Google

Exam Questions Professional-Cloud-Architect

Google Certified Professional - Cloud Architect (GCP)

NEW QUESTION 1
- (Exam Topic 1)
For this question, refer to the Mountkirk Games case study.
Mountkirk Games wants you to design their new testing strategy. How should the test coverage differ from their existing backends on the other platforms?

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

Answer: A

Explanation:
 From Scenario:
A few of their games were more popular than expected, and they had problems scaling their application servers, MySQL databases, and analytics tools.
Requirements for Game Analytics Platform include: Dynamically scale up or down based on game activity

NEW QUESTION 2
- (Exam Topic 1)
For this question, refer to the Mountkirk Games case study.
Mountkirk Games' gaming servers are not automatically scaling properly. Last month, they rolled out a new feature, which suddenly became very popular. A record number of users are trying to use the service, but many of them are getting 503 errors and very slow response times. What should they investigate first?

A. Verify that the database is online.
B. Verify that the project quota hasn't been exceeded.
C. Verify that the new feature code did not introduce any performance bugs.
D. Verify that the load-testing team is not running their tool against production.

Answer: B

Explanation:
* 503 is service unavailable error. If the database was online everyone would get the 503 error. https://cloud.google.com/docs/quota#capping_usage

NEW QUESTION 3
- (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 4
- (Exam 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 framewor
D. Focus on an API for the public.
E. Use Google App Engine with the Swagger (open API Specification) framewor
F. Focus on an API for the public.
G. Use Google Container Engine with a Django Python containe
H. Focus on an API for the public.

I. Use Google Container Engine with a Tomcat container with the Swagger (Open API Specification) framewor
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.-1712523
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 5
- (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 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)
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 8
- (Exam Topic 5)
You need to ensure reliability for your application and operations by supporting reliable task a scheduling for compute on GCP. Leveraging Google best practices, what should you do?

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

Answer: B

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

NEW QUESTION 9
- (Exam 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 10
- (Exam 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 10
- (Exam Topic 5)
Your company is using BigQuery as its enterprise data warehouse. Data is distributed over several Google Cloud projects. All queries on BigQuery need to be billed on a single project. You want to make sure that no query costs are incurred on the projects that contain the data. Users should be able to query the datasets, but not edit them.
How should you configure users' access roles?

A. Add all users to a grou
B. Grant the group the role of BigQuery user on the billing project and BigQuery dataViewer on the projects that contain the data.
C. Add all users to a grou
D. Grant the group the roles of BigQuery dataViewer on the billing project and BigQuery user on the projects that contain the data.
E. Add all users to a grou
F. Grant the group the roles of BigQuery jobUser on the billing project and BigQuery dataViewer on the projects that contain the data.
G. Add all users to a grou
H. Grant the group the roles of BigQuery dataViewer on the billing project and BigQuery jobUser on the projects that contain the data.

Answer: A

Explanation:
Reference: https://cloud.google.com/bigquery/docs/running-queries

NEW QUESTION 12
- (Exam 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 14
- (Exam Topic 5)
You have an App Engine application that needs to be updated. You want to test the update with production traffic before replacing the current application version. What should you do?

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

Answer: B

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

NEW QUESTION 16
- (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 18
- (Exam Topic 5)
You want your Google Kubernetes Engine cluster to automatically add or remove nodes based on CPUload. What should you do?

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

Answer: B

NEW QUESTION 21
- (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 23
- (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 28
- (Exam Topic 5)
A development manager is building a new application He asks you to review his requirements and identify what cloud technologies he can use to meet them. The application must
* 1. Be based on open-source technology for cloud portability
* 2. Dynamically scale compute capacity based on demand
* 3. Support continuous software delivery
* 4. Run multiple segregated copies of the same application stack
* 5. Deploy application bundles using dynamic templates
* 6. Route network traffic to specific services based on URL
Which combination of technologies will meet all of his requirements?

A. Google Container Engine, Jenkins, and Helm
B. Google Container Engine and Cloud Load Balancing
C. Google Compute Engine and Cloud Deployment Manager
D. Google Compute Engine, Jenkins, and Cloud Load Balancing

Answer: A

Explanation:
Helm for managing Kubernetes
Kubernetes can base on the URL to route traffic to different location (path) https://cloud.google.com/kubernetes-engine/docs/tutorials/http-balancer eg.apiVersion: networking.k8s.io/v1beta1
kind: Ingress metadata:
name: fanout-ingress spec:
rules:
- http: paths:
- path: /* backend: serviceName: web servicePort: 8080
- path: /v2/* backend: serviceName: web2 servicePort: 8080

NEW QUESTION 29
- (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 34
- (Exam Topic 5)
Your company creates rendering software which users can download from the company website. Your company has customers all over the world. You want to minimize latency for all your customers. You want to follow Google-recommended practices.
How should you store the files?

A. Save the files in a Multi-Regional Cloud Storage bucket.
B. Save the files in a Regional Cloud Storage bucket, one bucket per zone of the region.
C. Save the files in multiple Regional Cloud Storage buckets, one bucket per zone per region.
D. Save the files in multiple Multi-Regional Cloud Storage buckets, one bucket per multi-region.

Answer: A

Explanation:
https://cloud.google.com/storage/docs/locations#location-mr

NEW QUESTION 35
- (Exam 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 kubect1 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 kubect1 delete deployment/ echo-deployment and kubect1 create –f <yaml-file>
E. Update the service yaml file which the new container imag
F. Use kubect1 delete service/echoservice and kubect1 create –f <yaml-file>

Answer: A

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

NEW QUESTION 40
- (Exam 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 rul
F. Configure SSL proxy balancing.

Answer: B

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

NEW QUESTION 44
- (Exam Topic 5)
You have an application that makes HTTP requests to Cloud Storage. Occasionally the requests fail with HTTP status codes of 5xx and 429.
How should you handle these types of errors?

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

Answer: A

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

NEW QUESTION 47
- (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 51
- (Exam Topic 5)
You are developing a globally scaled frontend for a legacy streaming backend data API. This API expects events in strict chronological order with no repeat data for proper processing.
Which products should you deploy to ensure guaranteed-once FIFO (first-in, first-out) delivery of data?

A. Cloud Pub/Sub alone
B. Cloud Pub/Sub to Cloud DataFlow
C. Cloud Pub/Sub to Stackdriver
D. Cloud Pub/Sub to Cloud SQL

Answer: B

Explanation:
Reference https://cloud.google.com/pubsub/docs/ordering

NEW QUESTION 52

- (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 55
- (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 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 62
- (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 65
- (Exam 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 Stackdriver Monitoring 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 inCloud Datala

Answer: A

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

NEW QUESTION 69
- (Exam Topic 5)
Your applications will be writing their logs to BigQuery for analysis. Each application should have its own table.
Any logs older than 45 days should be removed. You want to optimize storage and follow Google recommended practices. What should you do?

A. Configure the expiration time for your tables at 45 days
B. Make the tables time-partitioned, and configure the partition expiration at 45 days
C. Rely on BigQuery's default behavior to prune application logs older than 45 days
D. Create a script that uses the BigQuery command line tool (bq) to remove records older than 45 days

Answer: B

Explanation:
https://cloud.google.com/bigquery/docs/managing-partitioned-tables

NEW QUESTION 71
- (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 74
- (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 75
- (Exam Topic 5)
You are running a cluster on Kubernetes Engine to serve a web application. Users are reporting that a specific part of the application is not responding anymore. You notice that all pods of your deployment keep restarting after 2 seconds. The application writes logs to standard output. You want to inspect the logs to find the cause of the issue. Which approach can you take?

A. Review the Stackdriver logs for each Compute Engine instance that is serving as a node in the cluster.
B. Review the Stackdriver logs for the specific Kubernetes Engine container that is serving the unresponsive part of the application.
C. Connect to the cluster using gcloud credentials and connect to a container in one of the pods to read the logs.
D. Review the Serial Port logs for each Compute Engine instance that is serving as a node in the cluster.

Answer: B

NEW QUESTION 76
- (Exam Topic 5)

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

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

Answer: B

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

NEW QUESTION 78
- (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 81
- (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 85
- (Exam Topic 5)
You are designing a mobile chat application. You want to ensure people cannot spoof chat messages, by providing a message were sent by a specific user. What should you do

A. Tag messages client side with the originating user identifier and the destination user.
B. Encrypt the message client side using block-based encryption with a shared key.
C. Use public key infrastructure (PKI) to encrypt the message client side using the originating user's private key.
D. Use a trusted certificate authority to enable SSL connectivity between the client application and the server.

Answer: C

NEW QUESTION 90
- (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 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 95
- (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 96
- (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

At com.google.appengine.runtime.Request.pr

At

Sun.securityutil.manifestENtryVerifier.ver

At java . net . URLClassLoader . defineCla

At sun . reflect. GeneratedMethodAccessors

At

Sun.reflect . DelegatingMethodAccesorImp1.

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 100
- (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 105
- (Exam 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 autoscalin
B. Use an Ingress resource to loadbalance the HTTPS traffic.
C. Use the Horizontal Pod Autoscaler and enable cluster autoscaling on the Kubernetes cluste
D. Use a Service resource of type LoadBalancer to load-balance the HTTPS traffic.
E. Enable autoscaling on the Compute Engine instance grou
F. Use an Ingress resource to load balance the HTTPS traffic.
G. Enable autoscaling on the Compute Engine instance grou
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 106
- (Exam Topic 5)
Your company has decided to build a backup replica of their on-premises user authentication PostgreSQL database on Google Cloud Platform. The database is 4 TB, and large updates are frequent. Replication requires private address space communication. Which networking approach should you use?

A. Google Cloud Dedicated Interconnect
B. Google Cloud VPN connected to the data center network
C. A NAT and TLS translation gateway installed on-premises
D. A Google Compute Engine instance with a VPN server installed connected to the data center network

Answer: A

Explanation:

https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations
Google Cloud Dedicated Interconnect provides direct physical connections and RFC 1918 communication between your on-premises network and Google's network. Dedicated Interconnect enables you to transfer large amounts of data between networks, which can be more cost effective than purchasing additional bandwidth over the public Internet or using VPN tunnels.
Benefits:
≫ Traffic between your on-premises network and your VPC network doesn't traverse the public Internet.
Traffic traverses a dedicated connection with fewer hops, meaning there are less points of failure where traffic might get dropped or disrupted.
≫ Your VPC network's internal (RFC 1918) IP addresses are directly accessible from your on-premises network. You don't need to use a NAT device or VPN tunnel to reach internal IP addresses. Currently, you can only reach internal IP addresses over a dedicated connection. To reach Google external IP addresses, you must use a separate connection.
≫ You can scale your connection to Google based on your needs. Connection capacity is delivered over one or more 10 Gbps Ethernet connections, with a maximum of eight connections (80 Gbps total per interconnect).
≫ The cost of egress traffic from your VPC network to your on-premises network is reduced. A dedicated connection is generally the least expensive method if you have a high-volume of traffic to and from Google's network.
References: https://cloud.google.com/interconnect/docs/details/dedicated

NEW QUESTION 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 111
- (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 116
- (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 121
- (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 122
- (Exam 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/pci-dss-compliance-ingcp#

NEW QUESTION 126
- (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 129
- (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 131
- (Exam 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 132
- (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 134
- (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 139
- (Exam 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 Storag
B. Use Cloud Datalab to explore and clean your data.
C. Upload your files into Cloud Storag
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 144
- (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 145
......

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

Professional-Cloud-Architect Practice Exam Features:

* Professional-Cloud-Architect Questions and Answers Updated Frequently

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

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

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

100% Actual & Verified — Instant Download, Please Click
Order The Professional-Cloud-Architect Practice Test Here