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 1)
For this question, refer to the Mountkirk Games case study.
Mountkirk Games has deployed their new backend on Google Cloud Platform (GCP). You want to create a thorough testing process for new versions of the backend before they are released to the public. You want the testing environment to scale in an economical way. How should you design the process?

A. Create a scalable environment in GCP for simulating production load.
B. Use the existing infrastructure to test the GCP-based backend at scale.
C. Build stress tests into each component of your application using resources internal to GCP to simulate load.
D. Create a set of static environments in GCP to test different levels of load — for example, high, medium, and low.

Answer: A

Explanation:
From scenario: Requirements for Game Backend Platform
? Dynamically scale up or down based on game activity
? Connect to a managed NoSQL database service
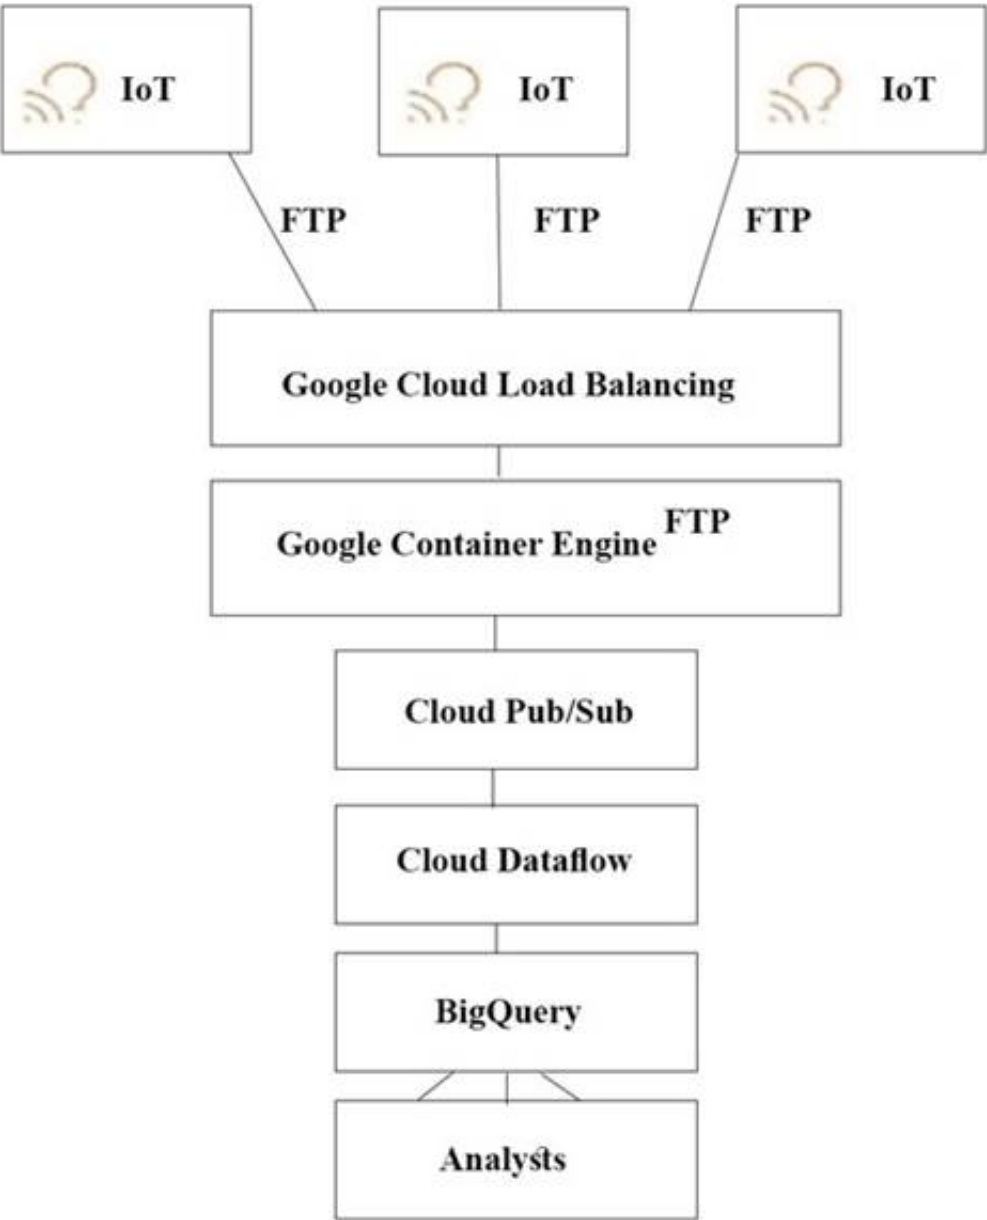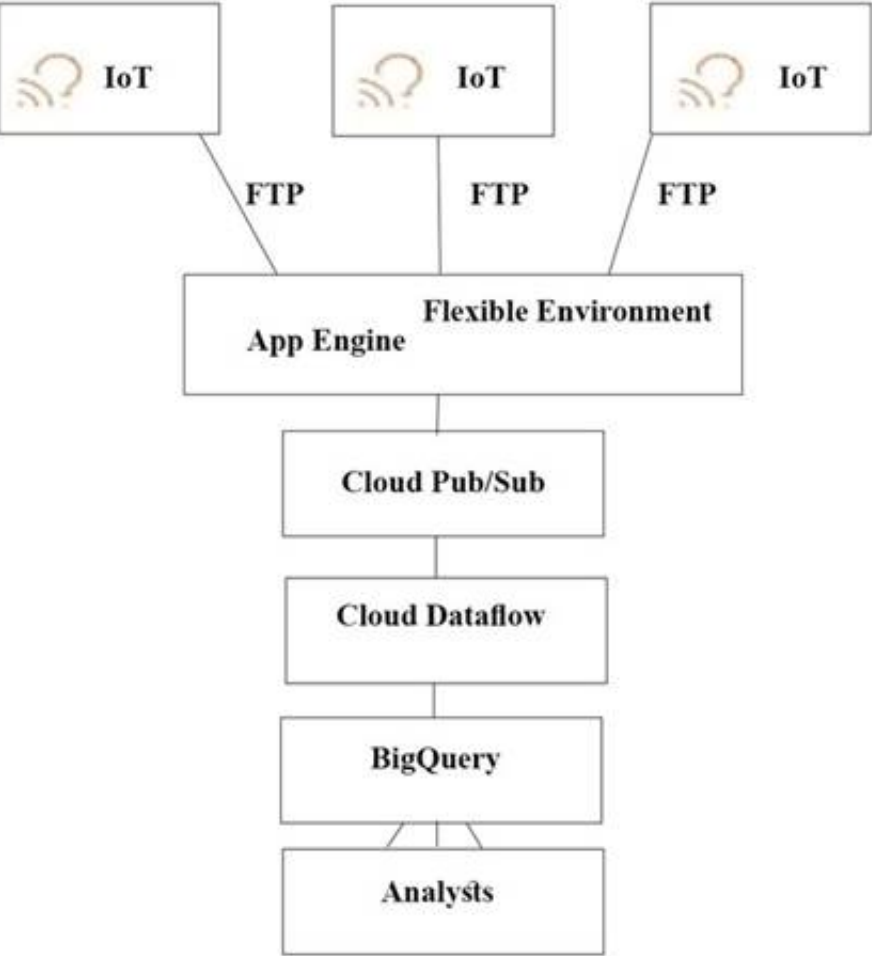? Run customize Linux distro

NEW QUESTION 3
- (Topic 2)
For this question, refer to the TerramEarth case study.
TerramEarth's CTO wants to use the raw data from connected vehicles to help identify approximately when a vehicle in the development team to focus their failure. You want to allow analysts to centrally query the vehicle data. Which architecture should you recommend?
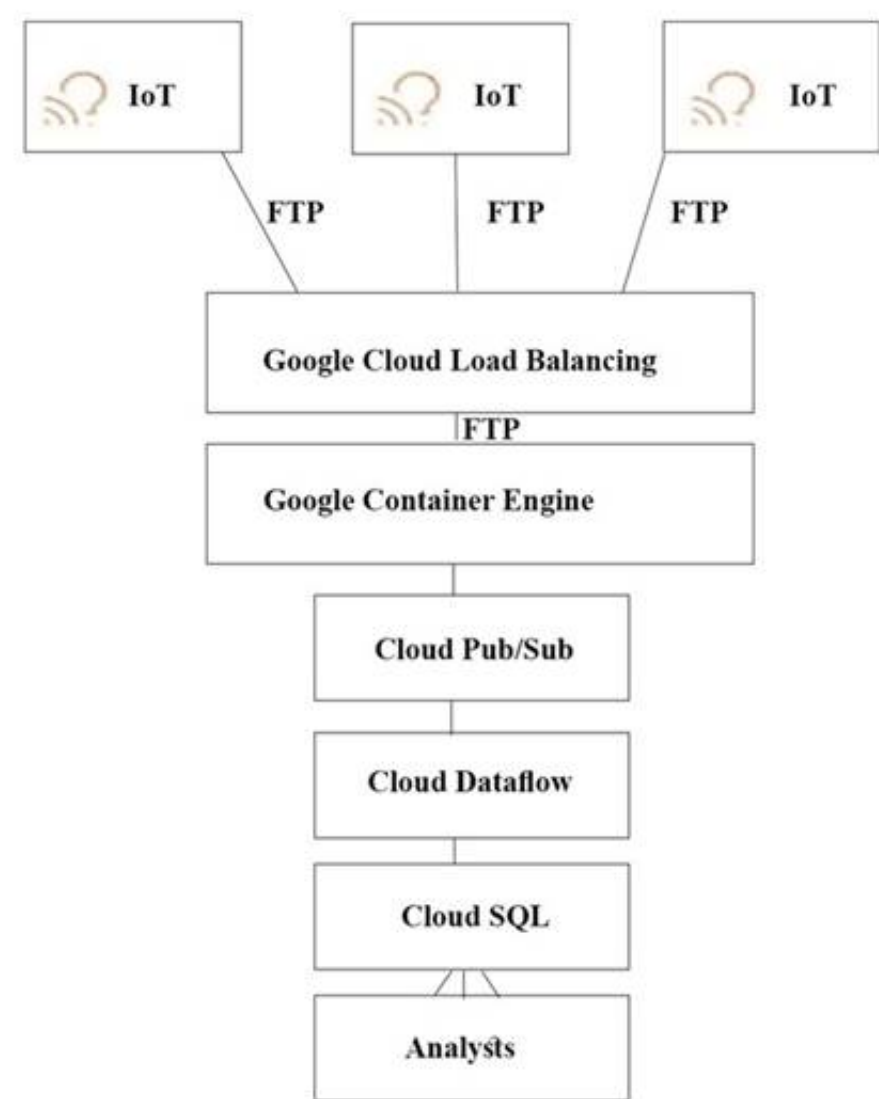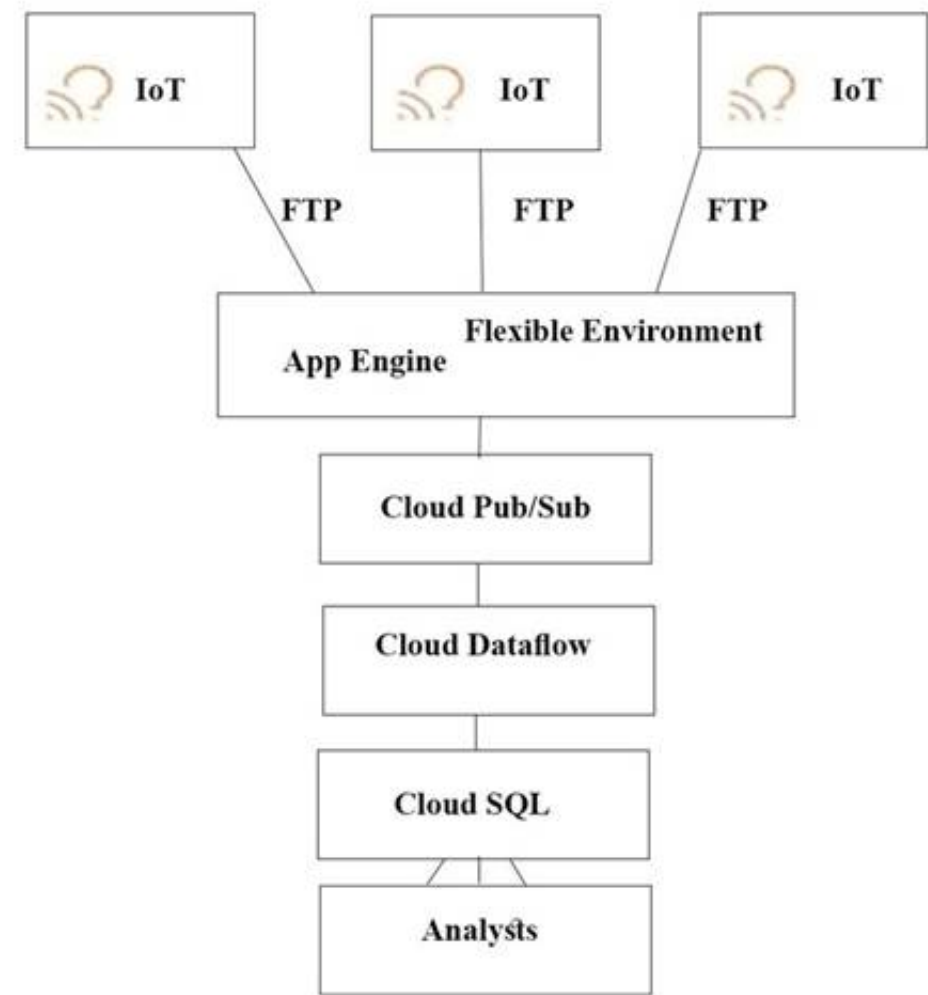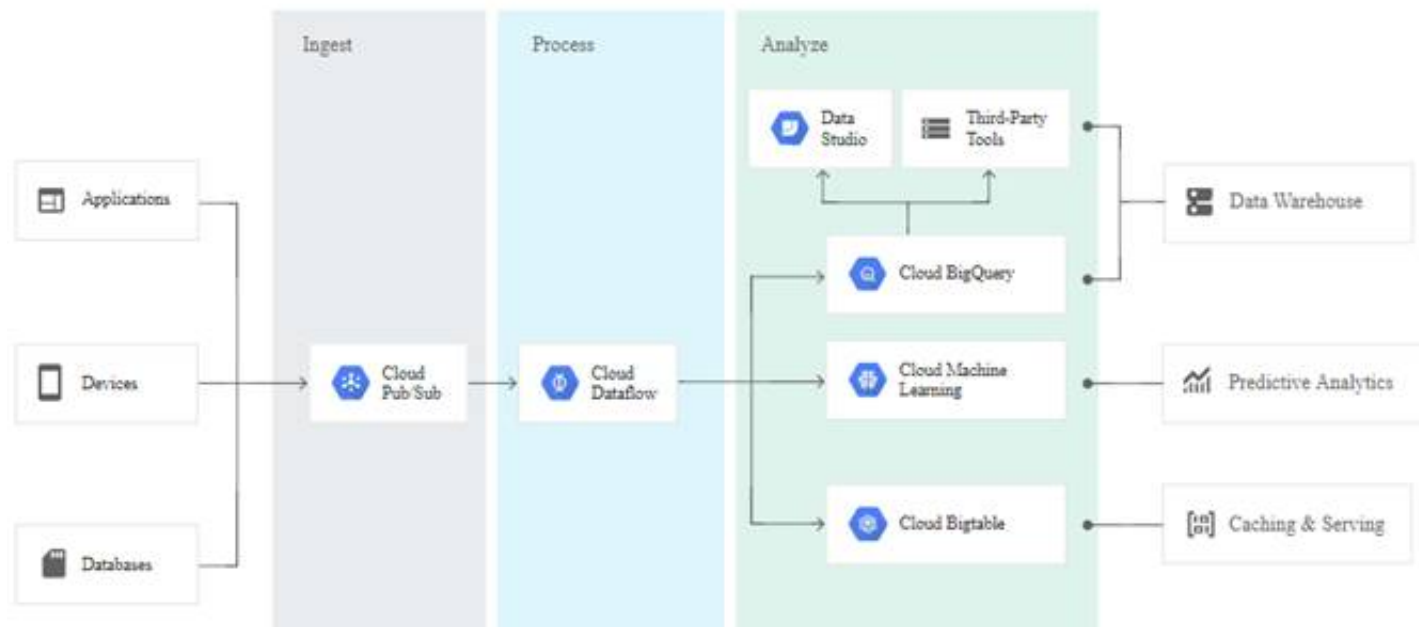A)

B)

C)

IoT   IoT   IoT

FTP   FTP   FTP

Google Cloud Load Balancing

FTP

Google Container Engine

Cloud Pub/Sub

Cloud Dataflow

Cloud SQL

Analysts

D)

IoT   IoT   IoT

FTP   FTP   FTP

App Engine   Flexible Environment

Cloud Pub/Sub

Cloud Dataflow

Cloud SQL

Analysts

A. Option A
B. Option B
C. Option C
D. Option D

Answer: A

Explanation:
 https://cloud.google.com/solutions/iot/ https://cloud.google.com/solutions/designing-connected-vehicle-platform https://cloud.google.com/solutions/designing-connected-vehicle-platform#data_ingestion
http://www.eweek.com/big-data-and-analytics/google-touts-value-of-cloud-iot-core-for- analyzing-connected-car-data
https://cloud.google.com/solutions/iot/
The push endpoint can be a load balancer.
A container cluster can be used. Cloud Pub/Sub for Stream Analytics

References: https://cloud.google.com/pubsub/ https://cloud.google.com/solutions/iot/ https://cloud.google.com/solutions/designing-connected-vehicle-platform
https://cloud.google.com/solutions/designing-connected-vehicle-platform#data_ingestion http://www.eweek.com/big-data-and-analytics/google-touts-value-of-cloud-iot-core-for-analyzing-connected-car-data
https://cloud.google.com/solutions/iot/

NEW QUESTION 4

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

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

Answer: B

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

NEW QUESTION 5

- (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 6

- (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 7
- (Topic 3)
For this question, refer to the JencoMart case study
A few days after JencoMart migrates the user credentials database to Google Cloud Platform and shuts down the old server, the new database server stops responding to SSH connections. It is still serving database requests to the application servers correctly. What three steps should you take to diagnose the problem? Choose 3 answers

A. Delete the virtual machine (VM) and disks and create a new one.
B. Delete the instance, attach the disk to a new VM, and investigate.
C. Take a snapshot of the disk and connect to a new machine to investigate.
D. Check inbound firewall rules for the network the machine is connected to.
E. Connect the machine to another network with very simple firewall rules and investigate.
F. Print the Serial Console output for the instance for troubleshooting, activate the interactive console, and investigate.

Answer: CDF

Explanation:
 https://cloud.google.com/compute/docs/troubleshooting/troubleshooting-ssh
D: Handling "Unable to connect on port 22" error message Possible causes include:
There is no firewall rule allowing SSH access on the port. SSH access on port 22 is enabled on all Compute Engine instances by default. If you have disabled access, SSH from the Browser will not work. If you run sshd on a port other than 22, you need to enable the access to that port with a custom firewall rule.
The firewall rule allowing SSH access is enabled, but is not configured to allow connections from GCP Console services. Source IP addresses for browser-based SSH sessions are dynamically allocated by GCP Console and can vary from session to session.
References: https://cloud.google.com/compute/docs/ssh-in-browser https://cloud.google.com/compute/docs/ssh-in-browser

NEW QUESTION 8
- (Topic 4)
For this question, refer to the Dress4Win case study.
Dress4Win would like to become familiar with deploying applications to the cloud by successfully deploying some applications quickly, as is. They have asked for your recommendation. What should you advise?

A. Identify self-contained applications with external dependencies as a first move to the cloud.
B. Identify enterprise applications with internal dependencies and recommend these as a first move to the cloud.
C. Suggest moving their in-house databases to the cloud and continue serving requests to on-premise applications.
D. Recommend moving their message queuing servers to the cloud and continue handling requests to on-premise applications.

Answer: A

Explanation:
 https://cloud.google.com/blog/products/gcp/the-five-phases-of-migrating-to-google-cloud-platform

NEW QUESTION 9
- (Topic 4)
For this question, refer to the Dress4Win case study.
The Dress4Win security team has disabled external SSH access into production virtual machines (VMs) on Google Cloud Platform (GCP). The operations team needs to remotely manage the VMs, build and push Docker containers, and manage Google Cloud Storage objects. What can they do?

A. Grant the operations engineers access to use Google Cloud Shell.
B. Configure a VPN connection to GCP to allow SSH access to the cloud VMs.
C. Develop a new access request process that grants temporary SSH access to cloud VMs when an operations engineer needs to perform a task.
D. Have the development team build an API service that allows the operations team to execute specific remote procedure calls to accomplish their tasks.

Answer: A

NEW QUESTION 10
- (Topic 4)
For this question, refer to the Dress4Win case study.
Dress4Win has configured a new uptime check with Google Stackdriver for several of their legacy services. The Stackdriver dashboard is not reporting the services as healthy. What should they do?

A. Install the Stackdriver agent on all of the legacy web servers.
B. In the Cloud Platform Console download the list of the uptime servers' IP addresses and create an inbound firewall rule
C. Configure their load balancer to pass through the User-Agent HTTP header when the value matches GoogleStackdriverMonitoring-UptimeChecks (https://cloud.google.com/monitoring)
D. Configure their legacy web servers to allow requests that contain user-Agent HTTP header when the value matches GoogleStackdriverMonitoring—UptimeChecks (https://cloud.google.com/monitoring)

Answer: B

NEW QUESTION 10
- (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 15
- (Topic 5)
You are responsible for the Google Cloud environment in your company Multiple departments need access to their own projects and the members within each department will have the same project responsibilities You want to structure your Google Cloud environment for minimal maintenance and maximum overview of 1AM permissions as each department's projects start and end You want to follow Google-recommended practices What should you do?

A. Create a Google Group per department and add all department members to their respective groups Create a folder per departmentand grant the respective group the required 1AM permissions at the folder level Add the projects under the respective folders
B. Grant all department members the required 1AM permissions for their respective projects
C. Create a Google Group per department and add all department members to theirrespective groups Grant each group the required I AM permissions for their respective projects
D. Create a folder per department and grant the respective members of the department the required 1AM permissions at the folder leve
E. Structure all projects for each department under the respective folders

Answer: A

Explanation:
This option follows the Google-recommended practices for structuring a Google Cloud environment for minimal maintenance and maximum overview of IAM permissions. By creating a Google Group per department and adding all department members to their respective groups, you can simplify user management and avoid granting IAM permissions to individual users. By creating a folder per department and granting the respective group the required IAM permissions at the folder level, you can enforce consistent policies across all projects within each department and avoid granting IAM permissions at the project level. By adding the projects under the respective folders, you can organize your resources hierarchically and leverage inheritance of IAM policies from folders to projects. The other options are not optimal for this scenario, because they either require granting IAM permissions to individual users (B, C), or do not use Google Groups to manage users (D). References:
? https://cloud.google.com/architecture/framework/system-design
? https://cloud.google.com/architecture/identity/best-practices-for-planning
? https://cloud.google.com/resource-manager/docs/creating-managing-folders

NEW QUESTION 16
- (Topic 5)
Your company has an application running on Compute Engine mat allows users to play their favorite music. There are a fixed number of instances Files are stored in Cloud Storage and data is streamed directly to users. Users are reporting that they sometimes need to attempt to play popular songs multiple times before they are successful. You need to improve the performance of the application. What should you do?
A.
* 1. Copy popular songs into CloudSQL as a blob
* 2. Update application code to retrieve data from CloudSQL when Cloud Storage is overloaded
B.
* 1. Create a managed instance group with Compute Engine instances
* 2. Create a global toad balancer and configure ii with two backbends
* Managed instance group
* Cloud Storage bucket
* 3. Enable Cloud CDN on the bucket backend
C.
* 1. Mount the Cloud Storage bucket using gcsfuse on all backend Compute Engine instances
* 2. Serve muse files directly from the backend Compute Engine instance
D.
* 1. Create a Cloud Filestore NFS volume and attach it to the backend Compute Engine instances
* 2. Download popular songs in Cloud Filestore
* 3. Serve music Wes directly from the backend Compute Engine instance

A.

Answer: B

NEW QUESTION 21
- (Topic 5)
Your company has a Google Cloud project that uses BigQuery for data warehousing on a pay-per-use basis. You want to monitor queries in real time to discover the most costly queries and which users spend the most. What should you do?
A.
* 1. Create a Cloud Logging sink to export BigQuery data access logs to Cloud Storage.
* 2. Develop a Dataflow pipeline to compute the cost of queries split by users.
B.
* 1. Create a Cloud Logging sink to export BigQuery data access logs to BigQuery.
* 2. Perform a BigQuery query on the generated table to extract the information you need.

C.
* 1. Activate billing export into BigQuery.
* 2. Perform a BigQuery query on the billing table to extract the information you need.
D.
* 1. In the BigQuery dataset that contains all the tables to be queried, add a label for each user that can launch a query.
* 2. Open the Billing page of the project.
* 3. Select Reports.
* 4. Select BigQuery as the product and filter by the user you want to check.

A.

Answer: C

Explanation:
 https://cloud.google.com/blog/products/data-analytics/taking-a-practical-approach-to-bigquery-cost-monitoring

NEW QUESTION 24
- (Topic 5)
Your company has just recently activated Cloud Identity to manage users. The Google Cloud Organization has been configured as wed. The security learn needs to secure protects that will be part of the Organization. They want to prohibit IAM users outside the domain from gaining permissions from now on. What should they do?

A. Configure an organization policy to restrict identities by domain
B. Configure an organization policy to block creation of service accounts
C. Configure Cloud Scheduler o trigger a Cloud Function every hour that removes all users that don't belong to the Cloud identity domain from all projects.
D. Create a technical user (e g . crawler@yourdomain com), and give it the protect owner rote at root organization level Write a bash script that• Lists all me IAM rules of all projects within the organization• Deletes all users that do not belong to the company domainCreate a Compute Engine instance m a project within the Organization and configure gcloud to be executed with technical user credentials Configure a cron job that executes the bash script every hour.

Answer: A

Explanation:
 https://cloud.google.com/resource-manager/docs/organization- policy/restricting-domains

NEW QUESTION 29
- (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 31
- (Topic 5)
You have deployed an application on Anthos clusters (formerly Anthos GKE). According to the SRE practices at your company you need to be alerted if the request latency is above a certain threshold for a specified amount of time. What should you do?

A. Enable the Cloud Trace API on your project and use Cloud Monitoring Alerts to send an alert based on the Cloud Trace metrics
B. Configure Anthos Config Management on your cluster and create a yaml file that defines the SLO and alerting policy you want to deploy in your cluster
C. Use Cloud Profiler to follow up the request latenc
D. Create a custom metric in Cloud Monitoring based on the results of Cloud Profiler, and create an Alerting Policy in case this metric exceeds the threshold
E. Install Anthos Service Mesh on your cluste
F. Use the Google Cloud Console to define a Service Level Objective (SLO)

Answer: D

Explanation:
 https://cloud.google.com/service-mesh/docs/overview https://cloud.google.com/service-mesh/docs/observability/slo-overview

NEW QUESTION 33
- (Topic 5)
You are building a continuous deployment pipeline for a project stored in a Git source repository and want to ensure that code changes can be verified deploying to production. What should you do?

A. Use Spinnaker to deploy builds to production using the red/black deployment strategy so that changes can easily be rolled back.
B. Use Spinnaker to deploy builds to production and run tests on production deployments.
C. Use Jenkins to build the staging branches and the master branc
D. Build and deploy changes to production for 10% of users before doing a complete rollout.

E. Use Jenkins to monitor tags in the repositor
F. Deploy staging tags to a staging environment for testing.After testing, tag the repository for production and deploy that to the production environment.

Answer: D

Explanation:
 Reference: https://github.com/GoogleCloudPlatform/continuous-deployment-on- kubernetes/blob/master/README.md

NEW QUESTION 38
- (Topic 5)
Your solution is producing performance bugs in production that you did not see in staging and test environments. You want to adjust your test and deployment procedures to avoid this problem in the future. What should you do?

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

Answer: C

NEW QUESTION 41
- (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 42
- (Topic 5)
Your customer wants to do resilience testing of their authentication layer. This consists of a regional managed instance group serving a public REST API that reads from and writes to a Cloud SQL instance.
What should you do?

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

Answer: C

NEW QUESTION 44
- (Topic 5)
You deploy your custom Java application to Google App Engine. It fails to deploy and gives you the following stack trace.

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

What should you do?

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

Answer: B

NEW QUESTION 47
- (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 51
- (Topic 5)
Your company runs several databases on a single MySQL instance. They need to take backups of a specific database at regular intervals. The backup activity needs to complete as quickly as possible and cannot be allowed to impact disk performance. How should you configure the storage?

A. Configure a cron job to use the gcloud tool to take regular backups using persistent disk snapshots.
B. Mount a Local SSD volume as the backup locatio
C. After the backup is complete, use gsutil to move the backup to Google Cloud Storage.
D. Use gcsfuse to mount a Google Cloud Storage bucket as a volume directly on the instance and write backups to the mounted location using mysqldump
E. Mount additional persistent disk volumes onto each virtual machine (VM) instance in a RAID10 array and use LVM to create snapshots to send to Cloud Storage.

Answer: B

Explanation:
https://cloud.google.com/compute/docs/instances/sql-server/best-practices

NEW QUESTION 56
- (Topic 5)
Your marketing department wants to send out a promotional email campaign. The development team wants to minimize direct operation management. They project a wide range of possible customer responses, from 100 to 500,000 click-throughs per day. The link leads to a simple website that explains the promotion and collects user information and preferences. Which infrastructure should you recommend? (CHOOSE TWO)

A. Use Google App Engine to serve the website and Google Cloud Datastore to store user data.
B. Use a Google Container Engine cluster to serve the website and store data to persistent disk.
C. Use a managed instance group to serve the website and Google Cloud Bigtable to store user data.
D. Use a single compute Engine virtual machine (VM) to host a web server, backed by Google Cloud SQL.

Answer: AC

Explanation:
 Reference: https://cloud.google.com/storage-options/
References: https://cloud.google.com/storage-options/

NEW QUESTION 60
- (Topic 5)
You need to deploy a stateful workload on Google Cloud. The workload can scale horizontally, but each instance needs to read and write to the same POSIX filesystem. At high load, the stateful workload needs to support up to 100 MB/s of writes. What should you do?

A. Use a persistent disk for each instance.
B. Use a regional persistent disk for each instance.
C. Create a Cloud Filestore instance and mount it in each instance.
D. Create a Cloud Storage bucket and mount it in each instance using gcsfuse.

Answer: C

Explanation:
 https://cloud.google.com/storage/docs/gcs-fuse#notes
Cloud Filestore: Cloud Filestore is a scalable and highly available shared file service fully managed by Google. Cloud Filestore provides persistent storage ideal for shared workloads. It is best suited for enterprise applications requiring persistent, durable, shared storage which is accessed by NFS or requires a POSIX compliant file system.
Reference: https://cloud.google.com/storage/docs/gcs-fuse

NEW QUESTION 65
- (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 70
- (Topic 5)
You need to reduce the number of unplanned rollbacks of erroneous production deployments in your company's web hosting platform. Improvement to the QA/Test processes accomplished an 80% reduction. Which additional two approaches can you take to further reduce the rollbacks? Choose 2 answers

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

Answer: AC

NEW QUESTION 75
- (Topic 5)
You are working with a data warehousing team that performs data analysis. The team needs to process data from external partners, but the data contains personally identifiable information (PII). You need to process and store the data without storing any of the PII data. What should you do?

A. Create a Dataflow pipeline to retrieve the data from the external source
B. As part of the pipeline use the Cloud Data Loss Prevention (Cloud DLP) API to remove any PII data Store the result in BigQuery
C. Create a Dataflow pipeline to retrieve the data from the external source
D. As part of the pipeline store all non-PII data in BigQuery and store all PII data in a Cloud Storage bucket that has a retention policy set.
E. Ask the external partners to upload an data on Cloud Storage Configure Bucket Lock for the bucket Create a Dataflow pipeline to read the data from the bucket

As part of the pipeline, use the Cloud Data Loss Prevention (Cloud DIP) API to remove any Pll data Store the result in BigQuery
F. Ask the external partners to import ail data in your BigQuery dataset Create a dataflow pipeline to copy the data into a new table As part of the Dataflow bucket skip all data in columns that have Pll data

Answer: A

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

NEW QUESTION 79
- (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 82
- (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 86
- (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 91
- (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 95
- (Topic 5)
Your company uses Google Kubernetes Engine (GKE) as a platform for all workloads. Your company has a single large GKE cluster that contains batch, stateful, and stateless workloads. The GKE cluster is configured with a single node pool with 200 nodes. Your company needs to reduce the cost of this cluster but does not want to compromise availability. What should you do?

A. Create a second GKE cluster for the batch workloads onl
B. Allocate the 200 original nodes across both clusters.
C. Configure a HorizontalPodAutoscaler for all stateless workloads and for all compatible stateful workload
D. Configure the cluster to use node auto scaling.
E. Configure CPU and memory limits on the namespaces in the cluste
F. Configure all Pods to have a CPU and memory limits.
G. Change the node pool to use spot VMs.

Answer: C

Explanation:
 One way to reduce the cost of a Google Kubernetes Engine (GKE) cluster without compromising availability is to use horizontal pod autoscalers (HPA) and node auto scaling. HPA allows you to automatically scale the number of Pods in a deployment based on the resource usage of the Pods. By configuring HPA for stateless workloads and for compatible stateful workloads, you can ensure that the number of Pods is automatically adjusted based on the actual resource usage, which can help to reduce costs. Node auto scaling allows you to automatically add or remove nodes from the node pool based on the resource usage of the cluster. By configuring node auto scaling, you can ensure that the cluster has the minimum number of nodes needed to meet the resource requirements of the workloads, which can also help to reduce costs.

NEW QUESTION 99
- (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 101
- (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 102
- (Topic 5)
You are configuring the cloud network architecture for a newly created project m Google Cloud that will host applications in Compote Engine Compute Engine virtual machine instances will be created in two different subnets (sub-a and sub-b) within a single region
• Instances in sub-a win have public IP addresses
• Instances in sub-b will have only private IP addresses
To download updated packages, instances must connect to a public repository outside the boundaries of Google Cloud You need to allow sub-b to access the external repository. What should you do?

A. Enable Private Google Access on sub-b
B. Configure Cloud NAT and select sub b m the NAT mapping section
C. Configure a bastion host instance in sub a to connect to instances in sub-b
D. Enable Identity Aware Proxy for TCP forwarding for instances in sub-b

Answer: B

Explanation:
? Cloud NAT (network address translation) lets Google Cloud virtual machine (VM) instances without external IP addresses and private Google Kubernetes Engine (GKE) clusters send outbound packets to the internet and receive any corresponding established inbound response packets1. By configuring Cloud NAT and selecting sub-b in the NAT mapping section, you can allow instances in sub-b to access the external repository without exposing them to the internet1.

NEW QUESTION 106
- (Topic 5)
Your company and one of its partners each nave a Google Cloud protect in separate organizations. Your company s protect (prj-a) runs in Virtual Private Cloud (vpc-a). The partner's project (prj-b) runs in vpc-b. There are two instances running on vpc-a and one instance running on vpc-b Subnets denned in both VPCs are not overlapping. You need to ensure that all instances communicate with each other via internal IPs minimizing latency and maximizing throughput. What should you do?

A. Set up a network peering between vpc-a and vpc-b
B. Set up a VPN between vpc-a and vpc-b using Cloud VPN
C. Configure IAP TCP forwarding on the instance in vpc b and then launch the following gcloud command from one of the instance in vpc-gcloud: `gcloud compute start-iap-tunnel INSTANCE_NAME_IN_VPC_B 22 \ --local-host-port=localhost:22`
* 1. Create an additional instance in vpc-a* 2. Create an additional instance n vpc-b* 3. Instal OpenVPN in newly created instances* 4. Configure a VPN tunnel between vpc-a and vpc-b with the help of OpenVPN

Answer: C

NEW QUESTION 110
- (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 112
- (Topic 5)
You have an outage in your Compute Engine managed instance group: all instance keep restarting after 5 seconds. You have a health check configured, but autoscaling is disabled. Your colleague, who is a Linux expert, offered to look into the issue. You need to make sure that he can access the VMs. What should you do?

A. Grant your colleague the IAM role of project Viewer
B. Perform a rolling restart on the instance group
C. Disable the health check for the instance grou
D. Add his SSH key to the project-wide SSH keys
E. Disable autoscaling for the instance grou
F. Add his SSH key to the project-wide SSH Keys

Answer: C

Explanation:
https://cloud.google.com/compute/docs/instance-groups/autohealing-instances-in-migs Health checks used for autohealing should be conservative so they don't preemptively delete and recreate your instances. When an autohealer health check is too aggressive, the autohealer might mistake busy instances for failed instances and unnecessarily restart them, reducing availability

NEW QUESTION 113
- (Topic 5)
Your customer support tool logs all email and chat conversations to Cloud Bigtable for retention and analysis.
What is the recommended approach for sanitizing this data of personally identifiable information or payment
card information before initial storage?

A. Hash all data using SHA256
B. Encrypt all data using elliptic curve cryptography
C. De-identify the data with the Cloud Data Loss Prevention API
D. Use regular expressions to find and redact phone numbers, email addresses, and credit card numbers

Answer: A

Explanation:
 Reference: https://cloud.google.com/solutions/pci-dss-compliance-ingcp#

NEW QUESTION 118
- (Topic 5)
A small number of API requests to your microservices-based application take a very long time. You know that each request to the API can traverse many services. You want to know which service takes the longest in those cases. What should you do?

A. Set timeouts on your application so that you can fail requests faster.

B. Send custom metrics for each of your requests to Stackdriver Monitoring.
C. Use Stackdriver Monitoring to look for insights that show when your API latencies are high.
D. Instrument your application with Stackdnver Trace in order to break down the request latencies at each microservice.

Answer: D

Explanation:
https://cloud.google.com/trace/docs/overview

NEW QUESTION 122
- (Topic 5)
For this question, refer to the TerramEarth case study. You are building a microservice- based application for TerramEarth. The application is based on Docker containers. You want to follow Google-recommended practices to build the application continuously and store the build artifacts. What should you do?
A.
* 1. Configure a trigger in Cloud Build for new source changes.
* 2. Invoke Cloud Build to build one container image, and tag the image with the label 'latest.'
* 3. Push the image to the Artifact Registry.
B.
* 1. Configure a trigger in Cloud Build for new source changes.
* 2. Invoke Cloud Build to build container images for each microservice, and tag them using the code commit hash.
* 3. Push the images to the Artifact Registry.
C.
* 1 Create a Scheduler job to check the repo every minute.
* 2. For any new change, invoke Cloud Build to build container images for the microservices.
* 3. Tag the images using the current timestamp, and push them to the Artifact Registry.
D.
* 1. Configure a trigger in Cloud Build for new source changes.
* 2. The trigger invokes build jobs and build container images for the microservices.
* 3. Tag the images with a version number, and push them to Cloud Storage.

A.

Answer: C

NEW QUESTION 126
- (Topic 5)
Your company plans to migrate a multi-petabyte data set to the cloud. The data set must be available 24hrs a day. Your business analysts have experience only with using a SQL interface. How should you store the data to optimize it for ease of analysis?

A. Load data into Google BigQuery.
B. Insert data into Google Cloud SQL.
C. Put flat files into Google Cloud Storage.
D. Stream data into Google Cloud Datastore.

Answer: A

Explanation:
Google Big Query is for multi peta byte storage , HA(High availability) which means 24 hours, SQL interface .
https://medium.com/google-cloud/the-12-components-of-google-bigquery-c2b49829a7c7 https://cloud.google.com/solutions/bigquery-data-warehouse
https://cloud.google.com/bigquery/
BigQuery is Google's serverless, highly scalable, low cost enterprise data warehouse designed to make all your data analysts productive. Because there is no infrastructure to manage, you can focus on analyzing data to find meaningful insights using familiar SQL and you don't need a database administrator.
BigQuery enables you to analyze all your data by creating a logical data warehouse over managed, columnar storage as well as data from object storage, and spreadsheets.
References: https://cloud.google.com/bigquery/

NEW QUESTION 127
- (Topic 5)
You want to automate the creation of a managed instance group and a startup script to install the OS package dependencies. You want to minimize the startup time for VMs in the instance group.
What should you do?

A. Use Terraform to create the managed instance group and a startup script to install the OS packagedependencies.
B. Create a custom VM image with all OS package dependencie
C. Use Deployment Manager to create the managed instance group with the VM image.
D. Use Puppet to create the managed instance group and install the OS package dependencies.
E. Use Deployment Manager to create the managed instance group and Ansible to install the OS package dependencies.

Answer: B

Explanation:
"Custom images are more deterministic and start more quickly than instances with startup scripts. However, startup scripts are more flexible and let you update the apps and settings in your instances more easily." https://cloud.google.com/compute/docs/instance- templates/create-instance-templates#using_custom_or_public_images_in_your_instance_templates

NEW QUESTION 128
- (Topic 5)
You need to deploy an application on Google Cloud that must run on a Debian Linux environment. The application requires extensive configuration in order to operate correctly. You want to ensure that you can install Debian distribution updates with minimal manual intervention whenever they become available. What

should you do?

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

Answer: B

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

NEW QUESTION 133
- (Topic 5)
You need to develop procedures to verify resilience of disaster recovery for remote recovery using GCP. Your production environment is hosted on-premises. You need to establish a secure, redundant connection between your on premises network and the GCP network.
What should you do?

A. Verify that Dedicated Interconnect can replicate files to GC
B. Verify that direct peering can establish asecure connection between your networks if Dedicated Interconnect fails.
C. Verify that Dedicated Interconnect can replicate files to GC
D. Verify that Cloud VPN can establish a secure connection between your networks if Dedicated Interconnect fails.
E. Verify that the Transfer Appliance can replicate files to GC
F. Verify that direct peering can establish asecure connection between your networks if the Transfer Appliance fails.
G. Verify that the Transfer Appliance can replicate files to GC
H. Verify that Cloud VPN can establish a secure connection between your networks if the Transfer Appliance fails.

Answer: B

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

NEW QUESTION 135
- (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 138
- (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 142

- (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 143
- (Topic 5)
You company has a Kubernetes application that pulls messages from Pub/Sub and stores them in Firestore. Because the application is simple, it was deployed as a single pod. The infrastructure team has analyzed Pub/Sub metrics and discovered that the application cannot process the messages in real time. Most of them wait for minutes before being processed. You need to scale the elaboration process that is I/O-intensive. What should you do?

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

Answer: A

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

NEW QUESTION 148
- (Topic 5)
You have developed a non-critical update to your application that is running in a managed instance group, and have created a new instance template with the update that you want to release. To prevent any possible impact to the application, you don't want to update any running instances. You want any new instances that are created by the managed instance group to contain the new update. What should you do?

A. Start a new rolling restart operation.
B. Start a new rolling replace operation.
C. Start a new rolling updat
D. Select the Proactive update mode.
E. Start a new rolling updat
F. Select the Opportunistic update mode.

Answer: D

Explanation:
 In certain scenarios, an opportunistic update is useful because you don't want to cause instability to the system if it can be avoided. For example, if you have a non- critical update that can be applied as necessary without any urgency and you have a MIG that is actively being autoscaled, perform an opportunistic update so that Compute Engine does not actively tear down your existing instances to apply the update. When resizing down, the autoscaler preferentially terminates instances with the old template as well as instances that are not yet in a RUNNING state.

NEW QUESTION 150
- (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 155

- (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 156
- (Topic 5)
You have found an error in your App Engine application caused by missing Cloud Datastore indexes. You have created a YAML file with the required indexes and want to deploy these new indexes to Cloud Datastore.
What should you do?

A. Point gcloud datastore create-indexes to your configuration file
B. Upload the configuration file the App Engine's default Cloud Storage bucket, and have App Engine detect the new indexes
C. In the GCP Console, use Datastore Admin to delete the current indexes and upload the new configuration file
D. Create an HTTP request to the built-in python module to send the index configuration file to your application

Answer: A

NEW QUESTION 157
- (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 158
- (Topic 5)
Your organization has decided to restrict the use of external IP addresses on instances to only approved instances. You want to enforce this requirement across all of your Virtual Private Clouds (VPCs). What should you do?

A. Remove the default route on all VPC
B. Move all approved instances into a new subnet that has a default route to an internet gateway.
C. Create a new VPC in custom mod
D. Create a new subnet for the approved instances, and set a default route to the internet gateway on this new subnet.
E. Implement a Cloud NAT solution to remove the need for external IP addresses entirely.
F. Set an Organization Policy with a constraint on constraints/compute.vmExternalIpAcces
G. List the approved instances in the allowedValues list.

Answer: D

Explanation:
Reference: https://cloud.google.com/compute/docs/ip-addresses/reserve-static-external-ip- address
https://cloud.google.com/compute/docs/ip-addresses/reserve-static-external-ip- address#disableexternalip
you might want to restrict external IP address so that only specific VM instances can use them. This option can help to prevent data exfiltration or maintain network isolation. Using an Organization Policy, you can restrict external IP addresses to specific VM instances with
constraints to control use of external IP addresses for your VM instances within an organization or a project.

NEW QUESTION 159
- (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 162
- (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 163
- (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 168
- (Topic 5)
You want to enable your running Google Kubernetes Engine cluster to scale as demand for your application changes.
What should you do?

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

Answer: C

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

NEW QUESTION 170
- (Topic 5)
You are monitoring Google Kubernetes Engine (GKE) clusters in a Cloud Monitoring workspace. As a Site Reliability Engineer (SRE), you need to triage incidents quickly. What should you do?

A. Navigate the predefined dashboards in the Cloud Monitoring workspace, and then add metrics and create alert policies.
B. Navigate the predefined dashboards in the Cloud Monitoring workspace, create custom metrics, and install alerting software on a Compute Engine instance.
C. Write a shell script that gathers metrics from GKE nodes, publish these metrics to a Pub/Sub topic, export the data to BigQuery, and make a Data Studio dashboard.
D. Create a custom dashboard in the Cloud Monitoring workspace for each incident, and then add metrics and create alert policies.

Answer: A

Explanation:
https://cloud.google.com/stackdriver/docs/solutions/gke/legacy-stackdriver/monitoring

NEW QUESTION 175

- (Topic 5)
You are using Cloud SQL as the database backend for a large CRM deployment. You want to scale as usage increases and ensure that you don't run out of storage, maintain 75% CPU usage cores, and keep replication lag below 60 seconds. What are the correct steps to meet your requirements?
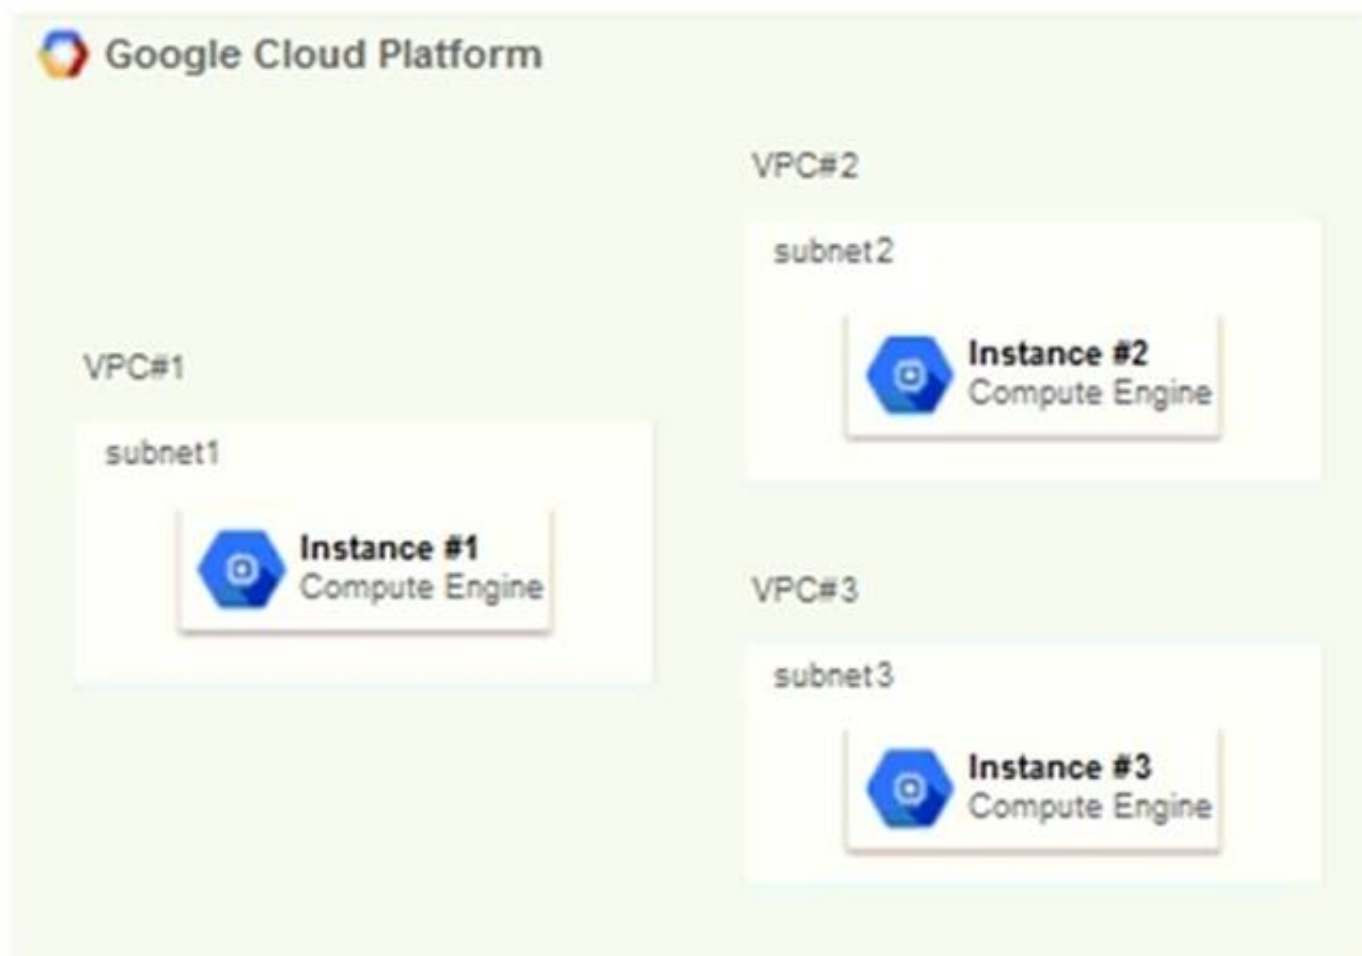
A. 1) Enable automatic storage increase for the instance.2) Create a Stackdriver alert when CPU usage exceeds 75%, and change the instance type to reduceCPU usage.3) Create a Stackdriver alert for replication lag, and shard the database to reduce replication time.
B. 1) Enable automatic storage increase for the instance.2) Change the instance type to a 32-core machine type to keep CPU usage below 75%.3) Create a Stackdriver alert for replication lag, and shard the database to reduce replication time.
C. 1) Create a Stackdriver alert when storage exceeds 75%, and increase the available storage on theinstance to create more space.2) Deploy memcached to reduce CPU load.3) Change the instance type to a 32-core machine type to reduce replication lag.
D. 1) Create a Stackdriver alert when storage exceeds 75%, and increase the available storage on theinstance to create more space.2) Deploy memcached to reduce CPU load.3) Create a Stackdriver alert for replication lag, and change the instance type to a 32-core machine type to reduce replication lag.

Answer: A

NEW QUESTION 179
- (Topic 5)
Your company has a project in Google Cloud with three Virtual Private Clouds (VPCs). There is a Compute Engine instance on each VPC. Network subnets do not overlap and must remain separated. The network configuration is shown below.

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

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

Answer: B

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

NEW QUESTION 180
- (Topic 5)
A production database virtual machine on Google Compute Engine has an ext4-formatted persistent disk for data files The database is about to run out of storage space How can you remediate the problem with the least amount of downtime?

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

Answer: A

Explanation:

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

NEW QUESTION 181
- (Topic 5)
You have an application that makes HTTP requests to Cloud Storage. Occasionally the requests fail with
HTTP status codes of 5xx and 429.
How should you handle these types of errors?

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

Answer: A

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

NEW QUESTION 182
- (Topic 5)
Your customer is receiving reports that their recently updated Google App Engine application is taking approximately 30 seconds to load for some of their users. This behavior was not reported before the update. What strategy should you take?

A. Work with your ISP to diagnose the problem.
B. Open a support ticket to ask for network capture and flow data to diagnose the problem, then roll back your application.
C. Roll back to an earlier known good release initially, then use Stackdriver Trace and logging to diagnose the problem in a development/test/staging environment.
D. Roll back to an earlier known good release, then push the release again at a quieter period to investigat
E. Then use Stackdriver Trace and logging to diagnose the problem.

Answer: C

Explanation:
Stackdriver Logging allows you to store, search, analyze, monitor, and alert on log data and events from Google Cloud Platform and Amazon Web Services (AWS). Our API also allows ingestion of any custom log data from any source. Stackdriver Logging is a fully managed service that performs at scale and can ingest application and system log data from thousands of VMs. Even better, you can analyze all that log data in real time.
References: https://cloud.google.com/logging/

NEW QUESTION 184
- (Topic 5)
Your company just finished a rapid lift and shift to Google Compute Engine for your compute needs. You have another 9 months to design and deploy a more cloud-native solution. Specifically, you want a system that is no-ops and auto-scaling. Which two compute products should you choose? Choose 2 answers

A. Compute Engine with containers
B. Google Kubernetes Engine with containers
C. Google App Engine Standard Environment
D. Compute Engine with custom instance types
E. Compute Engine with managed instance groups

Answer: BC

Explanation:
B: With Container Engine, Google will automatically deploy your cluster for you, update, patch, secure the nodes.
Kubernetes Engine's cluster autoscaler automatically resizes clusters based on the demands of the workloads you want to run.
C: Solutions like Datastore, BigQuery, AppEngine, etc are truly NoOps.
App Engine by default scales the number of instances running up and down to match the load, thus providing consistent performance for your app at all times while minimizing idle instances and thus reducing cost.
Note: At a high level, NoOps means that there is no infrastructure to build out and manage during usage of the platform. Typically, the compromise you make with NoOps is that you lose control of the underlying infrastructure.
References: https://www.quora.com/How-well-does-Google-Container-Engine-support-Google-Cloud-Platform%E2%80%99s-NoOps-claim

NEW QUESTION 187
- (Topic 5)
Your company has a support ticketing solution that uses App Engine Standard. The project that contains the App Engine application already has a Virtual Private Cloud(VPC) network fully
connected to the company's on-premises environment through a Cloud VPN tunnel. You want to enable App Engine application to communicate with a database that is running in
the company's on-premises environment. What should you do?

A. Configure private services access
B. Configure private Google access for on-premises hosts only
C. Configure serverless VPC access
D. Configure private Google access

Answer: A

Explanation:
https://cloud.google.com/appengine/docs/standard/python3/connecting-vpc https://cloud.google.com/appengine/docs/flexible/python/using-third-party-databases#on_premises

NEW QUESTION 188
- (Topic 5)
You want to create a private connection between your instances on Compute Engine and your on-premises data center. You require a connection of at least 20 Gbps. You want to follow Google-recommended practices.
How should you set up the connection?

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

Answer: A

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

NEW QUESTION 189
- (Topic 6)
For this question, refer to the Dress4Win case study. Considering the given business requirements, how would you automate the deployment of web and transactional data layers?

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

Answer: A

NEW QUESTION 191
- (Topic 6)
For this question, refer to the Dress4Win case study. You want to ensure that your on- premises architecture meets business requirements before you migrate your solution.
What change in the on-premises architecture should you make?

A. Replace RabbitMQ with Google Pub/Sub.
B. Downgrade MySQL to v5.7, which is supported by Cloud SQL for MySQL.
C. Resize compute resources to match predefined Compute Engine machine types.
D. Containerize the micro services and host them in Google Kubernetes Engine.

Answer: C

NEW QUESTION 194
- (Topic 7)
TerramEarth has a legacy web application that you cannot migrate to cloud. However, you still want to build a cloud-native way to monitor the application. If the application goes down, you want the URL to point to a "Site is unavailable" page as soon as possible. You also want your Ops team to receive a notification for the issue. You need to build a reliable solution for minimum cost
What should you do?

A. Create a scheduled job in Cloud Run to invoke a container every minut
B. The container will check the application URL If the application is down, switch the URL to the "Site is unavailable" page, and notify the Ops team.
C. Create a cron job on a Compute Engine VM that runs every minut
D. The cron job invokes a Python program to check the application URL If the application is down, switch the URL to the "Site is unavailable" page, and notify the Ops team.
E. Create a Cloud Monitoring uptime check to validate the application URL If it fails, put a message in a Pub/Sub queue that triggers a Cloud Function to switch the URL to the "Site is unavailable" page, and notify the Ops team.
F. Use Cloud Error Reporting to check the application URL If the application is down, switch the URL to the "Site is unavailable" page, and notify the Ops team.

Answer: C

Explanation:
https://cloud.google.com/blog/products/management-tools/how-to-use-pubsub-as-a-cloud-monitoring-notification-channel

NEW QUESTION 196
- (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 200
- (Topic 7)
For this question, refer to the TerramEarth case study. You need to implement a reliable, scalable GCP solution for the data warehouse for your company, TerramEarth. Considering the TerramEarth business and technical requirements, what should you do?

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

Answer: C

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

NEW QUESTION 202
- (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 203
- (Topic 8)
Your development teams release new versions of games running on Google Kubernetes Engine (GKE) daily.
You want to create service level indicators (SLIs) to evaluate the quality of the new versions from the user's perspective. What should you do?

A. Create CPU Utilization and Request Latency as service level indicators.
B. Create GKE CPU Utilization and Memory Utilization as service level indicators.
C. Create Request Latency and Error Rate as service level indicators.
D. Create Server Uptime and Error Rate as service level indicators.

Answer: C

NEW QUESTION 206
- (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 210
- (Topic 9)
For this question, refer to the Helicopter Racing League (HRL) case study. Your team is in charge of creating a
payment card data vault for card numbers used to bill tens of thousands of viewers, merchandise consumers, and season ticket holders. You need to implement a

custom card tokenization service that meets the following requirements:
• It must provide low latency at minimal cost.
• It must be able to identify duplicate credit cards and must not store plaintext card numbers.
• It should support annual key rotation.
Which storage approach should you adopt for your tokenization service?

A. Store the card data in Secret Manager after running a query to identify duplicates.
B. Encrypt the card data with a deterministic algorithm stored in Firestore using Datastore mode.
C. Encrypt the card data with a deterministic algorithm and shard it across multiple Memorystore instances.
D. Use column-level encryption to store the data in Cloud SQL.

Answer: B

NEW QUESTION 212
- (Topic 10)
For this question, refer to the EHR Healthcare case study. You are responsible for ensuring that EHR's use of Google Cloud will pass an upcoming privacy compliance audit. What should you do? (Choose two.)

A. Verify EHR's product usage against the list of compliant products on the Google Cloud compliance page.
B. Advise EHR to execute a Business Associate Agreement (BAA) with Google Cloud.
C. Use Firebase Authentication for EHR's user facing applications.
D. Implement Prometheus to detect and prevent security breaches on EHR's web-based applications.
E. Use GKE private clusters for all Kubernetes workloads.

Answer: AB

Explanation:
 https://cloud.google.com/security/compliance/hipaa

NEW QUESTION 216
- (Topic 10)
For this question, refer to the EHR Healthcare case study. You are a developer on the EHR customer portal team. Your team recently migrated the customer portal application to Google Cloud. The load has increased on the application servers, and now the application is logging many timeout errors. You recently incorporated Pub/Sub into the application architecture, and the application is not logging any Pub/Sub publishing errors. You want to improve publishing latency. What should you do?

A. Increase the Pub/Sub Total Timeout retry value.
B. Move from a Pub/Sub subscriber pull model to a push model.
C. Turn off Pub/Sub message batching.
D. Create a backup Pub/Sub message queue.

Answer: C

Explanation:
 https://cloud.google.com/pubsub/docs/publisher?hl=en#batching

NEW QUESTION 221
- (Topic 10)
For this question, refer to the EHR Healthcare case study. You are responsible for designing the Google Cloud network architecture for Google Kubernetes Engine. You want to follow Google best practices. Considering the EHR Healthcare business and technical requirements, what should you do to reduce the attack surface?

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

Answer: A

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

NEW QUESTION 225
- (Topic 10)
For this question, refer to the EHR Healthcare case study. You need to define the technical architecture for hybrid connectivity between EHR's on-premises systems and Google Cloud. You want to follow Google's recommended practices for production-level applications. Considering the EHR Healthcare business and technical requirements, what should you do?

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

Answer: D

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

NEW QUESTION 227
......

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

100% Actual & Verified — Instant Download, Please Click
[Order The Professional-Cloud-Architect Practice Test Here](#)