

Exam Questions Professional-Cloud-Architect

Google Certified Professional - Cloud Architect (GCP)

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



NEW QUESTION 1

- (Topic 1)

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

Mountkirk Games wants to set up a continuous delivery pipeline. Their architecture includes many small services that they want to be able to update and roll back quickly. Mountkirk Games has the following requirements:

- Services are deployed redundantly across multiple regions in the US and Europe.
- Only frontend services are exposed on the public internet.
- They can provide a single frontend IP for their fleet of services.
- Deployment artifacts are immutable. Which set of products should they use?

- A. Google Cloud Storage, Google Cloud Dataflow, Google Compute Engine
- B. Google Cloud Storage, Google App Engine, Google Network Load Balancer
- C. Google Kubernetes Registry, Google Container Engine, Google HTTP(S) Load Balancer
- D. Google Cloud Functions, Google Cloud Pub/Sub, Google Cloud Deployment Manager

Answer: C

NEW QUESTION 2

- (Topic 2)

For this question, refer to the TerramEarth case study.

TerramEarth's 20 million vehicles are scattered around the world. Based on the vehicle's location its telemetry data is stored in a Google Cloud Storage (GCS) regional bucket (US, Europe, or Asia). The CTO has asked you to run a report on the raw telemetry data to determine why vehicles are breaking down after 100 K miles. You want to run this job on all the data. What is the most cost-effective way to run this job?

- A. Move all the data into 1 zone, then launch a Cloud Dataproc cluster to run the job.
- B. Move all the data into 1 region, then launch a Google Cloud Dataproc cluster to run the job.
- C. Launch a cluster in each region to preprocess and compress the raw data, then move the data into a multi region bucket and use a Dataproc cluster to finish the job.
- D. Launch a cluster in each region to preprocess and compress the raw data, then move the data into a region bucket and use a Cloud Dataproc cluster to finish the job.

Answer: D

Explanation:

Storage guarantees 2 replicates which are geo diverse (100 miles apart) which can get better remote latency and availability.

More importantly, is that multiregional heavily leverages Edge caching and CDNs to provide the content to the end users.

All this redundancy and caching means that Multiregional comes with overhead to sync and ensure consistency between geo-diverse areas. As such, it's much better for write-once- read-many scenarios. This means frequently accessed (e.g. "hot" objects) around the world, such as website content, streaming videos, gaming or mobile applications.

References: <https://medium.com/google-cloud/google-cloud-storage-what-bucket-class-for-the-best-performance-5c847ac8f9f2>

NEW QUESTION 3

- (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 4

- (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 5

- (Topic 4)

For this question, refer to the Dress4Win case study.

As part of their new application experience, Dress4Win allows customers to upload images of themselves. The customer has exclusive control over who may view these images. Customers should be able to upload images with minimal latency and also be shown their images quickly on the main application page when they log in. Which configuration should Dress4Win use?

- A. Store image files in a Google Cloud Storage bucket
- B. Use Google Cloud Datastore to maintain metadata that maps each customer's ID and their image files.
- C. Store image files in a Google Cloud Storage bucket
- D. Add custom metadata to the uploaded images in Cloud Storage that contains the customer's unique ID.
- E. Use a distributed file system to store customers' image
- F. As storage needs increase, add more persistent disks and/or nodes
- G. Assign each customer a unique ID, which sets each file's owner attribute, ensuring privacy of images.
- H. Use a distributed file system to store customers' image
- I. As storage needs increase, add more persistent disks and/or nodes
- J. Use a Google Cloud SQL database to maintain metadata that maps each customer's ID to their image files.

Answer: A

NEW QUESTION 6

- (Topic 4)

For this question, refer to the Dress4Win case study.

The Dress4Win security team has disabled external SSH access into production virtual machines (VMs) on Google Cloud Platform (GCP). The operations team needs to remotely manage the VMs, build and push Docker containers, and manage Google Cloud Storage objects. What can they do?

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

Answer: A

NEW QUESTION 7

- (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 8

- (Topic 4)

For this question, refer to the Dress4Win case study.

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

- A. Google Cloud Storage Coldline to store the data, and gsutil to access the data.
- B. Google Cloud Storage Nearline to store the data, and gsutil to access the data.
- C. Google 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 9

- (Topic 4)

For this question, refer to the Dress4Win case study.

At Dress4Win, an operations engineer wants to create a low-cost solution to remotely archive copies of database backup files. The database files are compressed tar files stored in their current data center. How should he proceed?

- A. Create a cron script using gsutil to copy the files to a Coldline Storage bucket.
- B. Create a cron script using gsutil to copy the files to a Regional Storage bucket.
- C. Create a Cloud Storage Transfer Service Job to copy the files to a Coldline Storage bucket.
- D. Create a Cloud Storage Transfer Service job to copy the files to a Regional Storage bucket.

Answer: A

Explanation:

Follow these rules of thumb when deciding whether to use gsutil or Storage Transfer Service:

? When transferring data from an on-premises location, use gsutil.

? When transferring data from another cloud storage provider, use Storage Transfer Service.

? Otherwise, evaluate both tools with respect to your specific scenario.

Use this guidance as a starting point. The specific details of your transfer scenario will also help you determine which tool is more appropriate

<https://cloud.google.com/storage-transfer/docs/overview>

NEW QUESTION 10

- (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 10

- (Topic 5)

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

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

Answer: A

Explanation:

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

? <https://cloud.google.com/architecture/framework/system-design>

? <https://cloud.google.com/architecture/identity/best-practices-for-planning>

? <https://cloud.google.com/resource-manager/docs/creating-managing-folders>

NEW QUESTION 11

- (Topic 5)

Your company is running its application workloads on Compute Engine. The applications have been deployed in production, acceptance, and development environments. The production environment is business-critical and is used 24/7, while the acceptance and development environments are only critical during office hours. Your CFO has asked you to optimize these environments to achieve cost savings during idle times. What should you do?

- A. Create a shell script that uses the gcloud command to change the machine type of the development and acceptance instances to a smaller machine type outside of office hour.
- B. Schedule the shell script on one of the production instances to automate the task.
- C. Use Cloud Scheduler to trigger a Cloud Function that will stop the development and acceptance environments after office hours and start them just before office hours.
- D. Deploy the development and acceptance applications on a managed instance group and enable autoscaling.
- E. Use regular Compute Engine instances for the production environment, and use preemptible VMs for the acceptance and development environments.

Answer: B

Explanation:

Reference: <https://cloud.google.com/blog/products/it-ops/best-practices-for-optimizing-your-cloud-costs>

NEW QUESTION 16

- (Topic 5)

Your company has an application running on Compute Engine that allows users to play their favorite music. There are a fixed number of instances. Files are stored in Cloud Storage and data is streamed directly to users. Users are reporting that they sometimes need to attempt to play popular songs multiple times before they are successful. You need to improve the performance of the application. What should you do?

A.

- * 1. Copy popular songs into CloudSQL as a blob
- * 2. Update application code to retrieve data from CloudSQL when Cloud Storage is overloaded

B.

- * 1. Create a managed instance group with Compute Engine instances
- * 2. Create a global load balancer and configure it with two backends
- * Managed instance group
- * Cloud Storage bucket
- * 3. Enable Cloud CDN on the bucket backend

C.

- * 1. Mount the Cloud Storage bucket using gcsfuse on all backend Compute Engine instances
- * 2. Serve music files directly from the backend Compute Engine instance

D.

- * 1. Create a Cloud Filestore NFS volume and attach it to the backend Compute Engine instances
- * 2. Download popular songs in Cloud Filestore
- * 3. Serve music files directly from the backend Compute Engine instance

A.

Answer: B

NEW QUESTION 19

- (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 bucket
- 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 bucket
- 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 24

- (Topic 5)

Your company has a networking team and a development team. The development team runs applications on Compute Engine instances that contain sensitive data. The development team requires administrative permissions for Compute Engine. Your company requires all network resources to be managed by the networking team. The development team does not want the networking team to have access to the sensitive data on the instances. What should you do?

- A. * 1. Create a project with a standalone VPC and assign the Network Admin role to the networking team.* 2. Create a second project with a standalone VPC and assign the Compute Admin role to the development team.* 3. Use Cloud VPN to join the two VPCs.
- B. * 1. Create a project with a standalone Virtual Private Cloud (VPC), assign the Network Admin role to the networking team, and assign the Compute Admin role to the development team.
- C. * 1. Create a project with a Shared VPC and assign the Network Admin role to the networking team.* 2. Create a second project without a VPC, configure it as a Shared VPC service project, and assign the Compute Admin role to the development team.
- D. * 1. Create a project with a standalone VPC and assign the Network Admin role to the networking team.* 2. Create a second project with a standalone VPC and assign the Compute Admin role to the development team.* 3. Use VPC Peering to join the two VPCs.

Answer: C

Explanation:

In this scenario, a large organization has a central team that manages security and networking controls for the entire organization. Developers do not have permissions to make changes to any network or security settings defined by the security and networking team but they are granted permission to create resources such as virtual machines in shared subnets. To facilitate this the organization makes use of a shared VPC (Virtual Private Cloud). A shared VPC allows creation of a VPC network of RFC 1918 IP spaces that associated projects (service projects) can then use. Developers using the associated projects can create VM instances in the shared VPC network spaces. The organization's network and security admins can create subnets, VPNs, and firewall rules usable by all the projects in the VPC network. https://cloud.google.com/iam/docs/job-functions/networking#single_team_manages_security_network_for_organization

Reference: <https://cloud.google.com/vpc/docs/shared-vpc>

NEW QUESTION 29

- (Topic 5)

Your company is planning to upload several important files to Cloud Storage. After the upload is completed, they want to verify that the upload content is identical

to what they have on- premises. You want to minimize the cost and effort of performing this check. What should you do?

A.

- 1) Use gsutil -m to upload all the files to Cloud Storage.
- 2) Use gsutil cp to download the uploaded files
- 3) Use Linux diff to compare the content of the files

B.

- 1) Use gsutil -m to upload all the files to Cloud Storage.
- 2) Develop a custom Java application that computes CRC32C hashes
- 3) Use gsutil ls -L gs://[YOUR_BUCKET_NAME] to collect CRC32C hashes of the uploaded files
- 4) Compare the hashes

C.

- 1) Use Linux shasum to compute a digest of files you want to upload
- 2) Use gsutil -m to upload all the files to the Cloud Storage
- 3) Use gsutil cp to download the uploaded files
- 4) Use Linux shasum to compute a digest of the downloaded files 5. Compare the hashes

D.

- 1) Use gsutil -m to upload all the files to Cloud Storage.
- 2) Use gsutil hash -c FILE_NAME to generate CRC32C hashes of all on-premises files 3) Use gsutil ls -L gs://[YOUR_BUCKET_NAME] to collect CRC32C hashes of the uploaded files
- 4) Compare the hashes

A.

Answer: D

Explanation:

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

NEW QUESTION 33

- (Topic 5)

Your company has an application running on multiple Compute Engine instances. You need to ensure that the application can communicate with an on-premises service that requires high throughput via internal IPs, while minimizing latency. What should you do?

- A. Use OpenVPN to configure a VPN tunnel between the on-premises environment and Google Cloud.
- B. Configure a direct peering connection between the on-premises environment and Google Cloud.
- C. Use Cloud VPN to configure a VPN tunnel between the on-premises environment and Google Cloud.
- D. Configure a Cloud Dedicated Interconnect connection between the on-premises environment and Google Cloud.

Answer: D

Explanation:

Reference <https://cloud.google.com/architecture/setting-up-private-access-to-cloud-apis-through-vpn-tunnels>

NEW QUESTION 37

- (Topic 5)

Your company recently acquired a company that has infrastructure in Google Cloud. Each company has its own Google Cloud organization. Each company is using a Shared Virtual Private Cloud (VPC) to provide network connectivity for its applications. Some of the subnets used by both companies overlap. In order for both businesses to integrate, the applications need to have private network connectivity. These applications are not on overlapping subnets. You want to provide connectivity with minimal re-engineering. What should you do?

- A. Set up VPC peering and peer each Shared VPC together
- B. Configure SSH port forwarding on each application to provide connectivity between applications in the different Shared VPCs
- C. Migrate the projects from the acquired company into your company's Google Cloud organization. Relaunch the instances in your company's Shared VPC
- D. Set up a Cloud VPN gateway in each Shared VPC and peer Cloud VPNs

Answer: B

NEW QUESTION 41

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 42

- (Topic 5)

Your organization wants to control IAM policies for different departments independently, but centrally.

Which approach should you take?

- A. Multiple Organizations with multiple Folders
- B. Multiple Organizations, one for each department
- C. A single Organization with Folder for each department
- D. A single Organization with multiple projects, each with a central owner

Answer: C

Explanation:

Folders are nodes in the Cloud Platform Resource Hierarchy. A folder can contain projects, other folders, or a combination of both. You can use folders to group projects under an organization in a hierarchy. For example, your organization might contain multiple departments, each with its own set of GCP resources. Folders allow you to group these

resources on a per-department basis. Folders are used to group resources that share common IAM policies. While a folder can contain multiple folders or resources, a given folder or resource can have exactly one parent.

References: <https://cloud.google.com/resource-manager/docs/creating-managing-folders>

NEW QUESTION 43

- (Topic 5)

Your BigQuery project has several users. For audit purposes, you need to see how many queries each user ran in the last month.

- A. Connect Google Data Studio to BigQuery
- B. Create a dimension for the users and a metric for the amount of queries per user.
- C. In the BigQuery interface, execute a query on the JOBS table to get the required information.
- D. Use 'bq show' to list all job
- E. Per job, use 'bq ls' to list job information and get the required information.
- F. Use Cloud Audit Logging to view Cloud Audit Logs, and create a filter on the query operation to get the required information.

Answer: C

Explanation:

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

NEW QUESTION 46

- (Topic 5)

Your company has a Google Cloud project that uses BigQuery for data warehousing on a pay-per-use basis. You want to monitor queries in real time to discover the most costly queries and which users spend the most. What should you do?

- A.
 - * 1. Create a Cloud Logging sink to export BigQuery data access logs to Cloud Storage.
 - * 2. Develop a Dataflow pipeline to compute the cost of queries split by users.
- B.
 - * 1. Create a Cloud Logging sink to export BigQuery data access logs to BigQuery.
 - * 2. Perform a BigQuery query on the generated table to extract the information you need.
- C.
 - * 1. Activate billing export into BigQuery.
 - * 2. Perform a BigQuery query on the billing table to extract the information you need.
- D.
 - * 1. In the BigQuery dataset that contains all the tables to be queried, add a label for each user that can launch a query.
 - * 2. Open the Billing page of the project.
 - * 3. Select Reports.
 - * 4. Select BigQuery as the product and filter by the user you want to check.

A.

Answer: C

Explanation:

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

NEW QUESTION 49

- (Topic 5)

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

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

Answer: ABF

NEW QUESTION 53

- (Topic 5)

You are deploying a PHP App Engine Standard service with SQL as the backend. You want to minimize the number of queries to the database. What should you do?

- A. Set the memcache service level to dedicate

- B. Create a key from the hash of the query, and return database values from memcache before issuing a query to Cloud SQL.
- C. Set the memcache service level to dedicate
- D. Create a cron task that runs every minute to populate the cache with keys containing query results.
- E. Set the memcache service level to share
- F. Create a cron task that runs every minute to save all expected queries to a key called “cached-queries”.
- G. Set the memcache service level to share
- H. Create a key called “cached-queries”, and return database values from the key before using a query to Cloud SQL.

Answer: A

Explanation:

<https://cloud.google.com/appengine/docs/standard/php/memcache/using>

NEW QUESTION 55

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 59

- (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 62

- (Topic 5)

Your application needs to process credit card transactions. You want the smallest scope of

Payment Card Industry (PCI) compliance without compromising the ability to analyze transactional data and trends relating to which payment methods are used.

How should you design your architecture?

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

Answer: A

Explanation:

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

NEW QUESTION 64

- (Topic 5)

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

- A. Enable the Cloud Trace API on your project and use Cloud Monitoring Alerts to send an alert based on the Cloud Trace metrics
- B. Configure Anthos Config Management on your cluster and create a yaml file that defines the SLO and alerting policy you want to deploy in your cluster
- C. Use Cloud Profiler to follow up the request latenc

- D. Create a custom metric in Cloud Monitoring based on the results of Cloud Profiler, and create an Alerting Policy in case this metric exceeds the threshold
- E. Install Anthos Service Mesh on your cluster
- F. Use the Google Cloud Console to define a Service Level Objective (SLO)

Answer: D

Explanation:

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

NEW QUESTION 67

- (Topic 5)

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

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

Answer: C

NEW QUESTION 71

- (Topic 5)

Your company has decided to make a major revision of their API in order to create better experiences for their developers. They need to keep the old version of the API available and deployable, while allowing new customers and testers to try out the new API. They want to keep the same SSL and DNS records in place to serve both APIs. What should they do?

- A. Configure a new load balancer for the new version of the API.
- B. Reconfigure old clients to use a new endpoint for the new API.
- C. Have the old API forward traffic to the new API based on the path.
- D. Use separate backend pools for each API path behind the load balancer.

Answer: D

Explanation:

<https://cloud.google.com/endpoints/docs/openapi/lifecycle-management>

NEW QUESTION 76

- (Topic 5)

Your company is using Google Cloud. You have two folders under the Organization: Finance and Shopping. The members of the development team are in a Google Group. The development team group has been assigned the Project Owner role on the Organization. You want to prevent the development team from creating resources in projects in the Finance folder. What should you do?

- A. Assign the development team group the Project Viewer role on the Finance folder, and assign the development team group the Project Owner role on the Shopping folder.
- B. Assign the development team group only the Project Viewer role on the Finance folder.
- C. Assign the development team group the Project Owner role on the Shopping folder, and remove the development team group Project Owner role from the Organization.
- D. Assign the development team group only the Project Owner role on the Shopping folder.

Answer: C

Explanation:

<https://cloud.google.com/resource-manager/docs/cloud-platform-resource-hierarchy>

"Roles are always inherited, and there is no way to explicitly remove a permission for a lower-level resource that is granted at a higher level in the resource hierarchy. Given the above example, even if you were to remove the Project Editor role from Bob on the "Test GCP Project", he would still inherit that role from the "Dept Y" folder, so he would still have the permissions for that role on "Test GCP Project"."

Reference: <https://cloud.google.com/resource-manager/docs/creating-managing-folders>

NEW QUESTION 77

- (Topic 5)

Your customer runs a web service used by e-commerce sites to offer product recommendations to users. The company has begun experimenting with a machine learning model on Google Cloud Platform to improve the quality of results.

What should the customer do to improve their model's results over time?

- A. Export Cloud Machine Learning Engine performance metrics from Stackdriver to BigQuery, to be used to analyze the efficiency of the model.
- B. Build a roadmap to move the machine learning model training from Cloud GPUs to Cloud TPUs, which offer better results.
- C. Monitor Compute Engine announcements for availability of newer CPU architectures, and deploy the model to them as soon as they are available for additional performance.
- D. Save a history of recommendations and results of the recommendations in BigQuery, to be used as training data.

Answer: D

Explanation:

<https://cloud.google.com/solutions/building-a-serverless-ml-model>

NEW QUESTION 82

- (Topic 5)

Your customer wants to do resilience testing of their authentication layer. This consists of a regional managed instance group serving a public REST API that reads from and writes to a Cloud SQL instance.

What should you do?

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

Answer: C

NEW QUESTION 86

- (Topic 5)

A lead software engineer tells you that his new application design uses websockets and HTTP sessions that are not distributed across the web servers. You want to help him ensure his application will run properly on Google Cloud Platform. What should you do?

- A. Help the engineer to convert his websocket code to use HTTP streaming.
- B. Review the encryption requirements for websocket connections with the security team.
- C. Meet with the cloud operations team and the engineer to discuss load balancer options.
- D. Help the engineer redesign the application to use a distributed user session service that does not rely on websockets and HTTP sessions.

Answer: C

Explanation:

Google Cloud Platform (GCP) HTTP(S) load balancing provides global load balancing for HTTP(S) requests destined for your instances. The HTTP(S) load balancer has native support for the WebSocket protocol.

NEW QUESTION 88

- (Topic 5)

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

```
FROM ubuntu:16.04

COPY . /src

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

RUN pip install -r requirements.txt
```

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

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

Answer: CE

Explanation:

The speed of deployment can be changed by limiting the size of the uploaded app, limiting the complexity of the build necessary in the Dockerfile, if present, and by ensuring a fast and reliable internet connection.

Note: Alpine Linux is built around musl libc and busybox. This makes it smaller and more resource efficient than traditional GNU/Linux distributions. A container requires no more

than 8 MB and a minimal installation to disk requires around 130 MB of storage. Not only do you get a fully-fledged Linux environment but a large selection of packages from the repository.

References: <https://groups.google.com/forum/#!topic/google-appengine/hZMEkmmObDU> <https://www.alpinelinux.org/about/>

NEW QUESTION 92

- (Topic 5)

Your company has an application that is running on multiple instances of Compute Engine. It generates 1 TB per day of logs. For compliance reasons, the logs need to be kept for at least two years. The logs need to be available for active query for 30 days. After that, they just need to be retained for audit purposes. You want to implement a storage solution that is compliant, minimizes costs, and follows Google-recommended practices. What should you do?

- A.
 - * 1. Install the Cloud Ops agent on all instances.
 - * 2. Create a sink to export logs into a partitioned BigQuery table.
 - * 3. Set a time_partitioning_expiration of 30 days.
- B.
 - * 1. Install the Cloud Ops agent on all instances.
 - * 2. Create a sink to export logs into a regional Cloud Storage bucket.
 - * 3. Create an Object Lifecycle rule to move files into a Coldline Cloud Storage bucket after one month.
 - * 4. Configure a retention policy at the bucket level to create a lock.
- C.
 - * 1. Create a daily cron job, running on all instances, that uploads logs into a partitioned BigQuery table.
 - * 2. Set a time_partitioning_expiration of 30 days.
- D.

- * 1. Write a daily cron job, running on all instances, that uploads logs into a Cloud Storage bucket.
- * 2. Create a sink to export logs into a regional Cloud Storage bucket.
- * 3. Create an Object Lifecycle rule to move files into a Coldline Cloud Storage bucket after one month.

A.

Answer: B

Explanation:

The practice for managing logs generated on Compute Engine on Google Cloud is to install the Cloud Logging agent and send them to Cloud Logging. The sent logs will be aggregated into a Cloud Logging sink and exported to Cloud Storage. The reason for using Cloud Storage as the destination for the logs is that the requirement in question requires setting up a lifecycle based on the storage period. In this case, the log will be used for active queries for 30 days after it is saved, but after that, it needs to be stored for a longer period of time for auditing purposes. If the data is to be used for active queries, we can use BigQuery's Cloud Storage data query feature and move the data past 30 days to Coldline to build a cost-optimal solution. Therefore, the correct answer is as follows

- * 1. Install the Cloud Logging agent on all instances.

Create a sync that exports the logs to the region's Cloud Storage bucket.

- * 3. Create an Object Lifecycle rule to move the files to the Coldline Cloud Storage bucket after one month. * 4.
- * 4. set up a bucket-level retention policy using bucket locking."

NEW QUESTION 97

- (Topic 5)

Your organization requires that metrics from all applications be retained for 5 years for future analysis in possible legal proceedings. Which approach should you use?

- A. Grant the security team access to the logs in each Project.
- B. Configure Stackdriver Monitoring for all Projects, and export to BigQuery.
- C. Configure Stackdriver Monitoring for all Projects with the default retention policies.
- D. Configure Stackdriver Monitoring for all Projects, and export to Google Cloud Storage.

Answer: D

Explanation:

Overview of storage classes, price, and use cases <https://cloud.google.com/storage/docs/storage-classes>
Why export logs? <https://cloud.google.com/logging/docs/export/>
StackDriver Quotas and Limits for Monitoring <https://cloud.google.com/monitoring/quotas> The BigQuery pricing. <https://cloud.google.com/bigquery/pricing>

NEW QUESTION 100

- (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 102

- (Topic 5)

To reduce costs, the Director of Engineering has required all developers to move their development infrastructure resources from on-premises virtual machines (VMs) to Google Cloud Platform. These resources go through multiple start/stop events during the day and require state to persist. You have been asked to design the process of running a development environment in Google Cloud while providing cost visibility to the finance department. Which two steps should you take? Choose 2 answers

- A. Use the --no-auto-delete flag on all persistent disks and stop the VM.
- B. Use the -auto-delete flag on all persistent disks and terminate the VM.
- C. Apply VM CPU utilization label and include it in the BigQuery billing export.
- D. Use Google BigQuery billing export and labels to associate cost to groups.
- E. Store all state into local SSD, snapshot the persistent disks, and terminate the VM.
- F. Store all state in Google Cloud Storage, snapshot the persistent disks, and terminate the VM.

Answer: AD

Explanation:

<https://cloud.google.com/billing/docs/how-to/export-data-bigquery>

NEW QUESTION 106

- (Topic 5)

You need to design a solution for global load balancing based on the URL path being requested. You need to ensure operations reliability and end-to-end in-transit encryption based on Google best practices. What should you do?

- A. Create a cross-region load balancer with URL Maps.
- B. Create an HTTPS load balancer with URL maps.
- C. Create appropriate instance groups and instance
- D. Configure SSL proxy load balancing.
- E. Create a global forwarding rule
- F. Configure SSL proxy balancing.

Answer: B

Explanation:

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

NEW QUESTION 111

- (Topic 5)

You need to upload files from your on-premises environment to Cloud Storage. You want the files to be encrypted on Cloud Storage using customer-supplied encryption keys. What should you do?

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

Answer: A

Explanation:

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

NEW QUESTION 116

- (Topic 5)

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

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

Answer: AC

Explanation:

Reference: <https://cloud.google.com/storage-options/>

References: <https://cloud.google.com/storage-options/>

NEW QUESTION 118

- (Topic 5)

All compute Engine instances in your VPC should be able to connect to an Active Directory server on specific ports. Any other traffic emerging from your instances is not allowed. You want to enforce this using VPC firewall rules.

How should you configure the firewall rules?

- A. Create an egress rule with priority 1000 to deny all traffic for all instances
- B. Create another egress rule with priority 100 to allow the Active Directory traffic for all instances.
- C. Create an egress rule with priority 100 to deny all traffic for all instances
- D. Create another egress rule with priority 1000 to allow the Active Directory traffic for all instances.
- E. Create an egress rule with priority 1000 to allow the Active Directory traffic
- F. Rely on the implied deny egress rule with priority 100 to block all traffic for all instances.
- G. Create an egress rule with priority 100 to allow the Active Directory traffic
- H. Rely on the implied deny egress rule with priority 1000 to block all traffic for all instances.

Answer: B

Explanation:

<https://cloud.google.com/vpc/docs/firewalls>

NEW QUESTION 120

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 124

- (Topic 5)

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

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

Answer: AC

NEW QUESTION 126

- (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 127

- (Topic 5)

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

- A. Create a Dataflow pipeline to retrieve the data from the external source
- B. As part of the pipeline use the Cloud Data Loss Prevention (Cloud DLP) API to remove any PII data Store the result in BigQuery
- C. Create a Dataflow pipeline to retrieve the data from the external source
- D. As part of the pipeline store all non-PII data in BigQuery and store all PII data in a Cloud Storage bucket that has a retention policy set.
- E. Ask the external partners to upload an data on Cloud Storage Configure Bucket Lock for the bucket Create a Dataflow pipeline to read the data from the bucket As part of the pipeline, use the Cloud Data Loss Prevention (Cloud DIP) API to remove any PII data Store the result in BigQuery
- F. Ask the external partners to import ail data in your BigQuery dataset Create a dataflow pipeline to copy the data into a new table As part of the Dataflow bucket skip all data in columns that have PII data

Answer: A

Explanation:

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

NEW QUESTION 130

- (Topic 5)

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

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

Answer: C

Explanation:

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

NEW QUESTION 135

- (Topic 5)

You want to allow your operations learn to store togs from all the production protects in your Organization, without during logs from other projects All of the production projects are contained in a folder. You want to ensure that all logs for existing and new production projects are captured automatically. What should you do?

- A. Create an aggregated export on the Production folde

- B. Set the log sink to be a Cloud Storage bucket in an operations project
- C. Create an aggregated export on the Organization resource
- D. Set the log sink to be a Cloud Storage bucket in an operations project.
- E. Create log exports in the production project
- F. Set the log sinks to be a Cloud Storage bucket in an operations project.
- G. Create log exports in the production project
- H. Set the log sinks to be BigQuery datasets in the production projects and grant IAM access to the operations team to run queries on the datasets

Answer: A

Explanation:

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

NEW QUESTION 139

- (Topic 5)

Your team will start developing a new application using microservices architecture on Kubernetes Engine. As part of the development lifecycle, any code change that has been pushed to the remote develop branch on your GitHub repository should be built and tested automatically. When the build and test are successful, the relevant microservice will be deployed automatically in the development environment. You want to ensure that all code deployed in the development environment follows this process. What should you do?

- A. Have each developer install a pre-commit hook on their workstation that tests the code and builds the container when committing on the development branch
- B. After a successful commit, have the developer deploy the newly built container image on the development cluster.
- C. Install a post-commit hook on the remote git repository that tests the code and builds the container when code is pushed to the development branch
- D. After a successful commit, have the developer deploy the newly built container image on the development cluster.
- E. Create a Cloud Build trigger based on the development branch that tests the code, builds the container, and stores it in Container Registry
- F. Create a deployment pipeline that watches for new images and deploys the new image on the development cluster
- G. Ensure only the deployment tool has access to deploy new versions.
- H. Create a Cloud Build trigger based on the development branch to build a new container image and store it in Container Registry
- I. Rely on Vulnerability Scanning to ensure the code tests successfully
- J. As the final step of the Cloud Build process, deploy the new container image on the development cluster
- K. Ensure only Cloud Build has access to deploy new versions.

Answer: C

Explanation:

<https://cloud.google.com/container-registry/docs/overview>

Create a Cloud Build trigger based on the development branch that tests the code, builds the container, and stores it in Container Registry. Create a deployment pipeline that watches for new images and deploys the new image on the development cluster. Ensure only the deployment tool has access to deploy new versions.

NEW QUESTION 142

- (Topic 5)

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

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

Answer: D

Explanation:

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

? <https://cloud.google.com/kubernetes-engine/docs/how-to/updating-apps>

? <https://cloud.google.com/blog/products/containers-kubernetes/kubernetes-best-practices-setting-up-health-checks-with-readiness-and-liveness-probes>

NEW QUESTION 146

- (Topic 5)

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

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

Answer: A

NEW QUESTION 149

- (Topic 5)

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

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

Answer: B

Explanation:

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

NEW QUESTION 152

- (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 156

- (Topic 5)

Your company is running a stateless application on a Compute Engine instance. The application is used heavily during regular business hours and lightly outside of business hours. Users are reporting that the application is slow during peak hours. You need to optimize the application's performance. What should you do?

- A. Create a snapshot of the existing disk
- B. Create an instance template from the snapshot. Create an autoscaled managed instance group from the instance template.
- C. Create a snapshot of the existing disk
- D. Create a custom image from the snapshot
- E. Create an autoscaled managed instance group from the custom image.
- F. Create a custom image from the existing disk
- G. Create an instance template from the custom image
- H. Create an autoscaled managed instance group from the instance template.
- I. Create an instance template from the existing disk
- J. Create a custom image from the instance template. Create an autoscaled managed instance group from the custom image.

Answer: B

Explanation:

<https://cloud.google.com/compute/docs/instance-templates/create-instance-templates>

NEW QUESTION 161

- (Topic 5)

You are working at a financial institution that stores mortgage loan approval documents on Cloud Storage. Any change to these approval documents must be uploaded as a separate approval file, so you want to ensure that these documents cannot be deleted or overwritten for the next 5 years. What should you do?

- A. Create a retention policy on the bucket for the duration of 5 years
- B. Create a lock on the retention policy.
- C. Create the bucket with uniform bucket-level access, and grant a service account the role of Object Write
- D. Use the service account to upload new files.
- E. Use a customer-managed key for the encryption of the bucket
- F. Rotate the key after 5 years.
- G. Create the bucket with fine-grained access control, and grant a service account the role of Object Write
- H. Use the service account to upload new files.

Answer: A

Explanation:

Reference: <https://cloud.google.com/storage/docs/using-bucket-lock>

NEW QUESTION 162

- (Topic 5)

As part of implementing their disaster recovery plan, your company is trying to replicate their production

MySQL database from their private data center to their GCP project using a Google Cloud VPN connection. They are experiencing latency issues and a small amount of packet loss that is disrupting the replication. What should they do?

- A. Configure their replication to use UDP.
- B. Configure a Google Cloud Dedicated Interconnect.
- C. Restore their database daily using Google Cloud SQL.
- D. Add additional VPN connections and load balance them.
- E. Send the replicated transaction to Google Cloud Pub/Sub.

Answer: B

NEW QUESTION 167

- (Topic 5)

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

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

Answer: B

Explanation:

Relational databases were not designed to cope with the scale and agility challenges that face modern applications, nor were they built to take advantage of the commodity storage and processing power available today.

NoSQL fits well for:

– Developers are working with applications that create massive volumes of new, rapidly changing data types — structured, semi-structured, unstructured and polymorphic data.

NEW QUESTION 172

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 177

- (Topic 5)

Your company has a stateless web API that performs scientific calculations. The web API runs on a single Google Kubernetes Engine (GKE) cluster. The cluster is currently deployed in us-central1. Your company has expanded to offer your API to customers in Asia. You want to reduce the latency for the users in Asia. What should you do?

- A. Use a global HTTP(s) load balancer with Cloud CDN enabled
- B. Create a second GKE cluster in asia-southeast1, and expose both API's using a Service of type Load Balance
- C. Add the public IPs to the Cloud DNS zone
- D. Increase the memory and CPU allocated to the application in the cluster
- E. Create a second GKE cluster in asia-southeast1, and use kubemci to create a global HTTP(s) load balancer

Answer: D

Explanation:

https://cloud.google.com/kubernetes-engine/docs/concepts/multi-cluster-ingress#how_works

<https://github.com/GoogleCloudPlatform/k8s-multicloud-ingress> <https://cloud.google.com/blog/products/gcp/how-to-deploy-geographically-distributed-services-on-kubernetes-engine-with-kubemci>

NEW QUESTION 179

- (Topic 5)

You are designing an application for use only during business hours. For the minimum viable product release, you'd like to use a managed product that automatically "scales to zero" so you don't incur costs when there is no activity. Which primary compute resource should you choose?

- A. Cloud Functions
- B. Compute Engine
- C. Kubernetes Engine
- D. AppEngine flexible environment

Answer: A

Explanation:

<https://cloud.google.com/serverless-options>

NEW QUESTION 182

- (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 187

- (Topic 5)

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

Which of the following best reflects your recommendations for a cost-effective storage allocation?

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

Answer: D

Explanation:

<https://cloud.google.com/compute/docs/disks>

NEW QUESTION 189

- (Topic 5)

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

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

app = Flask(__name__)
sessions = {}

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

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

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

    return news.render(news_articles)

if __name__ == "__main__":
    app.run()
```

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

Answer: A

Explanation:

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

NEW QUESTION 191

- (Topic 5)

An application development team believes their current logging tool will not meet their needs for their new cloud-based product. They want a better tool to capture errors and help them analyze their historical log data. You want to help them find a solution that meets their needs, what should you do?

- A. Direct them to download and install the Google StackDriver logging agent.
- B. Send them a list of online resources about logging best practices.
- C. Help them define their requirements and assess viable logging tools.
- D. Help them upgrade their current tool to take advantage of any new features.

Answer: C

Explanation:

Help them define their requirements and assess viable logging tools. They know the requirements and the existing tools' problems. While it's true StackDriver Logging and Error Reporting possibly meet all their requirements, there might be other tools also meet their need. They need you to provide expertise to make assessment for new tools, specifically, logging tools that can "capture errors and help them analyze their historical log data".

References: <https://cloud.google.com/logging/docs/agent/installation>

NEW QUESTION 195

- (Topic 5)

Your company wants to start using Google Cloud resources but wants to retain their on- premises Active Directory domain controller for identity management. What should you do?

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

Answer: B

Explanation:

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

NEW QUESTION 200

- (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 203

- (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 create 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 are not enabled by default You have default service account but don't have the permission (scope) you can stop the instance, edit, change scope and restart it to enable the scope access. Of course, if you Run your script on a new virtual machine with the BigQuery access scope enabled, it also works

<https://cloud.google.com/compute/docs/access/service-accounts>

NEW QUESTION 208

- (Topic 5)

One of your primary business objectives is being able to trust the data stored in your application. You want to log all changes to the application data. How can you design your logging system to verify authenticity of your logs?

- A. Write the log concurrently in the cloud and on premises.
- B. Use a SQL database and limit who can modify the log table.
- C. Digitally sign each timestamp and log entry and store the signature.
- D. Create a JSON dump of each log entry and store it in Google Cloud Storage.

Answer: C

Explanation:

<https://cloud.google.com/storage/docs/access-logs>

References: <https://cloud.google.com/logging/docs/reference/tools/gcloud-logging>

NEW QUESTION 213

- (Topic 5)

You are designing a Data Warehouse on Google Cloud and want to store sensitive data in BigQuery. Your company requires you to generate encryption keys outside of Google Cloud. You need to implement a solution. What should you do?

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

Answer: D

Explanation:

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

NEW QUESTION 214

- (Topic 5)

A lead engineer wrote a custom tool that deploys virtual machines in the legacy data center. He wants to migrate the custom tool to the new cloud environment. You want to advocate for the adoption of Google Cloud Deployment Manager. What are two business risks of migrating to Cloud Deployment Manager? Choose 2 answers

- A. Cloud Deployment Manager uses Python.
- B. Cloud Deployment Manager APIs could be deprecated in the future.
- C. Cloud Deployment Manager is unfamiliar to the company's engineers.
- D. Cloud Deployment Manager requires a Google APIs service account to run.
- E. Cloud Deployment Manager can be used to permanently delete cloud resources.
- F. Cloud Deployment Manager only supports automation of Google Cloud resources.

Answer: CF

Explanation:

<https://cloud.google.com/deployment-manager/docs/deployments/deleting-deployments>

NEW QUESTION 219

- (Topic 5)

Your company has multiple on-premises systems that serve as sources for reporting. The data has not been maintained well and has become degraded over time. You want to use Google-recommended practices to detect anomalies in your company data. What should you do?

- A. Upload your files into Cloud Storage
- B. Use Cloud Datalab to explore and clean your data.
- C. Upload your files into Cloud Storage
- D. Use Cloud Dataprep to explore and clean your data.
- E. Connect Cloud Datalab to your on-premises system
- F. Use Cloud Datalab to explore and clean your data.
- G. Connect Cloud Dataprep to your on-premises system
- H. Use Cloud Dataprep to explore and clean your data.

Answer: B

Explanation:

<https://cloud.google.com/dataprep/>

NEW QUESTION 220

- (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 224

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 227

- (Topic 5)

Your web application uses Google Kubernetes Engine to manage several workloads. One workload requires a consistent set of hostnames even after pod scaling and relaunches.

Which feature of Kubernetes should you use to accomplish this?

- A. StatefulSets
- B. Role-based access control
- C. Container environment variables
- D. Persistent Volumes

Answer: A

Explanation:

<https://kubernetes.io/docs/tutorials/stateful-application/basic-stateful-set/>

NEW QUESTION 231

- (Topic 5)

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

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

Answer: B

NEW QUESTION 232

- (Topic 5)

Your company has an application running on App Engine that allows users to upload music files and share them with other people. You want to allow users to upload files directly into Cloud Storage from their browser session. The payload should not be passed through the backend. What should you do?

- A.
 - * 1. Set a CORS configuration in the target Cloud Storage bucket where the base URL of the App Engine application is an allowed origin.
 - * 2. Use the Cloud Storage Signed URL feature to generate a POST URL.
- B.
 - * 1. Set a CORS configuration in the target Cloud Storage bucket where the base URL of the App Engine application is an allowed origin.
 - * 2. Assign the Cloud Storage WRITER role to users who upload files.
- C.
 - * 1. Use the Cloud Storage Signed URL feature to generate a POST URL.
 - * 2. Use App Engine default credentials to sign requests against Cloud Storage.
- D.
 - * 1. Assign the Cloud Storage WRITER role to users who upload files.
 - * 2. Use App Engine default credentials to sign requests against Cloud Storage.

A.

Answer: B

NEW QUESTION 236

- (Topic 5)

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

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

Answer: B

Explanation:

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

NEW QUESTION 238

- (Topic 5)

Your company is forecasting a sharp increase in the number and size of Apache Spark and Hadoop jobs being run on your local datacenter. You want to utilize the cloud to help you scale this upcoming demand with the least amount of operations work and code change. Which product should you use?

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

Answer: B

Explanation:

Google Cloud Dataproc is a fast, easy-to-use, low-cost and fully managed service that lets you run the Apache Spark and Apache Hadoop ecosystem on Google Cloud Platform. Cloud Dataproc provisions big or small clusters rapidly, supports many popular job types, and is integrated with other Google Cloud Platform services, such as Google Cloud Storage and Stackdriver Logging, thus helping you reduce TCO.

References: <https://cloud.google.com/dataproc/docs/resources/faq>

NEW QUESTION 239

- (Topic 5)

The operations manager asks you for a list of recommended practices that she should consider when migrating a J2EE application to the cloud. Which three practices should you recommend? Choose 3 answers

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

Answer: AEF

Explanation:

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

NEW QUESTION 242

- (Topic 5)

You have found an error in your App Engine application caused by missing Cloud Datastore indexes. You have created a YAML file with the required indexes and want to deploy these new indexes to Cloud Datastore.

What should you do?

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

Answer: A

NEW QUESTION 247

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 251

- (Topic 5)

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

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

Answer: C

NEW QUESTION 254

- (Topic 5)

You are working at an institution that processes medical data. You are migrating several workloads onto Google Cloud. Company policies require all workloads to run on physically separated hardware, and workloads from different clients must also be separated You created a sole-tenant node group and added a node for each client. You need to deploy the workloads on these dedicated hosts. What should you do?

- A. Add the node group name as a network tag when creating Compute Engine instances in order to host each workload on the correct node group.
- B. Add the node name as a network tag when creating Compute Engine instances in order to host each workload on the correct node.
- C. Use node affinity labels based on the node group name when creating Compute Engine instances in order to host each workload on the correct node group
- D. Use node affinity labels based on the node name when creating Compute Engine instances in order to host each workload on the correct node.

Answer: C

Explanation:

https://cloud.google.com/compute/docs/nodes/provisioning-sole-tenant-vms#provision_a_sole-tenant_vm

https://cloud.google.com/compute/docs/nodes/provisioning-sole-tenant-vms#gcloud_2 When you create a VM, you request sole-tenancy by specifying node affinity or anti-affinity, referencing one or more node affinity labels. You specify custom node affinity labels when you create a node template, and Compute Engine automatically includes some default affinity labels on each node. By specifying affinity when you create a VM, you can schedule VMs together on a specific node or nodes in a node group. By specifying anti-affinity when you create a VM, you can ensure that certain VMs are not scheduled together on the same node or nodes in a node group.

NEW QUESTION 258

- (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 dat
- 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 dat
- D. Whenever a zonal outage occurs, use the instance template to spin up the application in another zone in the same regio
- E. Use the regional persistent disk for the application data.
- F. Create a snapshot schedule for the disk containing the application dat
- 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 dat
- I. Whenever a zonal outage occurs, use the instance template to spin up the application in another regio
- 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 259

- (Topic 5)

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

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

Answer: B

Explanation:

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

NEW QUESTION 263

- (Topic 5)

You want to establish a Compute Engine application in a single VPC across two regions. The application must communicate over VPN to an on-premises network. How should you deploy the VPN?

- A. Use VPC Network Peering between the VPC and the on-premises network.
- B. Expose the VPC to the on-premises network using IAM and VPC Sharing.
- C. Create a global Cloud VPN Gateway with VPN tunnels from each region to the on- premises peer gateway.
- D. Deploy Cloud VPN Gateway in each regio
- E. Ensure that each region has at least one VPN tunnel to the on-premises peer gateway.

Answer: C

Explanation:

<https://cloud.google.com/vpn/docs/how-to/creating-static-vpns>

NEW QUESTION 267

- (Topic 5)

Your customer wants to capture multiple GBs of aggregate real-time key performance indicators (KPIs) from their game servers running on Google Cloud Platform and monitor the KPIs with low latency. How should they capture the KPIs?

- A. Store time-series data from the game servers in Google Bigtable, and view it using Google Data Studio.
- B. Output custom metrics to Stackdriver from the game servers, and create a Dashboard in StackdriverMonitoring Console to view them.

- C. Schedule BigQuery load jobs to ingest analytics files uploaded to Cloud Storage every ten minutes, and visualize the results in Google Data Studio.
D. Insert the KPIs into Cloud Datastore entities, and run ad hoc analysis and visualizations of them in Cloud Datalab.

Answer: A

Explanation:

<https://cloud.google.com/monitoring/api/v3/metrics-details#metric-kinds>

NEW QUESTION 268

- (Topic 5)

Your architecture calls for the centralized collection of all admin activity and VM system logs within your project.

How should you collect these logs from both VMs and services?

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

Answer: B

Explanation:

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

NEW QUESTION 271

- (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 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 Kubernetes Engine cluster with the following command: `gcloud alpha container clusters update mycluster - --enable-autoscaling - --min-nodes=1 - --max-nodes=10`
D. Create a new Kubernetes Engine cluster with the following command: `gcloud alpha container clusters create 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 276

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 281

- (Topic 5)

You want your Google Kubernetes Engine cluster to automatically add or remove nodes based on CPU load. What should you do?

- A. Configure a HorizontalPodAutoscaler with a target CPU usage
B. Enable the Cluster Autoscaler from the GCP Console.
C. Configure a HorizontalPodAutoscaler with a target CPU usage
D. Enable autoscaling on the managed instance group for the cluster using the gcloud command.
E. Create a deployment and set the maxUnavailable and maxSurge properties
F. Enable the Cluster Autoscaler using the gcloud command.
G. Create a deployment and set the maxUnavailable and maxSurge properties
H. Enable autoscaling on the cluster managed instance group from the GCP Console.

Answer: B

NEW QUESTION 282

- (Topic 5)

You are using Cloud SQL as the database backend for a large CRM deployment. You want to scale as usage increases and ensure that you don't run out of

storage, maintain 75% CPU usage cores, and keep replication lag below 60 seconds. What are the correct steps to meet your requirements?

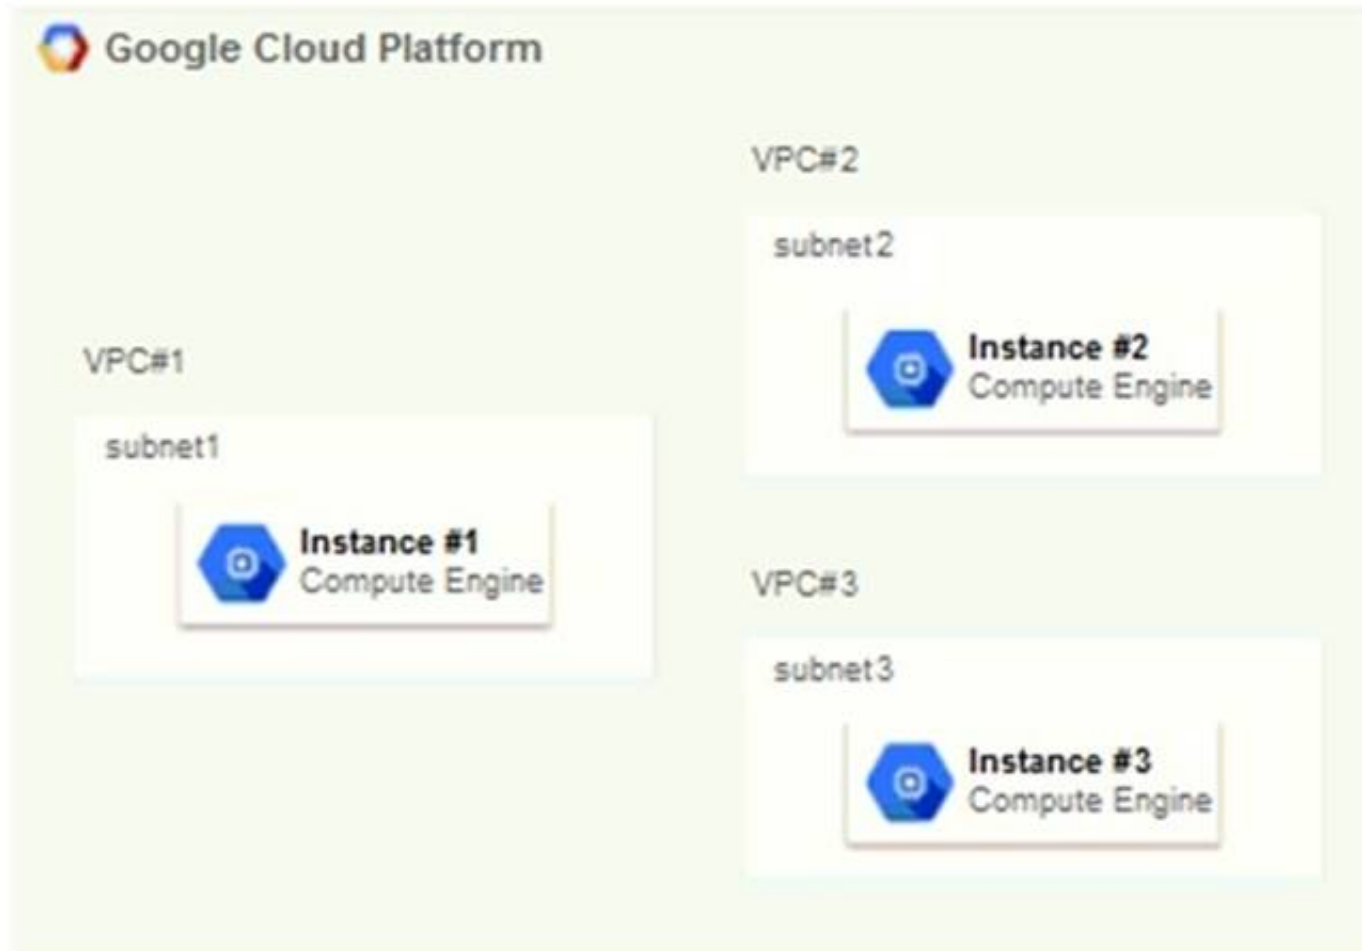
- A. 1) Enable automatic storage increase for the instance.2) Create a Stackdriver alert when CPU usage exceeds 75%, and change the instance type to reduce CPU usage.3) Create a Stackdriver alert for replication lag, and shard the database to reduce replication time.
 B. 1) Enable automatic storage increase for the instance.2) Change the instance type to a 32-core machine type to keep CPU usage below 75%.3) Create a Stackdriver alert for replication lag, and shard the database to reduce replication time.
 C. 1) Create a Stackdriver alert when storage exceeds 75%, and increase the available storage on the instance to create more space.2) Deploy memcached to reduce CPU load.3) Change the instance type to a 32-core machine type to reduce replication lag.
 D. 1) Create a Stackdriver alert when storage exceeds 75%, and increase the available storage on the instance to create more space.2) Deploy memcached to reduce CPU load.3) Create a Stackdriver alert for replication lag, and change the instance type to a 32-core machine type to reduce replication lag.

Answer: A

NEW QUESTION 286

- (Topic 5)

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



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

- A. Create a cloud router to advertise subnet #2 and subnet #3 to subnet #1.
 B. Add two additional NICs to Instance #1 with the following configuration:
 •NIC1 VPC: VPC #2 SUBNETWORK: subnet #2
 •NIC2 VPC: VPC #3 SUBNETWORK: subnet #3
 Update firewall rules to enable traffic between instances.
 C. Create two VPN tunnels via CloudVPN:
 •1 between VPC #1 and VPC #2.
 •1 between VPC #2 and VPC #3.
 Update firewall rules to enable traffic between the instances.
 D. Peer all three VPCs:
 •Peer VPC #1 with VPC #2.
 •Peer VPC #2 with VPC #3.
 Update firewall rules to enable traffic between the instances.

Answer: B

Explanation:

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

NEW QUESTION 291

- (Topic 5)

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

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

Answer: A

Explanation:

On Linux instances, connect to your instance and manually resize your partitions and file systems to use the additional disk space that you added. Extend the file system on the disk or the partition to use the added space. If you grew a partition on your disk, specify the partition. If your disk does not have a

partition table, specify only the disk ID.

```
sudo resize2fs /dev/[DISK_ID][PARTITION_NUMBER]
```

where [DISK_ID] is the device name and [PARTITION_NUMBER] is the partition number for the device where you are resizing the file system.

References: <https://cloud.google.com/compute/docs/disks/add-persistent-disk>

NEW QUESTION 296

- (Topic 5)

You are designing a large distributed application with 30 microservices. Each of your distributed microservices needs to connect to a database back-end. You want to store the credentials securely. Where should you store the credentials?

- A. In the source code
- B. In an environment variable
- C. In a secret management system
- D. In a config file that has restricted access through ACLs

Answer: C

Explanation:

https://cloud.google.com/docs/authentication/production#providing_credentials_to_your_application

NEW QUESTION 300

- (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 305

- (Topic 5)

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

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

Answer: C

Explanation:

<https://cloud.google.com/apis/design/versioning>

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

NEW QUESTION 310

- (Topic 5)

You have an application that makes HTTP requests to Cloud Storage. Occasionally the requests fail with HTTP status codes of 5xx and 429.

How should you handle these types of errors?

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

Answer: A

Explanation:

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

NEW QUESTION 311

- (Topic 5)

Your company's user-feedback portal comprises a standard LAMP stack replicated across two zones. It is deployed in the us-central1 region and uses autoscaled managed instance groups on all layers, except the database. Currently, only a small group of select customers have access to the portal. The portal meets a 99.99% availability SLA under these conditions. However, next quarter, your company will be making the portal available to all users, including unauthenticated users. You need to develop a resiliency testing strategy to ensure the system maintains the SLA once they introduce additional user load. What should you do?

- A. Capture existing users input, and replay captured user load until autoscale is triggered on all layer
- B. At the same time, terminate all resources in one of the zones.
- C. Create synthetic random user input, replay synthetic load until autoscale logic is triggered on at least one layer, and introduce "chaos" to the system by terminating random resources on both zones.
- D. Expose the new system to a larger group of users, and increase group ' size each day until autoscale logic is triggered on all layer
- E. At the same time, terminate random resources on both zones.
- F. Capture existing users input, and replay captured user load until resource utilization crosses 80%. Also, derive estimated number of users based on existing users usage of the app, and deploy enough resources to handle 200% of expected load.

Answer: A

NEW QUESTION 314

- (Topic 5)

You have deployed several instances on Compute Engine. As a security requirement, instances cannot have a public IP address. There is no VPN connection between Google

Cloud and your office, and you need to connect via SSH into a specific machine without violating the security requirements. What should you do?

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

Answer: C

Explanation:

https://cloud.google.com/iap/docs/using-tcp-forwarding#tunneling_with_ssh

Leveraging the BeyondCorp security model. "This January, we enhanced context-aware access capabilities in Cloud Identity-Aware Proxy (IAP) to help you protect SSH and RDP access to your virtual machines (VMs)—without needing to provide your VMs with public IP addresses, and without having to set up bastion hosts. "

<https://cloud.google.com/blog/products/identity-security/cloud-iap-enables-context-aware-access-to-vm-via-ssh-and-rdp-without-bastion-hosts>

Reference: <https://cloud.google.com/solutions/connecting-securely>

NEW QUESTION 315

- (Topic 5)

The database administration team has asked you to help them improve the performance of their new database server running on Google Compute Engine. The database is for importing and normalizing their performance statistics and is built with MySQL running on Debian Linux. They have an n1-standard-8 virtual machine with 80 GB of SSD persistent disk. What should they change to get better performance from this system?

- A. Increase the virtual machine's memory to 64 GB.
- B. Create a new virtual machine running PostgreSQL.
- C. Dynamically resize the SSD persistent disk to 500 GB.
- D. Migrate their performance metrics warehouse to BigQuery.
- E. Modify all of their batch jobs to use bulk inserts into the database.

Answer: C

NEW QUESTION 317

- (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 321

- (Topic 5)

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

- A. Compute Engine with containers

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

Answer: BC

Explanation:

B: With Container Engine, Google will automatically deploy your cluster for you, update, patch, secure the nodes.

Kubernetes Engine's cluster autoscaler automatically resizes clusters based on the demands of the workloads you want to run.

C: Solutions like Datastore, BigQuery, AppEngine, etc are truly NoOps.

App Engine by default scales the number of instances running up and down to match the load, thus providing consistent performance for your app at all times while minimizing idle instances and thus reducing cost.

Note: At a high level, NoOps means that there is no infrastructure to build out and manage during usage of the platform. Typically, the compromise you make with NoOps is that you lose control of the underlying infrastructure.

References: <https://www.quora.com/How-well-does-Google-Container-Engine-support-Google-Cloud-Platform%E2%80%99s-NoOps-claim>

NEW QUESTION 326

- (Topic 5)

Your company's test suite is a custom C++ application that runs tests throughout each day on Linux virtual machines. The full test suite takes several hours to complete, running on a limited number of on premises servers reserved for testing. Your company wants to move the testing infrastructure to the cloud, to reduce the amount of time it takes to fully test a change to the system, while changing the tests as little as possible. Which cloud infrastructure should you recommend?

- A. Google Compute Engine unmanaged instance groups and Network Load Balancer
- B. Google Compute Engine managed instance groups with auto-scaling
- C. Google Cloud Dataproc to run Apache Hadoop jobs to process each test
- D. Google App Engine with Google Stackdriver for logging

Answer: B

Explanation:

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

Google Compute Engine enables users to launch virtual machines (VMs) on demand. VMs can be launched from the standard images or custom images created by users.

Managed instance groups offer autoscaling capabilities that allow you to automatically add or remove instances from a managed instance group based on increases or decreases in load. Autoscaling helps your applications gracefully handle increases in traffic and reduces cost when the need for resources is lower.

NEW QUESTION 328

- (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 329

- (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 332

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 336

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

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

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

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

Money Back Guarantee

Professional-Cloud-Architect Practice Exam Features:

- * Professional-Cloud-Architect Questions and Answers Updated Frequently
- * Professional-Cloud-Architect Practice Questions Verified by Expert Senior Certified Staff
- * Professional-Cloud-Architect Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * Professional-Cloud-Architect Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year