

Google

Exam Questions Professional-Cloud-Architect

Google Certified Professional - Cloud Architect (GCP)



NEW QUESTION 1

- (Topic 1)

For this question, refer to the Mountkirk Games case study

Mountkirk Games needs to create a repeatable and configurable mechanism for deploying isolated application environments. Developers and testers can access each other's environments and resources, but they cannot access staging or production resources. The staging environment needs access to some services from production.

What should you do to isolate development environments from staging and production?

- A. Create a project for development and test and another for staging and production.
- B. Create a network for development and test and another for staging and production.
- C. Create one subnetwork for development and another for staging and production.
- D. Create one project for development, a second for staging and a third for production.

Answer: D

NEW QUESTION 2

- (Topic 1)

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

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

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

Answer: B

Explanation:

503 is service unavailable error. If the database was online everyone would get the 503 error.

https://cloud.google.com/docs/quota#capping_usage

NEW QUESTION 3

- (Topic 1)

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

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

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

Answer: A

Explanation:

From scenario: Requirements for Game Backend Platform

? Dynamically scale up or down based on game activity

? Connect to a managed NoSQL database service

? Run customize Linux distro

NEW QUESTION 4

- (Topic 2)

For this question, refer to the TerramEarth case study.

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 5

- (Topic 2)

For this question, refer to the TerramEarth case study.

TerramEarth plans to connect all 20 million vehicles in the field to the cloud. This increases the volume to 20 million 600 byte records a second for 40 TB an hour. How should you design the data ingestion?

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

Answer: B

Explanation:

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

NEW QUESTION 6

- (Topic 2)

For this question, refer to the TerramEarth case study.

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

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

Answer: D

Explanation:

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

NEW QUESTION 7

- (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 8

- (Topic 2)

For this question, refer to the TerramEarth case study.

The TerramEarth development team wants to create an API to meet the company's business requirements. You want the development team to focus their development effort on business value versus creating a custom framework. Which method should they use?

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

Answer: A

Explanation:

https://cloud.google.com/endpoints/docs/openapi/about-cloud-endpoints?hl=en_US&_ga=2.21787131.-1712523161.1522785064

<https://cloud.google.com/endpoints/docs/openapi/architecture-overview> <https://cloud.google.com/storage/docs/gsutil/commands/test>

Develop, deploy, protect and monitor your APIs with Google Cloud Endpoints. Using an Open API Specification or one of our API frameworks, Cloud Endpoints gives you the tools you need for every phase of API development.

From scenario: Business Requirements

Decrease unplanned vehicle downtime to less than 1 week, without increasing the cost of carrying surplus inventory

Support the dealer network with more data on how their customers use their equipment to better position new products and services

Have the ability to partner with different companies – especially with seed and fertilizer suppliers in the fast-growing agricultural business – to create compelling

joint offerings for their customers.

Reference: <https://cloud.google.com/certification/guides/cloud-architect/casestudy-terramearth>

NEW QUESTION 9

- (Topic 3)

For this question, refer to the JencoMart case study.

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

Which metrics should you track?

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


Answer: D

NEW QUESTION 10

- (Topic 3)

For this question, refer to the JencoMart case study.

The migration of JencoMart's application to Google Cloud Platform (GCP) is progressing too slowly. The infrastructure is shown in the diagram. You want to maximize throughput. What are three potential bottlenecks? (Choose 3 answers.)



- A. A single VPN tunnel, which limits throughput
- B. A tier of Google Cloud Storage that is not suited for this task
- C. A copy command that is not suited to operate over long distances
- D. Fewer virtual machines (VMs) in GCP than on-premises machines
- E. A separate storage layer outside the VMs, which is not suited for this task
- F. Complicated internet connectivity between the on-premises infrastructure and GCP

Answer: ADF

NEW QUESTION 10

- (Topic 3)

For this question, refer to the JencoMart case study

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

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

Answer: CDF

Explanation:

<https://cloud.google.com/compute/docs/troubleshooting/troubleshooting-ssh>

D: Handling "Unable to connect on port 22" error message Possible causes include:

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

NEW QUESTION 15

- (Topic 4)

For this question, refer to the Dress4Win case study.

Dress4Win would like to become familiar with deploying applications to the cloud by successfully deploying some applications quickly, as is. They have asked for your recommendation. What should you advise?

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

Answer: A

Explanation:

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

NEW QUESTION 19

- (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 24

- (Topic 4)

For this question, refer to the Dress4Win case study.

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

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

Answer: B

NEW QUESTION 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 34

- (Topic 5)

You want to store critical business information in Cloud Storage buckets. The information is regularly changed but previous versions need to be referenced on a regular basis. You want to ensure that there is a record of all changes to any information in these buckets. You want to ensure that accidental edits or deletions can be easily rolled back. Which feature should you enable?

- A. Bucket Lock
- B. Object Versioning
- C. Object change notification
- D. Object Lifecycle Management

Answer: B

NEW QUESTION 39

- (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 protects from the acquired company into your company's Google Cloud organization. Re-launch 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 40

- (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 41

- (Topic 5)

A development manager is building a new application. He asks you to review his requirements and identify what cloud technologies he can use to meet them. The application must

- * 1. Be based on open-source technology for cloud portability
- * 2. Dynamically scale compute capacity based on demand
- * 3. Support continuous software delivery
- * 4. Run multiple segregated copies of the same application stack
- * 5. Deploy application bundles using dynamic templates
- * 6. Route network traffic to specific services based on URL

Which combination of technologies will meet all of his requirements?

- A. Google Container Engine, Jenkins, and Helm
- B. Google Container Engine and Cloud Load Balancing
- C. Google Compute Engine and Cloud Deployment Manager
- D. Google Compute Engine, Jenkins, and Cloud Load Balancing

Answer: A

Explanation:

Helm for managing Kubernetes

Kubernetes can base on the URL to route traffic to different location (path)

<https://cloud.google.com/kubernetes-engine/docs/tutorials/http-balancer> eg. apiVersion: networking.k8s.io/v1beta1

kind: Ingress metadata:

name: fanout-ingress spec:

rules:

- http: paths:

- path: /* backend:

serviceName: web servicePort: 8080

- path: /v2/* backend: serviceName: web2 servicePort: 8080

NEW QUESTION 45

- (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 47

- (Topic 5)

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

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

Answer: D

Explanation:

In order to maintain a high request rate, avoid using sequential names. Using completely random object names will give you the best load distribution. Randomness after

a common prefix is effective under the prefix <https://cloud.google.com/storage/docs/request-rate>

NEW QUESTION 48

- (Topic 5)

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

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

Answer: C

Explanation:

From: <https://cloud.google.com/architecture/best-practices-for-building-containers>

Under: Tagging using the Git commit hash (bottom of page almost)

"In this case, a common way of handling version numbers is to use the Git commit SHA-1 hash (or a short version of it) as the version number. By design, the Git commit hash is immutable and references a specific version of your software.

You can use this commit hash as a version number for your software, but also as a tag for the Docker image built from this specific version of your software. Doing so makes Docker images traceable: because in this case the image tag is immutable, you instantly know which specific version of your software is running inside a given container."

NEW QUESTION 51

- (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 52

- (Topic 5)

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

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

Answer: C

Explanation:

The framework illustrated in the preceding diagram has four phases:

- Assess. In this phase, you assess your source environment, assess the workloads that you want to migrate to Google Cloud, and assess which VMs support each workload.
- Plan. In this phase, you create the basic infrastructure for Migrate for Compute Engine, such as provisioning the resource hierarchy and setting up network access.
- Deploy. In this phase, you migrate the VMs from the source environment to Compute Engine.
- Optimize. In this phase, you begin to take advantage of the cloud technologies and capabilities.

Reference: <https://cloud.google.com/architecture/migrating-vms-migrate-for-compute-engine-getting-started>

NEW QUESTION 57

- (Topic 5)

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

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

Answer: B

Explanation:

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

NEW QUESTION 61

- (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 share
- 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 62

- (Topic 5)

You are creating a solution to remove backup files older than 90 days from your backup Cloud Storage bucket. You want to optimize ongoing Cloud Storage spend. What should you do?

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

Answer: B

Explanation:

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

NEW QUESTION 65

- (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 Storage
- 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 69

- (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 71

- (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 76

- (Topic 5)

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

- A. Use Spinnaker to deploy builds to production using the red/black deployment strategy so that changes can easily be rolled back.
- B. Use Spinnaker to deploy builds to production and run tests on production deployments.
- C. Use Jenkins to build the staging branches and the master branch
- D. Build and deploy changes to production for 10% of users before doing a complete rollout.
- E. Use Jenkins to monitor tags in the repository
- F. Deploy staging tags to a staging environment for testing. After testing, tag the repository for production and deploy that to the production environment.

Answer: D

Explanation:

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

NEW QUESTION 79

- (Topic 5)

You want to optimize the performance of an accurate, real-time, weather-charting application. The data comes from 50,000 sensors sending 10 readings a second, in the format of a timestamp and sensor reading. Where should you store the data?

- A. Google BigQuery
- B. Google Cloud SQL
- C. Google Cloud Bigtable
- D. Google Cloud Storage

Answer: C

Explanation:

It is time-series data, So Big Table. <https://cloud.google.com/bigtable/docs/schema-design-time-series>

Google Cloud Bigtable is a scalable, fully-managed NoSQL wide-column database that is suitable for both real-time access and analytics workloads.

Good for:

- ? Low-latency read/write access
- ? High-throughput analytics
- ? Native time series support
- ? Common workloads:
- ? IoT, finance, adtech
- ? Personalization, recommendations
- ? Monitoring
- ? Geospatial datasets
- ? Graphs

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

NEW QUESTION 83

- (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 84

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 86

- (Topic 5)

Your operations team currently stores 10 TB of data in an object storage service from a third-party provider. They want to move this data to a Cloud Storage bucket as quickly as possible, following Google-recommended practices. They want to minimize the cost of this data migration. What approach should they use?

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

Answer: B

Explanation:

<https://cloud.google.com/architecture/migration-to-google-cloud-transferring-your-large-datasets#transfer-options>

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

NEW QUESTION 88

- (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 92

- (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 94

- (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 99

- (Topic 5)

You are deploying an application on App Engine that needs to integrate with an on-premises database. For security purposes, your on-premises database must not be accessible through the public Internet. What should you do?

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

Answer: D

Explanation:

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

NEW QUESTION 102

- (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 106

- (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 110

- (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 confi
- 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 115

- (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 119

- (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 122

- (Topic 5)

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

What should you do?

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

Answer: B

Explanation:

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

Cluster autoscaling

`--enable-autoscaling`

Enables autoscaling for a node pool.

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

Where:

`--max-nodes=MAX_NODES`

Maximum number of nodes in the node pool.

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

NEW QUESTION 123

- (Topic 5)

You are creating an App Engine application that uses Cloud Datastore as its persistence layer. You need to retrieve several root entities for which you have the identifiers. You want to minimize the overhead in operations performed by Cloud Datastore. What should you do?

- A. Create the Key object for each Entity and run a batch get operation
- B. Create the Key object for each Entity and run multiple get operations, one operation for each entity
- C. Use the identifiers to create a query filter and run a batch query operation
- D. Use the identifiers to create a query filter and run multiple query operations, one operation for each entity

Answer: C

Explanation:

<https://cloud.google.com/datastore/docs/concepts/entities#datastore-batch-upsert-nodejs>

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 131

- (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 136

- (Topic 5)

Your company operates nationally and plans to use GCP for multiple batch workloads, including some that are not time-critical. You also need to use GCP services that are HIPAA-certified and manage service costs.

How should you design to meet Google best practices?

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

Answer: B

Explanation:

<https://cloud.google.com/security/compliance/hipaa/>

NEW QUESTION 137

- (Topic 5)

You are deploying an application to Google Cloud. The application is part of a system. The application in Google Cloud must communicate over a private network with applications in a non-Google Cloud environment. The expected average throughput is 200 kbps. The business requires:

- 99.99% system availability
- cost optimization

You need to design the connectivity between the locations to meet the business requirements. What should you provision?

- A. A Classic Cloud VPN gateway connected with one tunnel to an on-premises VPN gateway.
- B. A Classic Cloud VPN gateway connected with two tunnels to an on-premises VPN gateway.
- C. An HA Cloud VPN gateway connected with two tunnels to an on-premises VPN gateway.
- D. Two HA Cloud VPN gateways connected to two on-premises VPN gateway
- E. Configure each HA CloudVPN gateway to have two tunnels, each connected to different on-premises VPN gateways.

Answer: C

Explanation:

https://cloud.google.com/network-connectivity/docs/vpn/concepts/topologies#configurations_that_support_9999_availability

NEW QUESTION 142

- (Topic 5)

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

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

Answer: B

NEW QUESTION 145

- (Topic 5)

Your development team has installed a new Linux kernel module on the batch servers in Google Compute Engine (GCE) virtual machines (VMs) to speed up the nightly batch process. Two days after the installation, 50% of web application deployed in the same nightly batch run. You want to collect details on the failure to pass back to the development team. Which three actions should you take? Choose 3 answers

- A. Use Stackdriver Logging to search for the module log entries.
- B. Read the debug GCE Activity log using the API or Cloud Console.
- C. Use gcloud or Cloud Console to connect to the serial console and observe the logs.
- D. Identify whether a live migration event of the failed server occurred, using in the activity log.
- E. Adjust the Google Stackdriver timeline to match the failure time, and observe the batch server metrics.
- F. Export a debug VM into an image, and run the image on a local server where kernel log messages will be displayed on the native screen.

Answer: ACE

Explanation:

<https://www.flexera.com/blog/cloud/2013/12/google-compute-engine-live-migration-passes-the-test/>
"With live migration, the virtual machines are moved without any downtime or noticeable service degradation"

NEW QUESTION 148

- (Topic 5)

Your company uses Google Kubernetes Engine (GKE) as a platform for all workloads. Your company has a single large GKE cluster that contains batch, stateful, and stateless workloads. The GKE cluster is configured with a single node pool with 200 nodes. Your company needs to reduce the cost of this cluster but does

not want to compromise availability. What should you do?

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

Answer: C

Explanation:

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

NEW QUESTION 152

- (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 156

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 159

- (Topic 5)

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

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

Answer: C

NEW QUESTION 163

- (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 year
- 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 bucke
- 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 166

- (Topic 5)

An application development team has come to you for advice. They are planning to write and deploy an HTTP(S) API using Go 1.12. The API will have a very unpredictable workload and must remain reliable during peaks in traffic. They want to minimize operational overhead for this application. What approach should you recommend?

- A. Use a Managed Instance Group when deploying to Compute Engine
- B. Develop an application with containers, and deploy to Google Kubernetes Engine (GKE)
- C. Develop the application for App Engine standard environment
- D. Develop the application for App Engine Flexible environment using a custom runtime

Answer: C

Explanation:

<https://cloud.google.com/appengine/docs/the-appengine-environments>

NEW QUESTION 167

- (Topic 5)

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

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

Answer: C

NEW QUESTION 170

- (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 171

- (Topic 5)

You set up an autoscaling instance group to serve web traffic for an upcoming launch. After configuring the instance group as a backend service to an HTTP(S) load balancer, you notice that virtual machine (VM) instances are being terminated and re-launched every minute. The instances do not have a public IP address. You have verified the appropriate web response is coming from each instance using the curl command. You want to ensure the backend is configured correctly. What should you do?

- A. Ensure that a firewall rule exists to allow source traffic on HTTP/HTTPS to reach the load balancer.
- B. Assign a public IP to each instance and configure a firewall rule to allow the load balancer to reach the instance public IP.
- C. Ensure that a firewall rule exists to allow load balancer health checks to reach the instances in the instance group.
- D. Create a tag on each instance with the name of the load balancer
- E. Configure a firewall rule with the name of the load balancer as the source and the instance tag as the destination.

Answer: C

Explanation:

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

The best practice when configuring a health check is to check health and serve traffic on the same port. However, it is possible to perform health checks on one port, but serve traffic on another. If you do use two different ports, ensure that firewall rules and services running on instances are configured appropriately. If you run health checks and serve traffic on the same port, but decide to switch ports at some point, be sure to update both the backend service and the health check. Backend services that do not have a valid global forwarding rule referencing it will not be health checked and will have no health status.

References: <https://cloud.google.com/compute/docs/load-balancing/http/backend-service>

NEW QUESTION 174

- (Topic 5)

Your company acquired a healthcare startup and must retain its customers' medical information for up to 4 more years, depending on when it was created. Your

corporate policy is to securely retain this data, and then delete it as soon as regulations allow. Which approach should you take?

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

Answer: C

Explanation:

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

NEW QUESTION 179

- (Topic 5)

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

- A. Set timeouts on your application so that you can fail requests faster.
- B. Send custom metrics for each of your requests to Stackdriver Monitoring.
- C. Use Stackdriver Monitoring to look for insights that show when your API latencies are high.
- D. Instrument your application with Stackdriver Trace in order to break down the request latencies at each microservice.

Answer: D

Explanation:

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

NEW QUESTION 181

- (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 185

- (Topic 5)

You need to migrate Hadoop jobs for your company's Data Science team without modifying the underlying infrastructure. You want to minimize costs and infrastructure management effort. What should you do?

- A. Create a Dataproc cluster using standard worker instances.
- B. Create a Dataproc cluster using preemptible worker instances.
- C. Manually deploy a Hadoop cluster on Compute Engine using standard instances.
- D. Manually deploy a Hadoop cluster on Compute Engine using preemptible instances.

Answer: B

Explanation:

Reference: <https://cloud.google.com/architecture/hadoop/hadoop-gcp-migration-jobs>

NEW QUESTION 187

- (Topic 5)

You need to develop procedures to test a disaster plan for a mission-critical application. You want to use Google-recommended practices and native capabilities within GCP. What should you do?

- A. Use Deployment Manager to automate service provisionin
- B. Use Activity Logs to monitor and debug your tests.
- C. Use Deployment Manager to automate provisionin
- D. Use Stackdriver to monitor and debug your tests.
- E. Use gcloud scripts to automate service provisionin
- F. Use Activity Logs monitor and debug your tests.

- G. Use automated scripts to automate service provisionin
- H. Use Activity Logs monitor and debug your tests.

Answer: B

Explanation:

<https://cloud.google.com/solutions/dr-scenarios-planning-guide>

NEW QUESTION 188

- (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 190

- (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 191

- (Topic 5)

You need to deploy an application on Google Cloud that must run on a Debian Linux environment. The application requires extensive configuration in order to operate correctly. You want to ensure that you can install Debian distribution updates with minimal manual intervention whenever they become available. What should you do?

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

Answer: B

Explanation:

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

NEW QUESTION 196

- (Topic 5)

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

- A. Use the Service Mesh visualization in the Cloud Console to inspect the telemetry between the microservices.

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

Answer: A

Explanation:

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

NEW QUESTION 199

- (Topic 5)

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

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

Answer: B

Explanation:

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

NEW QUESTION 200

- (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 204

- (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 Storang
- B. Use Cloud Datalab to explore and clean your data.
- C. Upload your files into Cloud Storang
- 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 exploreand clean your data.

Answer: B

Explanation:

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

NEW QUESTION 207

- (Topic 5)

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

- A. Inspect the logs and metrics from the instances in Cloud Logging and Cloud Monitoring.
- B. Restore a backup of the application database from a time before the application became slow.

- C. Deploy the applications on a managed instance group with autoscaling enable
- D. Add load balancer in front of the managed instance group, and have the users connect to the IP of the load balancer.
- E. Change the Compute Engine Instances behind the application to a machine type with more CPU and memory.

Answer: A

Explanation:

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

NEW QUESTION 212

- (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 213

- (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 218

- (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 220

- (Topic 5)

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

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

Answer: A

NEW QUESTION 224

- (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 228

- (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 229

- (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 234

- (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 to 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 237

- (Topic 5)

Your company has decided to build a backup replica of their on-premises user authentication PostgreSQL database on Google Cloud Platform. The database is 4 TB, and large updates are frequent. Replication requires private address space communication. Which networking approach should you use?

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

Answer: A

Explanation:

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

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

Benefits:

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

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

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

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

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

NEW QUESTION 238

- (Topic 5)

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

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

Answer: B

Explanation:

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

NEW QUESTION 239

- (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 region
- 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 240

- (Topic 5)

Your company wants to try out the cloud with low risk. They want to archive approximately 100 TB of their log data to the cloud and test the analytics features available to them there, while also retaining that data as a long-term disaster recovery backup. Which two steps should they take? Choose 2 answers

- A. Load logs into Google BigQuery.
- B. Load logs into Google Cloud SQL.
- C. Import logs into Google Stackdriver.
- D. Insert logs into Google Cloud Bigtable.
- E. Upload log files into Google Cloud Storage.

Answer: AE

NEW QUESTION 245

- (Topic 5)

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

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

-max-nodes=10 and redeploy your application

Answer: C

Explanation:

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

NEW QUESTION 250

- (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 253

- (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 254

- (Topic 5)

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



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

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

Answer: B

Explanation:

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

NEW QUESTION 255

- (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 259

- (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 262

- (Topic 5)

Your company has just recently activated Cloud Identity to manage users. The Google Cloud Organization has been configured as `wed`. The security team needs to secure the Organization. They want to prohibit IAM users outside the domain from gaining permissions from now on. What should they do?

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

Answer: A

Explanation:

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

NEW QUESTION 263

- (Topic 5)

During a high traffic portion of the day, one of your relational databases crashes, but the replica is never promoted to a master. You want to avoid this in the future. What should you do?

- A. Use a different database.
- B. Choose larger instances for your database.
- C. Create snapshots of your database more regularly.
- D. Implement routinely scheduled failovers of your databases.

Answer: D

Explanation:

<https://cloud.google.com/solutions/dr-scenarios-planning-guide>

NEW QUESTION 264

- (Topic 5)

Your company is using BigQuery as its enterprise data warehouse. Data is distributed over several Google Cloud projects. All queries on BigQuery need to be billed on a single project. You want to make sure that no query costs are incurred on the projects that contain the data. Users should be able to query the datasets, but not edit them.

How should you configure users' access roles?

- A. Add all users to a group
- B. Grant the group the role of BigQuery user on the billing project and BigQuerydataViewer on the projects that contain the data.
- C. Add all users to a group
- D. Grant the group the roles of BigQuery dataViewer on the billing project andBigQuery user on the projects that contain the data.
- E. Add all users to a group
- F. Grant the group the roles of BigQuery jobUser on the billing project and BigQuery dataViewer on the projects that contain the data.
- G. Add all users to a group
- H. Grant the group the roles of BigQuery dataViewer on the billing project andBigQuery jobUser on the projects that contain the data.

Answer: A

Explanation:

Reference: <https://cloud.google.com/bigquery/docs/running-queries>

NEW QUESTION 265

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 269

- (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 272

- (Topic 5)

You are designing a mobile chat application. You want to ensure people cannot spoof chat messages, by providing a message were sent by a specific user. What should you do

- A. Tag messages client side with the originating user identifier and the destination user.
- B. Encrypt the message client side using block-based encryption with a shared key.
- C. Use public key infrastructure (PKI) to encrypt the message client side using the originating user's privatekey.
- D. Use a trusted certificate authority to enable SSL connectivity between the client application and the server.

Answer: C

NEW QUESTION 277

- (Topic 5)

You want to create a private connection between your instances on Compute Engine and your on-premises data center. You require a connection of at least 20 Gbps. You want to follow Google-recommended practices.

How should you set up the connection?

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

Answer: A

Explanation:

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

NEW QUESTION 282

- (Topic 5)

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

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

Answer: B

Explanation:

Dataflow is for processing both the Batch and Stream.

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

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

NEW QUESTION 283

- (Topic 5)

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

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

Answer: D

Explanation:

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

NEW QUESTION 286

- (Topic 6)

For this question, refer to the Dress4Win case study. Considering the given business requirements, how would you automate the deployment of web and transactional data layers?

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

Answer: A

NEW QUESTION 289

- (Topic 6)

For this question, refer to the Dress4Win case study. Which of the compute services should be migrated as –is and would still be an optimized architecture for performance in the cloud?

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

Answer: C

NEW QUESTION 293

- (Topic 6)

For this question, refer to the Dress4Win case study. You want to ensure that your on-premises architecture meets business requirements before you migrate your solution.

What change in the on-premises architecture should you make?

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

Answer: C

NEW QUESTION 296

- (Topic 7)

You are migrating a Linux-based application from your private data center to Google Cloud. The TerramEarth security team sent you several recent Linux vulnerabilities published by Common Vulnerabilities and Exposures (CVE). You need assistance in understanding how these vulnerabilities could impact your migration. What should you do?

- A. Open a support case regarding the CVE and chat with the support engineer.
- B. Read the CVEs from the Google Cloud Status Dashboard to understand the impact.
- C. Read the CVEs from the Google Cloud Platform Security Bulletins to understand the impact
- D. Post a question regarding the CVE in Stack Overflow to get an explanation
- E. Post a question regarding the CVE in a Google Cloud discussion group to get an explanation

Answer: AC

Explanation:

<https://cloud.google.com/support/bulletins>

NEW QUESTION 301

- (Topic 7)

For this question, refer to the TerramEarth case study. To be compliant with European GDPR regulation, TerramEarth is required to delete data generated from its European customers after a period of 36 months when it contains personal data. In the new architecture, this data will be stored in both Cloud Storage and BigQuery. What should you do?

- A. Create a BigQuery table for the European data, and set the table retention period to 36 month
- B. For Cloud Storage, use gsutil to enable lifecycle management using a DELETE action with an Age condition of 36 months.
- C. Create a BigQuery table for the European data, and set the table retention period to 36 month
- D. For Cloud Storage, use gsutil to create a SetStorageClass to NONE action when with an Age condition of 36 months.
- E. Create a BigQuery time-partitioned table for the European data, and set the partition expiration period to 36 month
- F. For Cloud Storage, use gsutil to enable lifecycle management using a DELETE action with an Age condition of 36 months.
- G. Create a BigQuery time-partitioned table for the European data, and set the partition period to 36 month
- H. For Cloud Storage, use gsutil to create a SetStorageClass to NONE action with an Age condition of 36 months.

Answer: C

Explanation:

<https://cloud.google.com/bigquery/docs/managing-partitioned-tables#partition-expiration>

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

NEW QUESTION 305

- (Topic 8)

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

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

Answer: A

NEW QUESTION 307

- (Topic 8)

For this question, refer to the Mountkirk Games case study. You need to analyze and define the technical architecture for the compute workloads for your company, Mountkirk Games. Considering the Mountkirk Games business and technical requirements, what should you do?

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

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

Answer: D

NEW QUESTION 309

- (Topic 8)

You need to optimize batch file transfers into Cloud Storage for Mountkirk Games' new Google Cloud solution. The batch files contain game statistics that need to be staged in Cloud Storage and be processed by an extract transform load (ETL) tool. What should you do?

- A. Use gsutil to batch move files in sequence.
- B. Use gsutil to batch copy the files in parallel.
- C. Use gsutil to extract the files as the first part of ETL.
- D. Use gsutil to load the files as the last part of ETL.

Answer: B

Explanation:

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

NEW QUESTION 311

- (Topic 8)

For this question, refer to the Mountkirk Games case study. You need to analyze and define the technical architecture for the database workloads for your company, Mountkirk Games. Considering the business and technical requirements, what should you do?

- A. Use Cloud SQL for time series data, and use Cloud Bigtable for historical data queries.
- B. Use Cloud SQL to replace MySQL, and use Cloud Spanner for historical data queries.
- C. Use Cloud Bigtable to replace MySQL, and use BigQuery for historical data queries.
- D. Use Cloud Bigtable for time series data, use Cloud Spanner for transactional data, and use BigQuery for historical data queries.

Answer: D

Explanation:

<https://cloud.google.com/bigtable/docs/schema-design-time-series>

NEW QUESTION 314

- (Topic 8)

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

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

Answer: C

NEW QUESTION 318

- (Topic 8)

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

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

Answer: A

NEW QUESTION 323

- (Topic 9)

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

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

I. Set up Identity and Access Management (IAM) and Confidential Computing to trigger a Cloud Function.

Answer: A

NEW QUESTION 326

- (Topic 9)

For this question, refer to the Helicopter Racing League (HRL) case study. HRL wants better prediction accuracy from their ML prediction models. They want you to use Google's AI Platform so HRL can understand and interpret the predictions. What should you do?

- A. Use Explainable AI.
- B. Use Vision AI.
- C. Use Google Cloud's operations suite.
- D. Use Jupyter Notebooks.

Answer: A

Explanation:

Reference: <https://cloud.google.com/ai-platform/prediction/docs/ai-explanations/preparing-metadata>

NEW QUESTION 327

- (Topic 10)

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

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

Answer: AB

Explanation:

<https://cloud.google.com/security/compliance/hipaa>

NEW QUESTION 331

- (Topic 10)

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

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

Answer: A

Explanation:

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

NEW QUESTION 333

- (Topic 10)

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

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

Answer: D

Explanation:

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

NEW QUESTION 335

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

Professional-Cloud-Architect Practice Exam Features:

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

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