Exam Questions Professional-Cloud-Architect

Google Certified Professional - Cloud Architect (GCP)

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

NEW QUESTION 1
- (Exam Topic 5)
Your company just finished a rapid lift and shift to Google Compute Engine for your compute needs. You have another 9 months to design and deploy a more cloud-native solution. Specifically, you want a system that is no-ops and auto-scaling. Which two compute products should you choose? Choose 2 answers

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

Answer: BC

Explanation:
B: With Container Engine, Google will automatically deploy your cluster for you, update, patch, secure the nodes.
Kubernetes Engine's cluster autoscaler automatically resizes clusters based on the demands of the workloads you want to run.
C: Solutions like Datastore, BigQuery, AppEngine, etc are truly NoOps.
App Engine by default scales the number of instances running up and down to match the load, thus providing consistent performance for your app at all times while minimizing idle instances and thus reducing cost.
Note: At a high level, NoOps means that there is no infrastructure to build out and manage during usage of the platform. Typically, the compromise you make with NoOps is that you lose control of the underlying infrastructure.
References:
https://www.quora.com/How-well-does-Google-Container-Engine-support-Google-Cloud-Platform%E2%80%9

NEW QUESTION 2
- (Exam Topic 5)
One of your primary business objectives is being able to trust the data stored in your application. You want to log all changes to the application data. How can you design your logging system to verify authenticity of your logs?

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

Answer: C

Explanation:
https://cloud.google.com/storage/docs/access-logs
References: https://cloud.google.com/logging/docs/reference/tools/gcloud-logging

NEW QUESTION 3
- (Exam Topic 5)
You are using Cloud CDN to deliver static HTTP(S) website content hosted on a Compute Engine instance group. You want to improve the cache hit ratio. What should you do?

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

Answer: A

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

NEW QUESTION 4
- (Exam 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 5
- (Exam 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 6
- (Exam Topic 5)
You have been asked to select the storage system for the click-data of your company's large portfolio of websites. This data is streamed in from a custom website analytics package at a typical rate of 6,000 clicks per minute, with bursts of up to 8,500 clicks per second. It must been stored for future analysis by your data science and user experience teams. Which storage infrastructure should you choose?

A. Google Cloud SQL
B. Google Cloud Bigtable
C. Google Cloud Storage
D. Google cloud Datastore

Answer: C

Explanation:
https://cloud.google.com/bigquery/docs/loading-data-cloud-storage

NEW QUESTION 7
- (Exam Topic 5)
Your company's test suite is a custom C++ application that runs tests throughout each day on Linux virtual machines. The full test suite takes several hours to complete, running on a limited number of on premises servers reserved for testing. Your company wants to move the testing infrastructure to the cloud, to reduce the amount of time it takes to fully test a change to the system, while changing the tests as little as possible. Which cloud infrastructure should you recommend?

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

Answer: B

Explanation:
https://cloud.google.com/compute/docs/instance-groups/
Google Compute Engine enables users to launch virtual machines (VMs) on demand. VMs can be launched from the standard images or custom images created by users.
Managed instance groups offer autoscaling capabilities that allow you to automatically add or remove instances from a managed instance group based on increases or decreases in load. Autoscaling helps your applications gracefully handle increases in traffic and reduces cost when the need for resources is lower.

NEW QUESTION 8
- (Exam 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 bucke
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 bucke
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 9
- (Exam Topic 5)
You want to enable your running Google Container Engine cluster to scale as demand for your application changes.
What should you do?

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

Answer: B

Explanation:
https://cloud.google.com/kubernetes-engine/docs/concepts/cluster-autoscaler Cluster autoscaling
--enable-autoscaling
Enables autoscaling for a node pool.
Enables autoscaling in the node pool specified by --node-pool or the default node pool if --node-pool is not provided.
Where:
--max-nodes=MAX_NODES
Maximum number of nodes in the node pool.
Maximum number of nodes to which the node pool specified by --node-pool (or default node pool if unspecified) can scale.

NEW QUESTION 10
- (Exam 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 Stackdnver Trace in order to break down the request latencies at each microservice.

Answer: D

Explanation:
https://cloud.google.com/trace/docs/overview

NEW QUESTION 10
- (Exam 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 13
- (Exam Topic 5)
You are analyzing and defining business processes to support your startup's trial usage of GCP, and you don't yet know what consumer demand for your product will be. Your manager requires you to minimize GCP service costs and adhere to Google best practices. What should you do?

A. Utilize free tier and sustained use discount
B. Provision a staff position for service cost management.
C. Utilize free tier and sustained use discount
D. Provide training to the team about service cost management.
E. Utilize free tier and committed use discount
F. Provision a staff position for service cost management.
G. Utilize free tier and committed use discount
H. Provide training to the team about service cost management.

Answer: D

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

NEW QUESTION 17
- (Exam Topic 5)
Your company runs several databases on a single MySQL instance. They need to take backups of a specific database at regular intervals. The backup activity needs to complete as quickly as possible and cannot be allowed to impact disk performance. How should you configure the storage?

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

Answer: B

Explanation:

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

NEW QUESTION 19
- (Exam Topic 5)
You have an outage in your Compute Engine managed instance group: all instance keep restarting after 5 seconds. You have a health check configured, but autoscaling is disabled. Your colleague, who is a Linux expert, offered to look into the issue. You need to make sure that he can access the VMs. What should you do?

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

Answer: C

Explanation:
https://cloud.google.com/compute/docs/instance-groups/autohealing-instances-in-migs
Health checks used for autohealing should be conservative so they don't preemptively delete and recreate your instances. When an autohealer health check is too aggressive, the autohealer might mistake busy instances for failed instances and unnecessarily restart them, reducing availability

NEW QUESTION 21
- (Exam Topic 5)
You need to evaluate your team readiness for a new GCP project. You must perform the evaluation and create a skills gap plan incorporates the business goal of cost optimization. Your team has deployed two GCP projects successfully to date. What should you do?

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

Answer: B

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

NEW QUESTION 25
- (Exam 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 29
- (Exam 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 31
- (Exam 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 34
- (Exam Topic 5)
Your company is forecasting a sharp increase in the number and size of Apache Spark and Hadoop jobs being run on your local datacenter You want to utilize the cloud to help you scale this upcoming demand with the least amount of operations work and code change. Which product should you use?

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

Answer: B

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

NEW QUESTION 38
- (Exam Topic 5)
You have a Python web application with many dependencies that requires 0.1 CPU cores and 128 MB of memory to operate in production. You want to monitor and maximize machine utilization. You also to reliably deploy new versions of the application. Which set of steps should you take?

A. Perform the following:1) Create a managed instance group with f1-micro type machines.2) Use a startup script to clone the repository, check out the production branch, install the dependencies, and start the Python app.3) Restart the instances to automatically deploy new production releases.
B. Perform the following:1) Create a managed instance group with n1-standard-1 type machines.2) Build a Compute Engine image from the production branch that contains all of the dependencies and automatically starts the Python app.3) Rebuild the Compute Engine image, and update the instance template to deploy new production releases.
C. Perform the following:1) Create a Kubernetes Engine cluster with n1-standard-1 type machines.2) Build a Docker image from the production branch with all of the dependencies, and tag it with the version number.3) Create a Kubernetes Deployment with the imagePullPolicy set to "IfNotPresent" in the staging namespace, and then promote it to the production namespace after testing.
D. Perform the following:1) Create a Kubernetes Engine cluster with n1-standard-4 type machines.2) Build a Docker image from the master branch will all of the dependencies, and tag it with "latest".3) Create a Kubernetes Deployment in the default namespace with the imagePullPolicy set to "Always". Restart the pods to automatically deploy new production releases.

Answer: D

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

NEW QUESTION 41
- (Exam Topic 5)
You have created several preemptible Linux virtual machine instances using Google Compute Engine. You want to properly shut down your application before the virtual machines are preempted. What should you do?

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

Answer: C

NEW QUESTION 43
- (Exam 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=10and 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 45
- (Exam 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 Loggin
D. Consult logs for Kubernetes Engine and Cloud SQL.
E. In the GCP Console, navigate to Cloud SQ
F. Restore the latest backu
G. Use kubect1 to restart all pods.

Answer: C

NEW QUESTION 46
- (Exam Topic 5)
Your customer wants to do resilience testing of their authentication layer. This consists of a regional managed instance group serving a public REST API that reads from and writes to a Cloud SQL instance.
What should you do?

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

Answer: C

NEW QUESTION 50
- (Exam 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 dis
B. Create an instance template from the snapsho
C. Create an autoscaled managed instance group from the instance template.
D. Create a snapshot of the existing dis
E. Create a custom image from the snapsho
F. Create an autoscaled managed instance group from the custom image.
G. Create a custom image from the existing dis
H. Create an instance template from the custom image.Create an autoscaled managed instance group from the instance template.
I. Create an instance template from the existing dis
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 53
- (Exam 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 57
- (Exam 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 GCPConsole 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 59
- (Exam 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 63
......

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

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

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

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

Money Back Guarantee

Professional-Cloud-Architect Practice Exam Features:

* Professional-Cloud-Architect Questions and Answers Updated Frequently

* Professional-Cloud-Architect Practice Questions Verified by Expert Senior Certified Staff

* Professional-Cloud-Architect Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* Professional-Cloud-Architect Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year