

Professional-Cloud-Architect Dumps

Google Certified Professional - Cloud Architect (GCP)

<https://www.certleader.com/Professional-Cloud-Architect-dumps.html>



NEW QUESTION 1

- (Topic 1)

For this question, refer to the Mountkirk Games case study.

Mountkirk Games' gaming servers are not automatically scaling properly. Last month, they rolled out a new feature, which suddenly became very popular. A record number of users are trying to use the service, but many of them are getting 503 errors and very slow response times. What should they investigate first?

- A. Verify that the database is online.
- B. Verify that the project quota hasn't been exceeded.
- C. Verify that the new feature code did not introduce any performance bugs.
- D. Verify that the load-testing team is not running their tool against production.

Answer: B

Explanation:

503 is service unavailable error. If the database was online everyone would get the 503 error.
https://cloud.google.com/docs/quota#capping_usage

NEW QUESTION 2

- (Topic 1)

For this question, refer to the Mountkirk Games case study.

Mountkirk Games wants to set up a real-time analytics platform for their new game. The new platform must meet their technical requirements. Which combination of Google technologies will meet all of their requirements?

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

Answer: B

Explanation:

A real time requires Stream / Messaging so Pub/Sub, Analytics by Big Query.

Ingest millions of streaming events per second from anywhere in the world with Cloud Pub/Sub, powered by Google's unique, high-speed private network. Process the streams with Cloud Dataflow to ensure reliable, exactly-once, low-latency data transformation. Stream the transformed data into BigQuery, the cloud-native data warehousing service, for immediate analysis via SQL or popular visualization tools.

From scenario: They plan to deploy the game's backend on Google Compute Engine so they can capture streaming metrics, run intensive analytics.

Requirements for Game Analytics Platform

- ? Dynamically scale up or down based on game activity
- ? Process incoming data on the fly directly from the game servers
- ? Process data that arrives late because of slow mobile networks
- ? Allow SQL queries to access at least 10 TB of historical data
- ? Process files that are regularly uploaded by users' mobile devices
- ? Use only fully managed services

References: <https://cloud.google.com/solutions/big-data/stream-analytics/>

NEW QUESTION 3

- (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 4

- (Topic 2)

For this question, refer to the TerramEarth case study

Your development team has created a structured API to retrieve vehicle data. They want to allow third parties to develop tools for dealerships that use this vehicle event data. You want to support delegated authorization against this data. What should you do?

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

Answer: A

Explanation:

<https://cloud.google.com/appengine/docs/flexible/go/authorizing-apps>

https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#delegate_application_authorization_with_oauth2

Delegate application authorization with OAuth2

Cloud Platform APIs support OAuth 2.0, and scopes provide granular authorization over the methods that are supported. Cloud Platform supports both service-account and user-account OAuth, also called three-legged OAuth.

References: https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#delegate_application_authorization_with_oauth2

<https://cloud.google.com/appengine/docs/flexible/go/authorizing-apps>

NEW QUESTION 5

- (Topic 2)

For this question, refer to the TerramEarth case study.

To speed up data retrieval, more vehicles will be upgraded to cellular connections and be able to transmit data to the ETL process. The current FTP process is error-prone and restarts the data transfer from the start of the file when connections fail, which happens often. You want to improve the reliability of the solution and minimize data transfer time on the cellular connections. What should you do?

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

Answer: D

Explanation:

<https://cloud.google.com/storage/docs/locations>

NEW QUESTION 6

- (Topic 2)

Your agricultural division is experimenting with fully autonomous vehicles.

You want your architecture to promote strong security during vehicle operation. Which two architecture should you consider?

Choose 2 answers:

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

Answer: AC

NEW QUESTION 7

- (Topic 3)

For this question, refer to the JencoMart case study.

JencoMart has built a version of their application on Google Cloud Platform that serves traffic to Asia. You want to measure success against their business and technical goals.

Which metrics should you track?

- A. Error rates for requests from Asia
- B. Latency difference between US and Asia
- C. Total visits, error rates, and latency from Asia
- D. Total visits and average latency for users in Asia
- E. The number of character sets present in the database

Answer: D

NEW QUESTION 8

- (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 9

- (Topic 3)

For this question, refer to the JencoMart case study.

The JencoMart security team requires that all Google Cloud Platform infrastructure is deployed using a least privilege model with separation of duties for administration between production and development resources. What Google domain and project structure should you recommend?

- A. Create two G Suite accounts to manage users: one for development/test/staging and one for production.
- 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 10

- (Topic 4)

For this question, refer to the Dress4Win case study.

Dress4Win has end-to-end tests covering 100% of their endpoints. They want to ensure that the move to the cloud does not introduce any new bugs. Which additional testing methods should the developers employ to prevent an outage?

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

Answer: B

NEW QUESTION 10

- (Topic 4)

For this question, refer to the Dress4Win case study.

Dress4Win has asked you to recommend machine types they should deploy their application servers to. How should you proceed?

- A. Perform a mapping of the on-premises physical hardware cores and RAM to the nearest machine types in the cloud.
- B. Recommend that Dress4Win deploy application servers to machine types that offer the highest RAM to CPU ratio available.
- C. Recommend that Dress4Win deploy into production with the smallest instances available, monitor them over time, and scale the machine type up until the desired performance is reached.
- D. Identify the number of virtual cores and RAM associated with the application server virtual machines align them to a custom machine type in the cloud, monitor performance, and scale the machine types up until the desired performance is reached.

Answer: C

NEW QUESTION 15

- (Topic 4)

For this question, refer to the Dress4Win case study.

You want to ensure Dress4Win's sales and tax records remain available for infrequent viewing by auditors for at least 10 years. Cost optimization is your top priority. Which cloud services should you choose?

- A. Google Cloud Storage Coldline to store the data, and gsutil to access the data.
- B. Google Cloud Storage Nearline to store the data, and gsutil to access the data.
- C. Google Bigtable with US or EU as location to store the data, and gcloud to access the data.
- D. BigQuery to store the data, and a web server cluster in a managed instance group to access the data.
- 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 20

- (Topic 4)

For this question, refer to the Dress4Win case study.

As part of Dress4Win's plans to migrate to the cloud, they want to be able to set up a managed logging and monitoring system so they can handle spikes in their traffic load.

They want to ensure that:

- The infrastructure can be notified when it needs to scale up and down to handle the ebb and flow of usage throughout the day
 - Their administrators are notified automatically when their application reports errors.
 - They can filter their aggregated logs down in order to debug one piece of the application across many hosts
- Which Google StackDriver features should they use?

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

Answer: D

NEW QUESTION 24

- (Topic 5)

Your company captures all web traffic data in Google Analytics 260 and stores it in BigQuery. Each country has its own dataset. Each dataset has multiple tables. You want analysts from each country to be able to see and query only the data for their respective countries. How should you configure the access rights?

- A. Create a group per countr
- B. Add analysts to their respective country-group
- C. Create a single group 'all_analysts', and add all country-groups as member
- D. Grant the 'all-analysis' group the IAM role of BigQuery jobUse
- E. Share the appropriate dataset with view access with each respective analyst country-group.
- F. Create a group per countr
- G. Add analysts to their respective country-group
- H. Create a single group 'all_analysts', and add all country-groups as member
- I. Grant the 'all-analysis' group the IAM role of BigQuery jobUse
- J. Share the appropriate tables with view access with each respective analyst countrygroup.
- K. Create a group per countr
- L. Add analysts to their respective country-group
- M. Create a single group 'all_analysts', and add all country-groups as member
- N. Grant the 'all-analysis' group the IAM role of BigQuery dataViewe
- O. Share the appropriate dataset with view access with each respective analystcountry-group.
- P. Create a group per countr
- Q. Add analysts to their respective country-group
- R. Create a single group 'all_analysts', and add all country-groups as member
- S. Grant the 'all-analysis' group the IAM role of BigQuery dataViewe
- T. Share the appropriate table with view access with each respective analyst countrygroup.

Answer: A

NEW QUESTION 29

- (Topic 5)

Your company wants you to build a highly reliable web application with a few public APIs as the backend. You don't expect a lot of user traffic, but traffic could spike occasionally.

You want to leverage Cloud Load Balancing, and the solution must be cost-effective for users. What should you do?

- A. Store static content such as HTML and images in Cloud CD
- B. Host the APIs on App Engine and store the user data in Cloud SQL.
- C. Store static content such as HTML and images in a Cloud Storage bucke
- D. Host the APIs on a zonal Google Kubernetes Engine cluster with worker nodes in multiple zones, and save the user data in Cloud Spanner.
- E. Store static content such as HTML and images in Cloud CD
- F. Use Cloud Run to host the APIs and save the user data in Cloud SQL.
- G. Store static content such as HTML and images in a Cloud Storage bucke
- H. Use Cloud Functions to host the APIs and save the user data in Firestore.

Answer: D

Explanation:

<https://cloud.google.com/load-balancing/docs/https/setting-up-https-serverless#gcloud:-cloud-functions> <https://cloud.google.com/blog/products/networking/better-load-balancing-for-app-engine-cloud-run-and-functions>

NEW QUESTION 33

- (Topic 5)

Your company has a Google Workspace account and Google Cloud Organization. Some developers in the company have created Google Cloud projects outside of the Google Cloud Organization.

You want to create an Organization structure that allows developers to create projects, but prevents them from modifying production projects. You want to manage policies for all projects centrally and be able to set more restrictive policies for production projects.

You want to minimize disruption to users and developers when business needs change in the future. You want to follow Google-recommended practices. How should you design the Organization structure?

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

development projects outside of the Organization using the developer Google Workspace accounts* 4 Set the policies for all projects on the Organization* 5 Additionally set the production policies on the individual production projects

Answer: C

Explanation:

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

NEW QUESTION 38

- (Topic 5)

You need to ensure reliability for your application and operations by supporting reliable task scheduling for compute on GCP. Leveraging Google best practices, what should you do?

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

Answer: B

Explanation:

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

NEW QUESTION 43

- (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 bucket
- 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 bucket
- 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 45

- (Topic 5)

You are developing a globally scaled frontend for a legacy streaming backend data API.

This API expects

events in strict chronological order with no repeat data for proper processing.

Which products should you deploy to ensure guaranteed-once FIFO (first-in, first-out) delivery of data?

- A. Cloud Pub/Sub alone
- B. Cloud Pub/Sub to Cloud DataFlow
- C. Cloud Pub/Sub to Stackdriver
- D. Cloud Pub/Sub to Cloud SQL

Answer: B

Explanation:

Reference <https://cloud.google.com/pubsub/docs/ordering>

NEW QUESTION 50

- (Topic 5)

You are using Cloud CDN to deliver static HTTP(S) website content hosted on a Compute Engine instance group. You want to improve the cache hit ratio.

What should you do?

- A. Customize the cache keys to omit the protocol from the key.
- B. Shorten the expiration time of the cached objects.
- C. Make sure the HTTP(S) header "Cache-Region" points to the closest region of your users.
- D. Replicate the static content in a Cloud Storage bucket.
- E. Point CloudCDN toward a load balancer on that bucket.

Answer: A

Explanation:

Reference https://cloud.google.com/cdn/docs/bestpractices#using_custom_cache_keys_to_improve_cache_hit_ratio

NEW QUESTION 55

- (Topic 5)

Your operations team currently stores 10 TB of data in 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. Which approach should they use?

- A. Use the gsutil mv command to 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 57

- (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 62

- (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 on-premises database.
- B. Deploy your application on App Engine standard environment and use Cloud VPN to limit access to the on-premises 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 67

- (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 location.
- 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 69

- (Topic 5)

You want to enable your running Google Container Engine cluster to scale as demand for your application changes.

What should you do?

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

Answer: B

Explanation:

<https://cloud.google.com/kubernetes-engine/docs/concepts/cluster-autoscaler>

Cluster autoscaling

`--enable-autoscaling`

Enables autoscaling for a node pool.

Enables autoscaling in the node pool specified by `--node-pool` or the default node pool if `--node-pool` is not provided.

Where:

`--max-nodes=MAX_NODES`

Maximum number of nodes in the node pool.

Maximum number of nodes to which the node pool specified by `--node-pool` (or default node pool if unspecified) can scale.

NEW QUESTION 72

- (Topic 5)

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

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

Answer: CD

Explanation:

Backups help you restore lost data to your Cloud SQL instance. Additionally, if an instance is having a problem, you can restore it to a previous state by using the backup to overwrite it. Enable automated backups for any instance that contains necessary data. Backups protect your data from loss or damage.

Enabling automated backups, along with binary logging, is also required for some operations, such as clone and replica creation.

Reference: <https://cloud.google.com/sql/docs/mysql/backup-recovery/backups>

NEW QUESTION 73

- (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 GCP services 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 78

- (Topic 5)

Your company is migrating its on-premises data center into the cloud. As part of the migration, you want to integrate Kubernetes Engine for workload orchestration. Parts of your architecture must also be PCI DSS compliant.

Which of the following is most accurate?

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

Answer: D

Explanation:

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

NEW QUESTION 80

- (Topic 5)

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

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

Answer: B

NEW QUESTION 81

- (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 82

- (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 Logging
- D. Consult logs for Kubernetes Engine and Cloud SQL.
- E. In the GCP Console, navigate to Cloud SQ
- F. Restore the latest backup
- G. Use kubectl to restart all pods.

Answer: C

NEW QUESTION 85

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 90

- (Topic 5)

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

- A. Set up a network peering between vpc-a and vpc-b
- B. Set up a VPN between vpc-a and vpc-b using Cloud VPN
- C. Configure IAP TCP forwarding on the instance in vpc-b and then launch the following gcloud command from one of the instances in vpc-a: `gcloud compute start-l4o-tunnel INSTANCE_NAME_IPC_IP --local-host-port=localhost:22`
* 1. Create an additional instance in vpc-a * 2. Create an additional instance in vpc-b * 3. Install OpenVPN in newly created instances * 4. Configure a VPN tunnel between vpc-a and vpc-b with the help of OpenVPN

Answer: C

NEW QUESTION 93

- (Topic 5)

Your company is building a new architecture to support its data-centric business focus. You are responsible for setting up the network. Your company's mobile and web-facing applications will be deployed on-premises, and all data analysis will be conducted in GCP. The plan is to process and load 7 years of archived .csv files totaling 900 TB of data and then continue loading 10 TB of data daily. You currently have an existing 100-MB internet connection.

What actions will meet your company's needs?

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

Answer: B

Explanation:

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

NEW QUESTION 95

- (Topic 5)

You have created several preemptible Linux virtual machine instances using Google Compute Engine. You want to properly shut down your application before the virtual machines are preempted. What should you do?

- A. Create a shutdown script named `k99.shutdown` in the `/etc/rc.6.d/` directory.
- B. Create a shutdown script registered as a `xinetd` service in Linux and configure a `Stackdriver` endpoint check to call the service.
- C. Create a shutdown script and use it as the value for a new metadata entry with the key `shutdown-script` in the Cloud Platform Console when you create the new virtual machine instance.
- D. Create a shutdown script, registered as a `xinetd` service in Linux, and use the `gcloud compute instances add-metadata` command to specify the service URL as the value for a new metadata entry with the key `shutdown-script-url`

Answer: C

NEW QUESTION 96

- (Topic 5)

You need to deploy an application to Google Cloud. The application receives traffic via TCP and reads and writes data to the filesystem. The application does not support horizontal scaling. The application process requires full control over the data on the file system because concurrent access causes corruption. The business is willing to accept a downtime when an incident occurs, but the application must be available 24/7 to support their business operations. You need to design the architecture of this application on Google Cloud.

What should you do?

- A. Use a managed instance group with instances in multiple zones, use Cloud Filestore, and use an HTTP load balancer in front of the instances.
- B. Use a managed instance group with instances in multiple zones, use Cloud Filestore, and use a network load balancer in front of the instances.
- C. Use an unmanaged instance group with an active and standby instance in different zones, use a regional persistent disk, and use an HTTP load balancer in front of the instances.
- D. Use an unmanaged instance group with an active and standby instance in different zones, use a regional persistent disk, and use a network load balancer in front of the instances.

Answer: D

Explanation:

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

NEW QUESTION 97

- (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 Stackdriver Trace in order to break down the request latencies at each microservice.

Answer: D

Explanation:

<https://cloud.google.com/trace/docs/overview>

NEW QUESTION 99

- (Topic 5)

You are managing several projects on Google Cloud and need to interact on a daily basis with BigQuery, Bigtable and Kubernetes Engine using the `gcloud` CLI tool. You are travelling a lot and work on different workstations during the week. You want to avoid having to manage the `gcloud` CLI manually. What should you do?

- A. Use a package manager to install `gcloud` on your workstations instead of installing it manually.
- B. Create a Compute Engine instance and install `gcloud` on the instance. Connect to this instance via SSH to always use the same `gcloud` installation when interacting with Google Cloud.
- C. Install `gcloud` on all of your workstations. Run the command `gcloud components auto-update` on each workstation.

D. Use Google Cloud Shell in the Google Cloud Console to interact with Google Cloud

Answer: D

Explanation:

This option allows you to use the gcloud CLI tool without having to install or manage it manually on different workstations. Google Cloud Shell is a browser-based command-line tool that provides you with a temporary Compute Engine virtual machine instance preloaded with the Cloud SDK, including the gcloud CLI tool. You can access Google Cloud Shell from any web browser and use it to interact with BigQuery, Bigtable and Kubernetes Engine using the gcloud CLI tool. The other options are not optimal for this scenario, because they either require installing and updating the gcloud CLI tool on multiple workstations (A, C), or creating and maintaining a Compute Engine instance for the sole purpose of using the gcloud CLI tool (B). References:

? <https://cloud.google.com/shell/docs/overview>

? <https://cloud.google.com/sdk/gcloud/>

NEW QUESTION 101

- (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 105

- (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 data
- B. Lifecycle-managed Cloud Storage for log archives, thumbnails, and VM boot/data volumes.
- C. Memcache backed by Cloud Datastore for the customer session state data
- D. Lifecycle-managed CloudStorage for log archives, thumbnails, and VM boot/data volumes.
- E. Memcache backed by Cloud SQL for customer session state data
- 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 data
- 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 108

- (Topic 5)

Your company has an application deployed on Anthos clusters (formerly Anthos GKE) that is running multiple microservices. The cluster has both Anthos Service Mesh and Anthos Config Management configured. End users inform you that the application is responding very slowly. You want to identify the microservice that is causing the delay. What should you do?

- A. Use the Service Mesh visualization in the Cloud Console to inspect the telemetry between the microservices.
- B. Use Anthos Config Management to create a ClusterSelector selecting the relevant cluster
- C. On the Google Cloud Console page for Google Kubernetes Engine, view the Workloads and filter on the cluster
- D. Inspect the configurations of the filtered workloads.
- E. Use Anthos Config Management to create a namespaceSelector selecting the relevant cluster namespace
- F. On the Google Cloud Console page for Google Kubernetes Engine, visit the workloads and filter on the namespace
- G. Inspect the configurations of the filtered workloads.
- H. Reinstall istio using the default istio profile in order to collect request latency
- I. Evaluate the telemetry between the microservices in the Cloud Console.

Answer: A

Explanation:

The Anthos Service Mesh pages in the Google Cloud Console provide both summary and in-depth metrics, charts, and graphs that enable you to observe service behavior. You can monitor the overall health of your services, or drill down on a specific service to set a service level objective (SLO) or troubleshoot an issue.
<https://cloud.google.com/service-mesh/docs/observability/explore-dashboard> <https://cloud.google.com/anthos/service-mesh>

NEW QUESTION 111

- (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 a secure 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 a secure 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 115

- (Topic 5)

You write a Python script to connect to Google BigQuery from a Google Compute Engine virtual machine. The script is printing errors that it cannot connect to BigQuery. What should you do to fix the script?

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

Answer: B

Explanation:

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

NEW QUESTION 118

- (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 in 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 to 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 disk to the first instance
 - * 2. In case of a zone outage, force-attach the disk to the other instance

A.

Answer: D

NEW QUESTION 123

- (Topic 5)

Your team is developing a web application that will be deployed on Google Kubernetes Engine (GKE). Your CTO expects a successful launch and you need to ensure your application can handle the expected load of tens of thousands of users. You want to test the current deployment to ensure the latency of your application stays below a certain threshold. What should you do?

- A. Use a load testing tool to simulate the expected number of concurrent users and total requests to your application, and inspect the results.
- B. Enable autoscaling on the GKE cluster and enable horizontal pod autoscaling on your application deployment
- C. Send curl requests to your application, and validate if the auto-scaling works.

- D. Replicate the application over multiple GKE clusters in every Google Cloud region. Configure a global HTTP(S) load balancer to expose the different clusters over a single global IP address.
- E. Use Cloud Debugger in the development environment to understand the latency between the different microservices.

Answer: B

NEW QUESTION 126

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 128

- (Topic 5)

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

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

Answer: A

NEW QUESTION 130

- (Topic 5)

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

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

Answer: D

Explanation:

<https://cloud.google.com/architecture/prep-kubernetes-engine-for-prod#binary-authorization>

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

NEW QUESTION 135

- (Topic 5)

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

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

Answer: BE

Explanation:

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

NEW QUESTION 140

- (Topic 5)

Your company has decided to build a backup replica of their on-premises user authentication PostgreSQL database on Google Cloud Platform. The database is 4

TB, and large updates are frequent. Replication requires private address space communication. Which networking approach should you use?

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

Answer: A

Explanation:

<https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations>

Google Cloud Dedicated Interconnect provides direct physical connections and RFC 1918 communication between your on-premises network and Google's network. Dedicated Interconnect enables you to transfer large amounts of data between networks, which can be more cost effective than purchasing additional bandwidth over the public Internet or using VPN tunnels.

Benefits:

? Traffic between your on-premises network and your VPC network doesn't traverse the public Internet. Traffic traverses a dedicated connection with fewer hops, meaning there are less points of failure where traffic might get dropped or disrupted.

? Your VPC network's internal (RFC 1918) IP addresses are directly accessible from your on-premises network. You don't need to use a NAT device or VPN tunnel to reach internal IP addresses. Currently, you can only reach internal IP addresses over a dedicated connection. To reach Google external IP addresses, you must use a separate connection.

? You can scale your connection to Google based on your needs. Connection capacity is delivered over one or more 10 Gbps Ethernet connections, with a maximum of eight connections (80 Gbps total per interconnect).

? The cost of egress traffic from your VPC network to your on-premises network is reduced. A dedicated connection is generally the least expensive method if you have a high-volume of traffic to and from Google's network.

References: <https://cloud.google.com/interconnect/docs/details/dedicated>

NEW QUESTION 144

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 145

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 147

- (Topic 5)

Your company is designing its application landscape on Compute Engine. Whenever a zonal outage occurs, the application should be restored in another zone as quickly as possible with the latest application data. You need to design the solution to meet this requirement. What should you do?

- A. Create a snapshot schedule for the disk containing the application data
- B. Whenever a zonal outage occurs, use the latest snapshot to restore the disk in the same zone.
- C. Configure the Compute Engine instances with an instance template for the application, and use a regional persistent disk for the application data
- D. Whenever a zonal outage occurs, use the instance template to spin up the application in another zone in the same region
- E. Use the regional persistent disk for the application data.
- F. Create a snapshot schedule for the disk containing the application data
- G. Whenever a zonal outage occurs, use the latest snapshot to restore the disk in another zone within the same region.

- H. Configure the Compute Engine instances with an instance template for the application, and use a regional persistent disk for the application data
- I. Whenever a zonal outage occurs, use the instance template to spin up the application in another region
- J. Use the regional persistent disk for the application data,

Answer: B

Explanation:

Regional persistent disk is a storage option that provides synchronous replication of data between two zones in a region. Regional persistent disks can be a good building block to use when you implement HA services in Compute Engine. <https://cloud.google.com/compute/docs/disks/high-availability-regional-persistent-disk>

NEW QUESTION 150

- (Topic 5)

You want to enable your running Google Kubernetes Engine cluster to scale as demand for your application changes. What should you do?

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

Answer: C

Explanation:

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

NEW QUESTION 153

- (Topic 5)

Your company has developed a monolithic, 3-tier application to allow external users to upload and share files. The solution cannot be easily enhanced and lacks reliability. The development team would like to re-architect the application to adopt microservices and a fully managed service approach, but they need to convince their leadership that the effort is worthwhile. Which advantage(s) should they highlight to leadership?

- A. The new approach will be significantly less costly, make it easier to manage the underlying infrastructure, and automatically manage the CI/CD pipelines.
- B. The monolithic solution can be converted to a container with Docker
- C. The generated container can then be deployed into a Kubernetes cluster.
- D. The new approach will make it easier to decouple infrastructure from application, develop and release new features, manage the underlying infrastructure, manage CI/CD pipelines and perform A/B testing, and scale the solution if necessary.
- E. The process can be automated with Migrate for Compute Engine.

Answer: C

Explanation:

The new approach will make it easier to decouple infrastructure from an application, develop and release new features, manage the underlying infrastructure, manage CI/CD pipelines and perform A/B testing, and scale the solution if necessary.

NEW QUESTION 158

- (Topic 5)

Your company is moving 75 TB of data into Google Cloud. You want to use Cloud Storage and follow Google recommended practices. What should you do?

- A. Move your data onto a Transfer Appliance
- B. Use a Transfer Appliance Rehydrator to decrypt the data into Cloud Storage.
- C. Move your data onto a Transfer Appliance
- D. Use Cloud Dataprep to decrypt the data into Cloud Storage.
- E. Install gsutil on each server that contains data
- F. Use resumable transfers to upload the data into Cloud Storage.
- G. Install gsutil on each server containing data
- H. Use streaming transfers to upload the data into Cloud Storage.

Answer: A

Explanation:

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

NEW QUESTION 160

- (Topic 5)

Your company has just recently activated Cloud Identity to manage users. The Google Cloud Organization has been configured as well. The security team needs to secure projects 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 162

- (Topic 5)

The operations team in your company wants to save Cloud VPN log events (or one year). You need to configure the cloud infrastructure to save the logs. What should you do?

- A. Set up a filter in Cloud Logging and a topic in Pub/Sub to publish the logs
- B. Set up a Cloud Logging Dashboard titled Cloud VPN Logs, and then add a chart that queries for the VPN metrics over a one-year time period
- C. Enable the Compute Engine API and then enable logging on the firewall rules that match the traffic you want to save
- D. Set up a filter in Cloud Logging and a Cloud Storage bucket as an export target for the logs you want to save

Answer: D

NEW QUESTION 163

- (Topic 5)

Your customer is moving their corporate applications to Google Cloud Platform. The security team wants detailed visibility of all projects in the organization. You provision the Google Cloud Resource Manager and set up yourself as the org admin. What Google Cloud Identity and Access Management (Cloud IAM) roles should you give to the security team?

- A. Org viewer, project owner
- B. Org viewer, project viewer
- C. Org admin, project browser
- D. Project owner, network admin

Answer: B

Explanation:

<https://cloud.google.com/iam/docs/using-iam-securely>

NEW QUESTION 167

- (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 171

- (Topic 5)

You have an application deployed on Kubernetes Engine using a Deployment named echo-deployment. The deployment is exposed using a Service called echo-service. You need to perform an update to the application with minimal downtime to the application. What should you do?

- A. Use `kubectl set image deployment/echo-deployment <new-image>`
- B. Use the rolling update functionality of the Instance Group behind the Kubernetes cluster
- C. Update the deployment yaml file with the new container image
- D. Use `kubectl delete deployment/echo-deployment` and `kubectl create -f <yaml-file>`
- E. Update the service yaml file with the new container image
- F. Use `kubectl delete service/echoservice` and `kubectl create -f <yaml-file>`

Answer: A

Explanation:

https://cloud.google.com/kubernetes-engine/docs/how-to/updating-apps#updating_an_application

NEW QUESTION 172

- (Topic 5)

Your company has sensitive data in Cloud Storage buckets. Data analysts have Identity Access Management (IAM) permissions to read the buckets. You want to prevent data analysts from retrieving the data in the buckets from outside the office network. What should you do?

- A. * 1. Create a VPC Service Controls perimeter that includes the projects with the buckets.* 2. Create an access level with the CIDR of the office network.
- B. * 1. Create a firewall rule for all instances in the Virtual Private Cloud (VPC) network for source range.* 2. Use the Classless Inter-domain Routing (CIDR) of the office network.
- C. * 1. Create a Cloud Function to remove IAM permissions from the buckets, and another Cloud Function to add IAM permissions to the buckets.* 2. Schedule the Cloud Functions with Cloud Scheduler to add permissions at the start of business and remove permissions at the end of business.
- D. * 1. Create a Cloud VPN to the office network.* 2. Configure Private Google Access for on-premises hosts.

Answer: A

Explanation:

For all Google Cloud services secured with VPC Service Controls, you can ensure that: Resources within a perimeter are accessed only from clients within

authorized VPC networks using Private Google Access with either Google Cloud or on-premises. <https://cloud.google.com/vpc-service-controls/docs/overview>
<https://cloud.google.com/vpc-service-controls/docs/overview>. You create a service control across your VPC and any cloud bucket or any project resource to restrict access. Anything outside of it can't access the resources within service control perimeter

NEW QUESTION 174

- (Topic 5)

Your company has a Google Cloud project that uses BigQuery for data warehousing. There are some tables that contain personally identifiable information (PII). Only the compliance team may access the PII. 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 PII* 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 PII* 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 PII* 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 PII* 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 177

- (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 Deployment
- 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 Deployment
- 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 Pod
- F. Expose the Pod in the cluster using a Service, and use the Service DNS name to address the microservice from other microservices within the cluster.
- G. Deploy each microservice as a Pod
- 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 180

- (Topic 5)

You are analyzing and defining business processes to support your startup's trial usage of GCP, and you don't yet know what consumer demand for your product will be. Your manager requires you to minimize GCP service costs and adhere to Google best practices. What should you do?

- A. Utilize free tier and sustained use discount
- B. Provision a staff position for service cost management.
- C. Utilize free tier and sustained use discount
- D. Provide training to the team about service cost management.
- E. Utilize free tier and committed use discount
- F. Provision a staff position for service cost management.
- G. Utilize free tier and committed use discount
- H. Provide training to the team about service cost management.

Answer: D

Explanation:

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

NEW QUESTION 184

- (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 187

- (Topic 5)

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

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

Answer: B

Explanation:

Dataflow is for processing both the Batch and Stream.

Cloud Dataflow is a fully-managed service for transforming and enriching data in stream (real time) and batch (historical) modes with equal reliability and expressiveness -- no more complex workarounds or compromises needed.

References: <https://cloud.google.com/dataflow/>

NEW QUESTION 191

- (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 192

- (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 Engine
- B. Deploy a Cloud SQL server to replace MySQL
- 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 194

- (Topic 6)

For this question, refer to the Dress4Win case study. Dress4Win is expected to grow to 10 times its size in 1 year with a corresponding growth in data and traffic that mirrors the existing patterns of usage. The CIO has set the target of migrating production infrastructure to the cloud within the next 6 months. How will you configure the solution to scale for this growth without making major application changes and still maximize the ROI?

- A. Migrate the web application layer to App Engine, and MySQL to Cloud Datastore, and NAS to Cloud Storag
- B. Deploy RabbitMQ, and deploy Hadoop servers using Deployment Manager.
- C. Migrate RabbitMQ to Cloud Pub/Sub, Hadoop to BigQuery, and NAS to Compute Engine with Persistent Disk storag
- D. Deploy Tomcat, and deploy Nginx using Deployment Manager.
- E. Implement managed instance groups for Tomcat and Ngin
- F. Migrate MySQL to Cloud SQL, RabbitMQ to Cloud Pub/Sub, Hadoop to Cloud Dataproc, and NAS to Compute Engine with Persistent Disk storage.

- G. Implement managed instance groups for the Tomcat and Ngin
- H. Migrate MySQL to Cloud SQL, RabbitMQ to Cloud Pub/Sub, Hadoop to Cloud Dataproc, and NAS to Cloud Storage.

Answer: D

NEW QUESTION 199

- (Topic 6)

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

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

Answer: D

Explanation:

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

NEW QUESTION 201

- (Topic 7)

You have broken down a legacy monolithic application into a few containerized RESTful microservices. You want to run those microservices on Cloud Run. You also want to make sure the services are highly available with low latency to your customers. What should you do?

- A. Deploy Cloud Run services to multiple availability zone
- B. Create Cloud Endpoints that point to the service
- C. Create a global HTTP(S) Load Balancing instance and attach the Cloud Endpoints to its backend.
- D. Deploy Cloud Run services to multiple regions Create serverless network endpoint groups pointing to the service
- E. Add the serverless NEGs to a backend service that is used by a global HTTP(S) Load Balancing instance.
- F. Cloud Run services to multiple region
- G. In Cloud DNS, create a latency-based DNS name that points to the services.
- H. Deploy Cloud Run services to multiple availability zone
- I. Create a TCP/IP global load balance
- J. Add the Cloud Run Endpoints to its backend service.

Answer: B

Explanation:

<https://cloud.google.com/run/docs/multiple-regions>

NEW QUESTION 203

- (Topic 7)

For this question, refer to the TerraEarth case study. Considering the technical requirements, how should you reduce the unplanned vehicle downtime in GCP?

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

Answer: A

NEW QUESTION 207

- (Topic 7)

For this question, refer to the TerraEarth 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 208

- (Topic 8)

For this question, refer to the Mountkirk Games case study. Which managed storage option meets Mountkirk's technical requirement for storing game activity in a time series database service?

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

Answer: A

Explanation:

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

NEW QUESTION 212

- (Topic 8)

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

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

Answer: A

NEW QUESTION 213

- (Topic 8)

For this question, refer to the Mountkirk Games case study. Mountkirk Games wants you to design a way to test the analytics platform's resilience to changes in mobile network latency. What should you do?

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

Answer: D

NEW QUESTION 215

- (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 Engine instances.
- 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 220

- (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 policies
- F. Use preemptible Compute Engine instances.
- G. Create a global load balancer with managed instance groups and autoscaling policies
- H. Use non-preemptible Compute Engine instances.

Answer: D

NEW QUESTION 224

- (Topic 8)

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

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

Answer: C

NEW QUESTION 227

- (Topic 8)

You are implementing Firestore for Mountkirk Games. Mountkirk Games wants to give a new game programmatic access to a legacy game's Firestore database. Access should be as restricted as possible. What should you do?

- A. Create a service account (SA) in the legacy game's Google Cloud project, add this SA in the new game's IAM page, and then give it the Firebase Admin role in both projects
- B. Create a service account (SA) in the legacy game's Google Cloud project, add a second SA in the new game's IAM page, and then give the Organization Admin role to both SAs
- C. Create a service account (SA) in the legacy game's Google Cloud project, give it the Firebase Admin role, and then migrate the new game to the legacy game's project.
- D. Create a service account (SA) in the legacy game's Google Cloud project, give the SA the Organization Admin rule and then give it the Firebase Admin role in both projects

Answer: A

NEW QUESTION 229

- (Topic 9)

For this question, refer to the Helicopter Racing League (HRL) case study. The HRL development team releases a new version of their predictive capability application every Tuesday evening at 3

- A. a.
- B. UTC to a repository
- C. The security team at HRL has developed an in-house penetration test Cloud Function called Airwolf. The security team wants to run Airwolf against the predictive capability application as soon as it is released every Tuesday
- D. You need to set up Airwolf to run at the recurring weekly cadence
- E. What should you do?
- F. Set up Cloud Tasks and a Cloud Storage bucket that triggers a Cloud Function.
- G. Set up a Cloud Logging sink and a Cloud Storage bucket that triggers a Cloud Function.
- H. Configure the deployment job to notify a Pub/Sub queue that triggers a Cloud Function.
- I. Set up Identity and Access Management (IAM) and Confidential Computing to trigger a Cloud Function.

Answer: A

NEW QUESTION 234

- (Topic 9)

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

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

Answer: C

Explanation:

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

NEW QUESTION 239

- (Topic 9)

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

- A. `gcloud compute firewall rules update hlr-policy --priority 1000 --target tags=sourceip:fastly --allow tcp:443`
- B. `gcloud compute security policies rules update 1000 --security-policy hlr-policy --expression "evaluatePreconfiguredExpr('sourceip:fastly')" --action "allow"`

- C. gcloud compute firewall rules update sourceip:fastly \priority 1000 \allow tcp: 443
D. gcloud compute priority-policies rules update 1000 \security policy from fastly--src- ip-ranges"-- action " allow"

Answer: B

Explanation:

Reference: <https://cloud.google.com/load-balancing/docs/httpsD18912E1457D5D1DDCBD40AB3BF70D5D>

NEW QUESTION 241

- (Topic 10)

You need to upgrade the EHR connection to comply with their requirements. The new connection design must support business-critical needs and meet the same network and security policy requirements. What should you do?

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

Answer: D

NEW QUESTION 243

- (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 248

- (Topic 10)

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

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

Answer: C

Explanation:

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

NEW QUESTION 253

- (Topic 10)

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

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

Answer: A

Explanation:

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

NEW QUESTION 256

- (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 259

- (Topic 10)

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

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

Answer: A

Explanation:

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

NEW QUESTION 261

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your Professional-Cloud-Architect Exam with Our Prep Materials Via below:

<https://www.certleader.com/Professional-Cloud-Architect-dumps.html>