

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 wants you to design their new testing strategy. How should the test coverage differ from their existing backends on the other platforms?

- A. Tests should scale well beyond the prior approaches.
- B. Unit tests are no longer required, only end-to-end tests.
- C. Tests should be applied after the release is in the production environment.
- D. Tests should include directly testing the Google Cloud Platform (GCP) infrastructure.

Answer: A

Explanation:

From Scenario:

A few of their games were more popular than expected, and they had problems scaling their application servers, MySQL databases, and analytics tools.

Requirements for Game Analytics Platform include: Dynamically scale up or down based on game activity

NEW QUESTION 3

- (Topic 1)

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

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

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

Answer: B

Explanation:

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

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

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

Requirements for Game Analytics Platform

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

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

NEW QUESTION 4

- (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 5

- (Topic 2)

For this question refer to the TerramEarth case study.

Which of TerramEarth's legacy enterprise processes will experience significant change as a result of increased Google Cloud Platform adoption.

- A. Opex/capex allocation, LAN changes, capacity planning
- B. Capacity planning, TCO calculations, opex/capex allocation
- C. Capacity planning, utilization measurement, data center expansion
- D. Data Center expansion, TCO calculations, utilization measurement

Answer: B

Explanation:

Capacity planning, TCO calculations, opex/capex allocation From the case study, it can conclude that Management (CXO) all concern rapid provision of resources (infrastructure)

for growing as well as cost management, such as Cost optimization in Infrastructure, trade up front capital expenditures (Capex) for ongoing operating expenditures (Opex), and Total cost of ownership (TCO)

NEW QUESTION 6

- (Topic 2)

For this question, refer to the TerramEarth case study

You analyzed TerramEarth's business requirement to reduce downtime, and found that they can achieve a majority of time saving by reducing customers' wait time for parts You decided to focus on reduction of the 3 weeks aggregate reporting time Which modifications to the company's processes should you recommend?

- A. Migrate from CSV to binary format, migrate from FTP to SFTP transport, and develop machine learning analysis of metrics.
- B. Migrate from FTP to streaming transport, migrate from CSV to binary format, and develop machine learning analysis of metrics.
- C. Increase fleet cellular connectivity to 80%, migrate from FTP to streaming transport, and develop machine learning analysis of metrics.
- D. Migrate from FTP to SFTP transport, develop machine learning analysis of metrics, and increase dealer local inventory by a fixed factor.

Answer: C

Explanation:

The Avro binary format is the preferred format for loading compressed data. Avro data is faster to load because the data can be read in parallel, even when the data blocks are compressed.

Cloud Storage supports streaming transfers with the gsutil tool or boto library, based on HTTP chunked transfer encoding. Streaming data lets you stream data to and from your Cloud Storage account as soon as it becomes available without requiring that the data be first saved to a separate file. Streaming transfers are useful if you have a process that generates data and you do not want to buffer it locally before uploading it, or if you want to send the result from a computational pipeline directly into Cloud Storage.

References: <https://cloud.google.com/storage/docs/streaming> <https://cloud.google.com/bigquery/docs/loading-data>

NEW QUESTION 7

- (Topic 3)

For this question, refer to the JencoMart case study.

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

Which metrics should you track?

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

Answer: D

NEW QUESTION 8

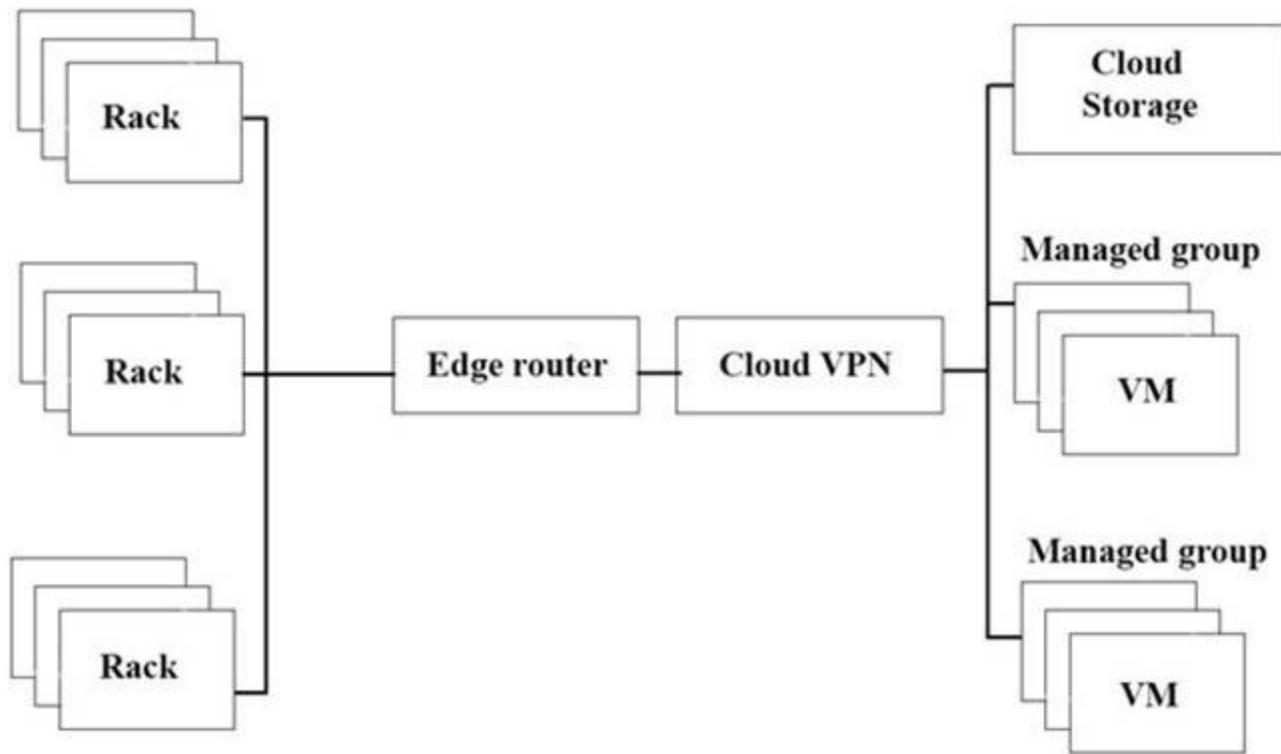
- (Topic 3)

For this question, refer to the JencoMart case study.

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.)

On-premises infrastructure

Google



- 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 9

- (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 10

- (Topic 4)

For this question, refer to the Dress4Win case study.

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

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

Answer: B

NEW QUESTION 10

- (Topic 4)

For this question, refer to the Dress4Win case study.

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

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

Answer: A

Explanation:

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

- ? When transferring data from an on-premises location, use gsutil.
- ? When transferring data from another cloud storage provider, use Storage Transfer Service.
- ? Otherwise, evaluate both tools with respect to your specific scenario.

Use this guidance as a starting point. The specific details of your transfer scenario will also help you determine which tool is more appropriate
<https://cloud.google.com/storage-transfer/docs/overview>

NEW QUESTION 14

- (Topic 4)

For this question, refer to the Dress4Win case study.

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

They want to ensure that:

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

Which Google StackDriver features should they use?

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

Answer: D

NEW QUESTION 17

- (Topic 4)

For this question, refer to the Dress4Win case study.

Dress4Win has asked you for advice on how to migrate their on-premises MySQL deployment to the cloud. They want to minimize downtime and performance impact to their on-premises solution during the migration. Which approach should you recommend?

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

Answer: B

NEW QUESTION 19

- (Topic 5)

Your company captures all web traffic data in Google Analytics 260 and stores it in BigQuery. Each country has its own dataset. Each dataset has multiple tables. You want analysts from each country

to be able to see and query only the data for their respective countries. How should you configure the access rights?

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

Answer: A

NEW QUESTION 20

- (Topic 5)

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

- A. Create a shell script that uses the gcloud command to change the machine type of the development and acceptance instances to a smaller machine type outside of office hour

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

Answer: B

Explanation:

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

NEW QUESTION 21

- (Topic 5)

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

A.

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

B.

- * 1. Create a managed instance group with Compute Engine instances
- * 2. Create a global load balancer and configure it with two backends

* Managed instance group

* Cloud Storage bucket

- * 3. Enable Cloud CDN on the bucket backend

C.

- * 1. Mount the Cloud Storage bucket using gcsfuse on all backend Compute Engine instances

- * 2. Serve music files directly from the backend Compute Engine instance

D.

- * 1. Create a Cloud Filestore NFS volume and attach it to the backend Compute Engine instances

- * 2. Download popular songs in Cloud Filestore

- * 3. Serve music files directly from the backend Compute Engine instance

A.

Answer: B

NEW QUESTION 24

- (Topic 5)

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

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

Answer: D

Explanation:

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

NEW QUESTION 26

- (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 30

- (Topic 5)

You have a Python web application with many dependencies that requires 0.1 CPU cores and 128 MB of memory to operate in production. You want to monitor and maximize machine utilization. You also want to reliably deploy new versions of the application. Which set of steps should you take?

- A. Perform the following: 1) Create a managed instance group with f1-micro type machines. 2) Use a startup script to clone the repository, check out the production branch, install the dependencies, and start the Python app. 3) Restart the instances to automatically deploy new production releases.
- B. Perform the following: 1) Create a managed instance group with n1-standard-1 type machines. 2) Build a Compute Engine image from the production branch that

contains all of the dependencies and automatically starts the Python app.3) Rebuild the Compute Engine image, and update the instance template to deploy new production releases.

C. Perform the following:1) Create a Kubernetes Engine cluster with n1-standard-1 type machines.2) Build a Docker image from the production branch with all of the dependencies, and tag it with the version number.3) Create a Kubernetes Deployment with the imagePullPolicy set to "IfNotPresent" in the staging namespace, and then promote it to the production namespace after testing.

D. Perform the following:1) Create a Kubernetes Engine (GKE) cluster with n1-standard-4 type machines.2) Build a Docker image from the master branch with all of the dependencies, and tag it with "latest".3) Create a Kubernetes Deployment in the default namespace with the imagePullPolicy set to "Always".Restart the pods to automatically deploy new production releases.

Answer: D

Explanation:

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

NEW QUESTION 34

- (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 37

- (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 41

- (Topic 5)

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

- A. Configure Private Google Access for on-premises only.
- B. Perform the following tasks:1) Create a service account.2) Give the BigQuery JobUser role and Storage Reader role to the service account.3) Remove all other IAM access from the project.
- C. Configure VPC Service Controls and configure Private Google Access.
- D. Configure Private Google Access.

Answer: C

Explanation:

<https://cloud.google.com/vpc-service-controls/docs/overview>

VPC Service Controls improves your ability to mitigate the risk of data exfiltration from Google Cloud services such as Cloud Storage and BigQuery.

NEW QUESTION 46

- (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 -1 gs://backups/** to find and remove items older than 90 days and schedule it with cron.

Answer: B

Explanation:

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

NEW QUESTION 48

- (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 53

- (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 56

- (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 57

- (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 62

- (Topic 5)

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

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

Answer: D

Explanation:

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

NEW QUESTION 65

- (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. When 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 67

- (Topic 5)

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

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

Answer: C

Explanation:

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

NEW QUESTION 69

- (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 74

- (Topic 5)

Your company runs several databases on a single MySQL instance. They need to take backups of a specific database at regular intervals. The backup activity needs to complete as quickly as possible and cannot be allowed to impact disk performance. How should you configure the storage?

- A. Configure a cron job to use the gcloud tool to take regular backups using persistent disk snapshots.
- B. Mount a Local SSD volume as the backup location.
- C. After the backup is complete, use gsutil to move the backup to Google Cloud Storage.
- D. Use gcsfuse to mount a Google Cloud Storage bucket as a volume directly on the instance and write backups to the mounted location using mysqldump.
- E. Mount additional persistent disk volumes onto each virtual machine (VM) instance in a RAID10 array and use LVM to create snapshots to send to Cloud Storage.

Answer: B

Explanation:

<https://cloud.google.com/compute/docs/instances/sql-server/best-practices>

NEW QUESTION 76

- (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 77

- (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 80

- (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-datastore-batch-upsert-nodejs>

NEW QUESTION 84

- (Topic 5)

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

How should you design to meet Google best practices?

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

Answer: B

Explanation:

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

NEW QUESTION 88

- (Topic 5)

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

Which of the following is most accurate?

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

Answer: D

Explanation:

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

NEW QUESTION 91

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 95

- (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 99

- (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 103

- (Topic 5)

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

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

Answer: D

Explanation:

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

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

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

NEW QUESTION 104

- (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 107

- (Topic 5)

You are configuring the cloud network architecture for a newly created project in Google Cloud that will host applications in Compute Engine. Compute Engine virtual machine instances will be created in two different subnets (sub-a and sub-b) within a single region.

- Instances in sub-a will have public IP addresses
- Instances in sub-b will have only private IP addresses

To download updated packages, instances must connect to a public repository outside the boundaries of Google Cloud. You need to allow sub-b to access the external repository. What should you do?

- A. Enable Private Google Access on sub-b
- B. Configure Cloud NAT and select sub-b in the NAT mapping section
- C. Configure a bastion host instance in sub-a to connect to instances in sub-b

D. Enable Identity Aware Proxy for TCP forwarding for instances in sub-b

Answer: B

Explanation:

? Cloud NAT (network address translation) lets Google Cloud virtual machine (VM) instances without external IP addresses and private Google Kubernetes Engine (GKE) clusters send outbound packets to the internet and receive any corresponding established inbound response packets¹. By configuring Cloud NAT and selecting sub-b in the NAT mapping section, you can allow instances in sub-b to access the external repository without exposing them to the internet¹.

NEW QUESTION 109

- (Topic 5)

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

What actions will meet your company's needs?

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

Answer: B

Explanation:

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

NEW QUESTION 112

- (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 bucket
- F. Rotate the key after 5 years.
- G. Create the bucket with fine-grained access control, and grant a service account the role of Object Write
- H. Use the service account to upload new files.

Answer: A

Explanation:

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

NEW QUESTION 116

- (Topic 5)

As part of implementing their disaster recovery plan, your company is trying to replicate their production MySQL database from their private data center to their GCP project using a Google Cloud VPN connection.

They are experiencing latency issues and a small amount of packet loss that is disrupting the replication. What should they do?

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

Answer: B

NEW QUESTION 120

- (Topic 5)

You have an outage in your Compute Engine managed instance group: all instances keep restarting after 5 seconds. You have a health check configured, but autoscaling is disabled. Your colleague, who is a Linux expert, offered to look into the issue. You need to make sure that he can access the VMs. What should you do?

- A. Grant your colleague the IAM role of project Viewer
- B. Perform a rolling restart on the instance group
- C. Disable the health check for the instance group
- D. Add his SSH key to the project-wide SSH keys
- E. Disable autoscaling for the instance group
- F. Add his SSH key to the project-wide SSH Keys

Answer: C

Explanation:

<https://cloud.google.com/compute/docs/instance-groups/autohealing-instances-in-migs> Health checks used for autohealing should be conservative so they don't

preemptively delete and recreate your instances. When an autohealer health check is too aggressive, the autohealer might mistake busy instances for failed instances and unnecessarily restart them, reducing availability

NEW QUESTION 123

- (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 127

- (Topic 5)

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 130

- (Topic 5)

Your web application must comply with the requirements of the European Union's General Data Protection Regulation (GDPR). You are responsible for the technical architecture of your web application. What should you do?

- A. Ensure that your web application only uses native features and services of Google Cloud Platform, because Google already has various certifications and provides "pass-on" compliance when you use native features.
- B. Enable the relevant GDPR compliance setting within the GCP Console for each of the services in use within your application.
- C. Ensure that Cloud Security Scanner is part of your test planning strategy in order to pick up any compliance gaps.
- D. Define a design for the security of data in your web application that meets GDPR requirements.

Answer: D

Explanation:

<https://cloud.google.com/security/gdpr/?tab=tab4>

Reference: <https://www.mobiloud.com/blog/gdpr-compliant-mobile-app/>

NEW QUESTION 131

- (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 136

- (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 138

- (Topic 5)

You want to automate the creation of a managed instance group and a startup script to install the OS package dependencies. You want to minimize the startup time for VMs in the instance group.

What should you do?

- A. Use Terraform to create the managed instance group and a startup script to install the OS package dependencies.
- B. Create a custom VM image with all OS package dependence
- C. Use Deployment Manager to create the managed instance group with the VM image.
- D. Use Puppet to create the managed instance group and install the OS package dependencies.
- E. Use Deployment Manager to create the managed instance group and Ansible to install the OS package dependencies.

Answer: B

Explanation:

"Custom images are more deterministic and start more quickly than instances with startup scripts. However, startup scripts are more flexible and let you update the apps and settings in your instances more easily." https://cloud.google.com/compute/docs/instance-templates/create-instance-templates#using_custom_or_public_images_in_your_instance_templates

NEW QUESTION 142

- (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 143

- (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 147

- (Topic 5)

You need to develop procedures to verify resilience of disaster recovery for remote recovery using GCP. Your production environment is hosted on-premises. You need to establish a secure, redundant connection between your on premises network and the GCP network.

What should you do?

- A. Verify that Dedicated Interconnect can replicate files to GC
- B. Verify that direct peering can establish a secure connection between your networks if Dedicated Interconnect fails.
- C. Verify that Dedicated Interconnect can replicate files to GC
- D. Verify that Cloud VPN can establish a secure connection between your networks if Dedicated Interconnect fails.
- E. Verify that the Transfer Appliance can replicate files to GC
- F. Verify that direct peering can establish a secure connection between your networks if the Transfer Appliance fails.
- G. Verify that the Transfer Appliance can replicate files to GC
- H. Verify that Cloud VPN can establish a secure connection between your networks if the Transfer Appliance fails.

Answer: B

Explanation:

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

NEW QUESTION 148

- (Topic 5)

A recent audit that a new network was created in Your GCP project. In this network, a GCE instance has an SSH port open the world. You want to discover this network's origin. What should you do?

- A. Search for Create VM entry in the Stackdriver alerting console.
- B. Navigate to the Activity page in the Home section
- C. Set category to Data Access and search for Create VM entry.
- D. In the logging section of the console, specify GCE Network as the logging section
- E. Search for the Create Instance entry.
- F. Connect to the GCE instance using project SSH Key
- G. Identify previous logins in system logs, and match these with the project owners list.

Answer: C

NEW QUESTION 152

- (Topic 5)

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

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

Answer: B

Explanation:

The error is most likely caused by the access scope issue. When create new instance, you have the default Compute engine default service account but most services access including BigQuery is not enable. Create an instance Most access are not enabled by default You have default service account but don't have the permission (scope) you can stop the instance, edit, change scope and restart it to enable the scope access. Of course, if you Run your script on a new virtual machine with the BigQuery access scope enabled, it also works

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

NEW QUESTION 153

- (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 156

- (Topic 5)

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

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

Answer: B

Explanation:

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

NEW QUESTION 161

- (Topic 5)

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

- A. Use a load testing tool to simulate the expected number of concurrent users and total requests to your application, and inspect the results.
- B. Enable autoscaling on the GKE cluster and enable horizontal pod autoscaling on your application deployment
- C. Send curl requests to your application, and validate if the auto scaling works.
- D. Replicate the application over multiple GKE clusters in every Google Cloud region. Configure a global HTTP(S) load balancer to expose the different clusters over a single global IP address.
- E. Use Cloud Debugger in the development environment to understand the latency between the different microservices.

Answer: B

NEW QUESTION 165

- (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 167

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 168

- (Topic 5)

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

Which feature of Kubernetes should you use to accomplish this?

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

Answer: A

Explanation:

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

NEW QUESTION 172

- (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 177

- (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 182

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 184

- (Topic 5)

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

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

Answer: C

NEW QUESTION 185

- (Topic 5)

Your organization has decided to restrict the use of external IP addresses on instances to only approved instances. You want to enforce this requirement across all of your Virtual Private Clouds (VPCs). What should you do?

- A. Remove the default route on all VPC
- B. Move all approved instances into a new subnet that has a default route to an internet gateway.
- C. Create a new VPC in custom mod
- D. Create a new subnet for the approved instances, and set a default route to the internet gateway on this new subnet.
- E. Implement a Cloud NAT solution to remove the need for external IP addresses entirely.
- F. Set an Organization Policy with a constraint on constraints/compute.vmExternallpAcces
- G. List the approved instances in the allowedValues list.

Answer: D

Explanation:

Reference: <https://cloud.google.com/compute/docs/ip-addresses/reserve-static-external-ip-address>
<https://cloud.google.com/compute/docs/ip-addresses/reserve-static-external-ip-address#disableexternalip>
you might want to restrict external IP address so that only specific VM instances can use them. This option can help to prevent data exfiltration or maintain network isolation. Using an Organization Policy, you can restrict external IP addresses to specific VM instances with constraints to control use of external IP addresses for your VM instances within an organization or a project.

NEW QUESTION 189

- (Topic 5)

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

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

Answer: C

Explanation:

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

NEW QUESTION 190

- (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

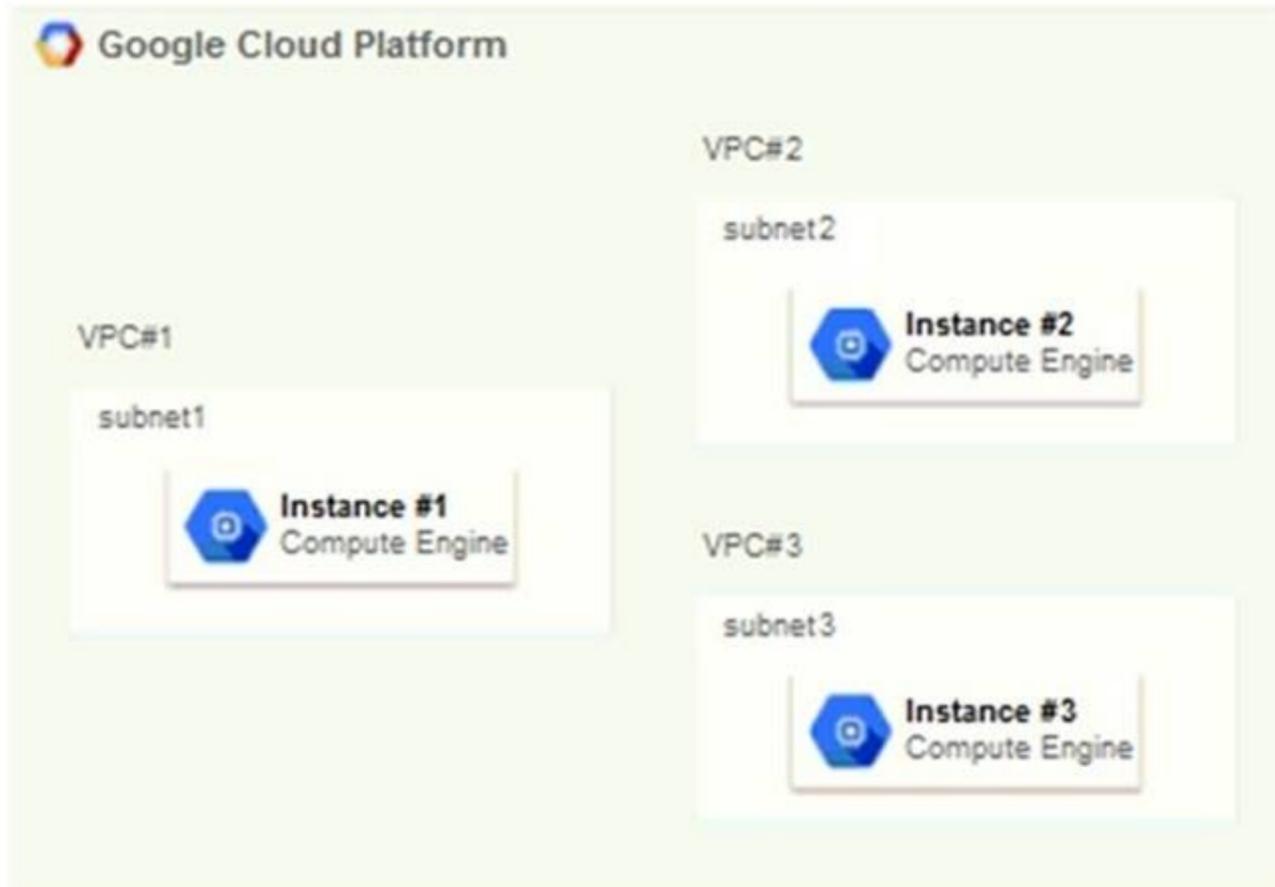
- A. Load logs into Google BigQuery.
- B. Load logs into Google Cloud SQL.
- C. Import logs into Google Stackdriver.
- D. Insert logs into Google Cloud Bigtable.
- E. Upload log files into Google Cloud Storage.

Answer: AE

NEW QUESTION 191

- (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 #3
 Update firewall rules to enable traffic between instances.
- C. Create two VPN tunnels via CloudVPN:
 - 1 between VPC #1 and VPC #2.
 - 1 between VPC #2 and VPC #3.
 Update firewall rules to enable traffic between the instances.
- D. Peer all three VPCs:
 - Peer VPC #1 with VPC #2.
 - Peer VPC #2 with VPC #3.
 Update firewall rules to enable traffic between the instances.

Answer: B

Explanation:

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

NEW QUESTION 196

- (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 200

- (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 203

- (Topic 5)

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

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

Answer: C

Explanation:

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

NEW QUESTION 207

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 212

- (Topic 5)

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

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

Answer: A

NEW QUESTION 217

- (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 218

- (Topic 5)

Your customer is receiving reports that their recently updated Google App Engine application is taking approximately 30 seconds to load for some of their users. This behavior was not reported before the update. What strategy should you take?

- A. Work with your ISP to diagnose the problem.
- B. Open a support ticket to ask for network capture and flow data to diagnose the problem, then roll back your application.
- C. Roll back to an earlier known good release initially, then use Stackdriver Trace and logging to diagnose the problem in a development/test/staging environment.
- D. Roll back to an earlier known good release, then push the release again at a quieter period to investigate.
- E. Then use Stackdriver Trace and logging to diagnose the problem.

Answer: C

Explanation:

Stackdriver Logging allows you to store, search, analyze, monitor, and alert on log data and events from Google Cloud Platform and Amazon Web Services (AWS). Our API also allows ingestion of any custom log data from any source. Stackdriver Logging is a fully managed service that performs at scale and can ingest application and system log data from thousands of VMs. Even better, you can analyze all that log data in real time.

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

NEW QUESTION 221

- (Topic 5)

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

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

Answer: B

NEW QUESTION 224

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 227

- (Topic 5)

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

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

Answer: A

Explanation:

For all Google Cloud services secured with VPC Service Controls, you can ensure that: Resources within a perimeter are accessed only from clients within authorized VPC networks using Private Google Access with either Google Cloud or on-premises. <https://cloud.google.com/vpc-service-controls/docs/overview>
<https://cloud.google.com/vpc-service-controls/docs/overview>. You create a service control across your VPC and any cloud bucket or any project resource to restrict access. Anything outside of it can't access the resources within service control perimeter

NEW QUESTION 228

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 230

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 232

- (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 236

- (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 238

- (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 239

- (Topic 6)

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

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

Answer: D

Explanation:

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

NEW QUESTION 244

- (Topic 7)

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

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

Answer: B

Explanation:

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

NEW QUESTION 247

- (Topic 7)

For this question, refer to the TerraEarth case study.

You start to build a new application that uses a few Cloud Functions for the backend. One use case requires a Cloud Function func_display to invoke another Cloud Function func_query. You want func_query only to accept invocations from func_display. You also want to follow Google's recommended best practices. What should you do?

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

Answer: B

Explanation:

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

NEW QUESTION 252

- (Topic 7)

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

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

I. Schedule a daily job to clean the data.

Answer: A

NEW QUESTION 257

- (Topic 8)

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

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

Answer: A

Explanation:

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

NEW QUESTION 259

- (Topic 8)

Your development teams release new versions of games running on Google Kubernetes Engine (GKE) daily. You want to create service level indicators (SLIs) to evaluate the quality of the new versions from the user's perspective. What should you do?

- A. Create CPU Utilization and Request Latency as service level indicators.
- B. Create GKE CPU Utilization and Memory Utilization as service level indicators.
- C. Create Request Latency and Error Rate as service level indicators.
- D. Create Server Uptime and Error Rate as service level indicators.

Answer: C

NEW QUESTION 260

- (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 264

- (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 269

- (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 273

- (Topic 9)

For this question, refer to the Helicopter Racing League (HRL) case study. A recent finance audit of cloud infrastructure noted an exceptionally high number of Compute Engine instances are allocated to do video encoding and transcoding. You suspect that these Virtual Machines are zombie machines that were not deleted after their workloads completed. You need to quickly get a list of which VM instances are idle. What should you do?

- A. Log into each Compute Engine instance and collect disk, CPU, memory, and network usage statistics for analysis.
- B. Use the gcloud compute instances list to list the virtual machine instances that have the idle: true label set.
- C. Use the gcloud recommender command to list the idle virtual machine instances.
- D. From the Google Console, identify which Compute Engine instances in the managed instance groups are no longer responding to health check probes.

Answer: C

Explanation:

Reference: <https://cloud.google.com/compute/docs/instances/viewing-and-applying-idle-vm-recommendations>

NEW QUESTION 278

- (Topic 10)

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

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

Answer: D

NEW QUESTION 279

- (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 283

- (Topic 10)

For this question, refer to the EHR Healthcare case study. EHR has single Dedicated Interconnect connection between their primary data center and Google's network. This connection satisfies EHR's network and security policies:

- On-premises servers without public IP addresses need to connect to cloud resources without public IP addresses
- Traffic flows from production network mgmt. servers to Compute Engine virtual machines should never traverse the public internet.

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

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

Answer: A

Explanation:

The case does not call out the throughput being an issue. However, to achieve 99.99%, you need to have 4 connections as per Google recommendations. However, in the options only A has the option to add an additional Interconnect connection. <https://cloud.google.com/network-connectivity/docs/interconnect/concepts/dedicated-overview#availability>

NEW QUESTION 285

- (Topic 10)

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

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

Answer: A

Explanation:

Binary Authorization to ensure only verified containers are deployed To ensure deployment are secure and and consistent, automatically scan images for

vulnerabilities with container analysis (https://cloud.google.com/docs/ci-cd/overview?hl=en&skip_cache=true)

NEW QUESTION 290

.....

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](#)