



Google

Exam Questions Professional-Cloud-Architect

Google Certified Professional - Cloud Architect (GCP)

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

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 1)

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

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

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

Answer: A

Explanation:

From scenario: Requirements for Game Backend Platform

? Dynamically scale up or down based on game activity

? Connect to a managed NoSQL database service

? Run customize Linux distro

NEW QUESTION 3

- (Topic 2)

For this question, refer to the TerramEarth case study.

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

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

Answer: D

Explanation:

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

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

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

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

NEW QUESTION 4

- (Topic 2)

For this question, refer to the TerramEarth case study.

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

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

Answer: B

Explanation:

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

NEW QUESTION 5

- (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 6

- (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 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.

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

- A. Store image files in a Google Cloud Storage bucket
- B. Use Google Cloud Datastore to maintain metadata that maps each customer's ID and their image files.

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

Answer: A

NEW QUESTION 9

- (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 10

- (Topic 4)

For this question, refer to the Dress4Win case study.

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

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

Answer: C

NEW QUESTION 10

- (Topic 5)

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

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

Answer: A

Explanation:

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

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

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

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

NEW QUESTION 14

- (Topic 5)

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

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

Answer: B

Explanation:

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

NEW QUESTION 16

- (Topic 5)

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

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

Answer: B

NEW QUESTION 17

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 20

- (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 23

- (Topic 5)

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

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

Answer: C

Explanation:

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

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

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

You can use this commit hash as a version number for your software, but also as a tag for the Docker image built from this specific version of your software. Doing

so makes Docker images traceable: because in this case the image tag is immutable, you instantly know which specific version of your software is running inside a given container."

NEW QUESTION 28

- (Topic 5)

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

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

Answer: C

Explanation:

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

NEW QUESTION 32

- (Topic 5)

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

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

A.

Answer: C

Explanation:

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

NEW QUESTION 37

- (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 38

- (Topic 5)

You are tasked with building an online analytical processing (OLAP) marketing analytics and reporting tool.

This requires a relational database that can operate on hundreds of terabytes of data. What is the Google recommended tool for such applications?

- A. Cloud Spanner, because it is globally distributed
- B. Cloud SQL, because it is a fully managed relational database
- C. Cloud Firestore, because it offers real-time synchronization across devices
- D. BigQuery, because it is designed for large-scale processing of tabular data

Answer: A

Explanation:

Reference: <https://cloud.google.com/files/BigQueryTechnicalWP.pdf>

NEW QUESTION 40

- (Topic 5)

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

should suddenly crash. What should you do?

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

Answer: B

Explanation:

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

NEW QUESTION 42

- (Topic 5)

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

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

Answer: D

Explanation:

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

NEW QUESTION 46

- (Topic 5)

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

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

Answer: C

Explanation:

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

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

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

NEW QUESTION 48

- (Topic 5)

Your organization has stored sensitive data in a Cloud Storage bucket. For regulatory reasons, your company must be able to rotate the encryption key used to encrypt the data in the bucket. The data will be processed in Dataproc. You want to follow Google-recommended practices for security. What should you do?

- A. Create a key with Cloud Key Management Service (KMS). Encrypt the data using the encrypt method of Cloud KMS.
- B. Create a key with Cloud Key Management Service (KMS). Set the encryption key on the bucket to the Cloud KMS key.
- C. Generate a GPG key pair
- D. Encrypt the data using the GPG key
- E. Upload the encrypted data to the bucket.
- F. Generate an AES-256 encryption key
- G. Encrypt the data in the bucket using the customer-supplied encryption keys feature.

Answer: B

Explanation:

<https://cloud.google.com/storage/docs/encryption/using-customer-managed-keys#add-object-key>

<https://cloud.google.com/storage/docs/encryption/using-customer-managed-keys>

NEW QUESTION 52

- (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 53

- (Topic 5)

You are implementing the infrastructure for a web service on Google Cloud. The web service needs to receive and store the data from 500,000 requests per second. The data will be queried later in real time, based on exact matches of a known set of attributes. There will be periods where the web service will not receive any requests. The business wants to keep costs low. Which web service platform and database should you use for the application?

- A. Cloud Run and BigQuery
- B. Cloud Run and Cloud Bigtable
- C. A Compute Engine autoscaling managed instance group and BigQuery
- D. A Compute Engine autoscaling managed instance group and Cloud Bigtable

Answer: B

Explanation:

<https://cloud.google.com/run/docs/about-instance-autoscaling> <https://cloud.google.com/blog/topics/developers-practitioners/bigtable-vs-bigquery-whats-difference>

NEW QUESTION 56

- (Topic 5)

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

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

Answer: C

Explanation:

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

NEW QUESTION 59

- (Topic 5)

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

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

Answer: D

Explanation:

Overview of storage classes, price, and use cases <https://cloud.google.com/storage/docs/storage-classes>

Why export logs? <https://cloud.google.com/logging/docs/export/>

StackDriver Quotas and Limits for Monitoring <https://cloud.google.com/monitoring/quotas> The BigQuery pricing. <https://cloud.google.com/bigquery/pricing>

NEW QUESTION 60

- (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 62

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 65

- (Topic 5)

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

How should you configure the firewall rules?

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

Answer: B

Explanation:

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

NEW QUESTION 69

- (Topic 5)

You need to deploy a stateful workload on Google Cloud. The workload can scale horizontally, but each instance needs to read and write to the same POSIX filesystem. At high load, the stateful workload needs to support up to 100 MB/s of writes. What should you do?

- A. Use a persistent disk for each instance.
- B. Use a regional persistent disk for each instance.
- C. Create a Cloud Filestore instance and mount it in each instance.
- D. Create a Cloud Storage bucket and mount it in each instance using gcsfuse.

Answer: C

Explanation:

<https://cloud.google.com/storage/docs/gcs-fuse#notes>

Cloud Filestore: Cloud Filestore is a scalable and highly available shared file service fully managed by Google. Cloud Filestore provides persistent storage ideal for shared workloads. It is best suited for enterprise applications requiring persistent, durable, shared storage which is accessed by NFS or requires a POSIX compliant file system.

Reference: <https://cloud.google.com/storage/docs/gcs-fuse>

NEW QUESTION 72

- (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 73

- (Topic 5)

You need to reduce the number of unplanned rollbacks of erroneous production deployments in your company's web hosting platform. Improvement to the QA/Test

processes accomplished an 80% reduction. Which additional two approaches can you take to further reduce the rollbacks? Choose 2 answers

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

Answer: AC

NEW QUESTION 74

- (Topic 5)

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

Which of the following is most accurate?

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

Answer: D

Explanation:

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

NEW QUESTION 78

- (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 80

- (Topic 5)

Your team will start developing a new application using microservices architecture on Kubernetes Engine. As part of the development lifecycle, any code change that has been pushed to the remote develop branch on your GitHub repository should be built and tested automatically. When the build and test are successful, the

relevant microservice will be deployed automatically in the development environment. You want to ensure that all code deployed in the development environment follows this process. What should you do?

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

Answer: C

Explanation:

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

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

NEW QUESTION 84

- (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 workloads
- 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 85

- (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 87

- (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 91

- (Topic 5)

As part of implementing their disaster recovery plan, your company is trying to replicate their production MySQL database from their private data center to their GCP project using a Google Cloud VPN connection. They are experiencing latency issues and a small amount of packet loss that is disrupting the replication. What should they do?

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

Answer: B

NEW QUESTION 93

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 97

- (Topic 5)

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 98

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 103

- (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 104

- (Topic 5)

Your company acquired a healthcare startup and must retain its customers' medical information for up to 4 more years, depending on when it was created. Your corporate policy is to securely retain this data, and then delete it as soon as regulations allow.

Which approach should you take?

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

Answer: C

Explanation:

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

NEW QUESTION 106

- (Topic 5)

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

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

Answer: D

Explanation:

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

NEW QUESTION 109

- (Topic 5)

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

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

Answer: D

Explanation:

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

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

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

NEW QUESTION 114

- (Topic 5)

Your company plans to migrate a multi-petabyte data set to the cloud. The data set must be available 24hrs a day. Your business analysts have experience only with using a SQL interface. How should you store the data to optimize it for ease of analysis?

- A. Load data into Google BigQuery.
- B. Insert data into Google Cloud SQL.
- C. Put flat files into Google Cloud Storage.
- D. Stream data into Google Cloud Datastore.

Answer: A

Explanation:

Google Big Query is for multi-peta-byte storage, HA (High availability) which means 24 hours, SQL interface.

<https://medium.com/google-cloud/the-12-components-of-google-bigquery-c2b49829a7c7> <https://cloud.google.com/solutions/bigquery-data-warehouse>

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

BigQuery is Google's serverless, highly scalable, low cost enterprise data warehouse designed to make all your data analysts productive. Because there is no infrastructure to manage, you can focus on analyzing data to find meaningful insights using familiar SQL and you don't need a database administrator. BigQuery enables you to analyze all your data by creating a logical data warehouse over managed, columnar storage as well as data from object storage, and spreadsheets.

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

NEW QUESTION 119

- (Topic 5)

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 123

- (Topic 5)

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

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

app = Flask(__name__)
sessions = {}

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

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

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

    return news.render(news_articles)

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

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

Answer: A

Explanation:

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

NEW QUESTION 124

- (Topic 5)

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

- A. Create a Compute Engine instance template using the most recent Debian image
- B. Create an instance from this template, and install and configure the application as part of the startup script
- C. Repeat this process whenever a new Google-managed Debian image becomes available.
- D. Create a Debian-based Compute Engine instance, install and configure the application, and use OS patch management to install available updates.

- E. Create an instance with the latest available Debian image
- F. Connect to the instance via SSH, and install and configure the application on the instance
- G. Repeat this process whenever a new Google-managed Debian image becomes available.
- H. Create a Docker container with Debian as the base image
- I. Install and configure the application as part of the Docker image creation process
- J. Host the container on Google Kubernetes Engine and restart the container whenever a new update is available.

Answer: B

Explanation:

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

NEW QUESTION 129

- (Topic 5)

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

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

Answer: CF

Explanation:

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

NEW QUESTION 130

- (Topic 5)

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

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

Answer: B

Explanation:

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

NEW QUESTION 133

- (Topic 5)

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

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

Answer: B

NEW QUESTION 137

- (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 139

- (Topic 5)

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

Which feature of Kubernetes should you use to accomplish this?

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

Answer: A

Explanation:

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

NEW QUESTION 144

- (Topic 5)

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

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

Answer: B

Explanation:

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

NEW QUESTION 146

- (Topic 5)

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 148

- (Topic 5)

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

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

Answer: AEF

Explanation:

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

NEW QUESTION 151

- (Topic 5)

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

What should you do?

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

Answer: A

NEW QUESTION 152

- (Topic 5)

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

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

Answer: C

NEW QUESTION 157

- (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 158

- (Topic 5)

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

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

Answer: B

Explanation:

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

NEW QUESTION 159

- (Topic 5)

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

- A. Store time-series data from the game servers in Google Bigtable, and view it using Google Data Studio.
- B. Output custom metrics to Stackdriver from the game servers, and create a Dashboard in StackdriverMonitoring Console to view them.
- C. Schedule BigQuery load jobs to ingest analytics files uploaded to Cloud Storage every ten minutes, and visualize the results in Google Data Studio.
- D. Insert the KPIs into Cloud Datastore entities, and run ad hoc analysis and visualizations of them in Cloud Datalab.

Answer: A

Explanation:

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

NEW QUESTION 160

- (Topic 5)

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

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

Answer: C

Explanation:

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

NEW QUESTION 161

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 165

- (Topic 5)

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

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

Answer: B

NEW QUESTION 167

- (Topic 5)

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

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

Answer: C

Explanation:

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

NEW QUESTION 170

- (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 175

- (Topic 5)

A development team at your company has created a dockerized HTTPS web application. You need to deploy the application on Google Kubernetes Engine (GKE) and make sure that the application scales automatically.

How should you deploy to GKE?

- A. Use the Horizontal Pod Autoscaler and enable cluster autoscaling.
- B. Use an Ingress resource to loadbalance the HTTPS traffic.
- C. Use the Horizontal Pod Autoscaler and enable cluster autoscaling on the Kubernetes cluster.
- D. Use a Service resource of type LoadBalancer to load-balance the HTTPS traffic.
- E. Enable autoscaling on the Compute Engine instance group.
- F. Use an Ingress resource to load balance the HTTPS traffic.
- G. Enable autoscaling on the Compute Engine instance group.
- H. Use a Service resource of type LoadBalancer to load-balance the HTTPS traffic.

Answer: B

Explanation:

<https://cloud.google.com/kubernetes-engine/docs/tutorials/http-balancer> <https://cloud.google.com/kubernetes-engine/docs/concepts/network-overview#ext-lb>

NEW QUESTION 179

- (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 182

- (Topic 5)

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

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

Answer: B

Explanation:

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

NEW QUESTION 184

- (Topic 5)

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

- A. Use `kubectl set image deployment/echo-deployment <new-image>`
- B. Use the rolling update functionality of the Instance Group behind the Kubernetes cluster

- C. Update the deployment yaml file with the new container imag
- D. Use kubectl delete deployment/echo-deployment and kubectl create -f <yaml-file>
- E. Update the service yaml file which the new container imag
- F. Use kubectl delete service/echoserviceand kubectl create -f <yaml-file>

Answer: A

Explanation:

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

NEW QUESTION 187

- (Topic 5)

Your company has a Google Cloud project that uses BigQuery for data warehousing. There are some tables that contain personally identifiable information (PII). Only the compliance team may access the PII. The other information in the tables must be available to the data science team. You want to minimize cost and the time it takes to assign appropriate access to the tables. What should you do?

- A. * 1 From the dataset where you have the source data, create views of tables that you want to share, excluding PII* 2 Assign an appropriate project-level IAM role to the members of the data science team 3 Assign access controls to the dataset that contains the view
- B. * 1 From the dataset where you have the source data, create materialized views of tables that you want to share, excluding PII* 2 Assign an appropriate project-level IAM role to the members of the data science team 3. Assign access controls to the dataset that contains the view.
- C. * 1 Create a dataset for the data science team* 2 Create views of tables that you want to share, excluding PII* 3 Assign an appropriate project-level IAM role to the members of the data science team 4 Assign access controls to the dataset that contains the view* 5 Authorize the view to access the source dataset
- D. * 1. Create a dataset for the data science team.* 2. Create materialized views of tables that you want to share, excluding PII* 3. Assign an appropriate project-level IAM role to the members of the data science team * 4 Assign access controls to the dataset that contains the view* 5 Authorize the view to access the source dataset

Answer: C

Explanation:

This option can help minimize cost and time by using views and authorized datasets. Views are virtual tables defined by a SQL query that can exclude PII columns from the source tables. Views do not incur storage costs and do not duplicate data. Authorized datasets are datasets that have access to another dataset's data without granting direct access to individual users or groups. By creating a dataset for the data science team and creating views of tables that exclude PII, you can share only the relevant information with the team. By assigning an appropriate project-level IAM role to the members of the data science team, you can grant them access to the BigQuery service and resources. By assigning access controls to the dataset that contains the view, you can grant them access to query the views. By authorizing the view to access the source dataset, you can enable the view to read data from the source tables without exposing PII. The other options are not optimal for this scenario, because they either use materialized views instead of views, which incur storage costs and duplicate data (B, D), or do not create a separate dataset for the data science team, which makes it harder to manage access controls (A). References:

? <https://cloud.google.com/bigquery/docs/views>

? <https://cloud.google.com/bigquery/docs/authorized-datasets>

NEW QUESTION 192

- (Topic 7)

For this question, refer to the TerraEarth case study. You are asked to design a new architecture for the ingestion of the data of the 200,000 vehicles that are connected to a cellular network. You want to follow Google-recommended practices.

Considering the technical requirements, which components should you use for the ingestion of the data?

- A. Google Kubernetes Engine with an SSL Ingress
- B. Cloud IoT Core with public/private key pairs
- C. Compute Engine with project-wide SSH keys
- D. Compute Engine with specific SSH keys

Answer: A

Explanation:

<https://cloud.google.com/solutions/iot-overview> <https://cloud.google.com/iot/quotas>

NEW QUESTION 196

- (Topic 8)

You need to implement a network ingress for a new game that meets the defined business and technical requirements. Mountkirk Games wants each regional game instance to be located in multiple Google Cloud regions. What should you do?

- A. Configure a global load balancer connected to a managed instance group running Compute Engine instances.
- B. Configure kubernetes with a global load balancer and Google Kubernetes Engine.
- C. Configure a global load balancer with Google Kubernetes Engine.
- D. Configure Ingress for Anthos with a global load balancer and Google Kubernetes Engine.

Answer: A

NEW QUESTION 198

- (Topic 8)

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

- A. Create network load balancer
- B. Use preemptible Compute Engine instances.
- C. Create network load balancer
- D. Use non-preemptible Compute Engine instances.

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

Answer: D

NEW QUESTION 200

- (Topic 8)

Mountkirk Games wants to limit the physical location of resources to their operating Google Cloud regions. What should you do?

- A. Configure an organizational policy which constrains where resources can be deployed.
- B. Configure IAM conditions to limit what resources can be configured.
- C. Configure the quotas for resources in the regions not being used to 0.
- D. Configure a custom alert in Cloud Monitoring so you can disable resources as they are created in other regions.

Answer: A

NEW QUESTION 202

- (Topic 8)

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

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

Answer: C

NEW QUESTION 207

- (Topic 8)

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

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

Answer: A

NEW QUESTION 209

- (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 211

- (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 213

.....

Relate Links

100% Pass Your Professional-Cloud-Architect Exam with Exambible Prep Materials

<https://www.exambible.com/Professional-Cloud-Architect-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>