

Google Professional Cloud Architect.43q

Passing Score: 800

Time Limit: 120 min



Website: <https://vceplus.com>

VCE to PDF Converter: <https://vceplus.com/vce-to-pdf/>

Facebook: <https://www.facebook.com/VCE.For.All.VN/>

Twitter : https://twitter.com/VCE_Plus

<https://vceplus.com/>

Google Certified Professional – Cloud Architect

Testlet 1

Company Overview

JencoMart is a global retailer with over 10,000 stores in 16 countries. The stores carry a range of goods, such as groceries, tires, and jewelry. One of the company's core values is excellent customer service. In addition, they recently introduced an environmental policy to reduce their carbon output by 50% over the next 5 years.

Company Background

JencoMart started as a general store in 1931, and has grown into one of the world's leading brands, known for great value and customer service. Over time, the company transitioned from only physical stores to a stores and online hybrid model, with 25% of sales online. Currently, JencoMart has little presence in Asia, but considers that market key for future growth.

Solution Concept

JencoMart wants to migrate several critical applications to the cloud but has not completed a technical review to determine their suitability for the cloud and the engineering required for migration. They currently host all of these applications on infrastructure that is at its end of life and is no longer supported.

Existing Technical Environment

JencoMart hosts all of its applications in 4 data centers: 3 in North American and 1 in Europe; most applications are dual-homed.

JencoMart understands the dependencies and resource usage metrics of their on-premises architecture.

Application: Customer loyalty portal

LAMP (Linux, Apache, MySQL and PHP) application served from the two JencoMart-owned U.S. data centers.

Database

- Oracle Database stores user profiles
 - 20 TB
 - Complex table structure
 - Well maintained, clean data
 - Strong backup strategy
- PostgreSQL database stores user credentials
 - Single-homed in US West
 - No redundancy
 - Backed up every 12 hours
 - 100% uptime service level agreement (SLA)
 - Authenticates all users

Compute

- 30 machines in US West Coast, each machine has:
 - Twin, dual core CPUs
 - 32 GB of RAM
 - Twin 250 GB HDD (RAID 1)
- 20 machines in US East Coast, each machine has:
 - Single, dual-core CPU
 - 24 GB of RAM
 - Twin 250 GB HDD (RAID 1)

Storage

- Access to shared 100 TB SAN in each location
- Tape backup every week

Business Requirements

- Optimize for capacity during peak periods and value during off-peak periods
 - Guarantee service availability and support
 - Reduce on-premises footprint and associated financial and environmental impact
 - Move to outsourcing model to avoid large upfront costs associated with infrastructure purchase
- Expand services into Asia

Technical Requirements

- Assess key application for cloud suitability
 - Modify applications for the cloud
 - Move applications to a new infrastructure
 - Leverage managed services wherever feasible
 - Sunset 20% of capacity in existing data centers
- Decrease latency in Asia

CEO Statement

JencoMart will continue to develop personal relationships with our customers as more people access the web. The future of our retail business is in the global market and the connection between online and in-store experiences. As a large, global company, we also have a responsibility to the environment through “green” initiatives and policies.

CTO Statement

The challenges of operating data centers prevent focus on key technologies critical to our long-term success. Migrating our data services to a public cloud infrastructure will allow us to focus on big data and machine learning to improve our service to customers. **CFO Statement**

Since its founding, JencoMart has invested heavily in our data services infrastructure. However, because of changing market trends, we need to outsource our infrastructure to ensure our long-term success. This model will allow us to respond to increasing customer demand during peak periods and reduce costs.

QUESTION 1

JencoMart has decided to migrate user profile storage to Google Cloud Datastore and the application servers to Google Compute Engine (GCE). During the migration, the existing infrastructure will need access to Datastore to upload the data.

What service account key-management strategy should you recommend?

<https://vceplus.com/>

- A. Provision service account keys for the on-premises infrastructure and for the GCE virtual machines (VMs)
- B. Authenticate the on-premises infrastructure with a user account and provision service account keys for the VMs
- C. Provision service account keys for the on-premises infrastructure and use Google Cloud Platform (GCP) managed keys for the VMs
- D. Deploy a custom authentication service on GCE/Google Container Engine (GKE) for the on-premises infrastructure and use GCP managed keys for the VMs

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Migrating data to Google Cloud Platform

Let's say that you have some data processing that happens on another cloud provider and you want to transfer the processed data to Google Cloud Platform. You can use a service account from the virtual machines on the external cloud to push the data to Google Cloud Platform. To do this, you must create and download a service account key when you create the service account and then use that key from the external process to call the Cloud Platform APIs.

References: https://cloud.google.com/iam/docs/understanding-service-accounts#migrating_data_to_google_cloud_platform

QUESTION 2

JencoMart has built a version of their application on Google Cloud Platform that serves traffic to Asia. You want to measure success against their business and technical goals.

Which metrics should you track?

- A. Error rates for requests from Asia
- B. Latency difference between US and Asia
- C. Total visits, error rates, and latency from Asia
- D. Total visits and average latency for users from Asia
- E. The number of character sets present in the database

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:


From scenario:

Business Requirements include: Expand services into Asia

Technical Requirements include: Decrease latency in Asia

QUESTION 3





The migration of JencoMart's application to Google Cloud Platform (GCP) is progressing too slowly. The infrastructure is shown in the diagram. You want to maximize throughput.

What are three potential bottlenecks? Choose 3 answers.

- A. A single VPN tunnel, which limits throughput
- B. A tier of Google Cloud Storage that is not suited for this task
- C. A copy command that is not suited to operate over long distances
- D. Fewer virtual machines (VMs) in GCP than on-premises machines
- E. A separate storage layer outside the VMs, which is not suited for this task
- F. Complicated internet connectivity between the on-premises infrastructure and GCP

Correct Answer: ACE

Section: [none]

Explanation

Explanation/Reference:

QUESTION 4

JencoMart wants to move their User Profiles database to Google Cloud Platform.

Which Google Database should they use?

- A. Cloud Spanner
- B. Google BigQuery
- C. Google Cloud SQL
- D. Google Cloud Datastore



<https://vceplus.com/>

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Common workloads for Google Cloud Datastore:

- User profiles
- Product catalogs

- Game state

References: <https://cloud.google.com/storage-options/>
<https://cloud.google.com/datastore/docs/concepts/overview> **Testlet 1**

Company Overview

Mountkirk Games makes online, session-based, multiplayer games for the most popular mobile platforms.

Company Background

Mountkirk Games builds all of their games with some server-side integration, and has historically used cloud providers to lease physical servers. A few of their games were more popular than expected, and they had problems scaling their application servers, MySQL databases, and analytics tools. Mountkirk's current model is to write game statistics to files and send them through an ETL tool that loads them into a centralized MySQL database for reporting.

Solution Concept

Mountkirk Games is building a new game, which they expect to be very popular. They plan to deploy the game's backend on Google Compute Engine so they can capture streaming metrics, run intensive analytics, and take advantage of its autoscaling server environment and integrate with a managed NoSQL database.

Technical Requirements

Requirements for Game Backend Platform

1. Dynamically scale up or down based on game activity
2. Connect to a managed NoSQL database service
3. Run customize Linux distro

Requirements for Game Analytics Platform

1. Dynamically scale up or down based on game activity
2. Process incoming data on the fly directly from the game servers
3. Process data that arrives late because of slow mobile networks
4. Allow SQL queries to access at least 10 TB of historical data
5. Process files that are regularly uploaded by users' mobile devices
6. Use only fully managed services

CEO Statement

Our last successful game did not scale well with our previous cloud provider, resulting in lower user adoption and affecting the game's reputation. Our investors want more key performance indicators (KPIs) to evaluate the speed and stability of the game, as well as other metrics that provide deeper insight into usage patterns so we can adapt the game to target users.

CTO Statement

Our current technology stack cannot provide the scale we need, so we want to replace MySQL and move to an environment that provides autoscaling, low latency load balancing, and frees us up from managing physical servers.

CFO Statement

We are not capturing enough user demographic data, usage metrics, and other KPIs. As a result, we do not engage the right users, we are not confident that our marketing is targeting the right users, and we are not selling enough premium Blast-Ups inside the games, which dramatically impacts our revenue.

QUESTION 1

Mountkirk Games has deployed their new backend on Google Cloud Platform (GCP). You want to create a through 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

Correct Answer: A

Section: [none]

Explanation



Explanation/Reference:

Explanation:

From scenario: Requirements for Game Backend Platform

1. Dynamically scale up or down based on game activity
2. Connect to a managed NoSQL database service
3. Run customize Linux distro

QUESTION 2

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?



<https://vceplus.com/>

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

Correct Answer: D

Section: [none]

Explanation



Explanation/Reference:

Explanation:

Google Cloud Functions is a serverless environment to build and connect cloud services.

Google Cloud Pub/Sub brings the scalability, flexibility, and reliability of enterprise message-oriented middleware to the cloud. By providing many-to-many, asynchronous messaging that decouples senders and receivers, it allows for secure and highly available communication between independently written applications. Google Cloud Pub/Sub delivers low-latency, durable messaging that helps developers quickly integrate systems hosted on the Google Cloud Platform and externally.

Incorrect Answers:

A: Cloud Dataflow is a fully-managed service for transforming and enriching data in stream (real time) and batch (historical) modes.

C: Store your private Docker container images on Cloud Platform for fast, scalable retrieval and deployment. Container Registry is a private Docker repository that works with popular continuous delivery systems. It runs on Cloud Platform to provide consistent uptime on an infrastructure protected by Google's security. You pay only for storage and internet egress you use, there is no per-image fee.

Reference: <https://cloud.google.com/load-balancing/> <https://cloud.google.com/solutions/ansible-with-spinnaker-tutorial> <http://blog.armory.io/what-is-immutable-infrastructure/> <https://cloud.google.com/compute/docs/load-balancing/http/>

QUESTION 3

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

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation: 503 is service unavailable error. If the database was online everyone would get the 503 error.

QUESTION 4

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

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

References: <https://cloud.google.com/appengine/docs/standard/go/creating-separate-dev-environments>

Question Set 1

QUESTION 1

A recent audit revealed that a new network was created in your GCP project. In this network, a GCE instance has an SSH port open to 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 section. Set category to Data Access and search for Create VM entry
- C. In the Logging section of the console, specify GCE Network as the logging section. Search for the Create Insert entry
- D. Connect to the GCE instance using project SSH keys. Identify previous logins in system logs, and match these with the project owners list

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Incorrect Answers:

A: To use the Stackdriver alerting console we must first set up alerting policies.

B: Data access logs only contain read-only operations.

Audit logs help you determine who did what, where, and when.

Cloud Audit Logging returns two types of logs:

- Admin activity logs
- Data access logs: Contains log entries for operations that perform read-only operations do not modify any data, such as get, list, and aggregated list methods.

QUESTION 2

You want to make a copy of a production Linux virtual machine in the US-Central region. You want to manage and replace the copy easily if there are changes on the production virtual machine. You will deploy the copy as a new instance in a different project in the US-East region.

What steps must you take?



<https://vceplus.com/>

- A. Use the Linux dd and netcat commands to copy and stream the root disk contents to a new virtual machine instance in the US-East region.
- B. Create a snapshot of the root disk and select the snapshot as the root disk when you create a new virtual machine instance in the US-East region.
- C. Create an image file from the root disk with Linux dd command, create a new virtual machine instance in the US-East region
- D. Create a snapshot of the root disk, create an image file in Google Cloud Storage from the snapshot, and create a new virtual machine instance in the US-East region using the image file the root disk.

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:



QUESTION 3

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 location. After the backup is complete, use gsutil to move the backup to Google Cloud Storage.
- C. 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.
- D. 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

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

References: <https://github.com/mvarrieur/MySQL-backup-to-Google-Cloud-Storage>
<https://cloud.google.com/storage/docs/gcs-fuse>

QUESTION 4

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 is 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

Correct Answer: BEF

Section: [none]

Explanation



Explanation/Reference:

QUESTION 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

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

QUESTION 6

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

Correct Answer: BE

Section: [none]

Explanation

Explanation/Reference:



QUESTION 7

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 enableautoscaling max-nodes=10`
- C. Update the existing Container Engine cluster with the following command: `gcloud alpha container clusters update mycluster - --enableautoscaling - --min-nodes=1 - --max-nodes=10`
- D. Create a new Container Engine cluster with the following command: `gcloud alpha container clusters create mycluster - --`

`enableautoscaling - --min-nodes=1 - --max-nodes=10` and
redploy your application

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

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.

Incorrect Answers:

C, D: Warning: Do not use Alpha Clusters or alpha features for production workloads.

Note: You can experiment with Kubernetes alpha features by creating an alpha cluster. Alpha clusters are short-lived clusters that run stable Kubernetes releases with all Kubernetes APIs and features enabled. Alpha clusters are designed for advanced users and early adopters to experiment with workloads that take advantage of new features before those features are production-ready. You can use Alpha clusters just like normal Kubernetes Engine clusters.

References: <https://cloud.google.com/sdk/gcloud/reference/container/clusters/create>

QUESTION 8

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-through 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 2 answers.

- 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, backend by Google Cloud SQL.

Correct Answer: AC


Section: [none]

Explanation

Explanation/Reference:

Explanation:





References: <https://cloud.google.com/storage-options/>

QUESTION 9

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 cloudnative 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 Container Engine with containers
- C. Google App Engine Standard Environment
- D. Compute Engine with custom instance types
- E. Compute Engine with managed instance groups

Correct Answer: BC

Section: [none]

Explanation

Explanation/Reference:

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%99s-NoOps-claim>

QUESTION 10

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

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Write a log entry. If the log does not exist, it is created. You can specify a severity for the log entry, and you can write a structured log entry by specifying --payloadtype=json and writing your message as a JSON string:

```
gcloud logging write LOG STRING gcloud logging write LOG  
JSON-STRING --payload-type=json
```

References: <https://cloud.google.com/logging/docs/reference/tools/gcloud-logging>

QUESTION 11

Your company has decided to make a major revision of their API in order to create better experiences for their developers. They need to keep the old version of the API available and deployable, while allowing new customers and testers to try out the new API. They want to keep the same SSL and DNS records in place to serve both APIs.

What should they do?

- A. Configure a new load balancer for the new version of the API
- B. Reconfigure old clients to use a new endpoint for the new API
- C. Have the old API forward traffic to the new API based on the path
- D. Use separate backend pools for each API path behind the load balancer

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

QUESTION 12

Your company plans to migrate a multi-petabyte data set to the cloud. The data set must be available 24hrs a day. Your business analysts have experience only with using a SQL interface.

How should you store the data to optimize it for ease of analysis?



<https://vceplus.com/>

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

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

BigQuery is Google's serverless, highly scalable, low cost enterprise data warehouse designed to make all your data analysts productive. Because there is no infrastructure to manage, you can focus on analyzing data to find meaningful insights using familiar SQL and you don't need a database administrator.

BigQuery enables you to analyze all your data by creating a logical data warehouse over managed, columnar storage as well as data from object storage, and spreadsheets.

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

QUESTION 13

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

Correct Answer: ADE

Section: [none]

Explanation

Explanation/Reference:

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

QUESTION 14

A news feed web service has the following code running on Google App Engine. During peak load, users report that they can see news articles they already viewed.

What is the most likely cause of this problem?



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

```
app = Flask(__name__)
sessions = {}
```

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

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

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

    return news.render(news_articles)
```

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

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

Section: [none]

Explanation

Explanation/Reference:

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

QUESTION 15

An application development team believes their current logging tool will not meet their needs for their new cloud-based product. They want a better 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

Correct Answer: A

Section: [none]

Explanation



Explanation/Reference:

Explanation:

The Stackdriver Logging agent streams logs from your VM instances and from selected third party software packages to Stackdriver Logging. Using the agent is optional but we recommend it. The agent runs under both Linux and Microsoft Windows.

Note: 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/docs/agent/installation>

QUESTION 16

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

Correct Answer: AC

Section: [none]

Explanation

Explanation/Reference:

QUESTION 17

To reduce costs, the Director of Engineering has required all developers to move their development infrastructure resources from on-premises virtual machines (VMs) to Google Cloud Platform. These resources go through multiple start/stop events during the day and require state to persist. You have been asked to design the process of running a development environment in Google Cloud while providing cost visibility to the finance department.

Which two steps should you take? Choose 2 answers.

- A. Use the -no-auto-delete flag on all persistent disks and stop the VM
- B. Use the -auto-delete flag on all persistent disks and terminate the VM
- C. Apply VM CPU utilization label and include it in the BigQuery billing export
- D. Use Google BigQuery billing export and labels to associate cost to groups
- E. Store all state into local SSD, snapshot the persistent disks, and terminate the VM
- F. Store all state in Google Cloud Storage, snapshot the persistent disks, and terminate the VM

Correct Answer: CE

Section: [none]

Explanation

Explanation/Reference:

Explanation:

C: Billing export to BigQuery enables you to export your daily usage and cost estimates automatically throughout the day to a BigQuery dataset you specify.

Labels applied to resources that generate usage metrics are forwarded to the billing system so that you can break down your billing charges based upon label criteria. For example, the Compute Engine service reports metrics on VM instances. If you deploy a project with 2,000 VMs, each of which is labeled distinctly, then only the first 1,000 label maps seen within the 1 hour window will be preserved.

E: You cannot stop an instance that has a local SSD attached. Instead, you must migrate your critical data off of the local SSD to a persistent disk or to another instance before you delete the instance completely.

You can stop an instance temporarily so you can come back to it at a later time. A stopped instance does not incur charges, but all of the resources that are attached to the instance will still be charged. Alternatively, if you are done using an instance, delete the instance and its resources to stop incurring charges.

References:

<https://cloud.google.com/billing/docs/how-to/export-data-bigquery>

<https://cloud.google.com/compute/docs/instances/stopping-or-deleting-an-instance>

QUESTION 18

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

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

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/>

QUESTION 19

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

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

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>

QUESTION 20

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

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

<https://www.sans.org/reading-room/whitepapers/compliance/ways-reduce-pci-dss-audit-scope-tokenizing-cardholder-data-33194>

QUESTION 21

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 have 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

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Google Cloud Bigtable is a scalable, fully-managed NoSQL wide-column database that is suitable for both real-time access and analytics workloads.

Good for:

- Low-latency read/write access
- High-throughput analytics
- Native time series support

Common workloads: ▪ IoT, finance, adtech

- Personalization, recommendations
- Monitoring
- Geospatial datasets ▪

Graphs

Incorrect Answers:

C: Google Cloud Storage is a scalable, fully-managed, highly reliable, and cost-efficient object / blob store.

Is good for:

- Images, pictures, and videos
- Objects and blobs ▪

Unstructured data

D: Google Cloud Datastore is a scalable, fully-managed NoSQL document database for your web and mobile applications.

Is good for:

- Semi-structured application data

- Hierarchical data ▪

- Durable key-value data ▪

Common workloads:

- User profiles

- Product catalogs

- Game state

References: <https://cloud.google.com/storage-options/>

QUESTION 22

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



<https://vceplus.com/>

- C. Schedule a cron script using `gsutil ls -lr gs://backups/**` to find and remove items older than 90 days
- D. Schedule a cron script using `gsutil ls -l gs://backups/**` to find and remove items older than 90 days and schedule it with cron

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

QUESTION 23

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

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

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>

QUESTION 24

The database administration team has asked you to help them improve the performance of their new database server running on Google Compute Engine. The database is for importing and normalizing their performance statistics and is built with MySQL running on Debian Linux. They have an n1-standard-8 virtual machine with 80 GB of SSD persistent disk.

What should they change to get better performance from this system?

- A. Increase the virtual machine's memory to 64 GB
- B. Create a new virtual machine running PostgreSQL
- C. Dynamically resize the SSD persistent disk to 500 GB
- D. Migrate their performance metrics warehouse to BigQuery
- E. Modify all of their batch jobs to use bulk inserts into the database

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

QUESTION 25

You want to optimize the performance of an accurate, real-time, weather-charting application. The data comes from 50,000 sensors sending 10 readings a second, in the format of a timestamp and sensor reading.

Where should you store the data?

- A. Google BigQuery
- B. Google Cloud SQL
- C. Google Cloud Bigtable
- D. Google Cloud Storage

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Google Cloud Bigtable is a scalable, fully-managed NoSQL wide-column database that is suitable for both real-time access and analytics workloads.

Good for:

- Low-latency read/write access
- High-throughput analytics
- Native time series support

Common workloads: ▪ IoT, finance, adtech

- Personalization, recommendations
- Monitoring
- Geospatial datasets
- Graphs

References: <https://cloud.google.com/storage-options/>

QUESTION 26

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 layers. At the same time, terminate all resources in one of the zones
- B. 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
- C. Expose the new system to a larger group of users, and increase group size each day until autoscale logic is triggered on all layers. At the same time, terminate random resources on both zones
- D. Capture existing users input, and replay captured user load until resource utilization crosses 80%. Also, derive estimated number of users based on existing user’s usage of the app, and deploy enough resources to handle 200% of expected load

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

QUESTION 27

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

Correct Answer: CE

Section: [none]

Explanation

Explanation/Reference:

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/>

QUESTION 28

Your solution is producing performance bugs in production that you did not see in staging and test environments. You want to adjust your test and deployment procedures to avoid this problem in the future.

What should you do?

- A. Deploy fewer changes to production
- B. Deploy smaller changes to production
- C. Increase the load on your test and staging environments
- D. Deploy changes to a small subset of users before rolling out to production

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

QUESTION 29

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

What should you do?

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

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

References: https://cloud.google.com/trace/docs/quickstart#find_a_trace

QUESTION 30

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



Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Take regular snapshots of your database system.

If your database system lives on a Compute Engine persistent disk, you can take snapshots of your system each time you upgrade. If your database system goes down or you need to roll back to a previous version, you can simply create a new persistent disk from your desired snapshot and make that disk the boot disk for a new Compute Engine instance. Note that, to avoid data corruption, this approach requires you to freeze the database system's disk while taking a snapshot.

Reference: <https://cloud.google.com/solutions/disaster-recovery-cookbook>

QUESTION 31

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

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Stackdriver Logging provides you with the ability to filter, search, and view logs from your cloud and open source application services. Allows you to define metrics based on log contents that are incorporated into dashboards and alerts. Enables you to export logs to BigQuery, Google Cloud Storage, and Pub/Sub.

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

QUESTION 32

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

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

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>

QUESTION 33

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

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

QUESTION 34

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

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

References: <https://cloud.google.com/kms/docs/secret-management>

QUESTION 35

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

Correct Answer: BF

Section: [none]

Explanation

Explanation/Reference:

Explanation:

What are two business risks of migrating to Cloud Deployment Manager?

Risk 1 Cloud Deployment Manager APIs could be deprecated in the future.

Risk 2 Cloud Deployment Manager only supports automation of Google Cloud resources.



<https://vceplus.com/>

