Exam Questions Professional-Cloud-Architect

Google Certified Professional - Cloud Architect (GCP)

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

NEW QUESTION 1
- (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 2
- (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 3
- (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 4
- (Topic 2)
For this question, refer to the TerramEarth case study.
The TerramEarth development team wants to create an API to meet the company's business requirements. You want the development team to focus their development effort on business value versus creating a custom framework. Which method should they use?

A. Use Google App Engine with Google Cloud Endpoint
B. Focus on an API for dealers and partners.
C. Use Google App Engine with a JAX-RS Jersey Java-based framewor
D. Focus on an API for the public.
E. Use Google App Engine with the Swagger (open API Specification) framewor
F. Focus on an API for the public.
G. Use Google Container Engine with a Django Python containe
H. Focus on an API for the public.
I. Use Google Container Engine with a Tomcat container with the Swagger (Open API Specification) framewor
J. Focus on an API for dealers and partners.

Answer: A

Explanation:

 https://cloud.google.com/endpoints/docs/openapi/about-cloud- endpoints?hl=en_US&_ga=2.21787131.-1712523161.1522785064
https://cloud.google.com/endpoints/docs/openapi/architecture-overview https://cloud.google.com/storage/docs/gsutil/commands/test
Develop, deploy, protect and monitor your APIs with Google Cloud Endpoints. Using an Open API Specification or one of our API frameworks, Cloud Endpoints gives you the tools you need for every phase of API development.
From scenario: Business Requirements
Decrease unplanned vehicle downtime to less than 1 week, without increasing the cost of carrying surplus inventory
Support the dealer network with more data on how their customers use their equipment to better position new products and services
Have the ability to partner with different companies – especially with seed and fertilizer suppliers in the fast-growing agricultural business – to create compelling joint offerings for their customers.
Reference: https://cloud.google.com/certification/guides/cloud-architect/casestudy- terramearth

NEW QUESTION 5
- (Topic 3)
For this question, refer to the JencoMart case study.
JencoMart wants to move their User Profiles database to Google Cloud Platform. Which Google Database should they use?

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

Answer: D

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

NEW QUESTION 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)
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 8
- (Topic 4)
For this question, refer to the Dress4Win case study.
You want to ensure Dress4Win's sales and tax records remain available for infrequent viewing by auditors for at least 10 years. Cost optimization is your top priority. Which cloud services should you choose?

A. Google Cloud Storage Coldline to store the data, and gsutil to access the data.
B. Google Cloud Storage Nearline to store the data, and gsutil to access the data.
C. Google Bigtabte with US or EU as location to store the data, and gcloud to access the data.
D. BigQuery to store the data, and a web server cluster in a managed instance group to access the dat

E. Google Cloud SQL mirrored across two distinct regions to store the data, and a Redis cluster in a managed instance group to access the data.

Answer: A

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

NEW QUESTION 9
- (Topic 4)
For this question, refer to the Dress4Win case study.
Dress4Win has asked you for advice on how to migrate their on-premises MySQL deployment to the cloud. They want to minimize downtime and performance impact to their on-premises solution during the migration. Which approach should you recommend?

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

Answer: B

NEW QUESTION 10
- (Topic 5)
Your company is planning to upload several important files to Cloud Storage. After the upload is completed, they want to verify that the upload content is identical to what they have on- premises. You want to minimize the cost and effort of performing this check. What should you do?
A.
1) Use gsutil -m to upload all the files to Cloud Storage.
2) Use gsutil cp to download the uploaded files
3) Use Linux diff to compare the content of the files
B.
1) Use gsutil -m to upload all the files to Cloud Storage.
2) Develop a custom Java application that computes CRC32C hashes
3) Use gsutil ls -L gs://[YOUR_BUCKET_NAME] to collect CRC32C hashes of the uploaded files
4) Compare the hashes
C.
1) Use Linux shasum to compute a digest of files you want to upload
2) Use gsutil -m to upload all the files to the Cloud Storage
3) Use gsutil cp to download the uploaded files
4) Use Linux shasum to compute a digest of the downloaded files 5.Compre the hashes
D.
1) Use gsutil -m to upload all the files to Cloud Storage.
2) Use gsutil hash -c FILE_NAME to generate CRC32C hashes of all on-premises files 3)Use gsutil ls -L gs://[YOUR_BUCKET_NAME] to collect CRC32C hashes of the uploaded files
4)Compare the hashes

A.

Answer: D

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

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

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

Answer: A

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

NEW QUESTION 15
- (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 17
- (Topic 5)
You are tasked with building an online analytical processing (OLAP) marketing analytics and reporting tool.
This requires a relational database that can operate on hundreds of terabytes of data. What is the Google recommended tool for such applications?

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

Answer: A

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

NEW QUESTION 18
- (Topic 5)
You are developing your microservices application on Google Kubernetes Engine. During testing, you want to validate the behavior of your application in case a specific microservice
should suddenly crash. What should you do?

A. Add a taint to one of the nodes of the Kubernetes cluste
B. For the specific microservice, configure a pod anti-affinity label that has the name of the tainted node as a value.
C. Use Istio's fault injection on the particular microservice whose faulty behavior you want to simulate.
D. Destroy one of the nodes of the Kubernetes cluster to observe the behavior.
E. Configure Istio's traffic management features to steer the traffic away from a crashing microservice.

Answer: B

Explanation:
Microservice runs on all nodes. The Micro service runs on Pod, Pod runs on Nodes. Nodes is nothing but Virtual machines. Once deployed the application microservices will get deployed across all Nodes. Destroying one node may not mimic the behaviour of microservice crashing as it may be running in other nodes. link: https://istio.io/latest/docs/tasks/traffic-management/fault-injection/

NEW QUESTION 22
- (Topic 5)
You are moving an application that uses MySQL from on-premises to Google Cloud. The application will run on Compute Engine and will use Cloud SQL. You want to cut over to the Compute Engine deployment of the application with minimal downtime and no data loss to your customers. You want to migrate the application with minimal modification. You also need to determine the cutover strategy. What should you do?

A. * 1. Set up Cloud VPN to provide private network connectivity between the Compute Engine application and the on-premises MySQL server.* 2. Stop the on-premises application.* 3. Create a mysqldump of the on-premises MySQL serve
B. * 4.Upload the dump to a Cloud Storage bucket.* 5. Import the dump into Cloud SQL.* 6. Modify the source code of the application to write queries to both databases and read from its local database.* 7. Start the Compute Engine applicatio
C. * 8.Stop the on-premises application.
D. * 1. Set up Cloud SQL proxy and MySQL prox
E. * 2.Create a mysqldump of the on-premises MySQL serve
F. * 3.Upload the dump to a Cloud Storage bucket.* 4.Import the dump into Cloud SQ
G. * 5.Stop the on-premises applicatio
H. * 6.Start the Compute Engine application.
I. * 1. Set up Cloud VPN to provide private network connectivity between the Compute Engine application and the on-premises MySQL server.* 2. Stop the on-premises application.* 3. Start the Compute Engine application, configured to read and write to the on-premises MySQL server.* 4. Create the replication configuration in Cloud SQL.* 5. Configure the source database server to accept connections from the Cloud SQL replic
J. * 6.Finalize the Cloud SQL replica configuration.* 7.When replication has been completed, stop the Compute Engine applicatio
K. * 8.Promote the Cloud SQL replica to a standalone instance.* 9.Restart the Compute Engine application, configured to read and write to the Cloud SQL standalone instance.
L. * 1. Stop the on-premises application.* 2.Create a mysqldump of the on-premises MySQL serve
M. * 3.Upload the dump to a Cloud Storage bucket.* 4. Import the dump into Cloud SQL.* 5. Start the application on Compute Engine.

Answer: C

Explanation:
External replica promotion migration In the migration strategy of external replica promotion, you create an external database replica and synchronize the existing data to that replica. This can happen with minimal downtime to the existing database. When you have a replica database, the two databases have different roles that are referred to in this document as primary and replica. After the data is synchronized, you promote the replica to be the primary in order to move the management layer with minimal impact to database uptime. In Cloud SQL, an easy way to accomplish the external replica promotion is to use the automated migration workflow. This process automates many of the steps that are needed for this type of migration.
https://cloud.google.com/architecture/migrating-mysql-to-cloudsql-concept
- The best option for migrating your MySQL database is to use an external replica promotion. In this strategy, you create a replica database and set your existing database as the primary. You wait until the two databases are in sync, and you then promote your MySQL replica database to be the primary. This process minimizes database downtime related to the database migration. - https://cloud.google.com/architecture/migrating-mysql-to-cloudsql-concept#external_replica_promotion_migration

NEW QUESTION 23
- (Topic 5)
You are creating a solution to remove backup files older than 90 days from your backup Cloud Storage bucket. You want to optimize ongoing Cloud Storage spend. What should you do?

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

Answer: B

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

NEW QUESTION 24
- (Topic 5)
Your application needs to process credit card transactions. You want the smallest scope of
Payment Card Industry (PCI) compliance without compromising the ability to analyze transactional data and trends relating to which payment methods are used. How should you design your architecture?

A. Create a tokenizer service and store only tokenized data.
B. Create separate projects that only process credit card data.
C. Create separate subnetworks and isolate the components that process credit card data.
D. Streamline the audit discovery phase by labeling all of the virtual machines (VMs) that process PCI data.
E. Enable Logging export to Google BigQuery and use ACLs and views to scope the data shared with the auditor.

Answer: A

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

NEW QUESTION 27
- (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 30
- (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 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 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 39
- (Topic 5)
You deploy your custom java application to google app engine. It fails to deploy and gives you the following stack trace:

```
Java.lang.securityException : SHA1 digest

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

At

Sun.securityutil.manifestENtryVerifier.ver

At java . net . URLClassLoader . defineCla

At sun . reflect. GeneratedMethodAccessors

At

Sun.reflect . DelegatingMethodAccesorImp1.

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

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

Answer: B

NEW QUESTION 43

- (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 48

- (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 51

- (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 54

- (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 56

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

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

Answer: A

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

NEW QUESTION 58
- (Topic 5)
Your company wants to track whether someone is present in a meeting room reserved for a scheduled meeting. There are 1000 meeting rooms across 5 offices on 3 continents. Each room is equipped with a motion sensor that reports its status every second. The data from the motion detector includes only a sensor ID and several different discrete items of information. Analysts will use this data, together with information about account owners and office locations. Which database type should you use?

A. Flat file
B. NoSQL
C. Relational
D. Blobstore

Answer: B

Explanation:
Relational databases were not designed to cope with the scale and agility challenges that face modern applications, nor were they built to take advantage of the commodity storage and processing power available today.
NoSQL fits well for:
⎯ Developers are working with applications that create massive volumes of new, rapidly changing data types — structured, semi-structured, unstructured and polymorphic data.

NEW QUESTION 59
- (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 60
- (Topic 5)
Your company acquired a healthcare startup and must retain its customers' medical information for up to 4 more years, depending on when it was created. Your corporate policy is to securely retain this data, and then delete it as soon as regulations allow.
Which approach should you take?

A. Store the data in Google Drive and manually delete records as they expire.
B. Anonymize the data using the Cloud Data Loss Prevention API and store it indefinitely.
C. Store the data using the Cloud Storage and use lifecycle management to delete files when they expire.
D. Store the data in Cloud Storage and run a nightly batch script that deletes all expired datA.

Answer: C

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

NEW QUESTION 61
- (Topic 5)
You have been engaged by your client to lead the migration of their application infrastructure to GCP. One of their current problems is that the on-premises high performance SAN is requiring frequent and expensive upgrades to keep up with the variety of workloads that are identified as follows: 20TB of log archives retained for legal reasons; 500 GB of VM boot/data volumes and templates; 500 GB of image thumbnails; 200 GB of customer session state data that allows customers to restart sessions even if off-line for several days.
Which of the following best reflects your recommendations for a cost-effective storage allocation?

A. Local SSD for customer session state dat
B. Lifecycle-managed Cloud Storage for logarchives, thumbnails, and VM boot/data volumes.
C. Memcache backed by Cloud Datastore for the customer session state dat

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

Answer: D

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

NEW QUESTION 62
- (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 65
- (Topic 5)
Your company wants to start using Google Cloud resources but wants to retain their on- premises Active
Directory domain controller for identity management. What should you do?

A. Use the Admin Directory API to authenticate against the Active Directory domain controller.
B. Use Google Cloud Directory Sync to synchronize Active Directory usernames with cloud identities andconfigure SAML SSO.
C. Use Cloud Identity-Aware Proxy configured to use the on-premises Active Directory domain controller as an identity provider.
D. Use Compute Engine to create an Active Directory (AD) domain controller that is a replica of the onpremises AD domain controller using Google Cloud Directory Sync.

Answer: B

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

NEW QUESTION 69
- (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 70
- (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 72
- (Topic 5)
You need to evaluate your team readiness for a new GCP project. You must perform the evaluation and create a skills gap plan incorporates the business goal of cost optimization. Your team has deployed two GCP projects successfully to date. What should you do?

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

Answer: B

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

NEW QUESTION 76
- (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 81
- (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 86
- (Topic 5)
You are managing several internal applications that are deployed on Compute Engine. Business users inform you that an application has become very slow over the past few days. You want to find the underlying cause in order to solve the problem. What should you do first?

A. Inspect the logs and metrics from the instances in Cloud Logging and Cloud Monitoring.
B. Restore a backup of the application database from a time before the application became slow.
C. Deploy the applications on a managed instance group with autoscaling enable
D. Add aload balancer in front of the managed instance group, and have the users connect to the IP of the load balancer.
E. Change the Compute Engine Instances behind the application to a machine type with more CPU and memory.

Answer: A

Explanation:
When an application becomes slow, the first step you should take is to gather information about the underlying cause of the problem. One way to do this is by inspecting the logs and metrics from the instances where the application is deployed. Google Cloud Platform (GCP) provides tools such as Cloud Logging and Cloud Monitoring that can help you to collect and analyze this information. By reviewing the logs and metrics from the instances, you may be able to identify issues such as resource shortages (e.g. CPU, memory, or disk), network problems, or application errors that are causing the performance issues. Once you have identified the underlying cause of the problem, you can take steps to resolve it.

NEW QUESTION 90
- (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 91
- (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 95
- (Topic 5)
You have an application that will run on Compute Engine. You need to design an architecture that takes into account a disaster recovery plan that requires your

application to fail over to another region in case of a regional outage. What should you do?

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

Answer: C

NEW QUESTION 98
- (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 100
- (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 105
- (Topic 5)
Your architecture calls for the centralized collection of all admin activity and VM system logs within your
project.
How should you collect these logs from both VMs and services?

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

Answer: B

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

NEW QUESTION 106
- (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 110
- (Topic 5)
Your company provides a recommendation engine for retail customers. You are providing retail customers with an API where they can submit a user ID and the API returns a list of recommendations for that user. You are responsible for the API lifecycle and want to ensure stability for your customers in case the API makes backward-incompatible changes. You want to follow Google-recommended practices. What should you do?

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

Answer: C

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

NEW QUESTION 111
- (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 113
- (Topic 5)
A development team at your company has created a dockerized HTTPS web application. You need to deploy the application on Google Kubernetes Engine (GKE) and make sure that the application scales automatically.
How should you deploy to GKE?

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

Answer: B

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

NEW QUESTION 114
- (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 118
- (Topic 5)
Your web application has several VM instances running within a VPC. You want to restrict communications between instances to only the paths and ports you authorize, but you don't want to rely on static IP addresses or subnets because the app can autoscale. How should you restrict communications?

A. Use separate VPCs to restrict traffic
B. Use firewall rules based on network tags attached to the compute instances
C. Use Cloud DNS and only allow connections from authorized hostnames
D. Use service accounts and configure the web application particular service accounts to have access

Answer: B

NEW QUESTION 122
- (Topic 5)
Your company has a Google Cloud project that uses BIgQuery 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 123
- (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 128
- (Topic 5)
Your company has just acquired another company, and you have been asked to integrate their existing Google Cloud environment into your company's data center. Upon investigation, you discover that some of the RFC 1918 IP ranges being used in the new company's Virtual Private Cloud (VPC) overlap with your data center IP space. What should you do to enable connectivity and make sure that there are no routing conflicts when connectivity is established?

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

Answer: A

Explanation:
To connect two networks together we need (1) either VPN or interconnect and (2) peering. When there is peering, you cannot have conflicting IP addresses. You can use either Cloud VPN or Cloud Interconnect to securely connect your on-premises network to your VPC network. (https://cloud.google.com/vpc/docs/vpc-peering#transit-network) At the time of peering, Google Cloud checks to see if there are any subnet IP ranges that overlap subnet IP ranges in the other network. If there is any overlap, peering is not established. (https://cloud.google.com/vpc/docs/vpc-peering#considerations) NAT is used to translate private to public IP and vice versa, however because we are connecting 2 networks together, they become private IPs. So it is not applicable.

NEW QUESTION 129
- (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 133
- (Topic 5)
You want to make a copy of a production Linux virtual machine in the US-Central region. You want to manage and replace the copy easily if there are changes on the production
virtual machine. You will deploy the copy as a new instances in a different project in the US-East region. What steps must you take?

A. Use the Linux dd and netcat command to copy and stream the root disk contents to a new virtual machine instance in the US-East region.
B. Create a snapshot of the root disk and select the snapshot as the root disk when you create a new virtual machine instance in the US-East region.
C. Create an image file from the root disk with Linux dd command, create a new disk from the image file, and use it to create a new virtual machine instance in the US-East region
D. Create a snapshot of the root disk, create an image file in Google Cloud Storage from the snapshot, and create a new virtual machine instance in the US-East region using the image file for the root disk.

Answer: D

Explanation:
https://stackoverflow.com/questions/36441423/migrate-google-compute-engine-instance-to-a-different-region

NEW QUESTION 135
- (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 137
- (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 142
- (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 146
- (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 148
- (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 152
- (Topic 7)
For this question, refer to the TerramEarth case study. You are asked to design a new architecture for the
ingestion of the data of the 200,000 vehicles that are connected to a cellular network. You want to follow
Google-recommended practices.
Considering the technical requirements, which components should you use for the ingestion of the data?

A. Google Kubernetes Engine with an SSL Ingress
B. Cloud IoT Core with public/private key pairs
C. Compute Engine with project-wide SSH keys
D. Compute Engine with specific SSH keys

Answer: A

Explanation:
https://cloud.google.com/solutions/iot-overview https://cloud.google.com/iot/quotas

NEW QUESTION 154
- (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 155
- (Topic 8)
For this question, refer to the Mountkirk Games case study. You need to analyze and define the technical architecture for the compute workloads for your company, Mountkirk Games. Considering the Mountkirk Games business and technical requirements, what should you do?

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

Answer: D

NEW QUESTION 158
- (Topic 8)
Mountkirk Games wants to limit the physical location of resources to their operating Google Cloud regions.
What should you do?

A. Configure an organizational policy which constrains where resources can be deployed.
B. Configure IAM conditions to limit what resources can be configured.
C. Configure the quotas for resources in the regions not being used to 0.
D. Configure a custom alert in Cloud Monitoring so you can disable resources as they are created in otherregions.

Answer: A

NEW QUESTION 163
- (Topic 8)
Your development team has created a mobile game app. You want to test the new mobile app on Android and iOS devices with a variety of configurations. You need to ensure that testing is efficient and cost-effective. What should you do?

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

Answer: C

NEW QUESTION 166
- (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 171
......

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

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

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

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

Money Back Guarantee

Professional-Cloud-Architect Practice Exam Features:

* Professional-Cloud-Architect Questions and Answers Updated Frequently

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

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

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