Professional-Cloud-Architect Dumps

Google Certified Professional - Cloud Architect (GCP)

https://www.certleader.com/Professional-Cloud-Architect-dumps.html
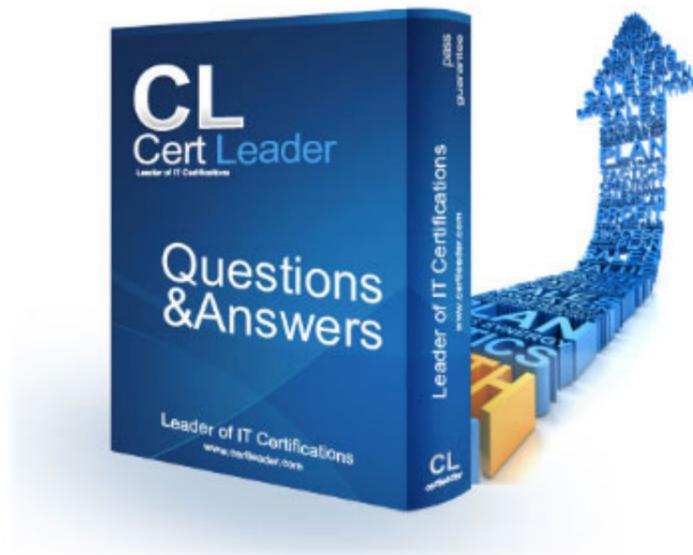
NEW QUESTION 1
- (Exam Topic 1)
For this question, refer to the Mountkirk Games case study.
Mountkirk Games has deployed their new backend on Google Cloud Platform (GCP). You want to create a thorough testing process for new versions of the backend before they are released to the public. You want the 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 2
- (Exam Topic 1)
For this question, refer to the Mountkirk Games case study.
Mountkirk Games wants to set up a real-time analytics platform for their new game. The new platform must meet their technical requirements. Which combination of Google technologies will meet all of their requirements?

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

Answer: B

Explanation:
A real time requires Stream / Messaging so Pub/Sub, Analytics by Big Query.
Ingest millions of streaming events per second from anywhere in the world with Cloud Pub/Sub, powered by Google's unique, high-speed private network. Process the streams with Cloud Dataflow to ensure reliable, exactly-once, low-latency data transformation. Stream the transformed data into BigQuery, the cloud-native data warehousing service, for immediate analysis via SQL or popular visualization tools.
From scenario: They plan to deploy the game's backend on Google Compute Engine so they can capture streaming metrics, run intensive analytics.
Requirements for Game Analytics Platform

‣ Dynamically scale up or down based on game activity

‣ Process incoming data on the fly directly from the game servers

‣ Process data that arrives late because of slow mobile networks

‣ Allow SQL queries to access at least 10 TB of historical data

‣ Process files that are regularly uploaded by users' mobile devices

‣ Use only fully managed services
References: https://cloud.google.com/solutions/big-data/stream-analytics/

NEW QUESTION 3
- (Exam 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 4
- (Exam Topic 2)
Your agricultural division is experimenting with fully autonomous vehicles.
You want your architecture to promote strong security during vehicle operation. Which two architecture should you consider?
Choose 2 answers:

A. Treat every micro service call between modules on the vehicle as untrusted.
B. Require IPv6 for connectivity to ensure a secure address space.
C. Use a trusted platform module (TPM) and verify firmware and binaries on boot.
D. Use a functional programming language to isolate code execution cycles.
E. Use multiple connectivity subsystems for redundancy.

F. Enclose the vehicle's drive electronics in a Faraday cage to isolate chips.

Answer: AC

NEW QUESTION 5
- (Exam Topic 4)
The current Dress4win system architecture has high latency to some customers because it is located in one data center.
As of a future evaluation and optimizing for performance in the cloud, Dresss4win wants to distribute it's system architecture to multiple locations when Google cloud platform. Which approach should they use?

A. Use regional managed instance groups and a global load balancer to increase performance because the regional managed instance group can grow instances in each region separately based on traffic.
B. Use a global load balancer with a set of virtual machines that forward the requests to a closer group ofvirtual machines managed by your operations team.
C. Use regional managed instance groups and a global load balancer to increase reliability by providing automatic failover between zones in different regions.
D. Use a global load balancer with a set of virtual machines that forward the requests to a closer group of virtual machines as part of a separate managed instance groups.

Answer: A

NEW QUESTION 6
- (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 7
- (Exam Topic 5)
Your customer is receiving reports that their recently updated Google App Engine application is taking approximately 30 seconds to load for some of their users. This behavior was not reported before the update. What strategy should you take?

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

Answer: C

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

NEW QUESTION 8
- (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 9
- (Exam 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 10
- (Exam Topic 5)
Your company wants to try out the cloud with low risk. They want to archive approximately 100 TB of their log data to the cloud and test the analytics features available to them there, while also retaining that data as a long-term disaster recovery backup. Which two steps should they take? Choose 2 answers

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

Answer: AE

NEW QUESTION 10
- (Exam 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 branc
D. Build and deploy changes to production for 10% of users before doing a complete rollout.
E. Use Jenkins to monitor tags in the repositor
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 13
- (Exam 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 error for
com/Altostrat/CloakedServlet.class
        at com.google.appengine.runtime.Request.process
-d36f818a24b8cf1d (Request.java)
        at
sun.security.util.ManifestEntryVerifier.verify
(ManifestEntryVerifier.java:210)
        at java.util.jar.JarVerifier.processEntry
(JarVerifier.java:218)
        at java.util.jar.JarVerifier.update
(JarVerifier.java:205)
        at
java.util.jar.JarVerifiersVerifierStream.read
(JarVerifier.java:428)
        at sun.misc.Resource.getBytes
(Resource.java:124)
        at java.net.URL.ClassLoader.defineClass
(URLClassLoader.java:273)
        at sun.reflect.GeneratedMethodAccessor5.invoke
(Unknown Source)
        at
sun.reflect.DelegatingMethodAccessorImpl.invoke
(DelegatingMethodAccessorImpl.java:43)
        at java.lang.reflect.Method.invoke
(Method.java:616)
        at java.lang.ClassLoader.loadClass
(ClassLoader.java:266)
```

What should you do?

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

Answer: B

NEW QUESTION 16
- (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 17
- (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 18

- (Exam Topic 5)
You are using Cloud Shell and need to install a custom utility for use in a few weeks. Where can you store the file so it is in the default execution path and persists across sessions?

A. ~/bin
B. Cloud Storage
C. /google/scripts
D. /usr/local/bin

Answer: D

Explanation:
https://medium.com/google-cloud/no-localhost-no-problem-using-google-cloud-shell-as-my-full-time-developm

NEW QUESTION 22
- (Exam Topic 5)
You are designing a large distributed application with 30 microservices. Each of your distributed microservices needs to connect to a database back-end. You want to store the credentials securely. Where should you store the credentials?

A. In the source code
B. In an environment variable
C. In a secret management system
D. In a config file that has restricted access through ACLs

Answer: C

Explanation:
https://cloud.google.com/docs/authentication/production#providing_credentials_to_your_application

NEW QUESTION 26
- (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 28
- (Exam Topic 5)
You created a pipeline that can deploy your source code changes to your infrastructure in instance groups for self healing.
One of the changes negatively affects your key performance indicator. You are not sure how to fix it and investigation could take up to a week. What should you do

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

Answer: C

NEW QUESTION 29
- (Exam Topic 5)
Auditors visit your teams every 12 months and ask to review all the Google Cloud Identity and Access Management (Cloud IAM) policy changes in the previous 12 months. You want to streamline and expedite the analysis and audit process. What should you do?

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

Answer: D

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

NEW QUESTION 30
- (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 34
- (Exam Topic 5)
You want to establish a Compute Engine application in a single VPC across two regions. The application must communicate over VPN to an on-premises network. How should you deploy the VPN?

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

Answer: C

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

NEW QUESTION 39
- (Exam 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 42
- (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 45
- (Exam Topic 5)
Your company is migrating its on-premises data center into the cloud. As part of the migration, you want to integrate Kubernetes Engine for workload orchestration. Parts of your architecture must also be PCI DSScompliant.
Which of the following is most accurate?

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

Answer: D

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

NEW QUESTION 50
- (Exam Topic 5)
Your company places a high value on being responsive and meeting customer needs quickly. Their primary business objectives are release speed and agility. You want to reduce the chance of security errors being accidentally introduced. Which two actions can you take? Choose 2 answers

A. Ensure every code check-in is peer reviewed by a security SME.
B. Use source code security analyzers as part of the CI/CD pipeline.
C. Ensure you have stubs to unit test all interfaces between components.
D. Enable code signing and a trusted binary repository integrated with your CI/CD pipeline.
E. Run a vulnerability security scanner as part of your continuous-integration /continuous-delivery (CI/CD) pipeline.

Answer: BE

Explanation:
https://docs.microsoft.com/en-us/vsts/articles/security-validation-cicd-pipeline?view=vsts

NEW QUESTION 51
- (Exam 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 52
- (Exam Topic 5)
You set up an autoscaling instance group to serve web traffic for an upcoming launch. After configuring the instance group as a backend service to an HTTP(S) load balancer, you notice that virtual machine (VM) instances are being terminated and re-launched every minute. The instances do not have a public IP address. You have verified the appropriate web response is coming from each instance using the curl command. You want to ensure the backend is configured correctly. What should you do?

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

Answer: C

Explanation:
https://cloud.google.com/vpc/docs/using-firewalls
The best practice when configuration a health check is to check health and serve traffic on the same port. However, it is possible to perform health checks on one port, but serve traffic on another. If you do use two different ports, ensure that firewall rules and services running on instances are configured appropriately. If you run health checks and serve traffic on the same port, but decide to switch ports at some point, be sure to update both the backend service and the health check.
Backend services that do not have a valid global forwarding rule referencing it will not be health checked and will have no health status.
References: https://cloud.google.com/compute/docs/load-balancing/http/backend-service

NEW QUESTION 55
- (Exam Topic 5)
Your web application has several VM instances running within a VPC. You want to restrict communications between instances to only the paths and ports you authorize, but you don't want to rely on static IP addresses or subnets because the app can autoscale. How should you restrict communications?

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

Answer: B

NEW QUESTION 60
- (Exam Topic 5)
Your company's user-feedback portal comprises a standard LAMP stack replicated across two zones. It is deployed in the us-central1 region and uses autoscaled managed instance groups on all layers, except the database. Currently, only a small group of select customers have access to the portal. The portal meets a 99.99% availability SLA under these conditions However next quarter, your company will be making the portal available to all users, including unauthenticated users. You need to develop a resiliency testing strategy to ensure the system maintains the SLA once they introduce additional user load. What should you do?

A. Capture existing users input, and replay captured user load until autoscale is triggered on all layer

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

Answer: A

NEW QUESTION 64
- (Exam Topic 5)
Your organization wants to control IAM policies for different departments independently, but centrally. Which approach should you take?

A. Multiple Organizations with multiple Folders
B. Multiple Organizations, one for each department
C. A single Organization with Folder for each department
D. A single Organization with multiple projects, each with a central owner

Answer: C

Explanation:
Folders are nodes in the Cloud Platform Resource Hierarchy. A folder can contain projects, other folders, or a combination of both. You can use folders to group projects under an organization in a hierarchy. For example, your organization might contain multiple departments, each with its own set of GCP resources. Folders allow you to group these resources on a per-department basis. Folders are used to group resources that share common IAM policies. While a folder can contain multiple folders or resources, a given folder or resource can have exactly one parent.
References: https://cloud.google.com/resource-manager/docs/creating-managing-folders

NEW QUESTION 69
- (Exam Topic 5)
Your marketing department wants to send out a promotional email campaign. The development team wants to minimize direct operation management. They project a wide range of possible customer responses, from 100 to 500,000 click-throughs per day. The link leads to a simple website that explains the promotion and collects user information and preferences. Which infrastructure should you recommend? (CHOOSE TWO)

A. Use Google App Engine to serve the website and Google Cloud Datastore to store user data.
B. Use a Google Container Engine cluster to serve the website and store data to persistent disk.
C. Use a managed instance group to serve the website and Google Cloud Bigtable to store user data.
D. Use a single compute Engine virtual machine (VM) to host a web server, backed by Google Cloud SQL.

Answer: AC

Explanation:
Reference: https://cloud.google.com/storage-options/ References: https://cloud.google.com/storage-options/

NEW QUESTION 72
- (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 76
- (Exam 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 fil
B. Use gsutil to upload the files.
C. Supply the encryption key using gcloud confi
D. Use gsutil to upload the files to that bucket.
E. Use gsutil to upload the files, and use the flag --encryption-key to supply the encryption key.
F. Use gsutil to create a bucket, and use the flag --encryption-key to supply the encryption ke
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 78
- (Exam Topic 5)
You need to develop procedures to verify resilience of disaster recovery for remote recovery using GCP. Your production environment is hosted on-premises. You need to establish a secure, redundant connection between your on premises network and the GCP network.
What should you do?

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

Answer: B

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

NEW QUESTION 80
- (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 84
- (Exam Topic 5)
Google Cloud Platform resources are managed hierarchically using organization, folders, and projects. When Cloud Identity and Access Management (IAM) policies exist at these different levels, what is the effective policy at a particular node of the hierarchy?

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

Answer: B

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

NEW QUESTION 89
- (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 90
- (Exam Topic 5)
You have been engaged by your client to lead the migration of their application infrastructure to GCP. One of their current problems is that the on-premises high performance SAN is requiring frequent and expensive upgrades to keep up with the variety of workloads that are identified as follows: 20TB of log archives retained for legal reasons; 500 GB of VM boot/data volumes and templates; 500 GB of image thumbnails; 200 GB of customer session state data that allows customers to restart sessions even if off-line for several days.
Which of the following best reflects your recommendations for a cost-effective storage allocation?

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

Answer: D

Explanation:
https://cloud.google.com/compute/docs/disks

NEW QUESTION 93
- (Exam 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 BigQuer
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 Is' 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 97
- (Exam Topic 5)
A recent audit that a new network was created in Your GCP project. In this network, a GCE instance has an SSH port open the world. You want to discover this network's origin. What should you do?

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

Answer: C

NEW QUESTION 99
- (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 104
- (Exam 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 108
- (Exam Topic 5)
An application development team believes their current logging tool will not meet their needs for their new cloud-based product. They want a bettor tool to capture errors and help them analyze their historical log data. You want to help them find a solution that meets their needs, what should you do?

A. Direct them to download and install the Google StackDriver logging agent.
B. Send them a list of online resources about logging best practices.
C. Help them define their requirements and assess viable logging tools.
D. Help them upgrade their current tool to take advantage of any new features.

Answer: C

Explanation:
Help them define their requirements and assess viable logging tools. They know the requirements and the existing tools' problems. While it's true StackDriver Logging and Error Reporting possibly meet all their requirements, there might be other tools also meet their need. They need you to provide expertise to make assessment for new tools, specifically, logging tools that can "capture errors and help them analyze their historical log data".
References: https://cloud.google.com/logging/docs/agent/installation

NEW QUESTION 109
- (Exam Topic 5)
Your application needs to process credit card transactions. You want the smallest scope of Payment Card Industry (PCI) compliance without compromising the ability to analyze transactional data and trends relating to which payment methods are used. How should you design your architecture?

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

Answer: A

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

NEW QUESTION 114
- (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 118
- (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 123
- (Exam Topic 5)
You are deploying an application on App Engine that needs to integrate with an on-premises database. For security purposes, your on-premises database must not be accessible through the public Internet. What should you do?

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

Answer: D

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

NEW QUESTION 125
- (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 128
- (Exam Topic 5)
Your company pushes batches of sensitive transaction data from its application server VMs to Cloud Pub/Sub for processing and storage. What is the Google-recommended way for your application to authenticate to the required Google Cloud services?

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

Answer: A

NEW QUESTION 131
- (Exam 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 cluste
B. Use Deployment Manager to create the deployment.
C. Use gcloud to create a Kubernetes cluste
D. Use kubect1 to create the deployment.
E. Use kubect1 to create a Kubernetes cluste
F. Use Deployment Manager to create the deployment.
G. Use kubect1 to create a Kubernetes cluste
H. Use kubect1 to create the deployment.

Answer: B

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

NEW QUESTION 135
- (Exam 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 137
- (Exam Topic 5)
During a high traffic portion of the day, one of your relational databases crashes, but the replica is never promoted to a master. You want to avoid this in the future. What should you do?

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

Answer: D

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

NEW QUESTION 138
- (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 141
- (Exam Topic 5)
Your development team has installed a new Linux kernel module on the batch servers in Google Compute Engine (GCE) virtual machines (VMs) to speed up the nightly batch process. Two days after the installation, 50% of web application deployed in the same
nightly batch run. You want to collect details on the failure to pass back to the development team. Which three actions should you take? Choose 3 answers

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

Answer: ACE

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

NEW QUESTION 146
- (Exam Topic 5)
Your company wants to track whether someone is present in a meeting room reserved for a scheduled meeting. There are 1000 meeting rooms across 5 offices on 3 continents. Each room is equipped with a motion sensor that reports its status every second. The data from the motion detector includes only a sensor ID and several different discrete items of information. Analysts will use this data, together with information about account owners and office locations. Which database type should you use?

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

Answer: B

Explanation:
Relational databases were not designed to cope with the scale and agility challenges that face modern applications, nor were they built to take advantage of the commodity storage and processing power available today.
NoSQL fits well for:
Developers are working with applications that create massive volumes of new, rapidly changing data types — structured, semi-structured, unstructured and polymorphic data.

NEW QUESTION 150
- (Exam Topic 5)
Your company is building a new architecture to support its data-centric business focus. You are responsible for setting up the network. Your company's mobile and web-facing applications will be deployed on-premises, and all data analysis will be conducted in GCP. The plan is to process and load 7 years of archived .csv files totaling 900 TB of data and then continue loading 10 TB of data daily. You currently have an existing 100-MB internet connection.
What actions will meet your company's needs?

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

Answer: B

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

NEW QUESTION 154
......

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

    All transactions are protected by VeriSign!

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

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