

Professional-Cloud-Architect Dumps

Google Certified Professional - Cloud Architect (GCP)

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



NEW QUESTION 1

- (Topic 1)

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

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

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

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

Answer: C

NEW QUESTION 2

- (Topic 2)

For this question, refer to the TerramEarth case study

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

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

Answer: A

Explanation:

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

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

Delegate application authorization with OAuth2

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

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

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

NEW QUESTION 3

- (Topic 2)

For this question, refer to the TerramEarth case study.

TerramEarth has equipped unconnected trucks with servers and sensors to collect telemetry data. Next year they want to use the data to train machine learning models. They want to store this data in the cloud while reducing costs. What should they do?

- A. Have the vehicle's computer compress the data in hourly snapshots, and store it in a Google Cloud storage (GCS) Nearline bucket.
- B. Push the telemetry data in Real-time to a streaming dataflow job that compresses the data, and store it in Google BigQuery.
- C. Push the telemetry data in real-time to a streaming dataflow job that compresses the data, and store it in Cloud Bigtable.
- D. Have the vehicle's computer compress the data in hourly snapshots, and store it in a GCS Coldline bucket.

Answer: D

Explanation:

Coldline Storage is the best choice for data that you plan to access at most once a year, due to its slightly lower availability, 90-day minimum storage duration, costs for data access, and higher per-operation costs. For example:

Cold Data Storage - Infrequently accessed data, such as data stored for legal or regulatory reasons, can be stored at low cost as Coldline Storage, and be available when you need it. Disaster recovery - In the event of a disaster recovery event, recovery time is key. Cloud Storage provides low latency access to data stored as Coldline Storage.

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

NEW QUESTION 4

- (Topic 2)

For this question refer to the TerramEarth case study

Operational parameters such as oil pressure are adjustable on each of TerramEarth's vehicles to increase their efficiency, depending on their environmental conditions. Your primary goal is to increase the operating efficiency of all 20 million cellular and unconnected vehicles in the field How can you accomplish this goal?

- A. Have your engineers inspect the data for patterns, and then create an algorithm with rules that make operational adjustments automatically.
- B. Capture all operating data, train machine learning models that identify ideal operations, and run locally to make operational adjustments automatically.
- C. Implement a Google Cloud Dataflow streaming job with a sliding window, and use Google Cloud Messaging (GCM) to make operational adjustments automatically.
- D. Capture all operating data, train machine learning models that identify ideal operations, and host in Google Cloud Machine Learning (ML) Platform to make operational adjustments automatically.

Answer: B

NEW QUESTION 5

- (Topic 2)

For this question, refer to the TerramEarth case study.

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

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

Answer: A

Explanation:

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

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

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

From scenario: Business Requirements

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

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

Have the ability to partner with different companies – especially with seed and fertilizer suppliers in the fast-growing agricultural business – to create compelling joint offerings for their customers.

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

NEW QUESTION 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

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

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

Answer: CDF

Explanation:

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

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

There is no firewall rule allowing SSH access on the port. SSH access on port 22 is enabled on all Compute Engine instances by default. If you have disabled access, SSH from the Browser will not work. If you run sshd on a port other than 22, you need to enable the access to that port with a custom firewall rule.

The firewall rule allowing SSH access is enabled, but is not configured to allow connections from GCP Console services. Source IP addresses for browser-based SSH sessions are dynamically allocated by GCP Console and can vary from session to session.

References: <https://cloud.google.com/compute/docs/ssh-in-browser> <https://cloud.google.com/compute/docs/ssh-in-browser>

NEW QUESTION 8

- (Topic 4)

For this question, refer to the Dress4Win case study.

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

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

Answer: A

NEW QUESTION 9

- (Topic 4)

For this question, refer to the Dress4Win case study.

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

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

Answer: C

NEW QUESTION 10

- (Topic 4)

For this question, refer to the Dress4Win case study.

Dress4Win has configured a new uptime check with Google Stackdriver for several of their legacy services. The Stackdriver dashboard is not reporting the services as healthy. What should they do?

- A. Install the Stackdriver agent on all of the legacy web servers.
- B. In the Cloud Platform Console download the list of the uptime servers' IP addresses and create an inbound firewall rule
- C. Configure their load balancer to pass through the User-Agent HTTP header when the value matches GoogleStackdriverMonitoring-UptimeChecks (<https://cloud.google.com/monitoring>)
- D. Configure their legacy web servers to allow requests that contain user-Agent HTTP header when the value matches GoogleStackdriverMonitoring—UptimeChecks (<https://cloud.google.com/monitoring>)

Answer: B

NEW QUESTION 10

- (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 12

- (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 13

- (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 asingle group'all_analysts', and add all country-groups as member
- D. Grant the 'all-analysis' group the IAM role ofBigQuery 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 ofBigQuery 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 ofBigQuery 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 ofBigQuery dataViewe
- T. Share the appropriate table with view access with each respective analyst countrygroup.

Answer: A

NEW QUESTION 17

- (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 usepreemptible 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 20

- (Topic 5)

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 24

- (Topic 5)

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

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

Answer: B

Explanation:

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

NEW QUESTION 29

- (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 33

- (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 35

- (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 37

- (Topic 5)

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

- A. Configure an organization policy to restrict identities by domain
- B. Configure an organization policy to block creation of service accounts
- C. Configure Cloud Scheduler o trigger a Cloud Function every hour that removes all users that don't belong to the Cloud identity domain from all projects.
- D. Create a technical user (e g . crawler@yourdomain com), and give it the protect owner rote at root organization level Write a bash script that• Lists all me IAM rules of all projects within the organization• Deletes all users that do not belong to the company domainCreate a Compute Engine instance m a project within the Organization and configure gcloud to be executed with technical user credentials Configure a cron job that executes the bash script every hour.

Answer: A

Explanation:

<https://cloud.google.com/resource-manager/docs/organization-policy/restricting-domains>

NEW QUESTION 42

- (Topic 5)

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

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

Answer: ABF

NEW QUESTION 44

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 48

- (Topic 5)

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

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

Answer: B

Explanation:

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

NEW QUESTION 52

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 55

- (Topic 5)

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

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

Answer: C

NEW QUESTION 56

- (Topic 5)

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

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

Answer: D

Explanation:

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

NEW QUESTION 60

- (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 63

- (Topic 5)

Your operations team has asked you to help diagnose a performance issue in a production application that runs on Compute Engine. The application is dropping requests that reach it when under heavy load. The process list for affected instances shows a single application process that is consuming all available CPU, and autoscaling has reached the upper limit of instances. There is no abnormal load on any other related systems, including the database. You want to allow production traffic to be served again as quickly as possible. Which action should you recommend?

- A. Change the autoscaling metric to `agent.googleapis.com/memory/percent_used`.
- B. Restart the affected instances on a staggered schedule.
- C. SSH to each instance and restart the application process.
- D. Increase the maximum number of instances in the autoscaling group.

Answer: D

Explanation:

Reference: <https://cloud.google.com/blog/products/sap-google-cloud/best-practices-for-sap-app-server-autoscaling-on-google-cloud>

NEW QUESTION 68

- (Topic 5)

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

What should you do?

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

Answer: B

Explanation:

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

NEW QUESTION 71

- (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 73

- (Topic 5)

The development team has provided you with a Kubernetes Deployment file. You have no infrastructure yet and need to deploy the application. What should you do?

- A. Use gcloud to create a Kubernetes cluster
- B. Use Deployment Manager to create the deployment.
- C. Use gcloud to create a Kubernetes cluster
- D. Use kubectl to create the deployment.
- E. Use kubectl to create a Kubernetes cluster
- F. Use Deployment Manager to create the deployment.
- G. Use kubectl to create a Kubernetes cluster
- H. Use kubectl to create the deployment.

Answer: B

Explanation:

<https://cloud.google.com/kubernetes-engine/docs/how-to/creating-a-cluster>

NEW QUESTION 77

- (Topic 5)

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

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

Answer: CD

Explanation:

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

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

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

NEW QUESTION 81

- (Topic 5)

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

- A. Schedule a cron job with Cloud Scheduler
- B. The scheduled job queries the logs every minute for the relevant events.

- C. Export logs to BigQuery, and trigger a query in BigQuery to process the log data for the relevant events.
- D. Export logs to a Pub/Sub topic, and trigger Cloud Function with the relevant log events.
- E. Export logs to a Cloud Storage bucket, and trigger Cloud Run with the relevant log events.

Answer: C

Explanation:

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

NEW QUESTION 86

- (Topic 5)

You have an application that runs in Google Kubernetes Engine (GKE). Over the last 2 weeks, customers have reported that a specific part of the application returns errors very frequently. You currently have no logging or monitoring solution enabled on your GKE cluster. You want to diagnose the problem, but you have not been able to replicate the issue. You want to cause minimal disruption to the application. What should you do?

- A. * 1. Update your GKE cluster to use Cloud Operations for GKE.* 2. Use the GKE Monitoring dashboard to investigate logs from affected Pods.
- B. * 1. Create a new GKE cluster with Cloud Operations for GKE enabled.* 2. Migrate the affected Pods to the new cluster, and redirect traffic for those Pods to the new cluster.* 3. Use the GKE Monitoring dashboard to investigate logs from affected Pods.
- C. * 1. Update your GKE cluster to use Cloud Operations for GKE, and deploy Prometheus.* 2. Set an alert to trigger whenever the application returns an error.
- D. * 1. Create a new GKE cluster with Cloud Operations for GKE enabled, and deploy Prometheus.* 2. Migrate the affected Pods to the new cluster, and redirect traffic for those Pods to the new cluster.* 3. Set an alert to trigger whenever the application returns an error.

Answer: A

Explanation:

Reference: <https://cloud.google.com/blog/products/management-tools/using-logging-your-apps-running-kubernetes-engine>

NEW QUESTION 89

- (Topic 5)

You deploy your custom java application to google app engine. It fails to deploy and gives you the following stack trace:

```

Java.lang.securityException : SHA1 digest

At com.google.appengine.runtime.Request.pro

At

Sun.securityutil.manifestEntryVerifier.ver

At java . net . URLClassLoader . defineCla

At sun . reflect . GeneratedMethodAccessors

At

Sun.reflect . DelegatingMethodAccesorImpl.

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

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

Answer: B

NEW QUESTION 90

- (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 95

- (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 97

- (Topic 5)

Your company uses Google Kubernetes Engine (GKE) as a platform for all workloads. Your company has a single large GKE cluster that contains batch, stateful, and stateless workloads. The GKE cluster is configured with a single node pool with 200 nodes. Your company needs to reduce the cost of this cluster but does not want to compromise availability. What should you do?

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

Answer: C

Explanation:

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

NEW QUESTION 101

- (Topic 5)

You have deployed an application to Kubernetes Engine, and are using the Cloud SQL proxy container to make the Cloud SQL database available to the services running on Kubernetes. You are notified that the application is reporting database connection issues. Your company policies require a post-mortem. What should you do?

- A. Use gcloud sql instances restart.
- B. Validate that the Service Account used by the Cloud SQL proxy container still has the Cloud Build Editor role.
- C. In the GCP Console, navigate to Stackdriver Logging.
- D. Consult logs for Kubernetes Engine and Cloud SQL.
- E. In the GCP Console, navigate to Cloud SQL.
- F. Restore the latest backup.
- G. Use kubectl to restart all pods.

Answer: C

NEW QUESTION 105

- (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 110

- (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 packets. 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 internet.

NEW QUESTION 112

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 117

- (Topic 5)

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

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

Answer: B

Explanation:

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

NEW QUESTION 120

- (Topic 5)

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

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

Answer: C

NEW QUESTION 122

- (Topic 5)

Your company is running a stateless application on a Compute Engine instance. The application is used

heavily during regular business hours and lightly outside of business hours. Users are reporting that the application is slow during peak hours. You need to optimize the application's performance. What should you do?

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

Answer: B

Explanation:

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

NEW QUESTION 124

- (Topic 5)

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

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

Answer: B

Explanation:

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

NoSQL fits well for:

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

NEW QUESTION 127

- (Topic 5)

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

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

Answer: C

Explanation:

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

NEW QUESTION 129

- (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)

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

- A. Create a Dataproc cluster using standard worker instances.

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

Answer: B

Explanation:

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

NEW QUESTION 132

- (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 137

- (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 dependencies
- 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 138

- (Topic 5)

You created a pipeline that can deploy your source code changes to your infrastructure in instance groups for self healing. One of the changes negatively affects your key performance indicator. You are not sure how to fix it and investigation could take up to a week. What should you do

- A. Log in to a server, and iterate a fix locally
- B. Change the instance group template to the previous one, and delete all instances.
- C. Revert the source code change and rerun the deployment pipeline
- D. Log into the servers with the bad code change, and swap in the previous code

Answer: C

NEW QUESTION 143

- (Topic 5)

You have an App Engine application that needs to be updated. You want to test the update with production traffic before replacing the current application version. What should you do?

- A. Deploy the update using the Instance Group Updater to create a partial rollout, which allows for canary testing.
- B. Deploy the update as a new version in the App Engine application, and split traffic between the new and current versions.
- C. Deploy the update in a new VPC, and use Google's global HTTP load balancing to split traffic between the update and current applications.
- D. Deploy the update as a new App Engine application, and use Google's global HTTP load balancing to split traffic between the new and current applications.

Answer: B

Explanation:

<https://cloud.google.com/appengine/docs/standard/python/splitting-traffic>

NEW QUESTION 145

- (Topic 5)

Your company wants to start using Google Cloud resources but wants to retain their on-premises Active

Directory domain controller for identity management. What should you do?

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

Answer: B

Explanation:

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

NEW QUESTION 149

- (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 153

- (Topic 5)

Your company has an enterprise application running on Compute Engine that requires high availability and high performance. The application has been deployed on two instances in two zones in the same region in active-passive mode. The application writes data to a persistent disk in the case of a single zone outage that data should be immediately made available to the other instance in the other zone. You want to maximize performance while minimizing downtime and data loss. What should you do?

- A.
 - * 1. Attach a persistent SSD disk to the first instance
 - * 2. Create a snapshot every hour
 - * 3. In case of a zone outage, recreate a persistent SSD disk in the second instance where data is coming from the created snapshot
- B.
 - * 1. Create a Cloud Storage bucket
 - * 2. Mount the bucket into the first instance with gcs-fuse
 - * 3. In case of a zone outage, mount the Cloud Storage bucket to the second instance with gcs-fuse
- C.
 - * 1. Attach a local SSD to the first instance disk
 - * 2. Execute an rsync command every hour where the target is a persistent SSD disk attached to the second instance
 - * 3. In case of a zone outage, use the second instance
- D.
 - * 1. Attach a regional SSD persistent disk to the first instance
 - * 2. In case of a zone outage, force-attach the disk to the other instance

A.

Answer: D

NEW QUESTION 155

- (Topic 5)

Your team needs to create a Google Kubernetes Engine (GKE) cluster to host a newly built application that requires access to third-party services on the internet. Your company does not allow any Compute Engine instance to have a public IP address on Google Cloud. You need to create a deployment strategy that adheres to these guidelines. What should you do?

- A. Create a Compute Engine instance, and install a NAT Proxy on the instance
- B. Configure all workloads on GKE to pass through this proxy to access third-party services on the Internet
- C. Configure the GKE cluster as a private cluster, and configure Cloud NAT Gateway for the cluster subnet
- D. Configure the GKE cluster as a route-based cluster
- E. Configure Private Google Access on the Virtual Private Cloud (VPC)
- F. Configure the GKE cluster as a private cluster
- G. Configure Private Google Access on the Virtual Private Cloud (VPC)

Answer: B

Explanation:

A Cloud NAT gateway can perform NAT for nodes and Pods in a private cluster, which is a type of VPC-native cluster. The Cloud NAT gateway must be configured to apply to at least the following subnet IP address ranges for the subnet that your cluster uses:

Subnet primary IP address range (used by nodes)

Subnet secondary IP address range used for Pods in the cluster Subnet secondary IP address range used for Services in the cluster

The simplest way to provide NAT for an entire private cluster is to configure a Cloud NAT gateway to apply to all of the cluster's subnet's IP address ranges.
<https://cloud.google.com/nat/docs/overview>

NEW QUESTION 159

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 161

- (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 166

- (Topic 5)

Auditors visit your teams every 12 months and ask to review all the Google Cloud Identity and Access Management (Cloud IAM) policy changes in the previous 12 months. You want to streamline and expedite the analysis and audit process. What should you do?

- A. Create custom Google Stackdriver alerts and send them to the auditor.
- B. Enable Logging export to Google BigQuery and use ACLs and views to scope the data shared with the auditor.
- C. Use cloud functions to transfer log entries to Google Cloud SQL and use ACLs and views to limit an auditor's view.
- D. Enable Google Cloud Storage (GCS) log export to audit logs into a GCS bucket and delegate access to the bucket.

Answer: D

Explanation:

Export the logs to Google Cloud Storage bucket - Archive Storage, as it will not be used for 1 year, price for which is \$0.004 per GB per Month. The price for long term storage in BigQuery is \$0.01 per GB per Month (250% more). Also for analysis purpose, whenever Auditors are there (once per year), you can use BigQuery and use GCS bucket as external data source. BigQuery supports querying Cloud Storage data from these storage classes:
Standard Nearline Coldline Archive

NEW QUESTION 169

- (Topic 5)

You have developed a non-critical update to your application that is running in a managed instance group, and have created a new instance template with the update that you want to release. To prevent any possible impact to the application, you don't want to update any running instances. You want any new instances that are created by the managed instance group to contain the new update. What should you do?

- A. Start a new rolling restart operation.
- B. Start a new rolling replace operation.
- C. Start a new rolling update.
- D. Select the Proactive update mode.
- E. Start a new rolling update.
- F. Select the Opportunistic update mode.

Answer: D

Explanation:

In certain scenarios, an opportunistic update is useful because you don't want to cause instability to the system if it can be avoided. For example, if you have a non-critical update that can be applied as necessary without any urgency and you have a MIG that is actively being autoscaled, perform an opportunistic update so that Compute Engine does not actively tear down your existing instances to apply the update. When resizing down, the autoscaler preferentially terminates

instances with the old template as well as instances that are not yet in a RUNNING state.

NEW QUESTION 170

- (Topic 5)

You are using a single Cloud SQL instance to serve your application from a specific zone.

You want to

introduce high availability. What should you do?

- A. Create a read replica instance in a different region
- B. Create a failover replica instance in a different region
- C. Create a read replica instance in the same region, but in a different zone
- D. Create a failover replica instance in the same region, but in a different zone

Answer: B

Explanation:

<https://cloud.google.com/sql/docs/mysql/high-availability>

NEW QUESTION 172

- (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 serve traffic, and use the HTTP load balancing service to fail over to the standby instance in case of a disaster.

Answer: C

NEW QUESTION 173

- (Topic 5)

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 176

- (Topic 5)

Your company is designing its application landscape on Compute Engine. Whenever a

zonal outage occurs, the application should be restored in another zone as quickly as possible with the latest application data. You need to design the solution to meet this requirement. What should you do?

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

Answer: B

Explanation:

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

NEW QUESTION 178

- (Topic 5)

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

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

Answer: B

Explanation:

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

NEW QUESTION 183

- (Topic 5)

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

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

Answer: C

Explanation:

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

NEW QUESTION 186

- (Topic 5)

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

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

Answer: C

Explanation:

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

NEW QUESTION 189

- (Topic 5)

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 191

- (Topic 5)

You have an application that makes HTTP requests to Cloud Storage. Occasionally the requests fail with HTTP status codes of 5xx and 429. How should you handle these types of errors?

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

Answer: A

Explanation:

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

NEW QUESTION 193

- (Topic 5)

You have deployed several instances on Compute Engine. As a security requirement, instances cannot have a public IP address. There is no VPN connection between Google Cloud and your office, and you need to connect via SSH into a specific machine without violating the security requirements. What should you do?

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

Answer: C

Explanation:

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

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

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

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

NEW QUESTION 195

- (Topic 5)

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

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

Answer: C

NEW QUESTION 200

- (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 204

- (Topic 5)

Your organization has a 3-tier web application deployed in the same network on Google Cloud Platform. Each tier (web, API, and database) scales independently of the others. Network traffic should flow through the web to the API tier and then on to the database tier. Traffic should not flow between the web and the database tier. How should you configure the network?

- A. Add each tier to a different subnetwork.
- B. Set up software based firewalls on individual VMs.
- C. Add tags to each tier and set up routes to allow the desired traffic flow.
- D. Add tags to each tier and set up firewall rules to allow the desired traffic flow.

Answer: D

Explanation:

<https://aws.amazon.com/blogs/aws/building-three-tier-architectures-with-security-groups/>

Google Cloud Platform(GCP) enforces firewall rules through rules and tags. GCP rules and tags can be defined once and used across all regions.

References: <https://cloud.google.com/docs/compare/openstack/> <https://aws.amazon.com/it/blogs/aws/building-three-tier-architectures-with-security-groups/>

NEW QUESTION 206

- (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 209

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 211

- (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 213

- (Topic 7)

TerraEarth has a legacy web application that you cannot migrate to cloud. However, you still want to build a cloud-native way to monitor the application. If the application goes down, you want the URL to point to a "Site is unavailable" page as soon as possible. You also want your Ops team to receive a notification for the issue. You need to build a reliable solution for minimum cost. What should you do?

- A. Create a scheduled job in Cloud Run to invoke a container every minute

- B. The container will check the application URL. If the application is down, switch the URL to the "Site is unavailable" page, and notify the Ops team.
- C. Create a cron job on a Compute Engine VM that runs every minute.
- D. The cron job invokes a Python program to check the application URL. If the application is down, switch the URL to the "Site is unavailable" page, and notify the Ops team.
- E. Create a Cloud Monitoring uptime check to validate the application URL. If it fails, put a message in a Pub/Sub queue that triggers a Cloud Function to switch the URL to the "Site is unavailable" page, and notify the Ops team.
- F. Use Cloud Error Reporting to check the application URL. If the application is down, switch the URL to the "Site is unavailable" page, and notify the Ops team.

Answer: C

Explanation:

<https://cloud.google.com/blog/products/management-tools/how-to-use-pubsub-as-a-cloud-monitoring-notification-channel>

NEW QUESTION 218

- (Topic 7)

TerraEarth has about 1 petabyte (PB) of vehicle testing data in a private data center. You want to move the data to Cloud Storage for your machine learning team. Currently, a 1-Gbps interconnect link is available for you. The machine learning team wants to start using the data in a month. What should you do?

- A. Request Transfer Appliances from Google Cloud, export the data to appliances, and return the appliances to Google Cloud.
- B. Configure the Storage Transfer service from Google Cloud to send the data from your data center to Cloud Storage.
- C. Make sure there are no other users consuming the 1-Gbps link, and use multi-thread transfer to upload the data to Cloud Storage.
- D. Export files to an encrypted USB device, send the device to Google Cloud, and request an import of the data to Cloud Storage.

Answer: A

NEW QUESTION 219

- (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 220

- (Topic 8)

For this question, refer to the Mountkirk Games case study. Mountkirk Games wants to design their solution for the future in order to take advantage of cloud and technology improvements as they become available. Which two steps should they take? (Choose two.)

- A. Store as much analytics and game activity data as financially feasible today so it can be used to train machine learning models to predict user behavior in the future.
- B. Begin packaging their game backend artifacts in container images and running them on Kubernetes Engine to improve the availability to scale up or down based on game activity.
- C. Set up a CI/CD pipeline using Jenkins and Spinnaker to automate canary deployments and improve development velocity.
- D. Adopt a schema versioning tool to reduce downtime when adding new game features that require storing additional player data in the database.
- E. Implement a weekly rolling maintenance process for the Linux virtual machines so they can apply critical kernel patches and package updates and reduce the risk of 0-day vulnerabilities.

Answer: BC

NEW QUESTION 222

- (Topic 8)

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

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

Answer: D

NEW QUESTION 224

- (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 228

- (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 233

- (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 238

- (Topic 10)

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

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

Answer: C

Explanation:

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

NEW QUESTION 240

- (Topic 10)

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

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

Answer: A

Explanation:

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

NEW QUESTION 242

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

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

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