

Professional Cloud Architect

Number: GoogleCloudArchitect

Passing Score: 800

Time Limit: 120 min

File Version: 1

Professional Cloud Architect



Website: <https://vceplus.com>

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

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

Twitter : https://twitter.com/VCE_Plus

<https://vceplus.com/>

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

The JencoMart security team requires that all Google Cloud Platform infrastructure is deployed using a least privilege model with separation of duties for administration between production and development resources.



What Google domain and project structure should you recommend?

- A. Create two G Suite accounts to manage users: one for development/test/staging and one for production. Each account should contain one project for every application
- B. Create two G Suite accounts to manage users: one with a single project for all development applications and one with a single project for all production applications
- C. Create a single G Suite account to manage users with each stage of each application in its own project
- D. Create a single G Suite account to manage users with one project for the development/test/staging environment and one project for the production environment

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Note: The principle of least privilege and separation of duties are concepts that, although semantically different, are intrinsically related from the standpoint of security. The intent behind both is to prevent people from having higher privilege levels than they actually need

- Principle of Least Privilege: Users should only have the least amount of privileges required to perform their job and no more. This reduces authorization exploitation by limiting access to resources such as targets, jobs, or monitoring templates for which they are not authorized.
- Separation of Duties: Beyond limiting user privilege level, you also limit user duties, or the specific jobs they can perform. No user should be given responsibility for more than one related function. This limits the ability of a user to perform a malicious action and then cover up that action.

Reference: <https://cloud.google.com/kms/docs/separation-of-duties>

QUESTION 2

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.

Reference: <https://cloud.google.com/compute/docs/ssh-in-browser> <https://cloud.google.com/compute/docs/ssh-in-browser>

QUESTION 3

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.

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

QUESTION 4

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

Which metrics should you track?

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

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:


From scenario:

Business Requirements include: Expand services into Asia

Technical Requirements include: Decrease latency in Asia

QUESTION 5





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

What are three potential bottlenecks? Choose 3 answers.

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

Correct Answer: ACE

Section: [none]

Explanation

Explanation/Reference:

QUESTION 6

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

Reference: <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. 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

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 wants you to design their new testing strategy. How should the test coverage differ from their existing backends on the other platforms?

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

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

From Scenario:

A few of their games were more popular than expected, and they had problems scaling their application servers, MySQL databases, and analytics tools.

Requirements for Game Analytics Platform include: Dynamically scale up or down based on game activity

QUESTION 2

Mountkirk Games has deployed their new backend on Google Cloud Platform (GCP). You want to create a thorough testing process for new versions of the backend before they are released to the public. You want the testing environment to scale in an economical way. How should you design the process?

- A. Create a scalable environment in GCP for simulating production load
- B. Use the existing infrastructure to test the GCP-based backend at scale
- C. Build stress tests into each component of your application using resources internal to GCP to simulate load
- D. Create a set of static environments in GCP to test different levels of load – for example, high, medium, and low

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 3

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: C

Section: [none]

Explanation

Explanation/Reference:

QUESTION 4

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 5

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: D

Section: [none]

Explanation

Explanation/Reference:

QUESTION 6

Mountkirk Games wants to set up a real-time analytics platform for their new game. The new platform must meet their technical requirements.

Which combination of Google technologies will meet all of their requirements?

- A. Kubernetes Engine, Cloud Pub/Sub, and Cloud SQL
- B. Cloud Dataflow, Cloud Storage, Cloud Pub/Sub, and BigQuery
- C. Cloud SQL, Cloud Storage, Cloud Pub/Sub, and Cloud Dataflow

- D. Cloud Dataproc, Cloud Pub/Sub, Cloud SQL, and Cloud Dataflow
- E. Cloud Pub/Sub, Compute Engine, Cloud Storage, and Cloud Dataproc

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Ingest millions of streaming events per second from anywhere in the world with Cloud Pub/Sub, powered by Google's unique, high-speed private network. Process the streams with Cloud Dataflow to ensure reliable, exactly-once, low-latency data transformation. Stream the transformed data into BigQuery, the cloud-native data warehousing service, for immediate analysis via SQL or popular visualization tools.

From scenario: They plan to deploy the game's backend on Google Compute Engine so they can capture streaming metrics, run intensive analytics.

Requirements for Game Analytics Platform

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

Reference: <https://cloud.google.com/solutions/big-data/stream-analytics/>

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 to migrate from their current analytics and statistics reporting model to one that meets their technical requirements on Google Cloud Platform.

Which two steps should be part of their migration plan? (Choose two.)

- A. Evaluate the impact of migrating their current batch ETL code to Cloud Dataflow.
- B. Write a schema migration plan to denormalize data for better performance in BigQuery.
- C. Draw an architecture diagram that shows how to move from a single MySQL database to a MySQL cluster.
- D. Load 10 TB of analytics data from a previous game into a Cloud SQL instance, and run test queries against the full dataset to confirm that they complete successfully.

E. Integrate Cloud Armor to defend against possible SQL injection attacks in analytics files uploaded to Cloud Storage.

Correct Answer: AB

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 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 instances.

Correct Answer: D

Section: [none]

Explanation



Explanation/Reference:

QUESTION 3

For this question, refer to the Mountkirk Games case study. Mountkirk Games wants to design their solution for the future in order to take advantage of cloud and technology improvements as they become available. Which two steps should they take? (Choose two.)

- A. Store as much analytics and game activity data as financially feasible today so it can be used to train machine learning models to predict user behavior in the future.
- B. Begin packaging their game backend artifacts in container images and running them on Google Kubernetes Engine to improve the availability to scale up or down based on game activity.
- C. Set up a CI/CD pipeline using Jenkins and Spinnaker to automate canary deployments and improve development velocity.
- D. Adopt a schema versioning tool to reduce downtime when adding new game features that require storing additional player data in the database.
- E. Implement a weekly rolling maintenance process for the Linux virtual machines so they can apply critical kernel patches and package updates and reduce the risk of 0-day vulnerabilities.

Correct Answer: CE

Section: [none]

Explanation

Explanation/Reference:

QUESTION 4

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 5

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: C

Section: [none]

Explanation

Explanation/Reference:

QUESTION 6

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?

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

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

QUESTION 7

For this question, refer to the Mountkirk Games case study. You are in charge of the new Game Backend Platform architecture. The game communicates with the backend over a REST API.

You want to follow Google-recommended practices. How should you design the backend?

- A. Create an instance template for the backend. For every region, deploy it on a multi-zone managed instance group. Use an L4 load balancer.
- B. Create an instance template for the backend. For every region, deploy it on a single-zone managed instance group. Use an L4 load balancer.
- C. Create an instance template for the backend. For every region, deploy it on a multi-zone managed instance group. Use an L7 load balancer.
- D. Create an instance template for the backend. For every region, deploy it on a single-zone managed instance group. Use an L7 load balancer.

Correct Answer: A

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

TerramEarth's CTO wants to use the raw data from connected vehicles to help identify approximately when a vehicle in the field will have a catastrophic failure. You want to allow analysts to centrally query the vehicle data.


Which architecture should you recommend?




<https://vceplus.com/>



A.




B.




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

C.



D.



Correct Answer: A

Section: [none]

Explanation


Explanation/Reference:

Explanation:

The push endpoint can be a load balancer.

A container cluster can be used.

Cloud Pub/Sub for Stream Analytics



Reference: <https://cloud.google.com/pubsub/>

<https://cloud.google.com/solutions/iot/>

<https://cloud.google.com/solutions/designing-connected-vehicle-platform> https://cloud.google.com/solutions/designing-connected-vehicle-platform#data_ingestion

<http://www.eweek.com/big-data-and-analytics/google-touts-value-of-cloud-iot-core-for-analyzing-connected-car-data>
<https://cloud.google.com/solutions/iot/>

QUESTION 2

The TerramEarth development team wants to create an API to meet the company's business requirements. You want the development team to focus their development effort on business value versus creating a custom framework.

Which method should they use?

- A. Use Google App Engine with Google Cloud Endpoints. Focus on an API for dealers and partners
- B. Use Google App Engine with a JAX-RS Jersey Java-based framework. Focus on an API for the public
- C. Use Google App Engine with the Swagger (Open API Specification) framework. Focus on an API for the public
- D. Use Google Container Engine with a Django Python container. Focus on an API for the public
- E. Use Google Container Engine with a Tomcat container with the Swagger (Open API Specification) framework. Focus on an API for dealers and partners

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Develop, deploy, protect and monitor your APIs with Google Cloud Endpoints. Using an Open API Specification or one of our API frameworks, Cloud Endpoints gives you the tools you need for every phase of API development.

From scenario:

Business Requirements

Decrease unplanned vehicle downtime to less than 1 week, without increasing the cost of carrying surplus inventory

Support the dealer network with more data on how their customers use their equipment to better position new products and services

Have the ability to partner with different companies – especially with seed and fertilizer suppliers in the fast-growing agricultural business – to create compelling joint offerings for their customers.

Reference: <https://cloud.google.com/certification/guides/cloud-architect/casestudy-terramearth>

QUESTION 3

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.

Reference: 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 4

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:

QUESTION 5

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: C

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.

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

QUESTION 6

Which of TerramEarth's legacy enterprise processes will experience significant change as a result of increased Google Cloud Platform adoption?

- A. Opex/capex allocation, LAN changes, capacity planning
- B. Capacity planning, TCO calculations, opex/capex allocation
- C. Capacity planning, utilization measurement, data center expansion
- D. Data Center expansion, TCO calculations, utilization measurement

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

QUESTION 7

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: D

Section: [none]

Explanation

Explanation/Reference:



QUESTION 8

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: D

Section: [none]

Explanation

Explanation/Reference:

QUESTION 9

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.

Reference: <https://cloud.google.com/storage/docs/storage-classes>

QUESTION 10

Your agricultural division is experimenting with fully autonomous vehicles. You want your architecture to promote strong security during vehicle operation.

Which two architectures should you consider? (Choose two.)

- A. Treat every micro service call between modules on the vehicle as untrusted.
- B. Require IPv6 for connectivity to ensure a secure address space.

- C. Use a trusted platform module (TPM) and verify firmware and binaries on boot.
- D. Use a functional programming language to isolate code execution cycles.
- E. Use multiple connectivity subsystems for redundancy.
- F. Enclose the vehicle's drive electronics in a Faraday cage to isolate chips.

Correct Answer: AC

Section: [none]

Explanation

Explanation/Reference:

QUESTION 11

Operational parameters such as oil pressure are adjustable on each of TerramEarth's vehicles to increase their efficiency, depending on their environmental conditions. Your primary goal is to increase the operating efficiency of all 20 million cellular and unconnected vehicles in the field.

How can you accomplish this goal?

- A. Have you engineers inspect the data for patterns, and then create an algorithm with rules that make operational adjustments automatically
- B. Capture all operating data, train machine learning models that identify ideal operations, and run locally to make operational adjustments automatically
- C. Implement a Google Cloud Dataflow streaming job with a sliding window, and use Google Cloud Messaging (GCM) to make operational adjustments automatically
- D. Capture all operating data, train machine learning models that identify ideal operations, and host in Google Cloud Machine Learning (ML) Platform to make operational adjustments automatically

Correct Answer: B

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.

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, TerramEarth collects a total of about 9 TB/day from these connected vehicles.

Existing Technical Environment

TerramEarth's existing architecture is composed of Linux and Windows-based systems that reside in a single U.S. west coast based data center. These systems gzip CSV files from the field and upload via FTP, 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
- 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

Technical Requirements

- Expand beyond a single datacenter to decrease latency to the American Midwest and east coast
- Create a backup strategy
- Increase security of data transfer from equipment to the datacenter
- Improve data in the data warehouse
- Use customer and equipment data to anticipate customer needs

Application 1: Data ingest

A custom Python application reads uploaded datafiles from a single server, writes to the data warehouse.

Compute:

- Windows Server 2008 R2
- 16 CPUs
- 128 GB of RAM
- 10 TB local HDD storage

Application 2: Reporting

An off the shelf application that business analysts use to run a daily report to see what equipment needs repair. Only 2 analysts of a team of 10 (5 west coast, 5 east coast) can connect to the reporting application at a time.

Compute:

- Off the shelf application. License tied to number of physical CPUs
- Windows Server 2008 R2
- 16 CPUs
- 32 GB of RAM - 500 GB HDD

Data warehouse:

- A single PostgreSQL server
- RedHat Linux
- 64 CPUs
- 128 GB of RAM
- 4x 6TB HDD in RAID 0

Executive Statement

Our competitive advantage has always been in our 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. My goals are to build our skills while addressing immediate market needs through incremental innovations.

QUESTION 1

For this question, refer to the TerramEarth case study. To be compliant with European GDPR regulation, TerramEarth is required to delete data generated from its European customers after a period of 36 months when it contains personal data. In the new architecture, this data will be stored in both Cloud Storage and BigQuery. What should you do?

- A. Create a BigQuery table for the European data, and set the table retention period to 36 months. For Cloud Storage, use gsutil to enable lifecycle management using a DELETE action with an Age condition of 36 months.
- B. Create a BigQuery table for the European data, and set the table retention period to 36 months. For Cloud Storage, use gsutil to create a SetStorageClass to NONE action when with an Age condition of 36 months.
- C. Create a BigQuery time-partitioned table for the European data, and set the partition expiration period to 36 months. For Cloud Storage, use gsutil to enable lifecycle management using a DELETE action with an Age condition of 36 months.
- D. Create a BigQuery time-partitioned table for the European data, and set the partition period to 36 months. For Cloud Storage, use gsutil to create a SetStorageClass to NONE action with an Age condition of 36 months.

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

QUESTION 2

For this question, refer to the TerramEarth case study. TerramEarth has decided to store data files in Cloud Storage. You need to configure Cloud Storage lifecycle rule to store 1 year of data and minimize file storage cost. Which two actions should you take?

- A. Create a Cloud Storage lifecycle rule with Age: "30", Storage Class: "Standard", and Action: "Set to Coldline", and create a second GCS life-cycle rule with Age: "365", Storage Class: "Coldline", and Action: "Delete".
- B. Create a Cloud Storage lifecycle rule with Age: "30", Storage Class: "Coldline", and Action: "Set to Nearline", and create a second GCS life-cycle rule with Age: "91", Storage Class: "Coldline", and Action: "Set to Nearline".
- C. Create a Cloud Storage lifecycle rule with Age: "90", Storage Class: "Standard", and Action: "Set to Nearline", and create a second GCS life-cycle rule with Age: "91", Storage Class: "Nearline", and Action: "Set to Coldline".
- D. Create a Cloud Storage lifecycle rule with Age: "30", Storage Class: "Standard", and Action: "Set to Coldline", and create a second GCS life-cycle rule with Age: "365", Storage Class: "Nearline", and Action: "Delete".

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:



QUESTION 3

For this question, refer to the TerramEarth case study. You need to implement a reliable, scalable GCP solution for the data warehouse for your company, TerramEarth.

Considering the TerramEarth business and technical requirements, what should you do?

- A. Replace the existing data warehouse with BigQuery. Use table partitioning.
- B. Replace the existing data warehouse with a Compute Engine instance with 96 CPUs.
- C. Replace the existing data warehouse with BigQuery. Use federated data sources.
- D. Replace the existing data warehouse with a Compute Engine instance with 96 CPUs. Add an additional Compute Engine preemptible instance with 32 CPUs.

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

QUESTION 4

For this question, refer to the TerramEarth case study. A new architecture that writes all incoming data to BigQuery has been introduced. You notice that the data is dirty, and want to ensure data quality on an automated daily basis while managing cost.

What should you do?

- A. Set up a streaming Cloud Dataflow job, receiving data by the ingestion process. Clean the data in a Cloud Dataflow pipeline.
- B. Create a Cloud Function that reads data from BigQuery and cleans it. Trigger the Cloud Function from a Compute Engine instance.
- C. Create a SQL statement on the data in BigQuery, and save it as a view. Run the view daily, and save the result to a new table.
- D. Use Cloud Dataprep and configure the BigQuery tables as the source. Schedule a daily job to clean the data.

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

QUESTION 5

For this question, refer to the TerramEarth case study. Considering the technical requirements, how should you reduce the unplanned vehicle downtime in GCP?

- A. Use BigQuery as the data warehouse. Connect all vehicles to the network and stream data into BigQuery using Cloud Pub/Sub and Cloud Dataflow. Use Google Data Studio for analysis and reporting.
- B. Use BigQuery as the data warehouse. Connect all vehicles to the network and upload gzip files to a Multi-Regional Cloud Storage bucket using gcloud. Use Google Data Studio for analysis and reporting.
- C. Use Cloud Dataproc Hive as the data warehouse. Upload gzip files to a MultiRegional Cloud Storage bucket. Upload this data into BigQuery using gcloud. Use Google Data Studio for analysis and reporting.
- D. Use Cloud Dataproc Hive as the data warehouse. Directly stream data into partitioned Hive tables. Use Pig scripts to analyze data.

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

QUESTION 6

For this question, refer to the TerramEarth case study. You are asked to design a new architecture for the ingestion of the data of the 200,000 vehicles that are connected to a cellular network. You want to follow Google-recommended practices.

Considering the technical requirements, which components should you use for the ingestion of the data?

- A. Google Kubernetes Engine with an SSL Ingress
- B. Cloud IoT Core with public/private key pairs
- C. Compute Engine with project-wide SSH keys
- D. Compute Engine with specific SSH keys

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:



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
 - Tomcat - Java micro-services
 - Nginx - static content
 - Apache Beam - Batch processing
 - iSCSI for VM hosts
 - Fiber channel SAN - MySQL databases
 - NAS - image storage, logs, backups
 - Real-time trending calculations
 - 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

The Dress4Win security team has disabled external SSH access into production virtual machines (VMs) on Google Cloud Platform (GCP).

The operations team needs to remotely manage the VMs, build and push Docker containers, and manage Google Cloud Storage objects.

What can they do?

- A. Grant the operations engineer access to use Google Cloud Shell.
- B. Configure a VPN connection to GCP to allow SSH access to the cloud VMs.
- C. Develop a new access request process that grants temporary SSH access to cloud VMs when an operations engineer needs to perform a task.
- D. Have the development team build an API service that allows the operations team to execute specific remote procedure calls to accomplish their tasks.

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

QUESTION 2

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 3

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: C

Section: [none]

Explanation

Explanation/Reference:

QUESTION 4

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

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: D

Section: [none]

Explanation

Explanation/Reference:

QUESTION 5

Dress4Win would like to become familiar with deploying applications to the cloud by successfully deploying some applications quickly, as is. They have asked for your recommendation.

What should you advise?

- A. Identify self-contained applications with external dependencies as a first move to the cloud.
- B. Identify enterprise applications with internal dependencies and recommend these as a first move to the cloud.
- C. Suggest moving their in-house databases to the cloud and continue serving requests to on-premise applications.
- D. Recommend moving their message queuing servers to the cloud and continue handling requests to on-premise applications.

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

QUESTION 6

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 7

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: B

Section: [none]

Explanation

Explanation/Reference:

QUESTION 8

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 9

Dress4Win has end-to-end tests covering 100% of their endpoints.

They want to ensure that the move to the cloud does not introduce any new bugs.

Which additional testing methods should the developers employ to prevent an outage?

- A. They should enable Google Stackdriver Debugger on the application code to show errors in the code.
- B. They should add additional unit tests and production scale load tests on their cloud staging environment.
- C. They should run the end-to-end tests in the cloud staging environment to determine if the code is working as intended.
- D. They should add canary tests so developers can measure how much of an impact the new release causes to latency.

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

QUESTION 10

You want to ensure Dress4Win's sales and tax records remain available for infrequent viewing by auditors for at least 10 years.

Cost optimization is your top priority.

Which cloud services should you choose?

- A. Google Cloud Storage Coldline to store the data, and gsutil to access the data.
- B. Google Cloud Storage Nearline to store the data, and gsutil to access the data.
- C. Google Bigtable with US or EU as location to store the data, and gcloud to access the data.
- D. BigQuery to store the data, and a web server cluster in a managed instance group to access the data. Google Cloud SQL mirrored across two distinct regions to store the data, and a Redis cluster in a managed instance group to access the data.

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/storage/docs/storage-classes>

QUESTION 11

The current Dress4win system architecture has high latency to some customers because it is located in one data center.

As of a future evaluation and optimizing for performance in the cloud, Dress4win wants to distribute its system architecture to multiple locations when Google cloud platform.

Which approach should they use?

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

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:



Testlet 1

Company Overview

Dress4win is a web-based company that helps their users organize and manage their personal wardrobe using a webapp 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 freemium app model. The application has grown from a few servers in the founder's garage to several hundred servers and appliances in a colocated 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 moving their development and test environments. They are also 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. All servers run Ubuntu LTS v16.04.

Databases:

- MySQL. 1 server for user data, inventory, static data:
 - MySQL 5.8
 - 8 core CPUs
 - 128 GB of RAM
 - 2x 5 TB HDD (RAID 1)
- Redis 3 server cluster for metadata, social graph, caching. Each server is:
 - Redis 3.2
 - 4 core CPUs - 32GB of RAM

Compute:

- 40 Web Application servers providing micro-services based APIs and static content.
 - Tomcat - Java
 - Nginx
 - 4 core CPUs
 - 32 GB of RAM
- 20 Apache Hadoop/Spark servers:
 - Data analysis

- Real-time trending calculations
- 8 core CPUS
- 128 GB of RAM
- 4x 5 TB HDD (RAID 1)
- 3 RabbitMQ servers for messaging, social notifications, and events:
 - 8 core CPUs
 - 32GB of RAM
- Miscellaneous servers:
 - Jenkins, monitoring, bastion hosts, security scanners
 - 8 core CPUs - 32GB of RAM

Storage appliances: ▪

iSCSI for VM hosts

- Fiber channel SAN – MySQL databases
- 1 PB total storage; 400 TB available ▪ NAS – image storage, logs, backups
- 100 TB total storage; 35 TB available

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.

Technical Requirements

- Easily create non-production environment in the cloud.
 - Implement an automation framework for provisioning resources in cloud.
 - Implement a continuous deployment process for deploying applications to the on-premises datacenter or cloud. ▪
- Support failover of the production environment to cloud during an emergency.
- Encrypt data on the wire and at rest.
 - Support multiple private connections between the production data center and cloud environment.

Executive Statement

Our investors are concerned about our ability to scale and contain costs with our current infrastructure. They are also concerned that a competitor could use a public cloud platform to offset their up-front investment and free them to focus on developing better features. Our traffic patterns are highest in the mornings and weekend evenings; during other times, 80% of our capacity is sitting idle.

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 for a public cloud strategy achieves a cost reduction between 30% and 50% over our current model.



For this question, refer to the Dress4Win case stud

QUESTION 1

y. Dress4Win is expected to grow to 10 times its size in 1 year with a corresponding growth in data and traffic that mirrors the existing patterns of usage. The CIO has set the target of migrating production infrastructure to the cloud within the next 6 months. How will you configure the solution to scale for this growth without making major application changes and still maximize the ROI?

- A. Migrate the web application layer to App Engine, and MySQL to Cloud Datastore, and NAS to Cloud Storage. Deploy RabbitMQ, and deploy Hadoop servers using Deployment Manager.
- B. Migrate RabbitMQ to Cloud Pub/Sub, Hadoop to BigQuery, and NAS to Compute Engine with Persistent Disk storage. Deploy Tomcat, and deploy Nginx using Deployment Manager.
- C. Implement managed instance groups for Tomcat and Nginx. Migrate MySQL to Cloud SQL, RabbitMQ to Cloud Pub/Sub, Hadoop to Cloud Dataproc, and NAS to Compute Engine with Persistent Disk storage.
- D. Implement managed instance groups for the Tomcat and Nginx. Migrate MySQL to Cloud SQL, RabbitMQ to Cloud Pub/Sub, Hadoop to Cloud Dataproc, and NAS to Cloud Storage.

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:



QUESTION 2

For this question, refer to the Dress4Win case study. Considering the given business requirements, how would you automate the deployment of web and transactional data layers?

- A. Deploy Nginx and Tomcat using Cloud Deployment Manager to Compute Engine. Deploy a Cloud SQL server to replace MySQL. Deploy Jenkins using Cloud Deployment Manager.
- B. Deploy Nginx and Tomcat using Cloud Launcher. Deploy a MySQL server using Cloud Launcher. Deploy Jenkins to Compute Engine using Cloud Deployment Manager scripts.
- C. Migrate Nginx and Tomcat to App Engine. Deploy a Cloud Datastore server to replace the MySQL server in a high-availability configuration. Deploy Jenkins to Compute Engine using Cloud Launcher.
- D. Migrate Nginx and Tomcat to App Engine. Deploy a MySQL server using Cloud Launcher. Deploy Jenkins to Compute Engine using Cloud Launcher.

For this question, refer to the Dress4Win case stud

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

QUESTION 3

y. Which of the compute services should be migrated as-is and would still be an optimized architecture for performance in the cloud?

- A. Web applications deployed using App Engine standard environment
- B. RabbitMQ deployed using an unmanaged instance group
- C. Hadoop/Spark deployed using Cloud Dataproc Regional in High Availability mode
- D. Jenkins, monitoring, bastion hosts, security scanners services deployed on custom machine types

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:



QUESTION 4

For this question, refer to the Dress4Win case study. To be legally compliant during an audit, Dress4Win must be able to give insights in all administrative actions that modify the configuration or metadata of resources on Google Cloud.

What should you do?

- A. Use Stackdriver Trace to create a trace list analysis.
- B. Use Stackdriver Monitoring to create a dashboard on the project's activity.
- C. Enable Cloud Identity-Aware Proxy in all projects, and add the group of Administrators as a member.
- D. Use the Activity page in the GCP Console and Stackdriver Logging to provide the required insight.

Correct Answer: D

Section: [none]

Explanation

For this question, refer to the Dress4Win case stud

Explanation/Reference:



<https://vceplus.com/>

QUESTION 5

y. You are responsible for the security of data stored in Cloud Storage for your company, Dress4Win. You have already created a set of Google Groups and assigned the appropriate users to those groups. You should use Google best practices and implement the simplest design to meet the requirements.

Considering Dress4Win's business and technical requirements, what should you do?

- A. Assign custom IAM roles to the Google Groups you created in order to enforce security requirements. Encrypt data with a customer-supplied encryption key when storing files in Cloud Storage.
- B. Assign custom IAM roles to the Google Groups you created in order to enforce security requirements. Enable default storage encryption before storing files in Cloud Storage.
- C. Assign predefined IAM roles to the Google Groups you created in order to enforce security requirements. Utilize Google's default encryption at rest when storing files in Cloud Storage.
- D. Assign predefined IAM roles to the Google Groups you created in order to enforce security requirements. Ensure that the default Cloud KMS key is set before storing files in Cloud Storage.

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

QUESTION 6

For this question, refer to the Dress4Win case stud

For this question, refer to the Dress4Win case study. You want to ensure that your on-premises architecture meets business requirements before you migrate your solution.

What change in the on-premises architecture should you make?

- A. Replace RabbitMQ with Google Pub/Sub.
- B. Downgrade MySQL to v5.7, which is supported by Cloud SQL for MySQL.
- C. Resize compute resources to match predefined Compute Engine machine types.
- D. Containerize the micro-services and host them in Google Kubernetes Engine.

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:



Question Set 1

QUESTION 1

Your company's test suite is a custom C++ application that runs tests throughout each day on Linux virtual machines. The full test suite takes several hours to complete, running on a limited number of on-premises servers reserved for testing. Your company wants to move the testing infrastructure to the cloud, to reduce the amount of time it takes to fully test a change to the system, while changing the tests as little as possible.

Which cloud infrastructure should you recommend?

- A. Google Compute Engine unmanaged instance groups and Network Load Balancer
- B. Google Compute Engine managed instance groups with auto-scaling
- C. Google Cloud Dataproc to run Apache Hadoop jobs to process each test
- D. Google App Engine with Google StackDriver for logging

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Google Compute Engine enables users to launch virtual machines (VMs) on demand. VMs can be launched from the standard images or custom images created by users.

Managed instance groups offer autoscaling capabilities that allow you to automatically add or remove instances from a managed instance group based on increases or decreases in load. Autoscaling helps your applications gracefully handle increases in traffic and reduces cost when the need for resources is lower.

Incorrect Answers:

B: There is no mention of incoming IP data traffic for the custom C++ applications.

C: Apache Hadoop is not fit for testing C++ applications. Apache Hadoop is an open-source software framework used for distributed storage and processing of datasets of big data using the MapReduce programming model.

D: Google App Engine is intended to be used for web applications.

Google App Engine (often referred to as GAE or simply App Engine) is a web framework and cloud computing platform for developing and hosting web applications in Google-managed data centers.

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

QUESTION 2

A recent audit revealed that a new network was created in your GCP project. In this network, a GCE instance has an SSH port open to the world. You want to discover this network's origin.

What should you do?

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

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Incorrect Answers:

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

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

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

Cloud Audit Logging returns two types of logs:

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

QUESTION 3

Your company runs several databases on a single MySQL instance. They need to take backups of a specific database at regular intervals. The backup activity needs to complete as quickly as possible and cannot be allowed to impact disk performance.

How should you configure the storage?

- A. Configure a cron job to use the gcloud tool to take regular backups using persistent disk snapshots.
- B. Mount a Local SSD volume as the backup location. After the backup is complete, use gsutil to move the backup to Google Cloud Storage.
- C. Use gcsfuse to mount a Google Cloud Storage bucket as a volume directly on the instance and write backups to the mounted location using mysqldump.
- D. Mount additional persistent disk volumes onto each virtual machine (VM) instance in a RAID10 array and use LVM to create snapshots to send to Cloud Storage

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

QUESTION 4

You want to enable your running Google **Kubernetes** Engine cluster to scale as demand for your application changes.

What should you do?

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

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

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

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

Which database type should you use?

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

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Relational databases were not designed to cope with the scale and agility challenges that face modern applications, nor were they built to take advantage of the commodity storage and processing power available today. NoSQL fits well for:

- Developers are working with applications that create massive volumes of new, rapidly changing data types — structured, semi-structured, unstructured and polymorphic data.

Incorrect Answers:

D: The Blobstore API allows your application to serve data objects, called blobs, that are much larger than the size allowed for objects in the Datastore service. Blobs are useful for serving large files, such as video or image files, and for allowing users to upload large data files.

Reference: <https://www.mongodb.com/nosql-explained>

QUESTION 7

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. Reference:

<https://cloud.google.com/logging/>

QUESTION 8

A production database virtual machine on Google Compute Engine has an `ext4`-formatted persistent disk for data files. The database is about to run out of storage space.

How can you remediate the problem with the least amount of downtime?

- A. In the Cloud Platform Console, increase the size of the persistent disk and use the `resize2fs` command in Linux.
- B. Shut down the virtual machine, use the Cloud Platform Console to increase the persistent disk size, then restart the virtual machine
- C. In the Cloud Platform Console, increase the size of the persistent disk and verify the new space is ready to use with the `fdisk` command in Linux

- D. In the Cloud Platform Console, create a new persistent disk attached to the virtual machine, format and mount it, and configure the database service to move the files to the new disk
- E. In the Cloud Platform Console, create a snapshot of the persistent disk restore the snapshot to a new larger disk, unmount the old disk, mount the new disk and restart the database service

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

On Linux instances, connect to your instance and manually resize your partitions and file systems to use the additional disk space that you added.

Extend the file system on the disk or the partition to use the added space. If you grew a partition on your disk, specify the partition. If your disk does not have a partition table, specify only the disk ID.

sudo resize2fs /dev/[DISK_ID][PARTITION_NUMBER] where [DISK_ID] is the device name and [PARTITION_NUMBER] is the partition number for the device where you are resizing the file system. Reference: <https://cloud.google.com/compute/docs/disks/add-persistent-disk>

QUESTION 9

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. Reference:

<https://cloud.google.com/dataproc/docs/resources/faq>

QUESTION 10

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

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

QUESTION 11

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

What should you do?

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

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

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

QUESTION 12

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 13

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 14

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.

Reference: <https://cloud.google.com/docs/compare/openstack/> <https://aws.amazon.com/it/blogs/aws/building-three-tier-architectures-with-security-groups/>

QUESTION 15

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 16

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: AE

Section: [none]

Explanation

Explanation/Reference:

QUESTION 17

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)
```

- 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 18

Your customer support tool logs all email and chat conversations to Cloud Bigtable for retention and analysis. What is the recommended approach for sanitizing this data of personally identifiable information or payment card information before initial storage?

- A. Hash all data using SHA256
- B. Encrypt all data using elliptic curve cryptography
- C. De-identify the data with the Cloud Data Loss Prevention API
- D. Use regular expressions to find and redact phone numbers, email addresses, and credit card numbers

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Reference: https://cloud.google.com/solutions/pai-dss-compliance-in-gcp#using_data_loss_prevention_api_to_sanitize_data

QUESTION 19

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

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

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

QUESTION 20

You want to create a private connection between your instances on Compute Engine and your on-premises data center. You require a connection of at least 20 Gbps. You want to follow Google-recommended practices. How should you set up the connection?

- A. Create a VPC and connect it to your on-premises data center using Dedicated Interconnect.
- B. Create a VPC and connect it to your on-premises data center using a single Cloud VPN.
- C. Create a Cloud Content Delivery Network (Cloud CDN) and connect it to your on-premises data center using Dedicated Interconnect.
- D. Create a Cloud Content Delivery Network (Cloud CDN) and connect it to your on-premises datacenter using a single Cloud VPN.

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:



QUESTION 21

You are analyzing and defining business processes to support your startup's trial usage of GCP, and you don't yet know what consumer demand for your product will be. Your manager requires you to minimize GCP service costs and adhere to Google best practices. What should you do?

- A. Utilize free tier and sustained use discounts. Provision a staff position for service cost management.
- B. Utilize free tier and sustained use discounts. Provide training to the team about service cost management.
- C. Utilize free tier and committed use discounts. Provision a staff position for service cost management.
- D. Utilize free tier and committed use discounts. Provide training to the team about service cost management.

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

QUESTION 22

You are building a continuous deployment pipeline for a project stored in a Git source repository and want to ensure that code changes can be verified deploying to production. What should you do?

- A. Use Spinnaker to deploy builds to production using the red/black deployment strategy so that changes can easily be rolled back.
- B. Use Spinnaker to deploy builds to production and run tests on production deployments.
- C. Use Jenkins to build the staging branches and the master branch. Build and deploy changes to production for 10% of users before doing a complete rollout.
- D. Use Jenkins to monitor tags in the repository. Deploy staging tags to a staging environment for testing. After testing, tag the repository for production and deploy that to the production environment.

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Reference: <https://github.com/GoogleCloudPlatform/continuous-deployment-on-kubernetes/blob/master/README.md>

QUESTION 23

Your company is migrating its on-premises data center into the cloud. As part of the migration, you want to integrate Google Kubernetes Engine (GKE) for workload orchestration. Parts of your architecture must also be PCI DSS-compliant. Which of the following is most accurate?

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

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

QUESTION 24

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

- A. The effective policy is determined only by the policy set at the node
- B. The effective policy is the policy set at the node and restricted by the policies of its ancestors

- C. The effective policy is the union of the policy set at the node and policies inherited from its ancestors
- D. The effective policy is the intersection of the policy set at the node and policies inherited from its ancestors

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

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

QUESTION 25

You have an application that will run on Compute Engine. You need to design an architecture that takes into account a disaster recovery plan that requires your application to fail over to another region in case of a regional outage. What should you do?

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

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

QUESTION 26

You are deploying an application on App Engine that needs to integrate with an on-premises database. For security purposes, your on-premises database must not be accessible through the public internet. What should you do?

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

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

QUESTION 27

You need to set up Microsoft SQL Server on GCP. Management requires that there's no downtime in case of a data center outage in any of the zones within a GCP region. What should you do?

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

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:



QUESTION 28

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: B

Section: [none]

Explanation

Explanation/Reference:

QUESTION 29

You need to evaluate your team readiness for a new GCP project. You must perform the evaluation and create a skills gap plan which 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 30

You are designing an application for use only during business hours. For the minimum viable product release, you'd like to use a managed product that automatically "scales to zero" so you don't incur costs when there is no activity. Which primary compute resource should you choose?

- A. Cloud Functions
- B. Compute Engine
- C. Google Kubernetes Engine
- D. AppEngine flexible environment

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

QUESTION 31

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: A**

Section: [none]

Explanation

Explanation/Reference:

QUESTION 32

Your customer wants to capture multiple GBs of aggregate real-time key performance indicators (KPIs) from their game servers running on Google Cloud Platform and monitor the KPIs with low latency. How should they capture the KPIs?

- A. Store time-series data from the game servers in Google Bigtable, and view it using Google Data Studio.
- B. Output custom metrics to Stackdriver from the game servers, and create a Dashboard in Stackdriver Monitoring Console to view them.
- C. Schedule BigQuery load jobs to ingest analytics files uploaded to Cloud Storage every ten minutes, and visualize the results in Google Data Studio.
- D. Insert the KPIs into Cloud Datastore entities, and run ad hoc analysis and visualizations of them in Cloud Datalab.

Correct Answer: A

Section: [none]

Explanation



Explanation/Reference:

Reference: <https://cloud.google.com/solutions/data-lifecycle-cloud-platform>

QUESTION 33

You are using a single Cloud SQL instance to serve your application from a specific zone. You want to introduce high availability. What should you do?

- A. Create a read replica instance in a different region
- B. Create a failover replica instance in a different region
- C. Create a read replica instance in the same region, but in a different zone
- D. Create a failover replica instance in the same region, but in a different zone

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

QUESTION 34

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

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

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

QUESTION 35

You are using Cloud SQL as the database backend for a large CRM deployment. You want to scale as usage increases and ensure that you don't run out of storage, maintain 75% CPU usage cores, and keep replication lag below 60 seconds. What are the correct steps to meet your requirements?

- A.
 1. Enable automatic storage increase for the instance.
 2. Create a Stackdriver alert when CPU usage exceeds 75%, and change the instance type to reduce CPU usage.
 3. Create a Stackdriver alert for replication lag, and shard the database to reduce replication time.
- B.
 1. Enable automatic storage increase for the instance.
 2. Change the instance type to a 32-core machine type to keep CPU usage below 75%.
 3. Create a Stackdriver alert for replication lag, and deploy memcache to reduce load on the master.
- C.
 1. Create a Stackdriver alert when storage exceeds 75%, and increase the available storage on the instance to create more space.
 2. Deploy memcached to reduce CPU load.
 3. Change the instance type to a 32-core machine type to reduce replication lag.
- D.
 1. Create a Stackdriver alert when storage exceeds 75%, and increase the available storage on the instance to create more space.
 2. Deploy memcached to reduce CPU load.
 3. Create a Stackdriver alert for replication lag, and change the instance type to a 32-core machine type to reduce replication lag.

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

QUESTION 36

You have deployed an application to Google Kubernetes Engine (GKE), and are using the Cloud SQL proxy container to make the Cloud SQL database available to the services running on Kubernetes. You are notified that the application is reporting database connection issues. Your company policies require a post-mortem. What should you do?

- A. Use `gcloud sql instances restart`.
- B. Validate that the Service Account used by the Cloud SQL proxy container still has the Cloud Build Editor role.
- C. In the GCP Console, navigate to Stackdriver Logging. Consult logs for (GKE) and Cloud SQL.
- D. In the GCP Console, navigate to Cloud SQL. Restore the latest backup. Use `kubectl` to restart all pods.

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

QUESTION 37

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

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

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

QUESTION 38

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

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 39

Your company operates nationally and plans to use GCP for multiple batch workloads, including some that are not time-critical. You also need to use GCP services that are HIPAA-certified and manage service costs.

How should you design to meet Google best practices?

- A. Provision preemptible VMs to reduce cost. Discontinue use of all GCP services and APIs that are not HIPAA-compliant.
- B. Provision preemptible VMs to reduce cost. Disable and then discontinue use of all GCP services and APIs that are not HIPAA-compliant.
- C. Provision standard VMs in the same region to reduce cost. Discontinue use of all GCP services and APIs that are not HIPAA-compliant.
- D. Provision standard VMs to the same region to reduce cost. Disable and then discontinue use of all GCP services and APIs that are not HIPAA-compliant.

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

QUESTION 40

Your customer wants to do resilience testing of their authentication layer. This consists of a regional managed instance group serving a public REST API that reads from and writes to a Cloud SQL instance.

What should you do?

- A. Engage with a security company to run web scrapers that look your for users' authentication data om malicious websites and notify you if any is found.

- B. Deploy intrusion detection software to your virtual machines to detect and log unauthorized access.
- C. Schedule a disaster simulation exercise during which you can shut off all VMs in a zone to see how your application behaves.
- D. Configure a read replica for your Cloud SQL instance in a different zone than the master, and then manually trigger a failover while monitoring KPIs for our REST API.

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

QUESTION 41

You want to automate the creation of a managed instance group. The VMs have many OS package dependencies. You want to minimize the startup time for new VMs in the instance group.

What should you do?

- A. Use Terraform to create the managed instance group and a startup script to install the OS package dependencies.
- B. Create a custom VM image with all OS package dependencies. Use Deployment Manager to create the managed instance group with the VM image.
- C. Use Puppet to create the managed instance group and install the OS package dependencies.
- D. Use Deployment Manager to create the managed instance group and Ansible to install the OS package dependencies.

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

QUESTION 42

Your architecture calls for the centralized collection of all admin activity and VM system logs within your project.

How should you collect these logs from both VMs and services?

- A. All admin and VM system logs are automatically collected by Stackdriver.
- B. Stackdriver automatically collects admin activity logs for most services. The Stackdriver Logging agent must be installed on each instance to collect system logs.
- C. Launch a custom syslogd compute instance and configure your GCP project and VMs to forward all logs to it.

D. Install the Stackdriver Logging agent on a single compute instance and let it collect all audit and access logs for your environment.

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

QUESTION 43

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: B

Section: [none]

Explanation

Explanation/Reference:



QUESTION 44

Your customer runs a web service used by e-commerce sites to offer product recommendations to users. The company has begun experimenting with a machine learning model on Google Cloud Platform to improve the quality of results.

What should the customer do to improve their model's results over time?

- A. Export Cloud Machine Learning Engine performance metrics from Stackdriver to BigQuery, to be used to analyze the efficiency of the model.
- B. Build a roadmap to move the machine learning model training from Cloud GPUs to Cloud TPUs, which offer better results.
- C. Monitor Compute Engine announcements for availability of newer CPU architectures, and deploy the model to them as soon as they are available for additional performance.
- D. Save a history of recommendations and results of the recommendations in BigQuery, to be used as training data.

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

QUESTION 45

A development team at your company has created a dockerized HTTPS web application. You need to deploy the application on Google Kubernetes Engine (GKE) and make sure that the application scales automatically.

How should you deploy to GKE?

- A. Use the Horizontal Pod Autoscaler and enable cluster autoscaling. Use an Ingress resource to load-balance the HTTPS traffic.
- B. Use the Horizontal Pod Autoscaler and enable cluster autoscaling on the Kubernetes cluster. Use a Service resource of type LoadBalancer to load-balance the HTTPS traffic.
- C. Enable autoscaling on the Compute Engine instance group. Use an Ingress resource to load balance the HTTPS traffic.
- D. Enable autoscaling on the Compute Engine instance group. Use a Service resource of type LoadBalancer to load-balance the HTTPS traffic.

Correct Answer: B

Section: [none]

Explanation



Explanation/Reference:

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

QUESTION 46

Your company creates rendering software which users can download from the company website. Your company has customers all over the world. You want to minimize latency for all your customers. You want to follow Google-recommended practices.

How should you store the files?

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

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

QUESTION 47

Your company acquired a healthcare startup and must retain its customers' medical information for up to 4 more years, depending on when it was created. Your corporate policy is to securely retain this data, and then delete it as soon as regulations allow.

Which approach should you take?

- A. Store the data in Google Drive and manually delete records as they expire.
- B. Anonymize the data using the Cloud Data Loss Prevention API and store it indefinitely.
- C. Store the data in Cloud Storage and use lifecycle management to delete files when they expire.
- D. Store the data in Cloud Storage and run a nightly batch script that deletes all expired data.

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:



<https://vceplus.com/>