

Google Certified Professional - Cloud Architect.73q

Number: GoogleCloudArchitect

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://www.vceplus.com/>

Google Certified Professional – Cloud Architect (English)

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

A few days after JencoMart migrates the user credentials database to Google Cloud Platform and shuts down the old server, the new database server stops responding to SSH connections. It is still serving database requests to the application servers correctly.

What three steps should you take to diagnose the problem? Choose 3 answers.

- A. Delete the virtual machine (VM) and disks and create a new one
- B. Delete the instance, attach the disk to a new VM, and investigate
- C. Take a snapshot of the disk and connect to a new machine to investigate
- D. Check inbound firewall rules for the network the machine is connected to
- E. Connect the machine to another network with very simple firewall rules and investigate
- F. Print the Serial Console output for the instance for troubleshooting, activate the interactive console, and investigate

Correct Answer: CDF

Section: [none]

Explanation

Explanation/Reference:

Explanation:

D: Handling "Unable to connect on port 22" error message

Possible causes include:

- There is no firewall rule allowing SSH access on the port. SSH access on port 22 is enabled on all Compute Engine instances by default. If you have disabled access, SSH from the Browser will not work. If you run sshd on a port other than 22, you need to enable the access to that port with a custom firewall rule. ▪ The firewall rule allowing SSH access is enabled, but is not configured to allow connections from GCP Console services. Source IP addresses for browserbased SSH sessions are dynamically allocated by GCP Console and can vary from session to session.

F: Handling "Could not connect, retrying..." error

You can verify that the daemon is running by navigating to the serial console output page and looking for output lines prefixed with the accounts-from-metadata: string. If you are using a standard image but you do not see these output prefixes in the serial console output, the daemon might be stopped. Reboot the instance to restart the daemon.

References:

<https://cloud.google.com/compute/docs/ssh-in-browser> <https://cloud.google.com/compute/docs/ssh-in-browser> **QUESTION 2**

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?

- 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 Kubernetes 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 3

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.





<https://www.vceplus.com/> 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 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

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?

- 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

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>

Testlet 1

Company Overview

Mountkirk Games makes online, session-based, multiplayer games for mobile platforms. They build all of their games using some server-side integration. Historically, they have used cloud providers to lease physical servers.

Due to the unexpected popularity of some of their games, they have had problems scaling their global audience, application servers, MySQL databases, and analytics tools.

Their 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.

Business Requirements

- Increase to a global footprint.
 - Improve uptime – downtime is loss of players.
 - Increase efficiency of the cloud resources we use. ▪
- Reduce latency to all customers.

Technical Requirements

Requirements for Game Backend Platform

- Dynamically scale up or down based on game activity.
 - Connect to a transactional database service to manage user profiles and game state.
 - Store game activity in a timeseries database service for future analysis.
 - As the system scales, ensure that data is not lost due to processing backlogs. ▪
- Run hardened Linux distro.

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 queries to access at least 10 TB of historical data
- Process files that are regularly uploaded by users' mobile devices

Executive 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. Additionally, 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.

QUESTION 1

For this question, refer to the Mountkirk Games case study. Mountkirk Games wants you to design a way to test the analytics platform's resilience to changes in mobile network latency. What should you do?

- A. Deploy failure injection software to the game analytics platform that can inject additional latency to mobile client analytics traffic.
- B. Build a test client that can be run from a mobile phone emulator on a Compute Engine virtual machine, and run multiple copies in Google Cloud Platform regions all over the world to generate realistic traffic.
- C. Add the ability to introduce a random amount of delay before beginning to process analytics files uploaded from mobile devices.

- D. Create an opt-in beta of the game that runs on players' mobile devices and collects response times from analytics endpoints running in Google Cloud Platform regions all over the world.

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

QUESTION 2

For this question, refer to the Mountkirk Games case study. You need to analyze and define the technical architecture for the database workloads for your company, Mountkirk Games. Considering the business and technical requirements, what should you do?

- A. Use Cloud SQL for time series data, and use Cloud Bigtable for historical data queries.
- B. Use Cloud SQL to replace MySQL, and use Cloud Spanner for historical data queries.
- C. Use Cloud Bigtable to replace MySQL, and use BigQuery for historical data queries.
- D. Use Cloud Bigtable for time series data, use Cloud Spanner for transactional data, and use BigQuery for historical data queries.

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

QUESTION 3

For this question, refer to the Mountkirk Games case study. Which managed storage option meets Mountkirk's technical requirement for storing game activity in a time series database service?



<https://www.vceplus.com/>

- A. Cloud Bigtable
- B. Cloud Spanner
- C. BigQuery
- D. Cloud Datastore

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

QUESTION 4

For this question, refer to the Mountkirk Games case study. You need to analyze and define the technical architecture for the compute workloads for your company, Mountkirk Games. Considering the Mountkirk games business and technical requirements, what should you do?

- A. Create network load balancers. Use preemptible Compute Engine instances.
- B. Create network load balancers. Use non-preemptible Compute Engine instances.
- C. Create a global load balancer with managed instance groups and autoscaling policies. Use preemptible Compute Engine instances.
- D. Create a global load balancer with managed instance groups and autoscaling policies. Use non-preemptible Compute Engine instance.

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Testlet 1

Company Overview

TerramEarth manufactures heavy equipment for the mining and agricultural industries: about 80% of their business is from mining and 20% from agriculture. They currently have over 500 dealers and service centers in 100 countries. Their mission is to build products that make their customers more productive.

Company background


TerramEarth was formed in 1946, when several small, family owned companies combined to retool after World War II. The company cares about their employees and customers and considers them to be extended members of their family.

TerramEarth is proud of their ability to innovate on their core products and find new markets as their customers' needs change. For the past 20 years, trends in the industry have been largely toward increasing productivity by using larger vehicles with a human operator.

Solution Concept

There are 20 million TerramEarth vehicles in operation that collect 120 fields of data per second. Data is stored locally on the vehicle and can be accessed for analysis when a vehicle is serviced. The data is downloaded via a maintenance port. This same port can be used to adjust operational parameters, allowing the vehicles to be upgraded in the field with new computing modules.

Approximately 200,000 vehicles are connected to a cellular network, allowing TerramEarth to collect data directly. At a rate of 120 fields of data per second with 22 hours of operation per day, Terram Earth collects a total of about 9 TB/day from these connected vehicles. **Existing Technical Environment**



TerramEarth's existing architecture is composed of Linux-based systems that reside in a data center. These systems gzip CSV files from the field and upload via FTP, transform and aggregate them, and place the data in their data warehouse. Because this process takes time, aggregated reports are based on data that is 3 weeks old.

With this data, TerramEarth has been able to preemptively stock replacement parts and reduce unplanned downtime of their vehicles by 60%. However, because the data is stale, some customers are without their vehicles for up to 4 weeks while they wait for replacement parts.

Business Requirements

- Decrease unplanned vehicle downtime to less than 1 week, without increasing the cost of carrying surplus inventory
- Support the dealer network with more data on how their customers use their equipment to better position new products and services
- Have the ability to partner with different companies – especially with seed and fertilizer suppliers in the fast-growing agricultural business – to create compelling joint offerings for their customers.

CEO Statement

We have been successful in capitalizing on the trend toward larger vehicles to increase the productivity of our customers. Technological change is occurring rapidly, and TerramEarth has taken advantage of connected devices technology to provide our customers with better services, such as our intelligent farming equipment. With this technology, we have been able to increase farmers' yields by 25%, by using past trends to adjust how our vehicles operate. These advances have led to the rapid growth of our agricultural product line, which we expect will generate 50% of our revenues by 2020.

CTO Statement

Our competitive advantage has always been in the manufacturing process, with our ability to build better vehicles for lower cost than our competitors. However, new products with different approaches are constantly being developed, and I'm concerned that we lack the skills to undergo the next wave of transformations in our industry. Unfortunately, our CEO doesn't take technology obsolescence seriously and he considers the many new companies in our industry to be niche players. My goals are to build our skills while addressing immediate market needs through incremental innovations.

QUESTION 1

Your development team has created a structured API to retrieve vehicle data. They want to allow third parties to develop tools for dealerships that use this vehicle event data. You want to support delegated authorization against this data.

What should you do?

- A. Build or leverage an OAuth-compatible access control system
- B. Build SAML 2.0 SSO compatibility into your authentication system
- C. Restrict data access based on the source IP address of the partner systems
- D. Create secondary credentials for each dealer that can be given to the trusted third party

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Delegate application authorization with OAuth2

Cloud Platform APIs support OAuth 2.0, and scopes provide granular authorization over the methods that are supported. Cloud Platform supports both serviceaccount and user-account OAuth, also called three-legged OAuth.

References: https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#delegate_application_authorization_with_oauth2
<https://cloud.google.com/appengine/docs/flexible/go/authorizing-apps>

QUESTION 2

TerramEarth plans to connect all 20 million vehicles in the field to the cloud. This increases the volume to 20 million 600 byte records a second for 40 TB an hour.

How should you design the data ingestion?

- A. Vehicles write data directly to GCS
- B. Vehicles write data directly to Google Cloud Pub/Sub
- C. Vehicles stream data directly to Google BigQuery
- D. Vehicles continue to write data using the existing system (FTP)

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Streamed data is available for real-time analysis within a few seconds of the first streaming insertion into a table.

Instead of using a job to load data into BigQuery, you can choose to stream your data into BigQuery one record at a time by using the `tabledata().insertAll()` method. This approach enables querying data without the delay of running a load job.

References: <https://cloud.google.com/bigquery/streaming-data-into-bigquery>

QUESTION 3

You analyzed TerramEarth's business requirement to reduce downtime, and found that they can achieve a majority of time saving by reducing customer's wait time for parts. You decided to focus on reduction of the 3 weeks aggregate reporting time.

Which modifications to the company's processes should you recommend?

- A. Migrate from CSV to binary format, migrate from FTP to SFTP transport, and develop machine learning analysis of metrics
- B. Migrate from FTP to streaming transport, migrate from CSV to binary format, and develop machine learning analysis of metrics
- C. Increase fleet cellular connectivity to 80%, migrate from FTP to streaming transport, and develop machine learning analysis of metrics
- D. Migrate from FTP to SFTP transport, develop machine learning analysis of metrics, and increase dealer local inventory by a fixed factor

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

The Avro binary format is the preferred format for loading compressed data. Avro data is faster to load because the data can be read in parallel, even when the data blocks are compressed.

Cloud Storage supports streaming transfers with the gsutil tool or boto library, based on HTTP chunked transfer encoding. Streaming data lets you stream data to and from your Cloud Storage account as soon as it becomes available without requiring that the data be first saved to a separate file. Streaming transfers are useful if you have a process that generates data and you do not want to buffer it locally before uploading it, or if you want to send the result from a computational pipeline directly into Cloud Storage.

References: <https://cloud.google.com/storage/docs/streaming> <https://cloud.google.com/bigquery/docs/loading-data>

QUESTION 4

To speed up data retrieval, more vehicles will be upgraded to cellular connections and be able to transmit data to the ETL process. The current FTP process is error-prone and restarts the data transfer from the start of the file when connections fail, which happens often. You want to improve the reliability of the solution and minimize data transfer time on the cellular connections.

What should you do?

- A. Use one Google Container Engine cluster of FTP servers. Save the data to a Multi-Regional bucket. Run the ETL process using data in the bucket
- B. Use multiple Google Container Engine clusters running FTP servers located in different regions. Save the data to Multi-Regional buckets in US, EU, and Asia. Run the ETL process using the data in the bucket
- C. Directly transfer the files to different Google Cloud Multi-Regional Storage bucket locations in US, EU, and Asia using Google APIs over HTTP(S). Run the ETL process using the data in the bucket
- D. Directly transfer the files to a different Google Cloud Regional Storage bucket location in US, EU, and Asia using Google APIs over HTTP(S). Run the ETL process to retrieve the data from each Regional bucket

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

From scenario: They currently have over 500 dealers and service centers in 100 countries.

Multi-Regional Storage is appropriate for storing data that is frequently accessed ("hot" objects), such as serving website content, interactive workloads, or data supporting mobile and gaming applications. Multi-Regional Storage data has the most availability compared to other storage classes.

Multi-Regional Storage is geo-redundant, which means Cloud Storage stores your data redundantly in at least two geographic places separated by at least 100 miles within the multi-regional location of the bucket. Geo-redundancy occurs asynchronously, but Multi-Regional Storage data is redundant within at least one geographic place as soon as you upload it. Like all Cloud Storage data, it is also immediately accessible worldwide.

References: <https://cloud.google.com/storage/docs/storage-classes#multi-regional>

QUESTION 5

TerramEarth's 20 million vehicles are scattered around the world. Based on the vehicle's location, its telemetry data is stored in a Google Cloud Storage (GCS) regional bucket (US, Europe, or Asia). The CTO has asked you to run a report on the raw telemetry data to determine why vehicles are breaking down after 100 K miles. You want to run this job on all the data.

What is the most cost-effective way to run this job?

- A. Move all the data into 1 zone, then launch a Cloud Dataproc cluster to run the job
- B. Move all the data into 1 region, then launch a Google Cloud Dataproc cluster to run the job
- C. Launch a cluster in each region to preprocess and compress the raw data, then move the data into a multi-region bucket and use a Dataproc cluster to finish the job
- D. Launch a cluster in each region to preprocess and compress the raw data, then move the data into a region bucket and use a Cloud Dataproc cluster to finish the job

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Storage guarantees 2 replicates which are geo diverse (100 miles apart) which can get better remote latency and availability.

More importantly, is that multiregional heavily leverages Edge caching and CDNs to provide the content to the end users.

All this redundancy and caching means that Multiregional comes with overhead to sync and ensure consistency between geo-diverse areas. As such, it's much better for write-once-read-many scenarios. This means frequently accessed (e.g. "hot" objects) around the world, such as website content, streaming videos, gaming or mobile applications.

References: <https://medium.com/google-cloud/google-cloud-storage-what-bucket-class-for-the-best-performance-5c847ac8f9f2>

QUESTION 6

TerramEarth has equipped all connected trucks with servers and sensors to collect 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's 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, and store it in a GCS Coldline bucket

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

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>

Testlet 1

Company Overview

Dress4win is a web-based company that helps their users organize and manage their personal wardrobe using a website and mobile application. The company also cultivates an active social network that connects their users with designers and retailers. They monetize their services through advertising, e-commerce, referrals, and a premium app model.

Company Background

Dress4win's application has grown from a few servers in the founder's garage to several hundred servers and appliances in a collocated data center. However, the capacity of their infrastructure is now insufficient for the application's rapid growth. Because of this growth and the company's desire to innovate faster, Dress4win is committing to a full migration to a public cloud.

Solution Concept

For the first phase of their migration to the cloud, Dress4win is considering moving their development and test environments. They are also considering building a disaster recovery site, because their current infrastructure is at a single location. They are not sure which components of their architecture they can migrate as is and which components they need to change before migrating them.

Existing Technical Environment

The Dress4win application is served out of a single data center location.

- Databases:
 - MySQL - user data, inventory, static data
 - Redis - metadata, social graph, caching
 - Application servers:
 - Tomcat - Java micro-services
 - Nginx - static content
 - Apache Beam - Batch processing
 - Storage appliances:
 - iSCSI for VM hosts
 - Fiber channel SAN - MySQL databases
 - NAS - image storage, logs, backups
 - Apache Hadoop/Spark servers:
 - Data analysis
 - Real-time trending calculations
 - MQ servers:
 - Messaging
 - Social notifications
 - Events
- Miscellaneous servers:
 - Jenkins, monitoring, bastion hosts, security scanners

Business Requirements

- Build a reliable and reproducible environment with scaled parity of production.
 - Improve security by defining and adhering to a set of security and Identity and Access Management (IAM) best practices for cloud.
 - Improve business agility and speed of innovation through rapid provisioning of new resources.
 - Analyze and optimize architecture for performance in the cloud.
- Migrate fully to the cloud if all other requirements are met.

Technical Requirements

- Evaluate and choose an automation framework for provisioning resources in cloud.
- Support failover of the production environment to cloud during an emergency.
- Identify production services that can migrate to cloud to save capacity.
- Use managed services whenever possible.
- Encrypt data on the wire and at rest.
- Support multiple VPN connections between the production data center and cloud environment.

CEO Statement

Our investors are concerned about our ability to scale and contain costs with our current infrastructure. They are also concerned that a new competitor could use a public cloud platform to offset their up-front investment and freeing them to focus on developing better features.

CTO Statement

We have invested heavily in the current infrastructure, but much of the equipment is approaching the end of its useful life. We are consistently waiting weeks for new gear to be racked before we can start new projects. Our traffic patterns are highest in the mornings and weekend evenings; during other times, 80% of our capacity is sitting idle.

CFO Statement

Our capital expenditure is now exceeding our quarterly projections. Migrating to the cloud will likely cause an initial increase in spending, but we expect to fully transition before our next hardware refresh cycle. Our total cost of ownership (TCO) analysis over the next 5 years puts a cloud strategy between 30 to 50% lower than our current model.

QUESTION 1

At Dress4Win, an operations engineer wants to create a low-cost solution to remotely archive copies of database backup files. The database files are compressed tar files stored in their current data center. How should he proceed?

- A. Create a cron script using gsutil to copy the files to a Coldline Storage bucket.
- B. Create a cron script using gsutil to copy the files to a Regional Storage bucket.
- C. Create a Cloud Storage Transfer Service Job to copy the files to a Coldline Storage bucket.
- D. Create a Cloud Storage Transfer Service job to copy the files to a Regional Storage bucket.

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Follow these rules of thumb when deciding whether to use gsutil or Storage Transfer Service:

- * When transferring data from an on-premises location, use gsutil.
- * When transferring data from another cloud storage provider, use Storage Transfer Service.
- * Otherwise, evaluate both tools with respect to your specific scenario.

Use this guidance as a starting point.

The specific details of your transfer scenario will also help you determine which tool is more appropriate.

QUESTION 2

Dress4Win has asked you to recommend machine types they should deploy their application servers to.

How should you proceed?

- A. Perform a mapping of the on-premises physical hardware cores and RAM to the nearest machine types in the cloud.
- B. Recommend that Dress4Win deploy application servers to machine types that offer the highest RAM to CPU ratio available.

- C. Recommend that Dress4Win deploy into production with the smallest instances available, monitor them over time, and scale the machine type up until the desired performance is reached.
- D. Identify the number of virtual cores and RAM associated with the application server virtual machines align them to a custom machine type in the cloud, monitor performance, and scale the machine types up until the desired performance is reached.

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

QUESTION 3

As part of Dress4Win's plans to migrate to the cloud, they want to be able to set up a managed logging and monitoring system so they can handle spikes in their traffic load.

They want to ensure that:

- * The infrastructure can be notified when it needs to scale up and down to handle the ebb and flow of usage throughout the day
- * Their administrators are notified automatically when their application reports errors.
- * They can filter their aggregated logs down in order to debug one piece of the application across many hosts



<https://www.vceplus.com/>

Which Google StackDriver features should they use?

- A. Logging, Alerts, Insights, Debug
- B. Monitoring, Trace, Debug, Logging
- C. Monitoring, Logging, Alerts, Error Reporting
- D. Monitoring, Logging, Debug, Error Report

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

QUESTION 4

Dress4Win has asked you for advice on how to migrate their on-premises MySQL deployment to the cloud.

They want to minimize downtime and performance impact to their on-premises solution during the migration.

Which approach should you recommend?

- A. Create a dump of the on-premises MySQL master server, and then shut it down, upload it to the cloud environment, and load into a new MySQL cluster.
- B. Setup a MySQL replica server/slave in the cloud environment, and configure it for asynchronous replication from the MySQL master server on-premises until cutover.
- C. Create a new MySQL cluster in the cloud, configure applications to begin writing to both on premises and cloud MySQL masters, and destroy the original cluster at cutover.
- D. Create a dump of the MySQL replica server into the cloud environment, load it into: Google Cloud Datastore, and configure applications to read/write to Cloud Datastore at cutover.

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

QUESTION 5

Dress4Win has configured a new uptime check with Google Stackdriver for several of their legacy services. The Stackdriver dashboard is not reporting the services as healthy.

What should they do?

- A. Install the Stackdriver agent on all of the legacy web servers.
- B. In the Cloud Platform Console download the list of the uptime servers' IP addresses and create an inbound firewall rule
- C. Configure their load balancer to pass through the User-Agent HTTP header when the value matches GoogleStackdriverMonitoring-UptimeChecks (<https://cloud.google.com/monitoring>)
- D. Configure their legacy web servers to allow requests that contain user-Agent HTTP header when the value matches GoogleStackdriverMonitoring-UptimeChecks (<https://cloud.google.com/monitoring>)

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

QUESTION 6

As part of their new application experience, Dress4Win allows customers to upload images of themselves. The customer has exclusive control over who may view these images.

Customers should be able to upload images with minimal latency and also be shown their images quickly on the main application page when they log in.

Which configuration should Dress4Win use?

- A. Store image files in a Google Cloud Storage bucket. Use Google Cloud Datastore to maintain metadata that maps each customer's ID and their image files.
- B. Store image files in a Google Cloud Storage bucket. Add custom metadata to the uploaded images in Cloud Storage that contains the customer's unique ID.
- C. Use a distributed file system to store customers' images. As storage needs increase, add more persistent disks and/or nodes. Assign each customer a unique ID, which sets each file's owner attribute, ensuring privacy of images.
- D. Use a distributed file system to store customers' images. As storage needs increase, add more persistent disks and/or nodes. Use a Google Cloud SQL database to maintain metadata that maps each customer's ID to their image files.

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Question Set 1

QUESTION 1

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

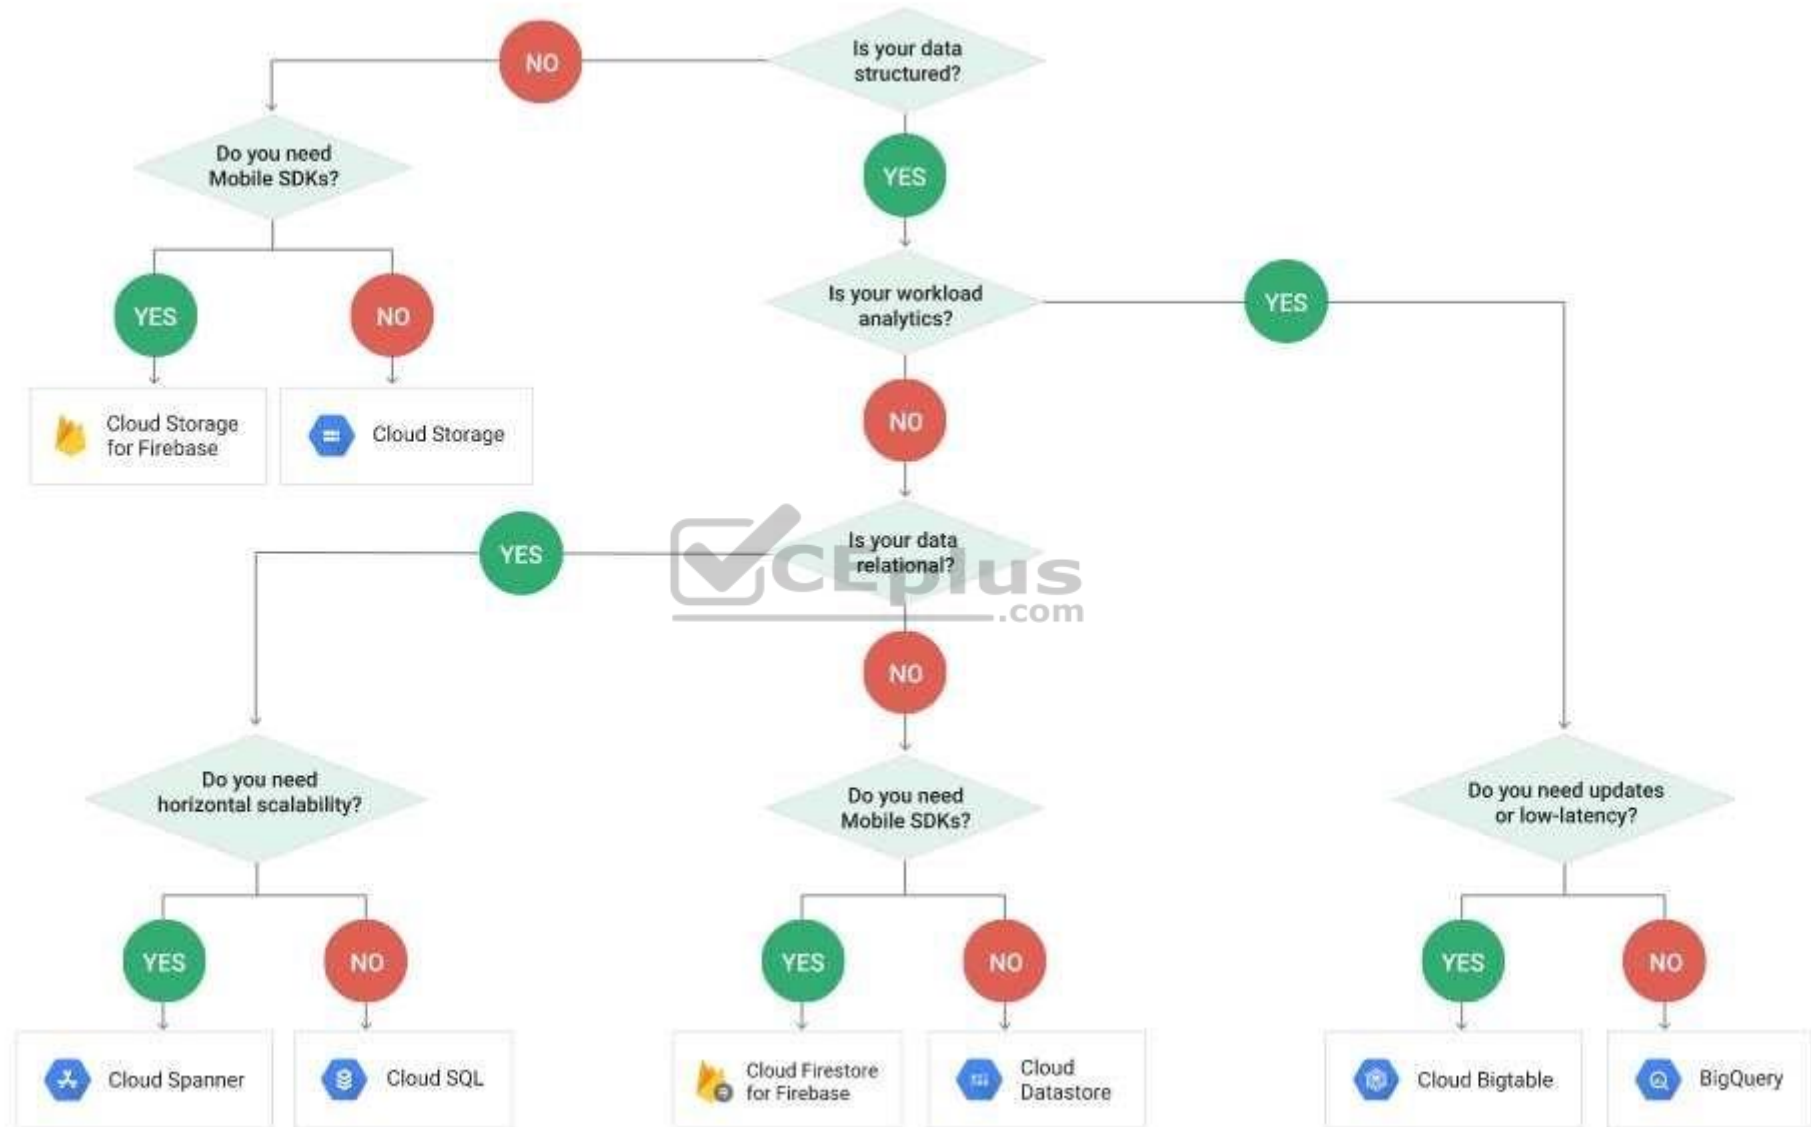
Section: [none]

Explanation

Explanation/Reference:

Explanation:





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

QUESTION 2

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 Kubernetes 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 3

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 4

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?

- 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 5

The operations manager asks you for a list of recommended practices that she should consider when migrating a J2EE application to the cloud.

Which three practices should you recommend? Choose 3 answers.

- A. Port the application code to run on Google App Engine
- B. Integrate Cloud Dataflow into the application to capture real-time metrics
- C. Instrument the application with a monitoring tool like Stackdriver Debugger
- D. Select an automation framework to reliably provision the cloud infrastructure
- E. Deploy a continuous integration tool with automated testing in a staging environment
- F. Migrate from MySQL to a managed NoSQL database like Google Cloud Datastore or Bigtable

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 6

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 7

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 8

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 9

Your customer is moving an existing corporate application to Google Cloud Platform from an on-premises data center. The business owners require minimal user disruption. There are strict security team requirements for storing passwords.

What authentication strategy should they use?

- A. Use G Suite Password Sync to replicate passwords into Google
- B. Federate authentication via SAML 2.0 to the existing Identity Provider
- C. Provision users in Google using the Google Cloud Directory Sync tool
- D. Ask users to set their Google password to match their corporate password

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Provision users to Google's directory

The global Directory is available to both Cloud Platform and G Suite resources and can be provisioned by a number of means. Provisioned users can take advantage of rich authentication features including single sign-on (SSO), OAuth, and two-factor verification. You can provision users automatically using one of the following tools and services:

Google Cloud Directory Sync (GCDS)

Google Admin SDK

A third-party connector

GCDS is a connector that can provision users and groups on your behalf for both Cloud Platform and G Suite. Using GCDS, you can automate the addition, modification, and deletion of users, groups, and non-employee contacts. You can synchronize the data from your LDAP directory server to your Cloud Platform domain by using LDAP queries. This synchronization is one-way: the data in your LDAP directory server is never modified.

References: <https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#authentication-and-identity>

QUESTION 10

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 liveprocessing 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

Correct Answer: B

Section: [none]

Explanation



Explanation/Reference:

Explanation:

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/>

QUESTION 11

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 12

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.



<https://www.vceplus.com/>

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 13

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:

Reference:

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

QUESTION 14

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 15

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 ls -l 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 16

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 Kubernetes 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 17

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 18

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 19

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 20

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 21

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 22

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 23

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:

QUESTION 24

A development manager is building a new application. He asks you to review his requirements and identify what cloud technologies he can use to meet them. The application must:

1. Be based on open-source technology for cloud portability
2. Dynamically scale compute capacity based on demand
3. Support continuous software delivery
4. Run multiple segregated copies of the same application stack
5. Deploy application bundles using dynamic templates
6. Route network traffic to specific services based on URL

Which combination of technologies will meet all of his requirements?

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

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Jenkins is an open-source automation server that lets you flexibly orchestrate your build, test, and deployment pipelines. Kubernetes Engine is a hosted version of Kubernetes, a powerful cluster manager and orchestration system for containers.

When you need to set up a continuous delivery (CD) pipeline, deploying Jenkins on Kubernetes Engine provides important benefits over a standard VM-based deployment

Incorrect Answers:

A: Helm is a tool for managing Kubernetes charts. Charts are packages of pre-configured Kubernetes resources.

Use Helm to:

- Find and use popular software packaged as Kubernetes charts
- Share your own applications as Kubernetes charts
- Create reproducible builds of your Kubernetes applications
- Intelligently manage your Kubernetes manifest files
- Manage releases of Helm packages

References: <https://cloud.google.com/solutions/jenkins-on-kubernetes-engine>

QUESTION 25

You have created several pre-emptible Linux virtual machine instances using Google Compute Engine. You want to properly shut down your application before the virtual machines are preempted.

What should you do?

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

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

A startup script, or a shutdown script, is specified through the metadata server, using startup script metadata keys.

Reference: <https://cloud.google.com/compute/docs/startupscript>

QUESTION 26

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

Correct Answer: D

Section: [none]

Explanation

**Explanation/Reference:**

Explanation:

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/>

QUESTION 27

Your development team has installed a new Linux kernel module on the batch servers in Google Compute Engine (GCE) virtual machines (VMs) to speed up the nightly batch process. Two days after the installation, 50% of the batch servers failed the nightly batch run. You want to collect details on the failure to pass back to the development team.

Which three actions should you take? Choose 3 answers.

- A. Use Stackdriver Logging to search for the module log entries
- B. Read the debug GCE Activity log using the API or Cloud Console
- C. Use gcloud or Cloud Console to connect to the serial console and observe the logs
- D. Identify whether a live migration event of the failed server occurred, using in the activity log

- E. Adjust the Google Stackdriver timeline to match the failure time, and observe the batch server metrics
- F. Export a debug VM into an image, and run the image on a local server where kernel log messages will be displayed on the native screen

Correct Answer: ACE

Section: [none]

Explanation

Explanation/Reference:

QUESTION 28

Your company wants to try out the cloud with low risk. They want to archive approximately 100 TB of their log data to the cloud and test the analytics features available to them there, while also retaining that data as a long-term disaster recovery backup.

Which two steps should you take? Choose 2 answers.

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



Correct Answer: AC

Section: [none]

Explanation

Explanation/Reference:

QUESTION 29

You created a pipeline that can deploy your source code changes to your infrastructure in instance groups for self-healing. One of the changes negatively affects your key performance indicator. You are not sure how to fix it, and investigation could take up to a week.

What should you do?

- A. Log in to a server, and iterate on the fix locally
- B. Revert the source code change, and rerun the deployment pipeline
- C. Log into the servers with the bad code change, and swap in the previous code

D. Change the instance group template to the previous one, and delete all instances

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

QUESTION 30

Your organization wants to control IAM policies for different departments independently, but centrally. Which approach should you take?

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

Correct Answer: C

Section: [none]

Explanation



Explanation/Reference:

Explanation:

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

References: <https://cloud.google.com/resource-manager/docs/creating-managing-folders>

QUESTION 31

You deploy your custom Java application to Google App Engine. It fails to deploy and gives you the following stack trace. What should you do?

```
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)
```




<https://www.vceplus.com/>

- 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 MD5 hash instead of SHA1

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:



QUESTION 32

You are designing a mobile chat application. You want to ensure people cannot spoof chat messages, by providing a message were sent by a specific user.

What should you do?

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

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Encrypting each block and tagging each message at the client side is an overhead on the application. Best method which has been adopted since years is contacting the SSL provider and use the public certificate to encrypt the traffic between client and the server.

QUESTION 33

As part of implementing their disaster recovery plan, your company is trying to replicate their production MySQL database from their private data center to their GCP project using a Google Cloud VPN connection. They are experiencing latency issues and a small amount of packet loss that is disrupting the replication. What should they do?

- A. Configure their replication to use UDP.
- B. Configure a Google Cloud Dedicated Interconnect.
- C. Restore their database daily using Google Cloud SQL.
- D. Add additional VPN connections and load balance them.
- E. Send the replicated transaction to Google Cloud Pub/Sub.

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:



QUESTION 34

You are working in a highly secured environment where public Internet access from the Compute Engine VMs is not allowed. You do not yet have a VPN connection to access an on-premises file server. You need to install specific software on a Compute Engine instance. How should you install the software?

- A. Upload the required installation files to Cloud Storage. Configure the VM on a subnet with a Private Google Access subnet. Assign only an internal IP address to the VM. Download the installation files to the VM using gsutil.
- B. Upload the required installation files to Cloud Storage and use firewall rules to block all traffic except the IP address range for Cloud Storage. Download the files to the VM using gsutil.
- C. Upload the required installation files to Cloud Source Repositories. Configure the VM on a subnet with a Private Google Access subnet. Assign only an internal IP address to the VM. Download the installation files to the VM using gcloud.
- D. Upload the required installation files to Cloud Source Repositories and use firewall rules to block all traffic except the IP address range for Cloud Source Repositories. Download the files to the VM using gsutil.

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

QUESTION 35

You have an application deployed on Kubernetes Engine using a Deployment named echo-deployment. The deployment is exposed using a Service called echoservice. You need to perform an update to the application with minimal downtime to the application. What should you do?

- A. Use `kubectl set image deployment/echo-deployment <new-image>`
- B. Use the rolling update functionality of the Instance Group behind the Kubernetes cluster
- C. Update the deployment yaml file with the new container image. Use `kubectl delete deployment/echo-deployment` and `kubectl create -f <yaml-file>`
- D. Update the service yaml file which the new container image. Use `kubectl delete service/echo-service` and `kubectl create -f <yaml-file>`

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/kubernetes-engine/docs/how-to/updating-apps>

QUESTION 36

Your company is using BigQuery as its enterprise data warehouse. Data is distributed over several Google Cloud projects. All queries on BigQuery need to be billed on a single project. You want to make sure that no query costs are incurred on the projects that contain the data. Users should be able to query the datasets, but not edit them.

How should you configure users' access roles?

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

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

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

QUESTION 37

You have developed an application using Cloud ML Engine that recognizes famous paintings from uploaded images. You want to test the application and allow specific people to upload images for the next 24 hours. Not all users have a Google Account. How should you have users upload images?

- A. Have users upload the images to Cloud Storage. Protect the bucket with a password that expires after 24 hours.
- B. Have users upload the images to Cloud Storage using a signed URL that expires after 24 hours.
- C. Create an App Engine web application where users can upload images. Configure App Engine to disable the application after 24 hours. Authenticate users via Cloud Identity.
- D. Create an App Engine web application where users can upload images for the next 24 hours. Authenticate users via Cloud Identity.

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

QUESTION 38

Your web application must comply with the requirements of the European Union's General Data Protection Regulation (GDPR). You are responsible for the technical architecture of your web application. What should you do?

- A. Ensure that your web application only uses native features and services of Google Cloud Platform, because Google already has various certifications and provides "pass-on" compliance when you use native features.
- B. Enable the relevant GDPR compliance setting within the GCP Console for each of the services in use within your application.
- C. Ensure that Cloud Security Scanner is part of your test planning strategy in order to pick up any compliance gaps.
- D. Define a design for the security of data in your web application that meets GDPR requirements.

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Reference: <https://www.mobiloud.com/blog/gdpr-compliant-mobile-app/>

QUESTION 39

The development team has provided you with a Kubernetes Deployment file. You have no infrastructure yet and need to deploy the application. What should you do?

- A. Use gcloud to create a Kubernetes cluster. Use Deployment Manager to create the deployment.
- B. Use gcloud to create a Kubernetes cluster. Use kubectl to create the deployment.
- C. Use kubectl to create a Kubernetes cluster. Use Deployment Manager to create the deployment.
- D. Use kubectl to create a Kubernetes cluster. Use kubectl to create the deployment.

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

QUESTION 40

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

- A. Allocate budget for team training. Set a deadline for the new GCP project.
- B. Allocate budget for team training. Create a roadmap for your team to achieve Google Cloud certification based on job role.
- C. Allocate budget to hire skilled external consultants. Set a deadline for the new GCP project.
- D. Allocate budget to hire skilled external consultants. Create a roadmap for your team to achieve Google Cloud certification based on job role.

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:



QUESTION 41

You are creating an App Engine application that uses Cloud Datastore as its persistence layer. You need to retrieve several root entities for which you have the identifiers. You want to minimize the overhead in operations performed by Cloud Datastore. What should you do?

- A. Create the Key object for each Entity and run a batch `get` operation
- B. Create the Key object for each Entity and run multiple `get` operations, one operation for each entity
- C. Use the identifiers to create a query filter and run a batch query operation
- D. Use the identifiers to create a query filter and run multiple query operations, one operation for each entity

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

QUESTION 42

You need to upload files from your on-premises environment to Cloud Storage. You want the files to be encrypted on Cloud Storage using customer-supplied encryption keys. What should you do?

- A. Supply the encryption key in a .boto configuration file. Use gsutil to upload the files.
- B. Supply the encryption key using gcloud config. Use gsutil to upload the files to that bucket.
- C. Use gsutil to upload the files, and use the flag --encryption-key to supply the encryption key.
- D. Use gsutil to create a bucket, and use the flag --encryption-key to supply the encryption key. Use gsutil to upload the files to that bucket.

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

QUESTION 43

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

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

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

QUESTION 44

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 on-premises AD domain controller using Google Cloud Directory Sync.

Correct Answer: C

Section: [none]

Explanation



Explanation/Reference:

Reference: <https://cloud.google.com/blog/products/identity-security/using-your-existing-identity-management-system-with-google-cloud-platform>

QUESTION 45

You want your Google Kubernetes Engine cluster to automatically add or remove nodes based on CPUload.

What should you do?

- A. Configure a HorizontalPodAutoscaler with a target CPU usage. Enable the Cluster Autoscaler from the GCP Console.
- B. Configure a HorizontalPodAutoscaler with a target CPU usage. Enable autoscaling on the managed instance group for the cluster using the gcloud command.
- C. Create a deployment and set the maxUnavailable and maxSurge properties. Enable the Cluster Autoscaler using the gcloud command.
- D. Create a deployment and set the maxUnavailable and maxSurge properties. Enable autoscaling on the cluster managed instance group from the GCP Console.

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:**QUESTION 46**

You need to develop procedures to verify resilience of disaster recovery for remote recovery using GCP. Your production environment is hosted on-premises. You need to establish a secure, redundant connection between your on premises network and the GCP network.

What should you do?

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

Correct Answer: C

Section: [none]

Explanation



Explanation/Reference:

QUESTION 47

You have been engaged by your client to lead the migration of their application infrastructure to GCP. One of their current problems is that the on-premises high performance SAN is requiring frequent and expensive upgrades to keep up with the variety of workloads that are identified as follows: 20TB of log archives retained for legal reasons; 500 GB of VM boot/data volumes and templates; 500 GB of image thumbnails; 200 GB of customer session state data that allows customers to restart sessions even if off-line for several days.

Which of the following best reflects your recommendations for a cost-effective storage allocation?

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

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

QUESTION 48

You are using Cloud CDN to deliver static HTTP(S) website content hosted on a Compute Engine instance group. You want to improve the cache hit ratio.

What should you do?

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

Correct Answer: A

Section: [none]

Explanation



Explanation/Reference:

Reference https://cloud.google.com/cdn/docs/best-practices#using_custom_cache_keys_to_improve_cache_hit_ratio

QUESTION 49

You have an App Engine application that needs to be updated. You want to test the update with production traffic before replacing the current application version.

What should you do?

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

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:



<https://www.vceplus.com/>

