Google

Exam Questions Professional-Cloud-Architect

Google Certified Professional - Cloud Architect (GCP)

NEW QUESTION 1
- (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 2
- (Topic 2)
For this question refer to the TerramEarth case study
Operational parameters such as oil pressure are adjustable on each of TerramEarth's vehicles to increase their efficiency, depending on their environmental conditions. Your primary goal is to increase the operating efficiency of all 20 million cellular and unconnected vehicles in the field How can you accomplish this goal?

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

Answer: B

NEW QUESTION 3
- (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 4
- (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 5
- (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 6
- (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 7
- (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 8
- (Topic 4)
Dress4win has end to end tests covering 100% of their endpoints.
They want to ensure that the move of cloud does not introduce any new bugs.
Which additional testing methods should the developers employ to prevent an outage?

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

Answer: B

NEW QUESTION 9
- (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 10
- (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 10
- (Topic 5)
The application reliability team at your company has added a debug feature to their backend service to send all server events to Google Cloud Storage for eventual analysis. The event records are at least 50 KB and at most 15 MB and are expected to peak at 3,000 events per second. You want to minimize data loss. Which process should you implement?

A. • Append metadata to file body.• Compress individual files.• Name files with serverName-Timestamp.• Create a new bucket if bucket is older than 1 hour and save individual files to the new bucke
B. Otherwise, save files to existing bucket
C. • Batch every 10,000 events with a single manifest file for metadata.• Compress event files and manifest file into a single archive file.• Name files using serverName-EventSequence.• Create a new bucket if bucket is older than 1 day and save the single archive file to the new bucke
D. Otherwise, save the single archive file to existing bucket.
E. • Compress individual files.• Name files with serverName-EventSequence.• Save files to one bucket• Set custom metadata headers for each object after saving.
F. • Append metadata to file body.• Compress individual files.• Name files with a random prefix pattern.• Save files to one bucket

Answer: D

Explanation:
In order to maintain a high request rate, avoid using sequential names. Using completely random object names will give you the best load distribution. Randomness after
a common prefix is effective under the prefix https://cloud.google.com/storage/docs/request-rate

NEW QUESTION 15
- (Topic 5)
Your company is developing a web-based application. You need to make sure that production deployments are linked to source code commits and are fully auditable. What should you do?

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

Answer: C

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

NEW QUESTION 17
- (Topic 5)
Your company is planning to perform a lift and shift migration of their Linux RHEL 6.5+ virtual machines. The virtual machines are running in an on-premises VMware environment. You want to migrate them to Compute Engine following Google- recommended practices. What should you do?

A. * 1. Define a migration plan based on the list of the applications and their dependencies.* 2. Migrate all virtual machines into Compute Engine individually with Migrate for Compute Engine.
B. * 1. Perform an assessment of virtual machines running in the current VMware environment.* 2. Create images of all disk
C. Import disks on Compute Engine.* 3. Create standard virtual machines where the boot disks are the ones you have imported.
D. * 1. Perform an assessment of virtual machines running in the current VMware environment.* 2. Define a migration plan, prepare a Migrate for Compute Engine migration RunBook, and execute the migration.
E. * 1. Perform an assessment of virtual machines running in the current VMware environment.* 2.Install a third-party agent on all selected virtual machine
F. 3.Migrate all virtual machines into Compute Engine.

Answer: C

Explanation:
The framework illustrated in the preceding diagram has four phases:
•Assess. In this phase, you assess your source environment, assess the workloads that you want to migrate to Google Cloud, and assess which VMs support each workload.
•Plan. In this phase, you create the basic infrastructure for Migrate for Compute Engine, such as provisioning the resource hierarchy and setting up network access.
•Deploy. In this phase, you migrate the VMs from the source environment to Compute Engine.
•Optimize. In this phase, you begin to take advantage of the cloud technologies and capabilities.
Reference: https://cloud.google.com/architecture/migrating-vms-migrate-for-compute-engine-getting-started

NEW QUESTION 21
- (Topic 5)
Your company is designing its data lake on Google Cloud and wants to develop different ingestion pipelines to collect unstructured data from different sources. After the data is stored in Google Cloud, it will be processed in several data pipelines to build a recommendation engine for end users on the website. The structure of the data retrieved from the source systems can change at any time. The data must be stored exactly as it was retrieved for reprocessing purposes in case the data structure is incompatible with the current processing pipelines. You need to design an architecture to support the use case after you retrieve the data. What should you do?

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

Answer: D

NEW QUESTION 25
- (Topic 5)
You are working in a highly secured environment where public Internet access from the Compute Engine VMs is not allowed. You do not yet have a VPN connection to access an on-premises file server. You need to install specific software on a Compute Engine instance. How should you install the software?

A. Upload the required installation files to Cloud Storag
B. Configure the VM on a subnet with a Private Google Access subne
C. Assign only an internal IP address to the V
D. Download the installation files to the VM using gsutil.
E. Upload the required installation files to Cloud Storage and use firewall rules to block all traffic except the IP address range for Cloud Storag
F. Download the files to the VM using gsutil.
G. Upload the required installation files to Cloud Source Repositorie
H. Configure the VM on a subnet with a Private Google Access subne
I. Assign only an internal IP address to the V
J. Download the installation files to the VM using gcloud.
K. Upload the required installation files to Cloud Source Repositories and use firewall rules to block all traffic except the IP address range for Cloud Source Repositorie
L. Download the files to the VM using gsutil.

Answer: A

Explanation:
https://cloud.google.com/vpc/docs/private-access-options#pga-supported

NEW QUESTION 29
- (Topic 5)
Your operations team currently stores 10 TB of data m an object storage service from a third-party provider. They want to move this data to a Cloud Storage

bucket as quickly as possible, following Google-recommended practices. They want to minimize the cost of this data migration. When approach should they use?

A. Use the gsutil mv command lo move the data
B. Use the Storage Transfer Service to move the data
C. Download the data to a Transfer Appliance and ship it to Google
D. Download the data to the on-premises data center and upload it to the Cloud Storage bucket

Answer: B

Explanation:
https://cloud.google.com/architecture/migration-to-google-cloud-transferring-your-large-datasets#transfer-options
https://cloud.google.com/storage-transfer-service

NEW QUESTION 31
- (Topic 5)
You are implementing the infrastructure for a web service on Google Cloud. The web service needs to receive and store the data from 500,000 requests per second. The data will be queried later in real time, based on exact matches of a known set of attributes. There will be periods where the web service will not receive any requests. The business wants to keep costs low. Which web service platform and database should you use for the application?

A. Cloud Run and BigQuery
B. Cloud Run and Cloud Bigtable
C. A Compute Engine autoscaling managed instance group and BigQuery
D. A Compute Engine autoscaling managed instance group and Cloud Bigtable

Answer: B

Explanation:
https://cloud.google.com/run/docs/about-instance-autoscaling https://cloud.google.com/blog/topics/developers-practitioners/bigtable-vs-bigquery-whats- difference

NEW QUESTION 35
- (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 36
- (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 39
- (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 40
- (Topic 5)
You are deploying an application on App Engine that needs to integrate with an on- premises database. For security purposes, your on-premises database must not be accessible through the public Internet. What should you do?

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

Answer: D

Explanation:

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

NEW QUESTION 43
- (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 44
- (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 Pll 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 Pll 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 Pll 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 47
- (Topic 5)
Your company sends all Google Cloud logs to Cloud Logging. Your security team wants to monitor the logs. You want to ensure that the security team can react quickly if an anomaly such as an unwanted firewall change or server breach is detected. You want to follow Google-recommended practices. What should you do?

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

Answer: C

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

NEW QUESTION 49
- (Topic 5)
Your company operates nationally and plans to use GCP for multiple batch workloads, including some that are not time-critical. You also need to use GCP services that are HIPAA-certified and manage service costs.
How should you design to meet Google best practices?

A. Provisioning preemptible VMs to reduce cos
B. Discontinue use of all GCP services and APIs that are not HIPAA-compliant.
C. Provisioning preemptible VMs to reduce cos
D. Disable and then discontinue use of all GCP and APIs that are not HIPAA-compliant.
E. Provision standard VMs in the same region to reduce cos
F. Discontinue use of all GCPservices and APIs that are not HIPAA-compliant.
G. Provision standard VMs to the same region to reduce cos
H. Disable and then discontinue use of all GCP services and APIs that are not HIPAA-compliant.

Answer: B

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

NEW QUESTION 51

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

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 54
- (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 57
- (Topic 5)
Your company has an application running as a Deployment in a Google Kubernetes Engine (GKE) cluster When releasing new versions of the application via a rolling deployment, the team has been causing outages The root cause of the outages is misconfigurations with parameters that are only used in production You want to put preventive measures for this in the platform to prevent outages What should you do?

A. Configure liveness and readiness probes in the Pod specification
B. Configure an uptime alert in Cloud Monitoring
C. Create a Scheduled Task to check whether the application is available
D. Configure health checks on the managed instance group

Answer: D

Explanation:
This option can help prevent outages caused by misconfigurations with parameters that are only used in production. Liveness and readiness probes are mechanisms to check the health and availability of the Pods and containers in a GKE cluster. Liveness probes determine if a container is still running, and if not, restart it. Readiness probes determine if a container is ready to serve requests, and if not, remove it from the load balancer. By configuring liveness and readiness probes in the Pod specification, you can ensure that your application can handle traffic and recover from failures gracefully during a rolling update. The other options are not optimal for this scenario, because they either do not prevent outages, but only alert or monitor them (B, C), or do not apply to GKE clusters, but to Compute Engine instances (D). References:
? https://cloud.google.com/kubernetes-engine/docs/how-to/updating-apps
? https://cloud.google.com/blog/products/containers-kubernetes/kubernetes-best- practices-setting-up-health-checks-with-readiness-and-liveness-probes

NEW QUESTION 62
- (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 65
- (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 67
- (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 72
- (Topic 5)
You have been asked to select the storage system for the click-data of your company's large portfolio of websites. This data is streamed in from a custom website analytics package at a typical rate of 6,000 clicks per minute, with bursts of up to 8,500 clicks per second. It must been stored for future analysis by your data science and user experience teams. Which storage infrastructure should you choose?

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

Answer: C

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

NEW QUESTION 77
- (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 78
- (Topic 5)
A news teed web service has the following code running on Google App Engine. During peak load, users report that they can see news articles they already viewed. What is the most likely cause of this problem?

```
import news
from flask import Flask, redirect, request
from flask.ext.api import status
from google.appengine.api import users

app = Flask(_name_)
sessions = {}

@app.route(“/”)
def homepage():
      user = users.get_current_user()
      if not user:
            return “Invalid login”,
status.HTTP_401_UNAUTHORIZED

      if user not in sessions:
            sessions[user] = {“viewed”: []}

      news_articles = news.get_new_news (user, sessions [user]
[“viewed”])
      sessions [user] [“viewed”] +- [n[“id”] for n
in news_articles]

      return news.render(news_articles)

if _name_ == “_main_”:
      app.run()
```

A. The session variable is local to just a single instance.
B. The session variable is being overwritten in Cloud Datastore.
C. The URL of the API needs to be modified to prevent caching.
D. The HTTP Expires header needs to be set to -1 to stop caching.

Answer: A

Explanation:
https://stackoverflow.com/questions/3164280/google-app-engine-cache-list-in-session-variable?rq=1

NEW QUESTION 81
- (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 85
- (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 87
- (Topic 5)
You team needs to create a Google Kubernetes Engine (GKE) cluster to host a newly built application that requires access to third-party services on the internet. Your company does not allow any Compute Engine instance to have a public IP address on Google Cloud. You need to create a deployment strategy that adheres to these guidelines. What should you do?

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

Answer: B

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

NEW QUESTION 89
- (Topic 5)
Your company has an application running on Google Cloud that is collecting data from thousands of physical devices that are globally distributed. Data is publish to Pub/Sub and streamed in real time into an SSO Cloud Bigtable cluster via a Dataflow pipeline. The operations team informs you that your Cloud Bigtable cluster has a hot-spot and queries are taking longer man expected You need to resolve the problem and prevent it from happening in the future What should you do?

A. Advise your clients to use HBase APIs instead of NodeJS APIs.
B. Review your RowKey strategy and ensure that keys are evenly spread across the alphabet.
C. Delete records older than 30 days.
D. Double the number of nodes you currently have.

Answer: B

NEW QUESTION 90
- (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 94

- (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 95
- (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 98
- (Topic 5)
Your company has just recently activated Cloud Identity to manage users. The Google Cloud Organization has been configured as wed. The security learn needs to secure protects that will be part of the Organization. They want to prohibit IAM users outside the domain from gaining permissions from now on. What should they do?

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

Answer: A

Explanation:
? An organization policy is a mechanism to configure constraints across your entire resource hierarchy1. By configuring an organization policy to restrict identities by domain, you can specify which domains are allowed or denied when granting IAM roles to users, groups, or service accounts2. This way, you can prohibit IAM users outside the domain from gaining permissions from now on2.

NEW QUESTION 103
- (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 108
- (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 111
- (Topic 5)
You are developing an application using different microservices that should remain internal to the cluster. You want to be able to configure each microservice with a specific number of replicas. You also want to be able to address a specific microservice from any other microservice in a uniform way, regardless of the number of replicas the microservice scales to. You need to implement this solution on Google Kubernetes Engine. What should you do?

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

Answer: A

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

NEW QUESTION 114
- (Topic 5)
You are using Cloud Shell and need to install a custom utility for use in a few weeks. Where can you store the file so it is in the default execution path and persists across sessions?

A. ~/bin
B. Cloud Storage
C. /google/scripts
D. /usr/local/bin

Answer: D

Explanation:
https://medium.com/google-cloud/no-localhost-no-problem-using-google-cloud-shell-as-my-full-time-development-environment-22d5a1942439

NEW QUESTION 118
- (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 123
- (Topic 6)
For this question, refer to the Dress4Win case study. Which of the compute services should be migrated as –is and would still be an optimized architecture for performance in the cloud?

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

Answer: C

NEW QUESTION 125
- (Topic 6)
For this question, refer to the Dress4Win case study. To be legally compliant during an audit, Dress4Win must be able to give insights in all administrative actions that modify the configuration or metadata of resources on Google Cloud.
What should you do?

A. Use Stackdriver Trace to create a trace list analysis.
B. Use Stackdriver Monitoring to create a dashboard on the project's activity.
C. Enable Cloud Identity-Aware Proxy in all projects, and add the group of Administrators as a member.
D. Use the Activity page in the GCP Console and Stackdriver Logging to provide the required insight.

Answer: A

Explanation:
https://cloud.google.com/logging/docs/audit/

NEW QUESTION 130
- (Topic 7)
For this question, refer to the TerramEarth case study.
You start to build a new application that uses a few Cloud Functions for the backend. One use case requires a Cloud Function func_display to invoke another Cloud Function func_query. You want func_query only to accept invocations from func_display. You also want to follow Google's recommended best practices.
What should you do?

A. Create a token and pass it in as an environment variable to func_displa
B. When invoking func_query, include the token in the request Pass the same token to func _query and reject the invocation if the tokens are different.
C. Make func_query 'Require authentication.' Create a unique service account and associate it to func_displa
D. Grant the service account invoker role for func_quer
E. Create an id token in func_display and include the token to the request when invoking func_query.
F. Make func _query 'Require authentication' and only accept internal traffi
G. Create those two functions in the same VP
H. Create an ingress firewall rule for func_query to only allow traffic from func_display.
I. Create those two functions in the same project and VP
J. Make func_query only accept internal traffi
K. Create an ingress firewall for func_query to only allow traffic from func_displa
L. Also, make sure both functions use the same service account.

Answer: B

Explanation:
https://cloud.google.com/functions/docs/securing/authenticating#authenticating_function_to_function_calls

NEW QUESTION 132
- (Topic 7)
For this question, refer to the TerramEarth case study. A new architecture that writes all incoming data to BigQuery has been introduced. You notice that the data is dirty, and want to ensure data quality on an automated daily basis while managing cost. What should you do?

A. Set up a streaming Cloud Dataflow job, receiving data by the ingestion proces
B. Clean the data in a Cloud Dataflow pipeline.
C. Create a Cloud Function that reads data from BigQuery and cleans i
D. Trigger i
E. Trigger the Cloud Function from a Compute Engine instance.
F. Create a SQL statement on the data in BigQuery, and save it as a vie
G. Run the view daily, and save the result to a new table.
H. Use Cloud Dataprep and configure the BigQuery tables as the sourc
I. Schedule a daily job to clean the data.

Answer: A

NEW QUESTION 137
- (Topic 8)
For this question, refer to the Mountkirk Games case study. Mountkirk Games wants to design their solution for the future in order to take advantage of cloud and technology improvements as they become available. Which two steps should they take? (Choose two.)

A. Store as much analytics and game activity data as financially feasible today so it can be used to train machine learning models to predict user behavior in the future.
B. Begin packaging their game backend artifacts in container images and running them on Kubernetes Engine to improve the availability to scale up or down based on game activity.
C. Set up a CI/CD pipeline using Jenkins and Spinnaker to automate canary deployments and improve development velocity.
D. Adopt a schema versioning tool to reduce downtime when adding new game features that require storing additional player data in the database.
E. Implement a weekly rolling maintenance process for the Linux virtual machines so theycan apply critical kernel patches and package updates and reduce the risk of 0-day vulnerabilities.

Answer: BC

NEW QUESTION 142
- (Topic 8)
You need to implement a network ingress for a new game that meets the defined business and technical requirements. Mountkirk Games wants each regional game instance to be located in multiple Google Cloud regions. What should you do?

A. Configure a global load balancer connected to a managed instance group running Compute Engineinstances.
B. Configure kubemci with a global load balancer and Google Kubernetes Engine.
C. Configure a global load balancer with Google Kubernetes Engine.
D. Configure Ingress for Anthos with a global load balancer and Google Kubernetes Engine.

Answer: A

NEW QUESTION 143
- (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 145
- (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 149
.....

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