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' 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 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 1)
For this question, refer to the Mountkirk Games case study
Mountkirk Games needs to create a repeatable and configurable mechanism for deploying isolated application environments. Developers and testers can access each other's environments and resources, but they cannot access staging or production resources. The staging environment needs access to some services from production.
What should you do to isolate development environments from staging and production?

A. Create a project for development and test and another for staging and production.
B. Create a network for development and test and another for staging and production.
C. Create one subnetwork for development and another for staging and production.
D. Create one project for development, a second for staging and a third for production.

Answer: D

NEW QUESTION 4
- (Exam Topic 1)
For this question, refer to the Mountkirk Games case study.
Mountkirk Games wants to set up a continuous delivery pipeline. Their architecture includes many small services that they want to be able to update and roll back quickly. Mountkirk Games has the following requirements:
• Services are deployed redundantly across multiple regions in the US and Europe.
• Only frontend services are exposed on the public internet.
• They can provide a single frontend IP for their fleet of services.
• Deployment artifacts are immutable. Which set of products should they use?

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

Answer: C

NEW QUESTION 5
- (Exam Topic 2)
For this question, refer to the TerramEarth case study.
TerramEarth has equipped unconnected trucks with servers and sensors to collet telemetry data. Next year they want to use the data to train machine learning models. They want to store this data in the cloud while reducing costs. What should they do?

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

Answer: D

Explanation:
Coldline Storage is the best choice for data that you plan to access at most once a year, due to its slightly lower availability, 90-day minimum storage duration, costs for data access, and higher per-operation costs. For example:
Cold Data Storage - Infrequently accessed data, such as data stored for legal or regulatory reasons, can be stored at low cost as Coldline Storage, and be available when you need it.
Disaster recovery - In the event of a disaster recovery event, recovery time is key. Cloud Storage provides low latency access to data stored as Coldline Storage.
References: https://cloud.google.com/storage/docs/storage-classes

NEW QUESTION 6
- (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 7
- (Exam Topic 5)
You have an application that will run on Compute Engine. You need to design an architecture that takes into account a disaster recovery plan that requires your application to fail over to another region in case of a regional outage. What should you do?

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

Answer: C

NEW QUESTION 8
- (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 9
- (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 10
- (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 10
- (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 14
- (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 15
- (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 19
- (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 22
- (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 24
- (Exam Topic 5)
Your customer is moving their corporate applications to Google Cloud Platform. The security team wants detailed visibility of all projects in the organization. You provision the Google Cloud Resource Manager and set up yourself as the org admin. What Google Cloud Identity and Access Management (Cloud IAM) roles should you give to the security team'?

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

Answer: B

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

NEW QUESTION 28
- (Exam Topic 5)
You want to enable your running Google Container Engine cluster to scale as demand for your application changes.
What should you do?

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

Answer: B

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

NEW QUESTION 32
- (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 35
- (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 38
- (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 39
- (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 43
- (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 46
- (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 50
- (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 54
- (Exam Topic 5)
A production database virtual machine on Google Compute Engine has an ext4-formatted persistent disk for data files The database is about to run out of storage space How can you remediate the problem with the least amount of downtime?

A. In the Cloud Platform Console, increase the size of the persistent disk and use the resize2fs command in Linux.
B. Shut down the virtual machine, use the Cloud Platform Console to increase the persistent disk size, then restart the virtual machine.
C. In the Cloud Platform Console, increase the size of the persistent disk and verify the new space is ready to use with the fdisk command in Linux.
D. In the Cloud Platform Console, create a new persistent disk attached to the virtual machine, format and mount it, and configure the database service to move the files to the new disk.
E. In the Cloud Platform Console, create a snapshot of the persistent disk, restore the snapshot to a new larger disk, unmount the old disk, mount the new disk, and restart the database service.

Answer: A

Explanation:
On Linux instances, connect to your instance and manually resize your partitions and file systems to use the additional disk space that you added.
Extend the file system on the disk or the partition to use the added space. If you grew a partition on your disk, specify the partition. If your disk does not have a partition table, specify only the disk ID.
sudo resize2fs /dev/[DISK_ID][PARTITION_NUMBER]
where [DISK_ID] is the device name and [PARTITION_NUMBER] is the partition number for the device where you are resizing the file system.
References: https://cloud.google.com/compute/docs/disks/add-persistent-disk

NEW QUESTION 56
- (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 59
- (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 63
- (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 66
- (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 70
- (Exam Topic 5)
You need to set up Microsoft SQL Server on GCP. Management requires that there's no downtime in case of a data center outage in any of the zones within a GCP region. What should you do?

A. Configure a Cloud SQL instance with high availability enabled.
B. Configure a Cloud Spanner instance with a regional instance configuration.
C. Set up SQL Server on Compute Engine, using Always On Availability Groups using Windows Failover Clusterin
D. Place nodes in different subnets.
E. Set up SQL Server Always On Availability Groups using Windows Failover Clusterin
F. Place nodes in different zones.

Answer: D

Explanation:
https://cloud.google.com/sql/docs/sqlserver/configure-ha

NEW QUESTION 72
- (Exam Topic 5)
Your company has successfully migrated to the cloud and wants to analyze their data stream to optimize operations. They do not have any existing code for this analysis, so they are exploring all their options. These options include a mix of batch and stream processing, as they are running some hourly jobs and live-processing some data as it comes in. Which technology should they use for this?

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

Answer: B

Explanation:
Dataflow is for processing both the Batch and Stream.
Cloud Dataflow is a fully-managed service for transforming and enriching data in stream (real time) and batch (historical) modes with equal reliability and expressiveness -- no more complex workarounds or compromises needed.
References: https://cloud.google.com/dataflow/

NEW QUESTION 73
- (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 74
- (Exam Topic 5)
You need to develop procedures to test a disaster plan for a mission-critical application. You want to use Google-recommended practices and native capabilities within GCP.
What should you do?

A. Use Deployment Manager to automate service provisionin
B. Use Activity Logs to monitor and debug your tests.
C. Use Deployment Manager to automate provisionin
D. Use Stackdriver to monitor and debug your tests.
E. Use gcloud scripts to automate service provisionin
F. Use Activity Logs monitor and debug your tests.
G. Use automated scripts to automate service provisionin
H. Use Activity Logs monitor and debug your tests.

Answer: B

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

NEW QUESTION 78
- (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 80
- (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 82
- (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 87
- (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 92
- (Exam Topic 5)

Your company wants to start using Google Cloud resources but wants to retain their on-premises Active Directory domain controller for identity management. What should you do?

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

Answer: B

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

NEW QUESTION 95
- (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 96
- (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 99
......

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