine

Answer: A

Explanation:
Reference: https://cloud.google.com/terms/services https://cloud.google.com/appengine/docs/standard/go/how-requests-are-handled

NEW QUESTION 176
- (Topic 5)
You are working at a sports association whose members range in age from 8 to 30. The association collects a large amount of health data, such as sustained injuries. You are storing this data in BigQuery. Current legislation requires you to delete such information upon request of the subject. You want to design a solution that can accommodate such a request. What should you do?

A. Use a unique identifier for each individua
B. Upon a deletion request, delete all rows from BigQuery with this identifier.
C. When ingesting new data in BigQuery, run the data through the Data Loss Prevention (DLP) API to identify any personal informatio
D. As part of the DLP scan, save the result to Data Catalo
E. Upon a deletion request, query Data Catalog to find the column with personal information.
F. Create a BigQuery view over the table that contains all dat
G. Upon a deletion request, exclude the rows that affect the subject's data from this vie
H. Use this view instead of the source table for all analysis tasks.
I. Use a unique identifier for each individua
J. Upon a deletion request, overwrite the column with the unique identifier with a salted SHA256 of its value.

Answer: B

Explanation:
Current legislation requires you to delete "SUCH" information upon request of the subject. " So from that point of view the question is not to delete the entire user records but specific data related to personal health data. With DLP you can use InfoTypes and InfoType detectors to specifically scan for those entries and how to act upon them (link https://cloud.google.com/dlp/docs/concepts-infotypes)
https://cloud.google.com/dlp#section-6

NEW QUESTION 177
- (Topic 5)
Your customer is moving an existing corporate application to Google Cloud Platform from an on-premises data center. The business owners require minimal user disruption. There are strict security team requirements for storing passwords. What authentication strategy should they use?

A. Use G Suite Password Sync to replicate passwords into Google.
B. Federate authentication via SAML 2.0 to the existing Identity Provider.
C. Provision users in Google using the Google Cloud Directory Sync tool.
D. Ask users to set their Google password to match their corporate password.

Answer: B

Explanation:
https://cloud.google.com/solutions/authenticating-corporate-users-in-a-hybrid-environment

NEW QUESTION 178
- (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 180
- (Topic 5)
Your company wants to try out the cloud with low risk. They want to archive approximately 100 TB of their log data to the cloud and test the analytics features available to them there, while also retaining that data as a long-term disaster recovery backup. Which two steps should they take? Choose 2 answers

A. Load logs into Google BigQuery.
B. Load logs into Google Cloud SQL.
C. Import logs into Google Stackdriver.
D. Insert logs into Google Cloud Bigtable.
E. Upload log files into Google Cloud Storage.

Answer: AE

NEW QUESTION 185
- (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 StackdriverMonitoring 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 188
- (Topic 5)
Your architecture calls for the centralized collection of all admin activity and VM system logs within your
project.
How should you collect these logs from both VMs and services?

A. All admin and VM system logs are automatically collected by Stackdriver.
B. Stackdriver automatically collects admin activity logs for most service
C. The Stackdriver Logging agentmust be installed on each instance to collect system logs.
D. Launch a custom syslogd compute instance and configure your GCP project and VMs to forward all logs to it.
E. Install the Stackdriver Logging agent on a single compute instance and let it collect all audit and access logs for your environment.

Answer: B

Explanation:
https://cloud.google.com/logging/docs/agent/default-logs

NEW QUESTION 193
- (Topic 5)
Your company provides a recommendation engine for retail customers. You are providing retail customers with an API where they can submit a user ID and the API returns a list of recommendations for that user. You are responsible for the API lifecycle and want to ensure stability for your customers in case the API makes backward-incompatible changes. You want to follow Google-recommended practices. What should you do?

A. Create a distribution list of all customers to inform them of an upcoming backward- incompatible change at least one month before replacing the old API with the new API.
B. Create an automated process to generate API documentation, and update the public API documentation as part of the CI/CD process when deploying an update to the API.
C. Use a versioning strategy for the APIs that increases the version number on every backward-incompatible change.
D. Use a versioning strategy for the APIs that adds the suffix “DEPRECATED” to the current API version number on every backward-incompatible chang
E. Use the current version number for the new API.

Answer: C

Explanation:
https://cloud.google.com/apis/design/versioning
All Google API interfaces must provide a major version number, which is encoded at the end of the protobuf package, and included as the first part of the URI path for REST APIs. If an API introduces a breaking change, such as removing or renaming a field, it must increment its API version number to ensure that existing user code does not suddenly break.

NEW QUESTION 198
- (Topic 5)
Your company is moving 75 TB of data into Google Cloud. You want to use Cloud Storage and follow Googlerecommended practices. What should you do?

A. Move your data onto a Transfer Applianc
B. Use a Transfer Appliance Rehydrator to decrypt the data into Cloud Storage.
C. Move your data onto a Transfer Applianc
D. Use Cloud Dataprep to decrypt the data into Cloud Storage.

E. Install gsutil on each server that contains dat
F. Use resumable transfers to upload the data into Cloud Storage.
G. Install gsutil on each server containing dat
H. Use streaming transfers to upload the data into CloudStorage.

Answer: A

Explanation:
https://cloud.google.com/transfer-appliance/docs/2.0/faq

NEW QUESTION 202
- (Topic 5)
You have deployed several instances on Compute Engine. As a security requirement, instances cannot have a public IP address. There is no VPN connection between Google
Cloud and your office, and you need to connect via SSH into a specific machine without violating the security requirements. What should you do?

A. Configure Cloud NAT on the subnet where the instance is hoste
B. Create an SSH connection to the Cloud NAT IP address to reach the instance.
C. Add all instances to an unmanaged instance grou
D. Configure TCP Proxy Load Balancing with the instance group as a backen
E. Connect to the instance using the TCP Proxy IP.
F. Configure Identity-Aware Proxy (IAP) for the instance and ensure that you have the role of IAP-secured Tunnel Use
G. Use the gcloud command line tool to ssh into the instance.
H. Create a bastion host in the network to SSH into the bastion host from your office locatio
I. From the bastion host, SSH into the desired instance.

Answer: C

Explanation:
https://cloud.google.com/iap/docs/using-tcp-forwarding#tunneling_with_ssh
Leveraging the BeyondCorp security model. "This January, we enhanced context-aware access capabilities in Cloud Identity-Aware Proxy (IAP) to help you protect SSH and RDP access to your virtual machines (VMs)—without needing to provide your VMs with public IP addresses, and without having to set up bastion hosts. "
https://cloud.google.com/blog/products/identity-security/cloud-iap-enables-context-aware- access-to-vms-via-ssh-and-rdp-without-bastion-hosts
Reference: https://cloud.google.com/solutions/connecting-securely

NEW QUESTION 206
- (Topic 5)
Your company is using BigQuery as its enterprise data warehouse. Data is distributed over several Google Cloud projects. All queries on BigQuery need to be billed on a single
project. You want to make sure that no query costs are incurred on the projects that contain the data. Users should be able to query the datasets, but not edit them.
How should you configure users' access roles?

A. Add all users to a grou
B. Grant the group the role of BigQuery user on the billing project and BigQuerydataViewer on the projects that contain the data.
C. Add all users to a grou
D. Grant the group the roles of BigQuery dataViewer on the billing project andBigQuery user on the projects that contain the data.
E. Add all users to a grou
F. Grant the group the roles of BigQuery jobUser on the billing project and BigQuery dataViewer on the projects that contain the data.
G. Add all users to a grou
H. Grant the group the roles of BigQuery dataViewer on the billing project andBigQuery jobUser on the projects that contain the data.

Answer: A

Explanation:
Reference: https://cloud.google.com/bigquery/docs/running-queries

NEW QUESTION 208
- (Topic 5)
Your company has a Google Cloud project that uses BlgQuery for data warehousing There are some tables that contain personally identifiable information (PI!) Only the compliance team may access the PH. The other information in the tables must be available to the data science team. You want to minimize cost and the time it takes to assign appropriate access to the tables What should you do?

A. * 1 From the dataset where you have the source data, create views of tables that you want to share, excluding Pll* 2 Assign an appropriate project-level IAM role to the members of the data science team 3 Assign access controls to the dataset that contains the view
B. * 1 From the dataset where you have the source data, create materialized views of tables that you want to share excluding Pll* 2 Assign an appropriate project-level IAM role to the members of the data science team 3. Assign access controls to the dataset that contains the view.
C. * 1 Create a dataset for the data science team* 2 Create views of tables that you want to share excluding Pll* 3 Assign an appropriate project-level IAM role to the members of the data science team 4 Assign access controls to the dataset that contains the view* 5 Authorize the view to access the source dataset
D. * 1. Create a dataset for the data science team.* 2. Create materialized views of tables that you want to share, excluding Pll* 3. Assign an appropriate project-level IAM role to the members of the data science team * 4 Assign access controls to the dataset that contains the view* 5 Authorize the view to access the source dataset

Answer: C

Explanation:
This option can help minimize cost and time by using views and authorized datasets. Views are virtual tables defined by a SQL query that can exclude PII columns from the source tables. Views do not incur storage costs and do not duplicate data. Authorized datasets are datasets that have access to another dataset's data without granting direct access to individual users or groups. By creating a dataset for the data science team and creating views of tables that exclude PII, you can share only the relevant information with the team. By assigning an appropriate project-level IAM role to the members of the data science

team, you can grant them access to the BigQuery service and resources. By assigning access controls to the dataset that contains the view, you can grant them access to query the views. By authorizing the view to access the source dataset, you can enable the view to read data from the source tables without exposing PII. The other options are not optimal for this scenario, because they either use materialized views instead of views, which incur storage costs and duplicate data (B, D), or do not create a separate dataset for the data science team, which makes it harder to manage access controls (A). References:
? https://cloud.google.com/bigquery/docs/views
? https://cloud.google.com/bigquery/docs/authorized-datasets

NEW QUESTION 209
- (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 211
- (Topic 5)
Your company has successfully migrated to the cloud and wants to analyze their data stream to optimize operations. They do not have any existing code for this analysis, so they are exploring all their options. These options include a mix of batch and stream processing, as they are running some hourly jobs and live-processing some data as it comes in. Which technology should they use for this?

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

Answer: B

Explanation:
Dataflow is for processing both the Batch and Stream.
Cloud Dataflow is a fully-managed service for transforming and enriching data in stream (real time) and batch (historical) modes with equal reliability and expressiveness -- no more complex workarounds or compromises needed.
References: https://cloud.google.com/dataflow/

NEW QUESTION 215
- (Topic 6)
For this question, refer to the Dress4Win case study. You are responsible for the security of data stored in
Cloud Storage for your company, Dress4Win. You have already created a set of Google Groups and assigned the appropriate users to those groups. You should use Google best practices and implement the simplest design to meet the requirements.
Considering Dress4Win's business and technical requirements, what should you do?

A. Assign custom IAM roles to the Google Groups you created in order to enforce security requirements.Encrypt data with a customer-supplied encryption key when storing files in Cloud Storage.
B. Assign custom IAM roles to the Google Groups you created in order to enforce security requirements.Enable default storage encryption before storing files in Cloud Storage.
C. Assign predefined IAM roles to the Google Groups you created in order to enforce security requirements.Utilize Google's default encryption at rest when storing files in Cloud Storage.
D. Assign predefined IAM roles to the Google Groups you created in order to enforce security requirement
E. Ensure that the default Cloud KMS key is set before storing files in Cloud Storage.

Answer: D

Explanation:
https://cloud.google.com/iam/docs/understanding-service-accounts

NEW QUESTION 220
- (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 223
- (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 Storagebucke
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 226
- (Topic 7)
TerramEarth has about 1 petabyte (PB) of vehicle testing data in a private data center. You want to move the data to Cloud Storage for your machine learning team. Currently, a 1- Gbps interconnect link is available for you. The machine learning team wants to start using
the data in a month. What should you do?

A. Request Transfer Appliances from Google Cloud, export the data to appliances, and return the appliances to Google Cloud.
B. Configure the Storage Transfer service from Google Cloud to send the data from your data center to Cloud Storage
C. Make sure there are no other users consuming the 1 Gbps link, and use multi-thread transfer to upload the data to Cloud Storage.
D. Export files to an encrypted USB device, send the device to Google Cloud, and request an import of the data to Cloud Storage

Answer: A

NEW QUESTION 230
- (Topic 8)
For this question, refer to the Mountkirk Games case study. Mountkirk Games wants to migrate from their current analytics and statistics reporting model to one that meets their technical requirements on Google Cloud Platform.
Which two steps should be part of their migration plan? (Choose two.)

A. Evaluate the impact of migrating their current batch ETL code to Cloud Dataflow.
B. Write a schema migration plan to denormalize data for better performance in BigQuery.
C. Draw an architecture diagram that shows how to move from a single MySQL database to a MySQL cluster.
D. Load 10 TB of analytics data from a previous game into a Cloud SQL instance, and run test queries against the full dataset to confirm that they complete successfully.
E. Integrate Cloud Armor to defend against possible SQL injection attacks in analytics files uploaded to Cloud Storage.

Answer: AB

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

NEW QUESTION 233
- (Topic 8)
Mountkirk Games wants you to secure the connectivity from the new gaming application platform to Google
Cloud. You want to streamline the process and follow Google-recommended practices. What should you do?

A. Configure Workload Identity and service accounts to be used by the application platform.
B. Use Kubernetes Secrets, which are obfuscated by defaul
C. Configure these Secrets to be used by theapplication platform.
D. Configure Kubernetes Secrets to store the secret, enable Application-Layer Secrets Encryption, and useCloud Key Management Service (Cloud KMS) to manage the encryption key
E. Configure these Secrets tobe used by the application platform.
F. Configure HashiCorp Vault on Compute Engine, and use customer managed encryption keys and CloudKey Management Service (Cloud KMS) to manage the encryption key
G. Configure these Secrets to be usedby the application platform.

Answer: A

NEW QUESTION 235
- (Topic 8)
For this question, refer to the Mountkirk Games case study. You need to analyze and define the technical architecture for the compute workloads for your company, Mountkirk Games. Considering the Mountkirk Games business and technical requirements, what should you do?

A. Create network load balancer
B. Use preemptible Compute Engine instances.
C. Create network load balancer
D. Use non-preemptible Compute Engine instances.
E. Create a global load balancer with managed instance groups and autoscaling policie
F. Use preemptible Compute Engine instances.
G. Create a global load balancer with managed instance groups and autoscaling policie
H. Use non-preemptible Compute Engine instances.

Answer: D

NEW QUESTION 237
- (Topic 8)
You need to optimize batch file transfers into Cloud Storage for Mountkirk Games' new Google Cloud solution. The batch files contain game statistics that need to be staged in Cloud Storage and be processed by an extract transform load (ETL) tool. What should you do?

A. Use gsutil to batch move files in sequence.
B. Use gsutil to batch copy the files in parallel.
C. Use gsutil to extract the files as the first part of ETL.
D. Use gsutil to load the files as the last part of ETL.

Answer: B

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

NEW QUESTION 238
- (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 thecontainers, and test the mobile app.
D. Upload your mobile app with different configurations to Firebase Hosting and test each configuration.

Answer: C

NEW QUESTION 241
- (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 databy season and event.
C. Use Cloud Spanner for its scalability and ability to version schemas with zero downtime.Split race datausing season as a primary key.
D. Use BigQuery for its scalability and ability to add columns to a schem
E. Partition race data based onseason.
F. Use Cloud SQL for its ability to automatically manage storage increases and compatibility with MySQ
G. Useseparate database instances for each season.

Answer: C

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

NEW QUESTION 242
- (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/httpsD18912E1457D5D1DDCBD40AB3BF70D5D

NEW QUESTION 247

- (Topic 10)
For this question, refer to the EHR Healthcare case study. In the past, configuration errors put public IP addresses on backend servers that should not have been accessible from the Internet. You need to ensure that no one can put external IP addresses on backend Compute Engine instances and that external IP addresses can only be configured on frontend Compute Engine instances. What should you do?

A. Create an Organizational Policy with a constraint to allow external IP addresses only on the frontend Compute Engine instances.
B. Revoke the compute.networkAdmin role from all users in the project with front end instances.
C. Create an Identity and Access Management (IAM) policy that maps the IT staff to the compute.networkAdmin role for the organization.
D. Create a custom Identity and Access Management (IAM) role named GCE_FRONTEND with the compute.addresses.create permission.

Answer: A

Explanation:
https://cloud.google.com/compute/docs/ip-addresses/reserve-static-external-ip-address#disableexternalip

NEW QUESTION 249
- (Topic 10)
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: D

NEW QUESTION 253
- (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 256
.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

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