

Google.Professional Cloud Architect.vJun-2024.by.Mark.143q

Number: Professional Cloud Architect.
Passing Score: 800
Time Limit: 120
File Version: 7.0

Exam Code: Professional Cloud Architect
Exam Name: Professional Cloud Architect on Google Cloud Platform



Mountkirk Games A Company Overview

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

Company Background

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

Solution Concept

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

Technical Requirements

Requirements for Game Backend Platform

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

Requirements for Game Analytics Platform

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



CEO Statement

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

CTO Statement

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

CFO Statement

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

QUESTION 1

Mountkirk Games 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:

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 through testing process for new versions of the backend before they are released to the public. You want the testing environment to scale in an economical way. How should you design the process?

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

Correct Answer: A

Section:

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:

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:

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:

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:

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

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

Mountkirk Games B

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: A, B

Section:

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:

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 Engine to improve the ability 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: C, E

Section:

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:

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:

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:

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:

Explanation:

JencoMart

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:

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.

References: <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: C, D, F

Section:

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 browser-based SSH sessions are dynamically allocated by GCP Console and can vary from session to session.

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

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

References:

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

QUESTION 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:

Explanation:

Migrating data to Google Cloud Platform

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

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

QUESTION 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:

Explanation:

From scenario:

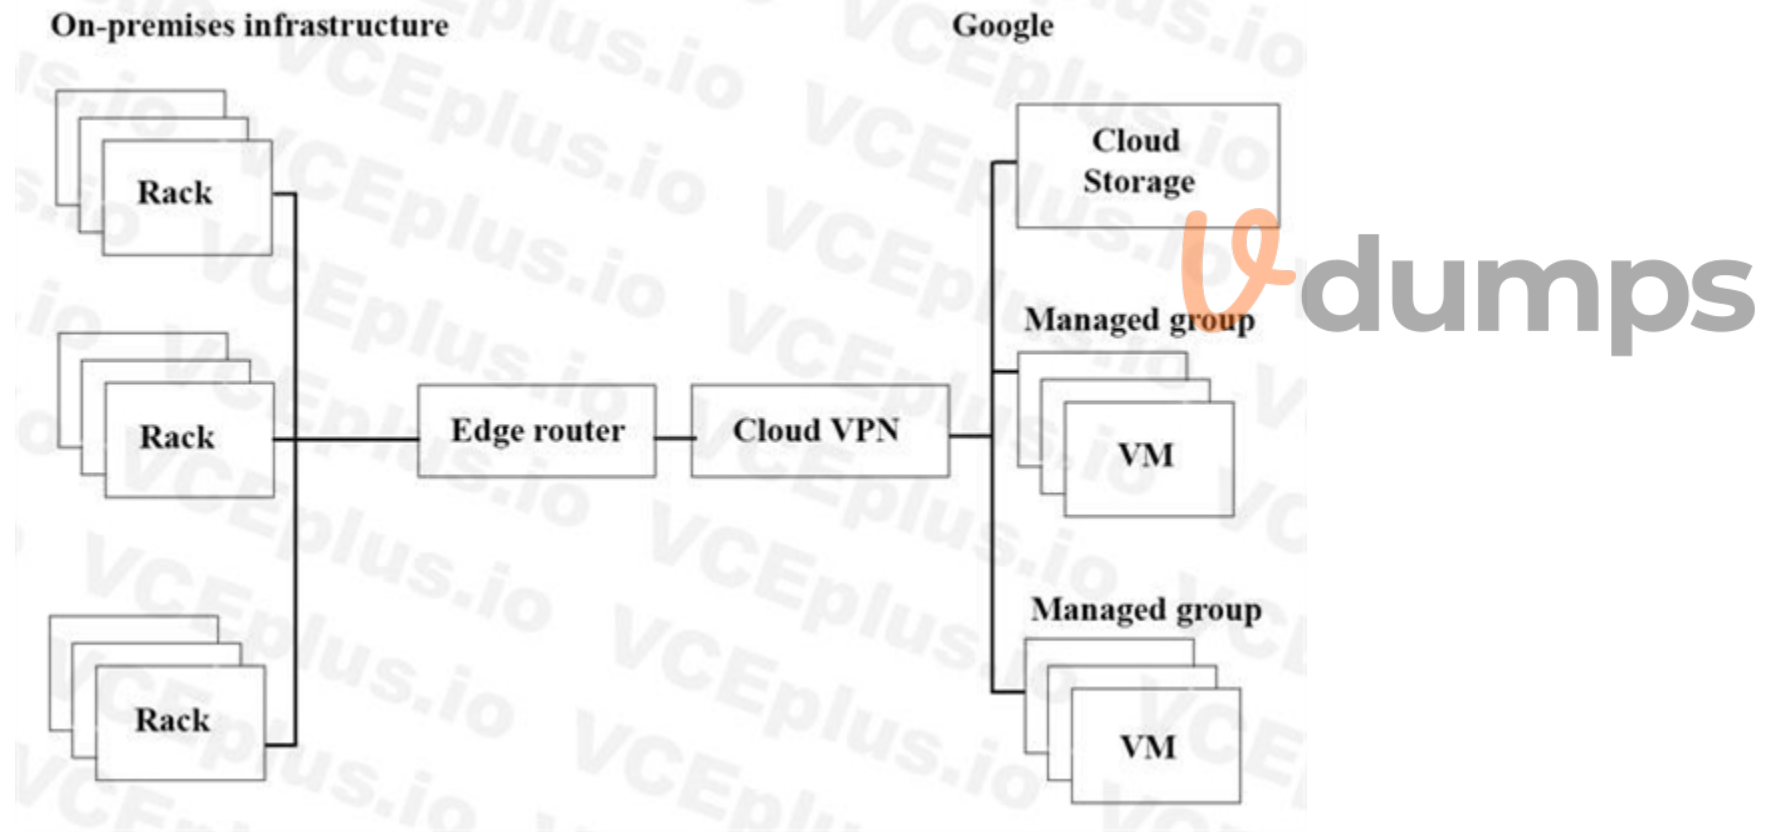
Business Requirements include: Expand services into Asia

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: A, C, E

Section:

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:

Explanation:

Common workloads for Google Cloud Datastore:

User profiles

Product catalogs

Game state

References: <https://cloud.google.com/storage-options/> <https://cloud.google.com/datastore/docs/concepts/overview>"f

Mountkirk Games C

Company overview

Mountkirk Games makes online, session-based, multiplayer games for mobile platforms. They have recently started expanding to other platforms after successfully migrating their on-premises environments to Google Cloud. Their most recent endeavor is to create a retro-style first-person shooter (FPS) game that allows hundreds of simultaneous players to join a geo-specific digital arena from multiple platforms and locations. A real-time digital banner will display a global leaderboard of all the top players across every active arena.

Solution concept

Mountkirk Games is building a new multiplayer game that they expect to be very popular. They plan to deploy the game's backend on Google Kubernetes Engine so they can scale rapidly and use Google's global load balancer to route players to the closest regional game arenas. In order to keep the global leader board in sync, they plan to use a multi-region Spanner cluster.

Existing technical environment

The existing environment was recently migrated to Google Cloud, and five games came across using lift-and-shift virtual machine migrations, with a few minor exceptions. Each new game exists in an isolated Google Cloud project nested below a folder that maintains most of the permissions and network policies. Legacy games with low traffic have been consolidated into a single project. There are also separate environments for development and testing.

Business requirements

Support multiple gaming platforms.

Support multiple regions.

Support rapid iteration of game features.

Minimize latency.

Optimize for dynamic scaling.

Use managed services and pooled resources.

Minimize costs.

Technical requirements

Dynamically scale based on game activity.

Publish scoring data on a near real-time global leaderboard.

Store game activity logs in structured files for future analysis.

Use GPU processing to render graphics server-side for multi-platform support.

Support eventual migration of legacy games to this new platform.

Executive statement

Our last game was the first time we used Google Cloud, and it was a tremendous success. We were able to analyze player behavior and game telemetry in ways that we never could before. This success allowed us to bet on a full migration to the cloud and to start building all-new games using cloud-native design principles. Our new game is our most ambitious to date and will open up doors for us to support more gaming platforms beyond mobile. Latency is our top priority, although cost management is the next most important challenge. As with our first cloud-based game, we have grown to expect the cloud to enable advanced analytics capabilities so we can rapidly iterate on our deployments of bug fixes and new functionality.

QUESTION 1

You need to optimize batch file transfers into Cloud Storage for Mountkirk Games' new Google Cloud solution. The batch files contain game statistics that need to be staged in Cloud Storage and be processed by an extract transform load (ETL) tool. What should you do?

- A. Use gsutil to batch move files in sequence.
- B. Use gsutil to batch copy the files in parallel.
- C. Use gsutil to extract the files as the first part of ETL.
- D. Use gsutil to load the files as the last part of ETL.

Correct Answer: B

Section:

Explanation:

Reference: <https://cloud.google.com/storage/docs/gsutil/commands/cp>

QUESTION 2

You are implementing Firestore for Mountkirk Games. Mountkirk Games wants to give a new game programmatic access to a legacy game's Firestore database. Access should be as restricted as possible. What should you do?

- A. Create a service account (SA) in the legacy game's Google Cloud project, add a second SA in the new game's IAM page, and then give the Organization Admin role to both SAs.
- B. Create a service account (SA) in the legacy game's Google Cloud project, give the SA the Organization Admin role, and then give it the Firebase Admin role in both projects.
- C. Create a service account (SA) in the legacy game's Google Cloud project, add this SA in the new game's IAM page, and then give it the Firebase Admin role in both projects.
- D. Create a service account (SA) in the legacy game's Google Cloud project, give it the Firebase Admin role, and then migrate the new game to the legacy game's project.

Correct Answer: C

Section:



QUESTION 3

Mountkirk Games wants to limit the physical location of resources to their operating Google Cloud regions. What should you do?

- A. Configure an organizational policy which constrains where resources can be deployed.
- B. Configure IAM conditions to limit what resources can be configured.
- C. Configure the quotas for resources in the regions not being used to 0.
- D. Configure a custom alert in Cloud Monitoring so you can disable resources as they are created in other regions.

Correct Answer: C

Section:

QUESTION 4

You need to implement a network ingress for a new game that meets the defined business and technical requirements. Mountkirk Games wants each regional game instance to be located in multiple Google Cloud regions. What should you do?

- A. Configure a global load balancer connected to a managed instance group running Compute Engine instances.
- B. Configure kubemci with a global load balancer and Google Kubernetes Engine.
- C. Configure a global load balancer with Google Kubernetes Engine.
- D. Configure Ingress for Anthos with a global load balancer and Google Kubernetes Engine.

Correct Answer: A

Section:

QUESTION 5

Your development teams release new versions of games running on Google Kubernetes Engine (GKE) daily. You want to create service level indicators (SLIs) to evaluate the quality of the new versions from the user's perspective. What should you do?

- A. Create CPU Utilization and Request Latency as service level indicators.
- B. Create GKE CPU Utilization and Memory Utilization as service level indicators.
- C. Create Request Latency and Error Rate as service level indicators.
- D. Create Server Uptime and Error Rate as service level indicators.

Correct Answer: A

Section:

QUESTION 6

Mountkirk Games wants you to secure the connectivity from the new gaming application platform to Google Cloud. You want to streamline the process and follow Google-recommended practices. What should you do?

- A. Configure Workload Identity and service accounts to be used by the application platform.
- B. Use Kubernetes Secrets, which are obfuscated by default. Configure these Secrets to be used by the application platform.
- C. Configure Kubernetes Secrets to store the secret, enable Application-Layer Secrets Encryption, and use Cloud Key Management Service (Cloud KMS) to manage the encryption keys. Configure these Secrets to be used by the application platform.
- D. Configure HashiCorp Vault on Compute Engine, and use customer managed encryption keys and Cloud Key Management Service (Cloud KMS) to manage the encryption keys. Configure these Secrets to be used by the application platform.

Correct Answer: A

Section:



QUESTION 7

Your development team has created a mobile game app. You want to test the new mobile app on Android and iOS devices with a variety of configurations. You need to ensure that testing is efficient and cost-effective. What should you do?

- A. Upload your mobile app to the Firebase Test Lab, and test the mobile app on Android and iOS devices.
- B. Create Android and iOS VMs on Google Cloud, install the mobile app on the VMs, and test the mobile app.
- C. Create Android and iOS containers on Google Kubernetes Engine (GKE), install the mobile app on the containers, and test the mobile app.
- D. Upload your mobile app with different configurations to Firebase Hosting and test each configuration.

Correct Answer: C

Section:

Explanation:

TerramEarth A

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.

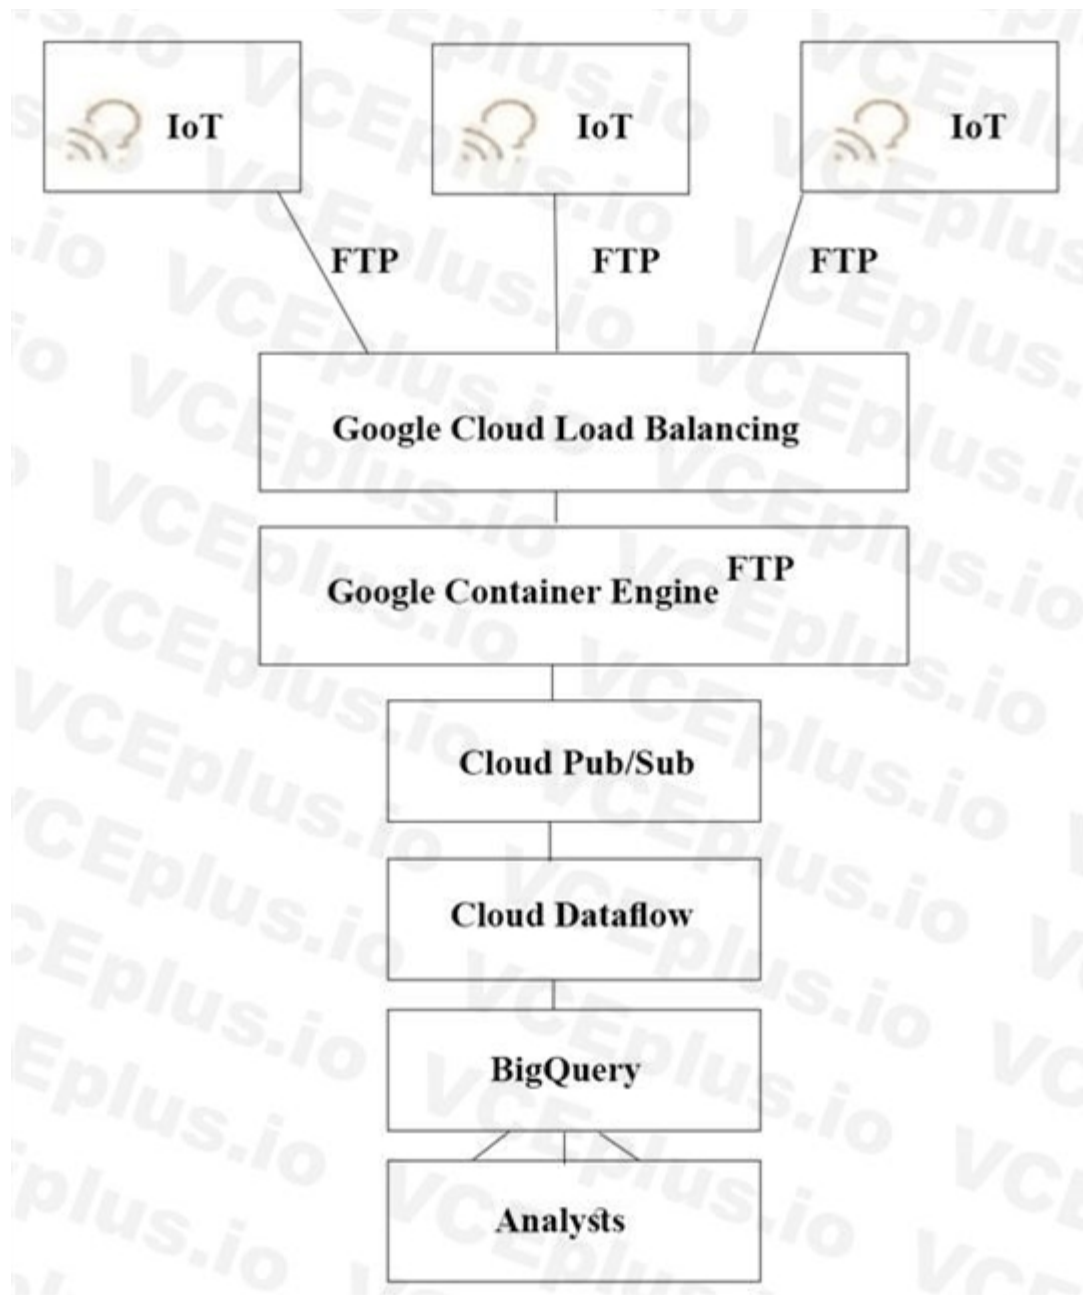
CTO Statement

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.

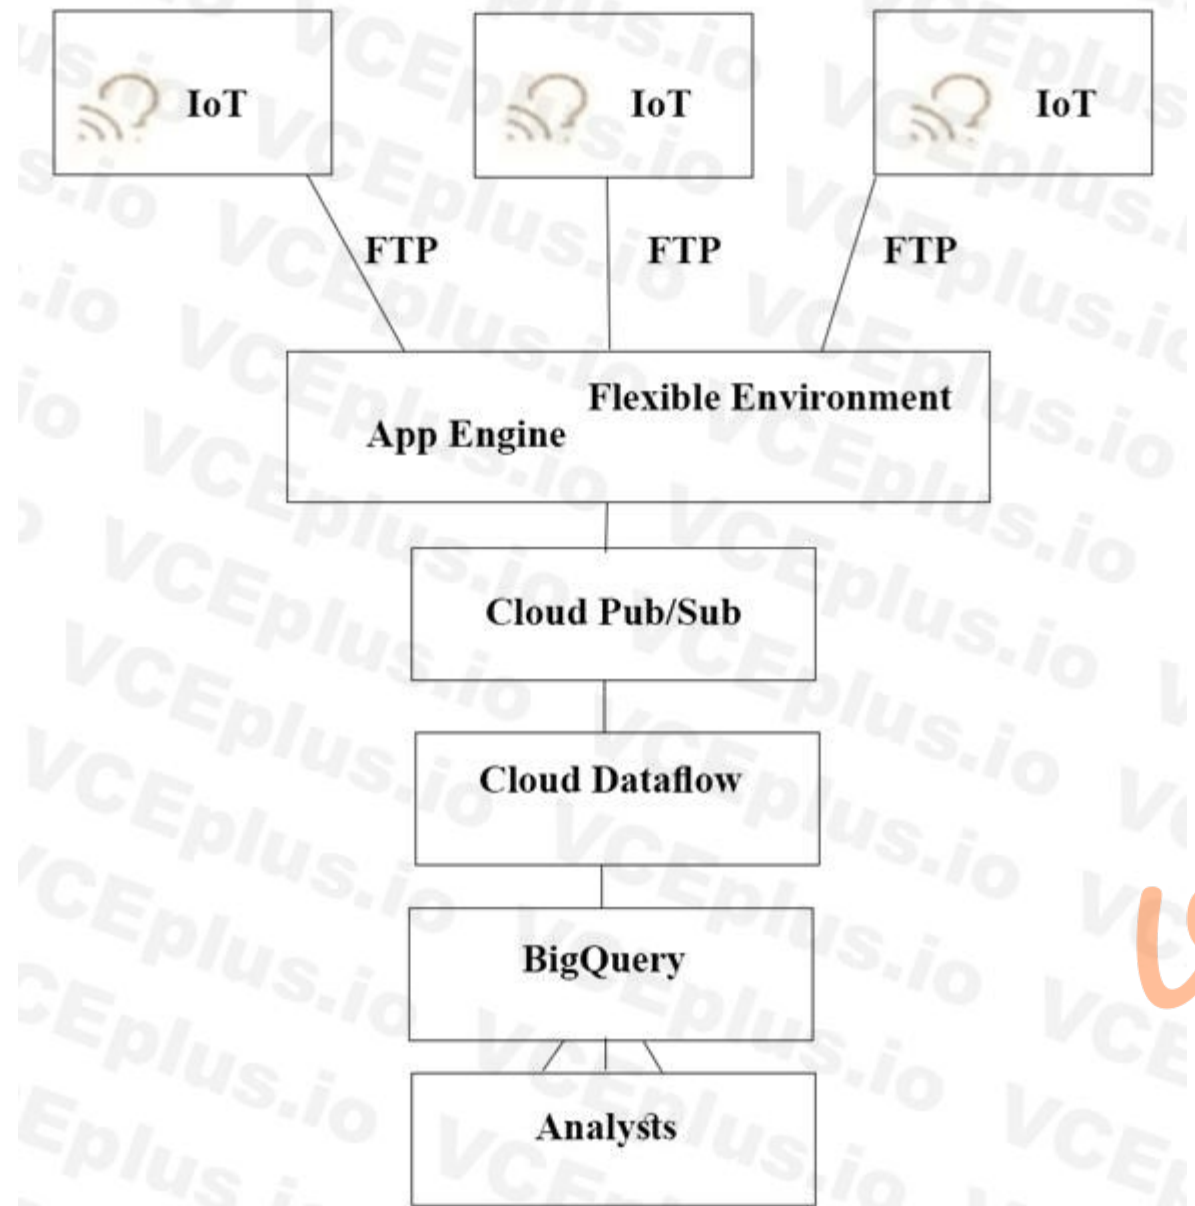
QUESTION 1

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?

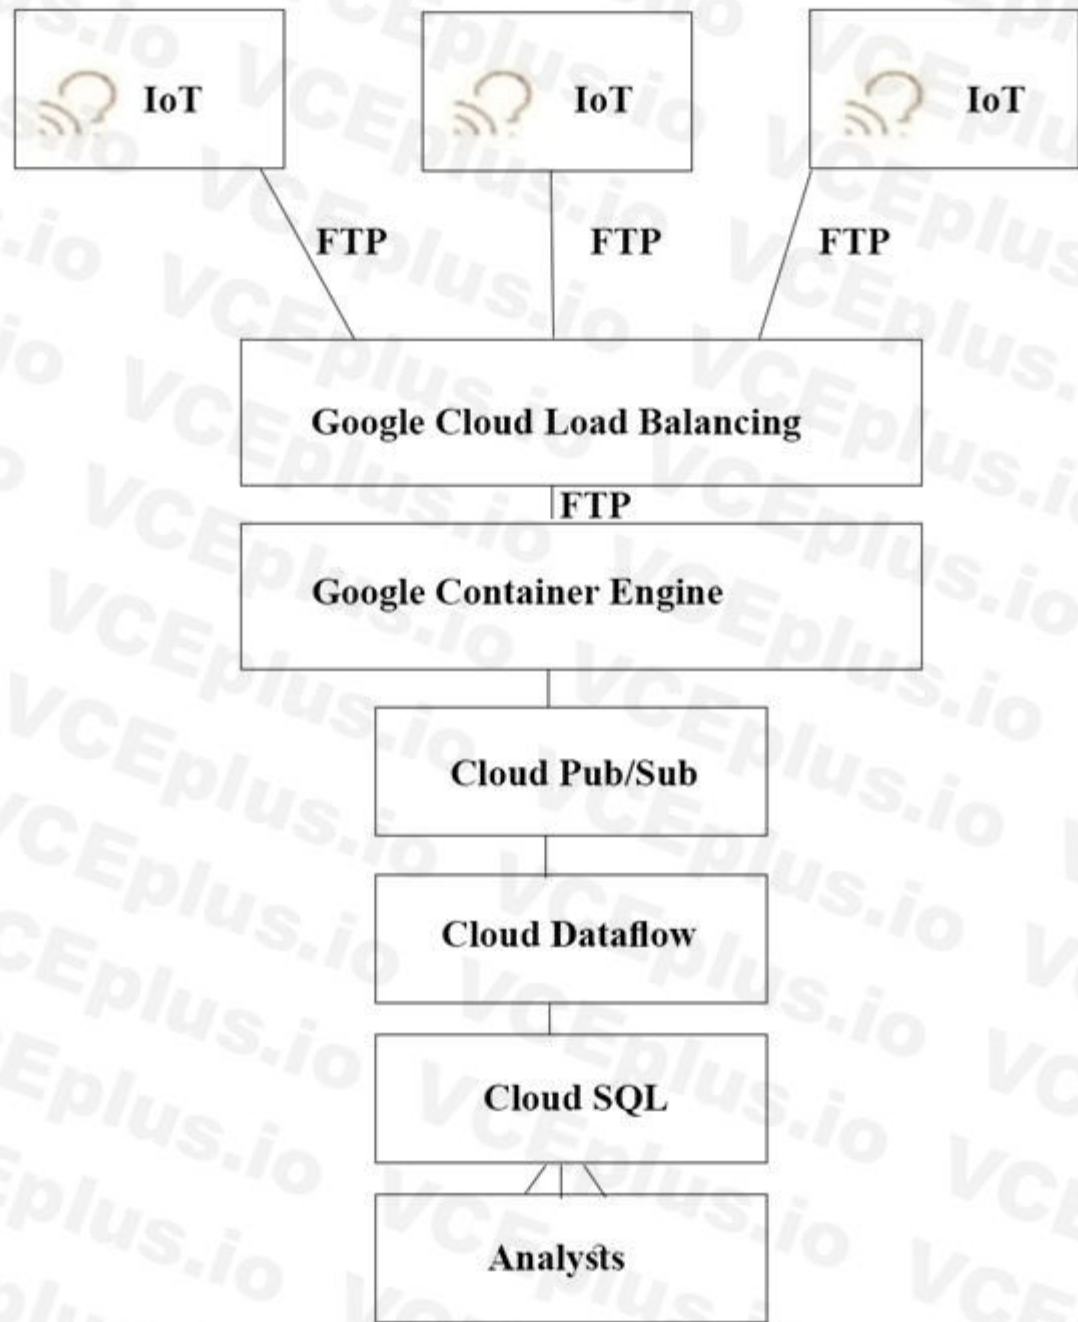
A.



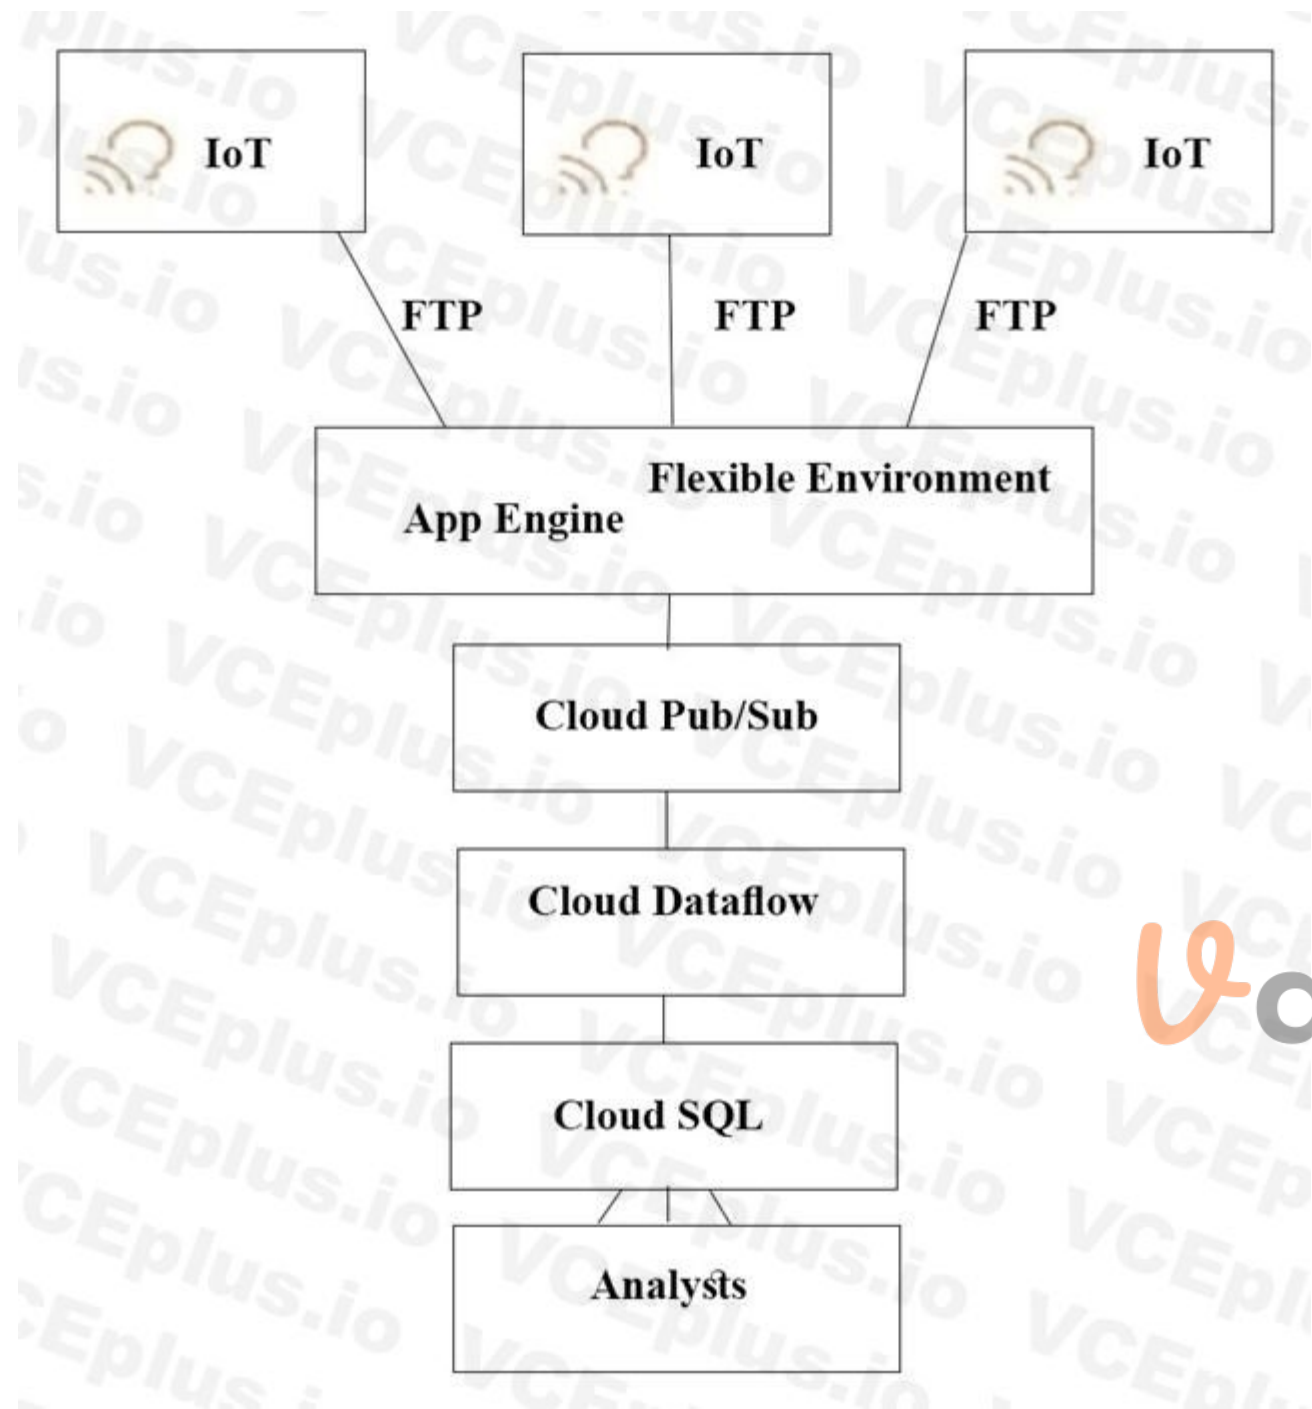
B.



C.



D.

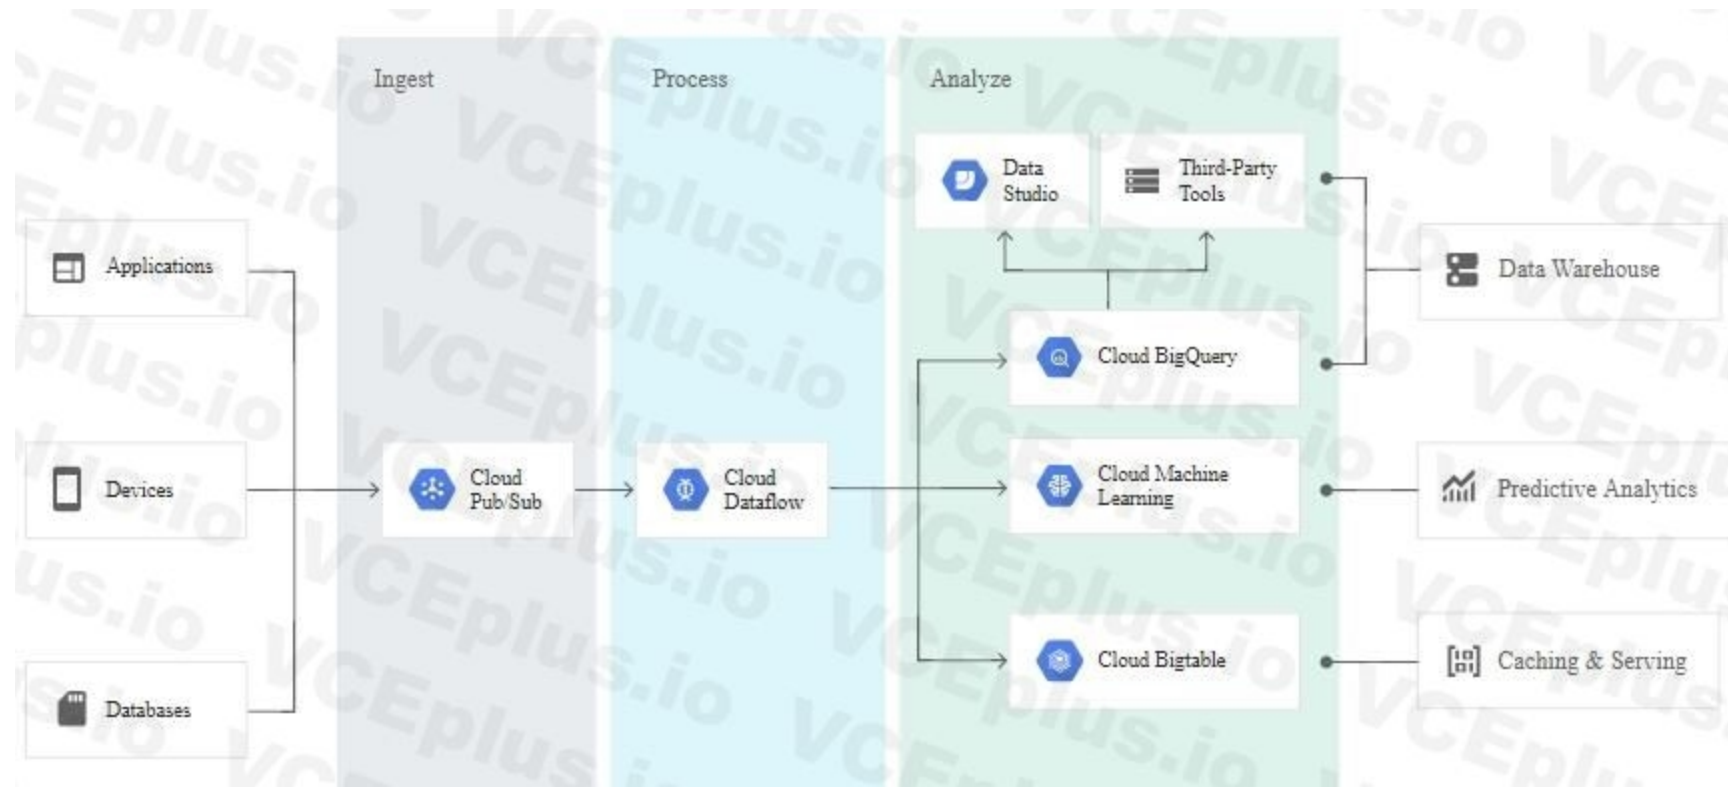


Correct Answer: A

Section:

Explanation:

The push endpoint can be a load balancer. A container cluster can be used. Cloud Pub/Sub for Stream Analytics



References: <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 TerraEarth 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:

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-terraearth>

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:

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 service-account and user-account OAuth, also called three-legged OAuth.

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

QUESTION 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:

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:

Explanation:

The Avro binary format is the preferred format for loading compressed data. Avro data is faster to load because the data can be read in parallel, even when the data blocks are compressed. Cloud Storage supports streaming transfers with the gsutil tool or boto library, based on HTTP chunked transfer encoding. Streaming data lets you stream data to and from your Cloud Storage account as soon as it becomes available without requiring that the data be first saved to a separate file. Streaming transfers are useful if you have a process that generates data and you do not want to buffer it locally before uploading it, or if you want to send the result from a computational pipeline directly into Cloud

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

QUESTION 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:

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:

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:

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:

Explanation:

Storage is the best choice for data that you plan to access at most once a year, due to its slightly lower availability, 90-day minimum storage duration, costs for data access, and higher per-operation costs. For example: Cold Data Storage - Infrequently accessed data, such as data stored for legal or regulatory reasons, can be stored at low cost as Coldline Storage, and be available when you need it.

Disaster recovery - In the event of a disaster recovery event, recovery time is key. Cloud Storage provides low latency access to data stored as Coldline Storage. References: <https://cloud.google.com/storage/docs/storage-classes>

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: A, C

Section:

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

Correct Answer: B

Section:

Explanation:

TerramEarth B

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 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. 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:

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:

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 pre-emptible instance with 32 CPUs.

Correct Answer: A

Section:

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:

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:

QUESTION 6

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 pre-emptible instance with 32 CPUs.

Correct Answer: A

Section:

QUESTION 7

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:

Explanation:

Exam G

QUESTION 1

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

What authentication strategy should they use?

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

Correct Answer: C

Section:

Explanation:

Provision users to Google's directory

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

You can provision users automatically using one of the following tools and services:

Google Cloud Directory Sync (GCDS)

Google Admin SDK

A third-party connector

GCDS is a connector that can provision users and groups on your behalf for both Cloud Platform and G Suite. Using GCDS, you can automate the addition, modification, and deletion of users, groups, and non-employee



contacts. You can synchronize the data from your LDAP directory server to your Cloud Platform domain by using LDAP queries. This synchronization is one-way: the data in your LDAP directory server is never modified.
References: <https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#authentication-and-identity>

QUESTION 2

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

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

Correct Answer: B

Section:

Explanation:

Section: [none]

Explanation:

Cloud Dataflow is a fully-managed service for transforming and enriching data in stream (real time) and batch (historical) modes with equal reliability and expressiveness -- no more complex workarounds or compromises needed.

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

QUESTION 3

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:

Explanation:

Stackdriver Logging allows you to store, search, analyze, monitor, and alert on log data and events from Google Cloud Platform and Amazon Web Services (AWS). Our API also allows ingestion of any custom log data from any source.

Stackdriver Logging is a fully managed service that performs at scale and can ingest application and system log data from thousands of VMs. Even better, you can analyze all that log data in real time. References: <https://cloud.google.com/logging/>

QUESTION 4

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:

Explanation:

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

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

`sudo resize2fs /dev/[DISK_ID][PARTITION_NUMBER]` where [DISK_ID] is the device name and [PARTITION_NUMBER] is the partition number for the device where you are resizing the file system. References:

<https://cloud.google.com/compute/docs/disks/add-persistent-disk>

QUESTION 5

Your application needs to process credit card transactions. You want the smallest scope of Payment Card Industry (PCI) compliance without compromising the ability to analyze transactional data and trends relating to which payment methods are used.

How should you design your architecture?

- A. Create a tokenizer service and store only tokenized data
- B. Create separate projects that only process credit card data
- C. Create separate subnetworks and isolate the components that process credit card data
- D. Streamline the audit discovery phase by labeling all of the virtual machines (VMs) that process PCI data
- E. Enable Logging export to Google BigQuery and use ACLs and views to scope the data shared with the auditor

Correct Answer: A

Section:

Explanation:

Reference:

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

QUESTION 6

You have been asked to select the storage system for the click-data of your company's large portfolio of websites. This data is streamed in from a custom website analytics package at a typical rate of 6,000 clicks per minute. With bursts of up to 8,500 clicks per second. It must have been stored for future analysis by your data science and user experience teams.

Which storage infrastructure should you choose?

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

Correct Answer: B

Section:

Explanation:

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

Low-latency read/write access

High-throughput analytics

Native time series support Common workloads: IoT, finance, adtech

Personalization, recommendations

Monitoring

Geospatial datasets Graphs

Incorrect Answers:

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

Images, pictures, and videos

Objects and blobs Unstructured data

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

Semi-structured application data

Hierarchical data Durable key-value data Common workloads:

User profiles

Product catalogs

Game state

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

QUESTION 7

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

What should you do?

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

Correct Answer: B

Section:

QUESTION 8

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:

Explanation:

Google Cloud Dataproc is a fast, easy-to-use, low-cost and fully managed service that lets you run the Apache Spark and Apache Hadoop ecosystem on Google Cloud Platform. Cloud Dataproc provisions big or small clusters rapidly, supports many popular job types, and is integrated with other Google Cloud Platform services, such as Google Cloud Storage and Stackdriver Logging, thus helping you reduce TCO. References:

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

QUESTION 9

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

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

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

Correct Answer: C

Section:

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:

Explanation:

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

Good for:

Low-latency read/write access

High-throughput analytics

Native time series support Common workloads: IoT, finance, adtech

Personalization, recommendations

Monitoring

Geospatial datasets

Graphs

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



QUESTION 11

Your company's user-feedback portal comprises a standard LAMP stack replicated across two zones. It is deployed in the us-central1 region and uses autoscaled managed instance groups on all layers, except the database. Currently, only a small group of select customers have access to the portal. The portal meets a 99,99% availability SLA under these conditions. However next quarter, your company will be making the portal available to all users, including unauthenticated users. You need to develop a resiliency testing strategy to ensure the system maintains the SLA once they introduce additional user load.

What should you do?

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

Correct Answer: B

Section:

QUESTION 12

One of the developers on your team deployed their application in Google Container Engine with the Dockerfile below. They report that their application deployments are taking too long.

```
FROM ubuntu:16.04
COPY . /src
RUN apt-get update && apt-get install -y python python-pip
RUN pip install -r requirements.txt
```

You want to optimize this Dockerfile for faster deployment times without adversely affecting the app's functionality. Which two actions should you take? Choose 2 answers.

- A. Remove Python after running pip
- B. Remove dependencies from requirements.txt
- C. Use a slimmed-down base image like Alpine Linux
- D. Use larger machine types for your Google Container Engine node pools
- E. Copy the source after the package dependencies (Python and pip) are installed

Correct Answer: C, E

Section:

Explanation:

The speed of deployment can be changed by limiting the size of the uploaded app, limiting the complexity of the build necessary in the Dockerfile, if present, and by ensuring a fast and reliable internet connection.

Note: Alpine Linux is built around musl libc and busybox. This makes it smaller and more resource efficient than traditional GNU/Linux distributions. A container requires no more than 8 MB and a minimal installation to disk requires around 130 MB of storage. Not only do you get a fully-fledged Linux environment but a large selection of packages from the repository.

References: <https://groups.google.com/forum/#!topic/google-appengine/hZMEkmmObDU> <https://www.alpinelinux.org/about/>

QUESTION 13

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

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



Correct Answer: D

Section:

QUESTION 14

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:

Explanation:

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

QUESTION 15

During a high traffic portion of the day, one of your relational databases crashes, but the replica is never promoted to a master. You want to avoid this in the future. What should you do?

- A. Use a different database
- B. Choose larger instances for your database
- C. Create snapshots of your database more regularly
- D. Implement routinely scheduled failovers of your databases

Correct Answer: D

Section:

QUESTION 16

Your organization requires that metrics from all applications be retained for 5 years for future analysis in possible legal proceedings. Which approach should you use?

- A. Grant the security team access to the logs in each Project
- B. Configure Stackdriver Monitoring for all Projects, and export to BigQuery
- C. Configure Stackdriver Monitoring for all Projects with the default retention policies
- D. Configure Stackdriver Monitoring for all Projects, and export to Google Cloud Storage

Correct Answer: B

Section:

Explanation:

Stackdriver Logging provides you with the ability to filter, search, and view logs from your cloud and open source application services. Allows you to define metrics based on log contents that are incorporated into dashboards and alerts.

Enables you to export logs to BigQuery, Google Cloud Storage, and Pub/Sub. References: <https://cloud.google.com/stackdriver/>

QUESTION 17

Your company has decided to build a backup replica of their on-premises user authentication PostgreSQL database on Google Cloud Platform. The database is 4 TB, and large updates are frequent. Replication requires private address space communication.

Which networking approach should you use?

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

Correct Answer: A

Section:

Explanation:

Google Cloud Dedicated Interconnect provides direct physical connections and RFC 1918 communication between your on-premises network and Google's network. Dedicated Interconnect enables you to transfer large amounts of data between networks, which can be more cost effective than purchasing additional bandwidth over the public Internet or using VPN tunnels.

Benefits:

Traffic between your on-premises network and your VPC network doesn't traverse the public Internet. Traffic traverses a dedicated connection with fewer hops, meaning there are less points of failure where traffic might get dropped or disrupted.

Your VPC network's internal (RFC 1918) IP addresses are directly accessible from your on-premises network. You don't need to use a NAT device or VPN tunnel to reach internal IP addresses. Currently, you can only reach internal IP addresses over a dedicated connection. To reach Google external IP addresses, you must use a separate connection.

You can scale your connection to Google based on your needs. Connection capacity is delivered over one or more 10 Gbps Ethernet connections, with a maximum of eight connections (80 Gbps total per interconnect).

The cost of egress traffic from your VPC network to your on-premises network is reduced. A dedicated connection is generally the least expensive method if you have a high-volume of traffic to and from Google's network.

References:

<https://cloud.google.com/interconnect/docs/details/dedicated>

QUESTION 18

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: B

Section:

QUESTION 19

You are designing a large distributed application with 30 microservices. Each of your distributed microservices needs to connect to a database back-end. You want to store the credentials securely. Where should you store the credentials?

- A. In the source code
- B. In an environment variable
- C. In a secret management system
- D. In a config file that has restricted access through ACLs

Correct Answer: C

Section:

Explanation:

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

**QUESTION 20**

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: B, F

Section:

QUESTION 21

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

- A. Be based on open-source technology for cloud portability
- B. Dynamically scale compute capacity based on demand

- C. Support continuous software delivery
- D. Run multiple segregated copies of the same application stack
- E. Deploy application bundles using dynamic templates
- F. Route network traffic to specific services based on URL
Which combination of technologies will meet all of his requirements?
- G. Google Kubernetes Engine, Jenkins, and Helm
- H. Google Kubernetes Engine and Cloud Load Balancing
- I. Google Kubernetes Engine and Cloud Deployment Manager
- J. Google Kubernetes Engine, Jenkins, and Cloud Load Balancing

Correct Answer: D

Section:

Explanation:

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

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

Incorrect Answers:

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

Use Helm to:

Find and use popular software packaged as Kubernetes charts

Share your own applications as Kubernetes charts

Create reproducible builds of your Kubernetes applications

Intelligently manage your Kubernetes manifest files

Manage releases of Helm packages

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



QUESTION 22

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

What should you do?

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

Correct Answer: C

Section:

Explanation:

A startup script, or a shutdown script, is specified through the metadata server, using startup script metadata keys. Reference: <https://cloud.google.com/compute/docs/startupscript>

QUESTION 23

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:

Explanation:

Google Cloud Platform(GCP) enforces firewall rules through rules and tags. GCP rules and tags can be defined once and used across all regions.

References: <https://cloud.google.com/docs/compare/openstack/>

<https://aws.amazon.com/it/blogs/aws/building-three-tier-architectures-with-security-groups/>

QUESTION 24

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: A, C, E

Section:



QUESTION 25

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: A, E

Section:

QUESTION 26

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

What should you do?

- A. Log in to a server, and iterate on the fix locally
- B. Revert the source code change, and rerun the deployment pipeline

- C. Log into the servers with the bad code change, and swap in the previous code
- D. Change the instance group template to the previous one, and delete all instances

Correct Answer: B

Section:

QUESTION 27

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

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

Correct Answer: C

Section:

Explanation:

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

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

QUESTION 28

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:

QUESTION 29

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

- A. Tag messages client side with the originating user identifier and the destination user.

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

Correct Answer: C

Section:

QUESTION 30

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

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

Correct Answer: B

Section:

QUESTION 31

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:

Explanation:

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

QUESTION 32

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:

QUESTION 33

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:

QUESTION 34

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: A

Section:

QUESTION 35

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:

Explanation:

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

QUESTION 36

You have an outage in your Compute Engine managed instance group: all instance keep restarting after 5 seconds. You have a health check configured, but autoscaling is disabled. Your colleague, who is a Linux expert, offered to look into the issue. You need to make sure that he can access the VMs. What should you do?

- A. Grant your colleague the IAM role of project Viewer
- B. Perform a rolling restart on the instance group
- C. Disable the health check for the instance group. Add his SSH key to the project-wide SSH keys
- D. Disable autoscaling for the instance group. Add his SSH key to the project-wide SSH Keys

Correct Answer: C

Section:

QUESTION 37

Your company is migrating its on-premises data center into the cloud. As part of the migration, you want to integrate Kubernetes Engine for workload orchestration. Parts of your architecture must also be PCI 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. Kubernetes Engine cannot be used under PCI DSS because it is considered shared hosting.
- C. Kubernetes Engine and GCP provide the tools you need to build a PCI DSS-compliant environment.
- D. All Google Cloud services are usable because Google Cloud Platform is certified PCI-compliant.

Correct Answer: C

Section:

QUESTION 38

Your company has multiple on-premises systems that serve as sources for reporting. The data has not been maintained well and has become degraded over time. You want to use Google-recommended practices to detect anomalies in your company data. What should you do?

- A. Upload your files into Cloud Storage. Use Cloud Datalab to explore and clean your data.
- B. Upload your files into Cloud Storage. Use Cloud Dataprep to explore and clean your data.
- C. Connect Cloud Datalab to your on-premises systems. Use Cloud Datalab to explore and clean your data.
- D. Connect Cloud Dataprep to your on-premises systems. Use Cloud Dataprep to explore and clean your data.

Correct Answer: B

Section:

QUESTION 39

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:

Explanation:

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

QUESTION 40

You are migrating your on-premises solution to Google Cloud in several phases. You will use Cloud VPN to maintain a connection between your on-premises systems and Google Cloud until the migration is completed. You want to make sure all your on-premise systems remain reachable during this period. How should you organize your networking in Google Cloud?

- A. Use the same IP range on Google Cloud as you use on-premises
- B. Use the same IP range on Google Cloud as you use on-premises for your primary IP range and use a secondary range that does not overlap with the range you use on-premises
- C. Use an IP range on Google Cloud that does not overlap with the range you use on-premises
- D. Use an IP range on Google Cloud that does not overlap with the range you use on-premises for your primary IP range and use a secondary range with the same IP range as you use on-premises



Correct Answer: D

Section:

QUESTION 41

You have found an error in your App Engine application caused by missing Cloud Datastore indexes. You have created a YAML file with the required indexes and want to deploy these new indexes to Cloud Datastore. What should you do?

- A. Point gcloud datastore create-indexes to your configuration file
- B. Upload the configuration file to the App Engine's default Cloud Storage bucket, and have App Engine detect the new indexes
- C. In the GCP Console, use Datastore Admin to delete the current indexes and upload the new configuration file
- D. Create an HTTP request to the built-in python module to send the index configuration file to your application

Correct Answer: A

Section:

QUESTION 42

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 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:

QUESTION 43

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 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:

QUESTION 44

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

- A. Upload the required installation files to Cloud Storage. Configure the VM on a subnet with a Private Google Access subnet. Assign only an internal IP address to the VM. Download the installation files to the VM using gsutil.

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

Correct Answer: B

Section:

QUESTION 45

Your company is moving 75 TB of data into Google Cloud. You want to use Cloud Storage and follow Google-recommended practices. What should you do?

- A. Move your data onto a Transfer Appliance. Use a Transfer Appliance Rehydrator to decrypt the data into Cloud Storage.
- B. Move your data onto a Transfer Appliance. Use Cloud Dataprep to decrypt the data into Cloud Storage.
- C. Install gsutil on each server that contains data. Use resumable transfers to upload the data into Cloud Storage.
- D. Install gsutil on each server containing data. Use streaming transfers to upload the data into Cloud Storage.

Correct Answer: A

Section:

QUESTION 46

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: A

Section:

QUESTION 47

You need to design a solution for global load balancing based on the URL path being requested. You need to ensure operations reliability and end-to-end in-transit encryption based on Google best practices. What should you do?

- A. Create a cross-region load balancer with URL Maps.
- B. Create an HTTPS load balancer with URL maps.
- C. Create appropriate instance groups and instances. Configure SSL proxy load balancing.
- D. Create a global forwarding rule. Configure SSL proxy load balancing.

Correct Answer: B

Section:

Explanation:

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

QUESTION 48

You have an application that makes HTTP requests to Cloud Storage. Occasionally the requests fail with HTTP status codes of 5xx and 429.

How should you handle these types of errors?

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

Correct Answer: B

Section:

Explanation:

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

QUESTION 49

You need to develop procedures to test a disaster plan for a mission-critical application. You want to use Google-recommended practices and native capabilities within GCP. What should you do?

- A. Use Deployment Manager to automate service provisioning. Use Activity Logs to monitor and debug your tests.
- B. Use Deployment Manager to automate provisioning. Use Stackdriver to monitor and debug your tests.
- C. Use gcloud scripts to automate service provisioning. Use Activity Logs to monitor and debug your tests.
- D. Use gcloud scripts to automate service provisioning. Use Stackdriver to monitor and debug your tests.

Correct Answer: B

Section:

QUESTION 50

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:

QUESTION 51

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:

QUESTION 52

You are deploying a PHP App Engine Standard service with Cloud SQL as the backend. You want to minimize the number of queries to the database. What should you do?

- A. Set the memcache service level to dedicated. Create a key from the hash of the query, and return database values from memcache before issuing a query to Cloud SQL.
- B. Set the memcache service level to dedicated. Create a cron task that runs every minute to populate the cache with keys containing query results.
- C. Set the memcache service level to shared. Create a cron task that runs every minute to save all expected queries to a key called "cached-queries".
- D. Set the memcache service level to shared. Create a key called "cached-queries", and return database values from the key before using a query to Cloud SQL.

Correct Answer: A

Section:

QUESTION 53

You need to ensure reliability for your application and operations by supporting reliable task scheduling for compute on GCP. Leveraging Google best practices, what should you do?

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

Correct Answer: B

Section:



QUESTION 54

Your company is building a new architecture to support its data-centric business focus. You are responsible for setting up the network. Your company's mobile and web-facing applications will be deployed on-premises, and all data analysis will be conducted in GCP. The plan is to process and load 7 years of archived .csv files totaling 900 TB of data and then continue loading 10 TB of data daily. You currently have an existing 100-MB internet connection.

What actions will meet your company's needs?

- A. Compress and upload both archived files and files uploaded daily using the gsutil -m option.
- B. Lease a Transfer Appliance, upload archived files to it, and send it to Google to transfer archived data to Cloud Storage. Establish a connection with Google using a Dedicated Interconnect or Direct Peering connection and use it to upload files daily.
- C. Lease a Transfer Appliance, upload archived files to it, and send it to Google to transfer archived data to Cloud Storage. Establish one Cloud VPN Tunnel to VPC networks over the public internet, and compress and upload files daily using the gsutil -m option.
- D. Lease a Transfer Appliance, upload archived files to it, and send it to Google to transfer archived data to Cloud Storage. Establish a Cloud VPN Tunnel to VPC networks over the public internet, and compress and upload files daily.

Correct Answer: B

Section:

QUESTION 55

You are developing a globally scaled frontend for a legacy streaming backend data API. This API expects events in strict chronological order with no repeat data for proper processing.

Which products should you deploy to ensure guaranteed-once FIFO (first-in, first-out) delivery of data?

- A. Cloud Pub/Sub alone

- B. Cloud Pub/Sub to Cloud Dataflow
- C. Cloud Pub/Sub to Stackdriver
- D. Cloud Pub/Sub to Cloud SQL

Correct Answer: D

Section:

Explanation:

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

QUESTION 56

Your company is planning to perform a lift and shift migration of their Linux RHEL 6.5+ virtual machines. The virtual machines are running in an on-premises VMware environment. You want to migrate them to Compute Engine following Google-recommended practices. What should you do?

- A. 1. Define a migration plan based on the list of the applications and their dependencies.
- B. Migrate all virtual machines into Compute Engine individually with Migrate for Compute Engine.
- C. 1. Perform an assessment of virtual machines running in the current VMware environment.
- D. Create images of all disks. Import disks on Compute Engine.
- E. Create standard virtual machines where the boot disks are the ones you have imported.
- F. 1. Perform an assessment of virtual machines running in the current VMware environment.
- G. Define a migration plan, prepare a Migrate for Compute Engine migration RunBook, and execute the migration.
- H. 1. Perform an assessment of virtual machines running in the current VMware environment.
- I. Install a third-party agent on all selected virtual machines.
- J. Migrate all virtual machines into Compute Engine.

Correct Answer: C

Section:

Explanation:

The framework illustrated in the preceding diagram has four phases:

- Assess. In this phase, you assess your source environment, assess the workloads that you want to migrate to Google Cloud, and assess which VMs support each workload.
- Plan. In this phase, you create the basic infrastructure for Migrate for Compute Engine, such as provisioning the resource hierarchy and setting up network access.
- Deploy. In this phase, you migrate the VMs from the source environment to Compute Engine.
- Optimize. In this phase, you begin to take advantage of the cloud technologies and capabilities.

Reference: <https://cloud.google.com/architecture/migrating-vms-migrate-for-compute-engine-getting-started>

QUESTION 57

You need to deploy an application to Google Cloud. The application receives traffic via TCP and reads and writes data to the filesystem. The application does not support horizontal scaling. The application process requires full control over the data on the file system because concurrent access causes corruption. The business is willing to accept a downtime when an incident occurs, but the application must be available 24/7 to support their business operations. You need to design the architecture of this application on Google Cloud. What should you do?

- A. Use a managed instance group with instances in multiple zones, use Cloud Filestore, and use an HTTP load balancer in front of the instances.
- B. Use a managed instance group with instances in multiple zones, use Cloud Filestore, and use a network load balancer in front of the instances.
- C. Use an unmanaged instance group with an active and standby instance in different zones, use a regional persistent disk, and use an HTTP load balancer in front of the instances.
- D. Use an unmanaged instance group with an active and standby instance in different zones, use a regional persistent disk, and use a network load balancer in front of the instances.

Correct Answer: D

Section:

Explanation:

Reference: <https://cloud.google.com/compute/docs/instance-groups>

QUESTION 58

Your company has an application running on multiple Compute Engine instances. You need to ensure that the application can communicate with an on-premises service that requires high throughput via internal IPs, while minimizing latency. What should you do?

- A. Use OpenVPN to configure a VPN tunnel between the on-premises environment and Google Cloud.
- B. Configure a direct peering connection between the on-premises environment and Google Cloud.
- C. Use Cloud VPN to configure a VPN tunnel between the on-premises environment and Google Cloud.
- D. Configure a Cloud Dedicated Interconnect connection between the on-premises environment and Google Cloud.

Correct Answer: C

Section:

Explanation:

Reference <https://cloud.google.com/architecture/setting-up-private-access-to-cloud-apis-through-vpn-tunnels>

QUESTION 59

You are managing an application deployed on Cloud Run for Anthos, and you need to define a strategy for deploying new versions of the application. You want to evaluate the new code with a subset of production traffic to decide whether to proceed with the rollout. What should you do?

- A. Deploy a new revision to Cloud Run with the new version. Configure traffic percentage between revisions.
- B. Deploy a new service to Cloud Run with the new version. Add a Cloud Load Balancing instance in front of both services.
- C. In the Google Cloud Console page for Cloud Run, set up continuous deployment using Cloud Build for the development branch. As part of the Cloud Build trigger, configure the substitution variable TRAFFIC_PERCENTAGE with the percentage of traffic you want directed to a new version.
- D. In the Google Cloud Console, configure Traffic Director with a new Service that points to the new version of the application on Cloud Run. Configure Traffic Director to send a small percentage of traffic to the new version of the application.

Correct Answer: C

Section:

QUESTION 60

You are monitoring Google Kubernetes Engine (GKE) clusters in a Cloud Monitoring workspace. As a Site Reliability Engineer (SRE), you need to triage incidents quickly. What should you do?

- A. Navigate the predefined dashboards in the Cloud Monitoring workspace, and then add metrics and create alert policies.
- B. Navigate the predefined dashboards in the Cloud Monitoring workspace, create custom metrics, and install alerting software on a Compute Engine instance.
- C. Write a shell script that gathers metrics from GKE nodes, publish these metrics to a Pub/Sub topic, export the data to BigQuery, and make a Data Studio dashboard.
- D. Create a custom dashboard in the Cloud Monitoring workspace for each incident, and then add metrics and create alert policies.

Correct Answer: D

Section:

Explanation:

Reference: <https://cloud.google.com/monitoring/charts/dashboards>

QUESTION 61

You are implementing a single Cloud SQL MySQL second-generation database that contains business-critical transaction data. You want to ensure that the minimum amount of data is lost in case of catastrophic failure. Which two features should you implement? (Choose two.)

- A. Sharding
- B. Read replicas
- C. Binary logging
- D. Automated backups
- E. Semisynchronous replication

Correct Answer: C, D

Section:

Explanation:

Section: [none]

Backups help you restore lost data to your Cloud SQL instance. Additionally, if an instance is having a problem, you can restore it to a previous state by using the backup to overwrite it. Enable automated backups for any instance that contains necessary data. Backups protect your data from loss or damage.

Enabling automated backups, along with binary logging, is also required for some operations, such as clone and replica creation. Reference: <https://cloud.google.com/sql/docs/mysql/backup-recovery/backups>

QUESTION 62

You are working at a sports association whose members range in age from 8 to 30. The association collects a large amount of health data, such as sustained injuries. You are storing this data in BigQuery. Current legislation requires you to delete such information upon request of the subject. You want to design a solution that can accommodate such a request. What should you do?

- A. Use a unique identifier for each individual. Upon a deletion request, delete all rows from BigQuery with this identifier.
- B. When ingesting new data in BigQuery, run the data through the Data Loss Prevention (DLP) API to identify any personal information. As part of the DLP scan, save the result to Data Catalog. Upon a deletion request, query Data Catalog to find the column with personal information.
- C. Create a BigQuery view over the table that contains all data. Upon a deletion request, exclude the rows that affect the subject's data from this view. Use this view instead of the source table for all analysis tasks.
- D. Use a unique identifier for each individual. Upon a deletion request, overwrite the column with the unique identifier with a salted SHA256 of its value.

Correct Answer: B

Section:

QUESTION 63

Your company has announced that they will be outsourcing operations functions. You want to allow developers to easily stage new versions of a cloud-based application in the production environment and allow the outsourced operations team to autonomously promote staged versions to production. You want to minimize the operational overhead of the solution. Which Google Cloud product should you migrate to?

- A. App Engine
- B. GKE On-Prem
- C. Compute Engine
- D. Google Kubernetes Engine

Correct Answer: D

Section:

Explanation:

Reference: <https://cloud.google.com/security/compliance/eba-outsourcing-mapping-gcp>

QUESTION 64

Your company is running its application workloads on Compute Engine. The applications have been deployed in production, acceptance, and development environments. The production environment is business-critical and is used 24/7, while the acceptance and development environments are only critical during office hours. Your CFO has asked you to optimize these environments to achieve cost savings during idle times. What should you do?

- A. Create a shell script that uses the `gcloud` command to change the machine type of the development and acceptance instances to a smaller machine type outside of office hours. Schedule the shell script on one of the production instances to automate the task.

- B. Use Cloud Scheduler to trigger a Cloud Function that will stop the development and acceptance environments after office hours and start them just before office hours.
- C. Deploy the development and acceptance applications on a managed instance group and enable autoscaling.
- D. Use regular Compute Engine instances for the production environment, and use preemptible VMs for the acceptance and development environments.

Correct Answer: D

Section:

Explanation:

Reference: <https://cloud.google.com/blog/products/it-ops/best-practices-for-optimizing-your-cloud-costs>

QUESTION 65

You are moving an application that uses MySQL from on-premises to Google Cloud. The application will run on Compute Engine and will use Cloud SQL. You want to cut over to the Compute Engine deployment of the application with minimal downtime and no data loss to your customers. You want to migrate the application with minimal modification. You also need to determine the cutover strategy. What should you do?

- A. 1. Set up Cloud VPN to provide private network connectivity between the Compute Engine application and the on-premises MySQL server.
- B. Stop the on-premises application.
- C. Create a mysqldump of the on-premises MySQL server.
- D. Upload the dump to a Cloud Storage bucket.
- E. Import the dump into Cloud SQL.
- F. Modify the source code of the application to write queries to both databases and read from its local database.
- G. Start the Compute Engine application.
- H. Stop the on-premises application.
- I. 1. Set up Cloud SQL proxy and MySQL proxy.
- J. Create a mysqldump of the on-premises MySQL server.
- K. Upload the dump to a Cloud Storage bucket.
- L. Import the dump into Cloud SQL.
- M. Stop the on-premises application.
- N. Start the Compute Engine application.
- O. 1. Set up Cloud VPN to provide private network connectivity between the Compute Engine application and the on-premises MySQL server.
- P. Stop the on-premises application.
- Q. Start the Compute Engine application, configured to read and write to the on-premises MySQL server.
- R. Create the replication configuration in Cloud SQL.
- S. Configure the source database server to accept connections from the Cloud SQL replica.
- T. Finalize the Cloud SQL replica configuration.
- U. When replication has been completed, stop the Compute Engine application.
- V. Promote the Cloud SQL replica to a standalone instance.
- W. Restart the Compute Engine application, configured to read and write to the Cloud SQL standalone instance.
- X. 1. Stop the on-premises application.
- Y. Create a mysqldump of the on-premises MySQL server.
- Z. Upload the dump to a Cloud Storage bucket.
- [. Import the dump into Cloud SQL.
- \. Start the application on Compute Engine.

Correct Answer: A

Section:



QUESTION 66

Your organization has decided to restrict the use of external IP addresses on instances to only approved instances. You want to enforce this requirement across all of your Virtual Private Clouds (VPCs). What should you do?

- A. Remove the default route on all VPCs. Move all approved instances into a new subnet that has a default route to an internet gateway.
- B. Create a new VPC in custom mode. Create a new subnet for the approved instances, and set a default route to the internet gateway on this new subnet.
- C. Implement a Cloud NAT solution to remove the need for external IP addresses entirely.
- D. Set an Organization Policy with a constraint on constraints/compute.vmExternallpAccess. List the approved instances in the allowedValues list.

Correct Answer: D

Section:

Explanation:

Reference: <https://cloud.google.com/compute/docs/ip-addresses/reserve-static-external-ip-address>

QUESTION 67

Your company uses the Firewall Insights feature in the Google Network Intelligence Center. You have several firewall rules applied to Compute Engine instances. You need to evaluate the efficiency of the applied firewall ruleset. When you bring up the Firewall Insights page in the Google Cloud Console, you notice that there are no log rows to display. What should you do to troubleshoot the issue?

- A. Enable Virtual Private Cloud (VPC) flow logging.
- B. Enable Firewall Rules Logging for the firewall rules you want to monitor.
- C. Verify that your user account is assigned the compute.networkAdmin Identity and Access Management (IAM) role.
- D. Install the Google Cloud SDK, and verify that there are no Firewall logs in the command line output.

Correct Answer: B

Section:

Explanation:

Reference: <https://cloud.google.com/network-intelligence-center/docs/firewall-insights/how-to/using-firewall-insights>

**QUESTION 68**

Your company has sensitive data in Cloud Storage buckets. Data analysts have Identity Access Management (IAM) permissions to read the buckets. You want to prevent data analysts from retrieving the data in the buckets from outside the office network. What should you do?

- A. 1. Create a VPC Service Controls perimeter that includes the projects with the buckets.
- B. Create an access level with the CIDR of the office network.
- C. 1. Create a firewall rule for all instances in the Virtual Private Cloud (VPC) network for source range.
- D. Use the Classless Inter-domain Routing (CIDR) of the office network.
- E. 1. Create a Cloud Function to remove IAM permissions from the buckets, and another Cloud Function to add IAM permissions to the buckets.
- F. Schedule the Cloud Functions with Cloud Scheduler to add permissions at the start of business and remove permissions at the end of business.
- G. 1. Create a Cloud VPN to the office network.
- H. Configure Private Google Access for on-premises hosts.

Correct Answer: C

Section:

QUESTION 69

You have developed a non-critical update to your application that is running in a managed instance group, and have created a new instance template with the update that you want to release. To prevent any possible impact to the application, you don't want to update any running instances. You want any new instances that are created by the managed instance group to contain the new update. What should you do?

- A. Start a new rolling restart operation.
- B. Start a new rolling replace operation.
- C. Start a new rolling update. Select the Proactive update mode.
- D. Start a new rolling update. Select the Opportunistic update mode.

Correct Answer: C

Section:

QUESTION 70

Your company is designing its application landscape on Compute Engine. Whenever a zonal outage occurs, the application should be restored in another zone as quickly as possible with the latest application data. You need to design the solution to meet this requirement. What should you do?

- A. Create a snapshot schedule for the disk containing the application data. Whenever a zonal outage occurs, use the latest snapshot to restore the disk in the same zone.
- B. Configure the Compute Engine instances with an instance template for the application, and use a regional persistent disk for the application data. Whenever a zonal outage occurs, use the instance template to spin up the application in another zone in the same region. Use the regional persistent disk for the application data.
- C. Create a snapshot schedule for the disk containing the application data. Whenever a zonal outage occurs, use the latest snapshot to restore the disk in another zone within the same region.
- D. Configure the Compute Engine instances with an instance template for the application, and use a regional persistent disk for the application data. Whenever a zonal outage occurs, use the instance template to spin up the application in another region. Use the regional persistent disk for the application data,

Correct Answer: D

Section:

QUESTION 71

Your company has just acquired another company, and you have been asked to integrate their existing Google Cloud environment into your company's data center. Upon investigation, you discover that some of the RFC 1918 IP ranges being used in the new company's Virtual Private Cloud (VPC) overlap with your data center IP space. What should you do to enable connectivity and make sure that there are no routing conflicts when connectivity is established?

- A. Create a Cloud VPN connection from the new VPC to the data center, create a Cloud Router, and apply new IP addresses so there is no overlapping IP space.
- B. Create a Cloud VPN connection from the new VPC to the data center, and create a Cloud NAT instance to perform NAT on the overlapping IP space.
- C. Create a Cloud VPN connection from the new VPC to the data center, create a Cloud Router, and apply a custom route advertisement to block the overlapping IP space.
- D. Create a Cloud VPN connection from the new VPC to the data center, and apply a firewall rule that blocks the overlapping IP space.

Correct Answer: A

Section:

QUESTION 72

You need to migrate Hadoop jobs for your company's Data Science team without modifying the underlying infrastructure. You want to minimize costs and infrastructure management effort. What should you do?

- A. Create a Dataproc cluster using standard worker instances.
- B. Create a Dataproc cluster using preemptible worker instances.
- C. Manually deploy a Hadoop cluster on Compute Engine using standard instances.
- D. Manually deploy a Hadoop cluster on Compute Engine using preemptible instances.

Correct Answer: A

Section:

Explanation:

Reference: <https://cloud.google.com/architecture/hadoop/hadoop-gcp-migration-jobs>

QUESTION 73

Your company has a project in Google Cloud with three Virtual Private Clouds (VPCs). There is a Compute Engine instance on each VPC. Network subnets do not overlap and must remain separated. The network configuration is shown below.



Instance #1 is an exception and must communicate directly with both Instance #2 and Instance #3 via internal IPs. How should you accomplish this?

- A. Create a cloud router to advertise subnet #2 and subnet #3 to subnet #1.
- B. Add two additional NICs to Instance #1 with the following configuration:
 - NIC1
 - ? VPC: VPC #2
 - ? SUBNETWORK: subnet #2
 - NIC2
 - ? VPC: VPC #3
 - ? SUBNETWORK: subnet #3Update firewall rules to enable traffic between instances.
- C. Create two VPN tunnels via CloudVPN:
 - 1 between VPC #1 and VPC #2.
 - 1 between VPC #2 and VPC #3.Update firewall rules to enable traffic between the instances.
- D. Peer all three VPCs:
 - Peer VPC #1 with VPC #2.
 - Peer VPC #2 with VPC #3.Update firewall rules to enable traffic between the instances.

Correct Answer: B

Section:

QUESTION 74

You need to deploy an application on Google Cloud that must run on a Debian Linux environment. The application requires extensive configuration in order to operate correctly. You want to ensure that you can install Debian distribution updates with minimal manual intervention whenever they become available. What should you do?

- A. Create a Compute Engine instance template using the most recent Debian image. Create an instance from this template, and install and configure the application as part of the startup script. Repeat this process whenever a new Google-managed Debian image becomes available.
- B. Create a Debian-based Compute Engine instance, install and configure the application, and use OS patch management to install available updates.
- C. Create an instance with the latest available Debian image. Connect to the instance via SSH, and install and configure the application on the instance. Repeat this process whenever a new Google-managed Debian image becomes available.
- D. Create a Docker container with Debian as the base image. Install and configure the application as part of the Docker image creation process. Host the container on Google Kubernetes Engine and restart the container whenever a new update is available.

Correct Answer: B

Section:

Explanation:

Reference: <https://cloud.google.com/compute/docs/os-patch-management>

QUESTION 75

You have an application that runs in Google Kubernetes Engine (GKE). Over the last 2 weeks, customers have reported that a specific part of the application returns errors very frequently. You currently have no logging or monitoring solution enabled on your GKE cluster. You want to diagnose the problem, but you have not been able to replicate the issue. You want to cause minimal disruption to the application. What should you do?

- A. 1. Update your GKE cluster to use Cloud Operations for GKE.
- B. Use the GKE Monitoring dashboard to investigate logs from affected Pods.
- C. 1. Create a new GKE cluster with Cloud Operations for GKE enabled.
- D. Migrate the affected Pods to the new cluster, and redirect traffic for those Pods to the new cluster.
- E. Use the GKE Monitoring dashboard to investigate logs from affected Pods.
- F. 1. Update your GKE cluster to use Cloud Operations for GKE, and deploy Prometheus.
- G. Set an alert to trigger whenever the application returns an error.
- H. 1. Create a new GKE cluster with Cloud Operations for GKE enabled, and deploy Prometheus.
- I. Migrate the affected Pods to the new cluster, and redirect traffic for those Pods to the new cluster.
- J. Set an alert to trigger whenever the application returns an error.

Correct Answer: C

Section:

Explanation:

Reference: <https://cloud.google.com/blog/products/management-tools/using-logging-your-apps-running-kubernetes-engine>

QUESTION 76

You need to deploy a stateful workload on Google Cloud. The workload can scale horizontally, but each instance needs to read and write to the same POSIX filesystem. At high load, the stateful workload needs to support up to 100 MB/s of writes. What should you do?

- A. Use a persistent disk for each instance.
- B. Use a regional persistent disk for each instance.
- C. Create a Cloud Filestore instance and mount it in each instance.
- D. Create a Cloud Storage bucket and mount it in each instance using gcsfuse.

Correct Answer: D

Section:

Explanation:

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

QUESTION 77

Your company has an application deployed on Anthos clusters (formerly Anthos GKE) that is running multiple microservices. The cluster has both Anthos Service Mesh and Anthos Config Management configured. End users inform you that the application is responding very slowly. You want to identify the microservice that is causing the delay. What should you do?

- A. Use the Service Mesh visualization in the Cloud Console to inspect the telemetry between the microservices.
- B. Use Anthos Config Management to create a ClusterSelector selecting the relevant cluster. On the Google Cloud Console page for Google Kubernetes Engine, view the Workloads and filter on the cluster. Inspect the configurations of the filtered workloads.
- C. Use Anthos Config Management to create a namespaceSelector selecting the relevant cluster namespace. On the Google Cloud Console page for Google Kubernetes Engine, visit the workloads and filter on the namespace. Inspect the configurations of the filtered workloads.
- D. Reinstall istio using the default istio profile in order to collect request latency. Evaluate the telemetry between the microservices in the Cloud Console.

Correct Answer: A

Section:

QUESTION 78

You are working at a financial institution that stores mortgage loan approval documents on Cloud Storage. Any change to these approval documents must be uploaded as a separate approval file, so you want to ensure that these documents cannot be deleted or overwritten for the next 5 years. What should you do?

- A. Create a retention policy on the bucket for the duration of 5 years. Create a lock on the retention policy.
- B. Create the bucket with uniform bucket-level access, and grant a service account the role of Object Writer. Use the service account to upload new files.
- C. Use a customer-managed key for the encryption of the bucket. Rotate the key after 5 years.
- D. Create the bucket with fine-grained access control, and grant a service account the role of Object Writer. Use the service account to upload new files.

Correct Answer: A

Section:

Explanation:

Reference: <https://cloud.google.com/storage/docs/using-bucket-lock>



QUESTION 79

Your company has just recently activated Cloud Identity to manage users. The Google Cloud Organization has been configured as well. The security team needs to secure projects that will be part of the Organization. They want to prohibit IAM users outside the domain from gaining permissions from now on. What should they do?

- A. Configure an organization policy to restrict identities by domain.
- B. Configure an organization policy to block creation of service accounts.
- C. Configure Cloud Scheduler to trigger a Cloud Function every hour that removes all users that don't belong to the Cloud Identity domain from all projects.
- D. Create a technical user (e.g., crawler@yourdomain.com), and give it the project owner role at root organization level.

Write a bash script that:

- Lists all the IAM rules of all projects within the organization.
- Deletes all users that do not belong to the company domain.

Create a Compute Engine instance in a project within the Organization and configure gcloud to be executed with technical user credentials. Configure a cron job that executes the bash script every hour.

Correct Answer: D

Section:

Explanation:

Reference: <https://sysdig.com/blog/gcp-security-best-practices/>

QUESTION 80

Your company has an application running on Google Cloud that is collecting data from thousands of physical devices that are globally distributed. Data is published to Pub/Sub and streamed in real time into an SSD Cloud

Bigtable cluster via a Dataflow pipeline. The operations team informs you that your Cloud Bigtable cluster has a hotspot, and queries are taking longer than expected. You need to resolve the problem and prevent it from happening in the future. What should you do?

- A. Advise your clients to use HBase APIs instead of NodeJS APIs.
- B. Delete records older than 30 days.
- C. Review your RowKey strategy and ensure that keys are evenly spread across the alphabet.
- D. Double the number of nodes you currently have.

Correct Answer: C

Section:

QUESTION 81

Your company has a Google Cloud project that uses BigQuery for data warehousing. There are some tables that contain personally identifiable information (PII). Only the compliance team may access the PII. The other information in the tables must be available to the data science team. You want to minimize cost and the time it takes to assign appropriate access to the tables. What should you do?

- A. 1. From the dataset where you have the source data, create views of tables that you want to share, excluding PII.
- B. Assign an appropriate project-level IAM role to the members of the data science team.
- C. Assign access controls to the dataset that contains the view.
- D. 1. From the dataset where you have the source data, create materialized views of tables that you want to share, excluding PII.
- E. Assign an appropriate project-level IAM role to the members of the data science team.
- F. Assign access controls to the dataset that contains the view.
- G. 1. Create a dataset for the data science team.
- H. Create views of tables that you want to share, excluding PII.
- I. Assign an appropriate project-level IAM role to the members of the data science team.
- J. Assign access controls to the dataset that contains the view.
- K. Authorize the view to access the source dataset.
- L. 1. Create a dataset for the data science team.
- M. Create materialized views of tables that you want to share, excluding PII.
- N. Assign an appropriate project-level IAM role to the members of the data science team.
- O. Assign access controls to the dataset that contains the view.
- P. Authorize the view to access the source dataset.



Correct Answer: C

Section:

Explanation:

Reference: https://cloud.google.com/blog/topics/developers-practitioners/bigquery-admin-reference-guide-data-governance?skip_cache=true

QUESTION 82

Your operations team currently stores 10 TB of data in an object storage service from a third-party provider. They want to move this data to a Cloud Storage bucket as quickly as possible, following Google recommended practices. They want to minimize the cost of this data migration. Which approach should they use?

- A. Use the gsutil mv command to move the data.
- B. Use the Storage Transfer Service to move the data.
- C. Download the data to a Transfer Appliance, and ship it to Google.
- D. Download the data to the on-premises data center, and upload it to the Cloud Storage bucket.

Correct Answer: B

Section:

Explanation:

Reference: <https://cloud.google.com/architecture/migration-to-google-cloud-transferring-your-large-datasets>

QUESTION 83

Your company and one of its partners each have a Google Cloud project in separate organizations.

Your company's project (prj-a) runs in Virtual Private Cloud (vpc-a). The partner's project (prj-b) runs in vpc-b. There are two instances running on vpc-a and one instance running on vpc-b. Subnets defined in both VPCs are not overlapping.

You need to ensure that all instances communicate with each other via internal IPs minimizing latency and maximizing throughput. What should you do?

```
gcloud compute start-lap-tunnel INSTANCE_NAME_IN_VPC_B 22 \
--local-host-port=localhost:22
```

- A. Set up a network peering between vpc-a and vpc-b
- B. Set up a VPN between vpc-a and vpc-b using Cloud VPN
- C. Configure IAP TCP forwarding on the instance in vpc-b and then launch the following gcloud command from one of the instances in vpc-a:
gcloud compute start-lap-tunnel INSTANCE_NAME_IN_VPC_B 22 \ --local-host-port=localhost:22
- D. Create an additional instance in vpc-a
- E. Create an additional instance in vpc-b
- F. Install OpenVPN in newly created instances
- G. Configure a VPN tunnel between vpc-a and vpc-b with the help of OpenVPN

Correct Answer: C

Section:

QUESTION 84

Your company has an application running on Google Cloud that is collecting data from thousands of physical devices that are globally distributed. Data is published to Pub/Sub and streamed in real time into an SSO Cloud Bigtable cluster via a Dataflow pipeline. The operations team informs you that your Cloud Bigtable cluster has a hot-spot and queries are taking longer than expected. You need to resolve the problem and prevent it from happening in the future. What should you do?

- A. Advise your clients to use HBase APIs instead of NodeJS APIs.
- B. Review your RowKey strategy and ensure that keys are evenly spread across the alphabet.
- C. Delete records older than 30 days.
- D. Double the number of nodes you currently have.

Correct Answer: B

Section:

QUESTION 85

Your company recently acquired a company that has infrastructure in Google Cloud. Each company has its own Google Cloud organization. Each company is using a Shared Virtual Private Cloud (VPC) to provide network connectivity for its applications. Some of the subnets used by both companies overlap. In order for both businesses to integrate, the applications need to have private network connectivity. These applications are not on overlapping subnets. You want to provide connectivity with minimal re-engineering. What should you do?

- A. Set up VPC peering and peer each Shared VPC together
- B. Configure SSH port forwarding on each application to provide connectivity between applications in the different Shared VPCs
- C. Migrate the projects from the acquired company into your company's Google Cloud organization. Relaunch the instances in your company's Shared VPC
- D. Set up a Cloud VPN gateway in each Shared VPC and peer Cloud VPNs

Correct Answer: B

Section:

QUESTION 86

Your operations team currently stores 10 TB of data in an object storage service from a third-party provider. They want to move this data to a Cloud Storage bucket as quickly as possible, following Google-recommended practices. They want to minimize the cost of this data migration. Which approach should they use?

- A. Use the `gsutil mv` command to move the data
- B. Use the Storage Transfer Service to move the data
- C. Download the data to a Transfer Appliance and ship it to Google
- D. Download the data to the on-premises data center and upload it to the Cloud Storage bucket

Correct Answer: B

Section:

Explanation:

<https://cloud.google.com/architecture/migration-to-google-cloud-transferring-your-largedatasets#transfer-options>

<https://cloud.google.com/storage-transfer-service>

QUESTION 87

You want to allow your operations team to store logs from all the production projects in your Organization, without duplicating logs from other projects. All of the production projects are contained in a folder. You want to ensure that all logs for existing and new production projects are captured automatically. What should you do?

- A. Create an aggregated export on the Production folder. Set the log sink to be a Cloud Storage bucket in an operations project
- B. Create an aggregated export on the Organization resource. Set the log sink to be a Cloud Storage bucket in an operations project.
- C. Create log exports in the production projects. Set the log sinks to be a Cloud Storage bucket in an operations project.
- D. Create log exports in the production projects. Set the log sinks to be BigQuery datasets in the production projects and grant IAM access to the operations team to run queries on the datasets

Correct Answer: B

Section:

QUESTION 88

You are configuring the cloud network architecture for a newly created project in Google Cloud that will host applications in Compute Engine. Compute Engine virtual machine instances will be created in two different subnets (sub-a and sub-b) within a single region.

- Instances in sub-a will have public IP addresses
- Instances in sub-b will have only private IP addresses

To download updated packages, instances must connect to a public repository outside the boundaries of Google Cloud. You need to allow sub-b to access the external repository. What should you do?

- A. Enable Private Google Access on sub-b
- B. Configure Cloud NAT and select sub-b in the NAT mapping section
- C. Configure a bastion host instance in sub-a to connect to instances in sub-b
- D. Enable Identity Aware Proxy for TCP forwarding for instances in sub-b

Correct Answer: B

Section:

QUESTION 89

Your company has just recently activated Cloud Identity to manage users. The Google Cloud Organization has been configured as usual. The security team needs to secure projects that will be part of the Organization. They want to prohibit IAM users outside the domain from gaining permissions from now on. What should they do?

- A. Configure an organization policy to restrict identities by domain
- B. Configure an organization policy to block creation of service accounts
- C. Configure Cloud Scheduler to trigger a Cloud Function every hour that removes all users that don't belong to the Cloud identity domain from all projects.

Correct Answer: A

Section:

QUESTION 90

Your company has just recently activated Cloud Identity to manage users. The Google Cloud Organization has been configured as wed. The security learn needs to secure protects that will be part of the Organization. They want to prohibit IAM users outside the domain from gaining permissions from now on. What should they do?

- A. Configure an organization policy to restrict identities by domain
- B. Configure an organization policy to block creation of service accounts
- C. Configure Cloud Scheduler o trigger a Cloud Function every hour that removes all users that don't belong to the Cloud identity domain from all projects.
- D. Create a technical user (e g . crawler@yourdomain com), and give it the protect owner role at root organization level Write a bash script that
 - Lists all me IAM rules of all projects within the organization
 - Deletes all users that do not belong to the company domain Create a Compute Engine instance m a project within the Organization and configure gcloud to be executed with technical user credentials Configure a cron job that executes the bash script every hour.

Correct Answer: A

Section:

QUESTION 91

You want to store critical business information in Cloud Storage buckets. The information is regularly changed but previous versions need to be referenced on a regular basis. You want to ensure that there is a record of all changes to any information in these buckets. You want to ensure that accidental edits or deletions can be easily roiled back. Which feature should you enable?

- A. Bucket Lock
- B. Object Versioning
- C. Object change notification
- D. Object Lifecycle Management

Correct Answer: B

Section:

QUESTION 92

You are working with a data warehousing team that performs data analysis. The team needs to process data from external partners, but the data contains personally identifiable information (PII). You need to process and store the data without storing any of the PII data. What should you do?

- A. Create a Dataflow pipeline to retrieve the data from the external sources. As part of the pipeline use the Cloud Data Loss Prevention (Cloud DLP) API to remove any PII data Store the result in BigQuery
- B. Create a Dataflow pipeline to retrieve the data from the external sources. As part of the pipeline store all non-PII data in BigQuery and store all PII data in a Cloud Storage bucket that has a retention policy set.
- C. Ask the external partners to upload an data on Cloud Storage Configure Bucket Lock for the bucket Create a Dataflow pipeline to read the data from the bucket As part of the pipeline, use the Cloud Data Loss Prevention (Cloud DIP) API to remove any PII data Store the result in BigQuery
- D. Ask the external partners to import ail data in your BigQuery dataset Create a dataflow pipeline to copy the data into a new table As part of the Dataflow bucket skip all data in columns that have PII data

Correct Answer: A

Section:

Explanation:

Create a Dataflow pipeline to retrieve the data from the external sources, he did not specify the way he is going to create it, it might be a pub/sub or external table or whatever.

QUESTION 93

Your company wants to migrate their 10-TB on-premises database export into Cloud Storage. You want to minimize the time it takes to complete this activity, the overall cost and database load. The bandwidth between the on-premises environment and Google Cloud is 1 Gbps. You want to follow Google-recommended practices. What should you do?

- A. Use the Data Transfer appliance to perform an offline migration
- B. Use a commercial partner ETL solution to extract the data from the on-premises database and upload it into Cloud Storage
- C. Develop a Dataflow job to read data directly from the database and write it into Cloud Storage
- D. Compress the data and upload it with `gsutil -m` to enable multi-threaded copy

Correct Answer: A

Section:

Explanation:

The Data Transfer appliance is a Google-provided hardware device that can be used to transfer large amounts of data from on-premises environments to Cloud Storage. It is suitable for scenarios where the bandwidth between the on-premises environment and Google Cloud is low or insufficient, and the data size is large. The Data Transfer appliance can minimize the time it takes to complete the migration, the overall cost and database load, by avoiding network bottlenecks and reducing bandwidth consumption. The Data Transfer appliance also encrypts the data at rest and in transit, ensuring data security and privacy. The other options are not optimal for this scenario, because they either require a high-bandwidth network connection (B, C, D), or incur additional costs and complexity (B, C).

Reference:

<https://cloud.google.com/data-transfer-appliance/docs/overview>

<https://cloud.google.com/blog/products/storage-data-transfer/introducing-storage-transfer-service-for-on-premises-data>

QUESTION 94

You are responsible for the Google Cloud environment in your company. Multiple departments need access to their own projects and the members within each department will have the same project responsibilities. You want to structure your Google Cloud environment for minimal maintenance and maximum overview of IAM permissions as each department's projects start and end. You want to follow Google-recommended practices. What should you do?

- A. Create a Google Group per department and add all department members to their respective groups. Create a folder per department and grant the respective group the required IAM permissions at the folder level. Add the projects under the respective folders.
- B. Grant all department members the required IAM permissions for their respective projects.
- C. Create a Google Group per department and add all department members to their respective groups. Grant each group the required IAM permissions for their respective projects.
- D. Create a folder per department and grant the respective members of the department the required IAM permissions at the folder level. Structure all projects for each department under the respective folders.

Correct Answer: A

Section:

Explanation:

This option follows the Google-recommended practices for structuring a Google Cloud environment for minimal maintenance and maximum overview of IAM permissions. By creating a Google Group per department and adding all department members to their respective groups, you can simplify user management and avoid granting IAM permissions to individual users. By creating a folder per department and granting the respective group the required IAM permissions at the folder level, you can enforce consistent policies across all projects within each department and avoid granting IAM permissions at the project level. By adding the projects under the respective folders, you can organize your resources hierarchically and leverage inheritance of IAM policies from folders to projects. The other options are not optimal for this scenario, because they either require granting IAM permissions to individual users (B, C), or do not use Google Groups to manage users (D).

Reference:

<https://cloud.google.com/architecture/framework/system-design>

<https://cloud.google.com/architecture/identity/best-practices-for-planning>

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

QUESTION 95

You are managing several projects on Google Cloud and need to interact on a daily basis with BigQuery, Bigtable and Kubernetes Engine using the gcloud CLI tool. You are travelling a lot and work on different workstations during the week. You want to avoid having to manage the gcloud CLI manually. What should you do?

- A. Use a package manager to install gcloud on your workstations instead of installing it manually
- B. Create a Compute Engine instance and install gcloud on the instance. Connect to this instance via SSH to always use the same gcloud installation when interacting with Google Cloud
- C. Install gcloud on all of your workstations. Run the command `gcloud components auto-update` on each workstation
- D. Use Google Cloud Shell in the Google Cloud Console to interact with Google Cloud

Correct Answer: D

Section:

Explanation:

This option allows you to use the gcloud CLI tool without having to install or manage it manually on different workstations. Google Cloud Shell is a browser-based command-line tool that provides you with a temporary Compute Engine virtual machine instance preloaded with the Cloud SDK, including the gcloud CLI tool. You can access Google Cloud Shell from any web browser and use it to interact with BigQuery, Bigtable and Kubernetes Engine using the gcloud CLI tool. The other options are not optimal for this scenario, because they either require installing and updating the gcloud CLI tool on multiple workstations (A, C), or creating and maintaining a Compute Engine instance for the sole purpose of using the gcloud CLI tool (B).

Reference:

<https://cloud.google.com/shell/docs/overview>

<https://cloud.google.com/sdk/gcloud/>

QUESTION 96

The operations team in your company wants to save Cloud VPN log events (or one year). You need to configure the cloud infrastructure to save the logs. What should you do?

- A. Set up a filter in Cloud Logging and a topic in Pub/Sub to publish the logs
- B. Set up a Cloud Logging Dashboard titled Cloud VPN Logs, and then add a chart that queries for the VPN metrics over a one-year time period
- C. Enable the Compute Engine API and then enable logging on the firewall rules that match the traffic you want to save
- D. Set up a filter in Cloud Logging and a Cloud Storage bucket as an export target for the logs you want to save

Correct Answer: D

Section:

QUESTION 97

Your company has an application running as a Deployment in a Google Kubernetes Engine (GKE) cluster. When releasing new versions of the application via a rolling deployment, the team has been causing outages. The root cause of the outages is misconfigurations with parameters that are only used in production. You want to put preventive measures for this in the platform to prevent outages. What should you do?

- A. Configure liveness and readiness probes in the Pod specification
- B. Configure an uptime alert in Cloud Monitoring
- C. Create a Scheduled Task to check whether the application is available
- D. Configure health checks on the managed instance group

Correct Answer: A

Section:

Explanation:

This option can help prevent outages caused by misconfigurations with parameters that are only used in production. Liveness and readiness probes are mechanisms to check the health and availability of the Pods and containers in a GKE cluster. Liveness probes determine if a container is still running, and if not, restart it. Readiness probes determine if a container is ready to serve requests, and if not, remove it from the load balancer. By configuring liveness and readiness probes in the Pod specification, you can ensure that your application can handle traffic and recover from failures gracefully during a rolling update. The other options are not optimal for this scenario, because they either do not prevent outages, but only alert or monitor them (B, C), or do not apply to GKE clusters, but to Compute Engine instances (D).

Reference:

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

<https://cloud.google.com/blog/products/containers-kubernetes/kubernetes-best-practices-setting-up-health-checks-with-readiness-and-liveness-probes>

QUESTION 98

Your company has a Google Cloud project that uses BigQuery for data warehousing. There are some tables that contain personally identifiable information (PII). Only the compliance team may access the PII. The other information in the tables must be available to the data science team. You want to minimize cost and the time it takes to assign appropriate access to the tables. What should you do?

- A. 1 From the dataset where you have the source data, create views of tables that you want to share, excluding PII
2 Assign an appropriate project-level IAM role to the members of the data science team
3 Assign access controls to the dataset that contains the view
- B. 1 From the dataset where you have the source data, create materialized views of tables that you want to share excluding PII
2 Assign an appropriate project-level IAM role to the members of the data science team
3. Assign access controls to the dataset that contains the view.
- C. 1 Create a dataset for the data science team
2 Create views of tables that you want to share excluding PII
3 Assign an appropriate project-level IAM role to the members of the data science team
4 Assign access controls to the dataset that contains the view
5 Authorize the view to access the source dataset
- D. 1. Create a dataset for the data science team.
2. Create materialized views of tables that you want to share, excluding PII
3. Assign an appropriate project-level IAM role to the members of the data science team
4 Assign access controls to the dataset that contains the view
5 Authorize the view to access the source dataset

Correct Answer: C

Section:

Explanation:

This option can help minimize cost and time by using views and authorized datasets. Views are virtual tables defined by a SQL query that can exclude PII columns from the source tables. Views do not incur storage costs and do not duplicate data. Authorized datasets are datasets that have access to another dataset's data without granting direct access to individual users or groups. By creating a dataset for the data science team and creating views of tables that exclude PII, you can share only the relevant information with the team. By assigning an appropriate project-level IAM role to the members of the data science team, you can grant them access to the BigQuery service and resources. By assigning access controls to the dataset that contains the view, you can grant them access to query the views. By authorizing the view to access the source dataset, you can enable the view to read data from the source tables without exposing PII. The other options are not optimal for this scenario, because they either use materialized views instead of views, which incur storage costs and duplicate data (B, D), or do not create a separate dataset for the data science team, which makes it harder to manage access controls (A).

Reference:

<https://cloud.google.com/bigquery/docs/views>

<https://cloud.google.com/bigquery/docs/authorized-datasets>

QUESTION 99

Your company has a Google Workspace account and Google Cloud Organization. Some developers in the company have created Google Cloud projects outside of the Google Cloud Organization.

You want to create an Organization structure that allows developers to create projects, but prevents them from modifying production projects. You want to manage policies for all projects centrally and be able to set more restrictive policies for production projects.

You want to minimize disruption to users and developers when business needs change in the future. You want to follow Google-recommended practices. How should you design the Organization structure?

- A. 1 Create a second Google Workspace account and Organization. 2 Grant all developers the Project Creator IAM role on the new Organization. 3 Move the developer projects into the new Organization. 4 Set the policies for all projects on both Organizations. 5 Additionally set the production policies on the original Organization.
- B. 1 Create a folder under the Organization resource named 'Production'. 2 Grant all developers the Project Creator IAM role on the Organization. 3. Move the developer projects into the Organization. 4 Set the policies for all projects on the Organization. 5 Additionally set the production policies on the 'Production' folder.
- C. 1 Create folders under the Organization resource named 'Development' and 'Production'. 2 Grant all developers the Project Creator IAM role on the 'Development' folder. 3. Move the developer projects into the 'Development' folder. 4 Set the policies for all projects on the Organization. 5 Additionally set the production policies on the 'Production' folder.



D. 1 Designate the Organization for production projects only 2 Ensure that developers do not have the Project Creator IAM role on the Organization 3 Create development projects outside of the Organization using the developer Google Workspace accounts 4 Set the policies for all projects on the Organization 5 Additionally set the production policies on the individual production projects

Correct Answer: C

Section:

Explanation:

This option can help create an organization structure that allows developers to create projects, but prevents them from modifying production projects. Folders are containers for projects and other folders within Google Cloud organizations. Folders allow resources to be structured hierarchically and inherit policies from their parent resources. By creating folders under the organization resource named "Development" and "Production", you can organize your projects by environment and apply different policies to them. By granting all developers the Project Creator IAM role on the "Development" folder, you can allow them to create projects under that folder, but not under the "Production" folder. By moving the developer projects into the "Development" folder, you can ensure that they are subject to the policies set on that folder. By setting the policies for all projects on the organization, you can manage policies centrally and efficiently. By additionally setting the production policies on the "Production" folder, you can enforce more restrictive policies for production projects and prevent developers from modifying them. The other options are not optimal for this scenario, because they either create a second Google Workspace account and organization, which increases complexity and cost (A), or do not use folders to organize projects by environment, which makes it harder to manage policies and permissions (B, D).

Reference:

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

<https://cloud.google.com/architecture/framework/system-design>

