

Q1)

You want to make a copy of a production Linux virtual machine in the US-Central region.

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

What steps must you take?

✓ Create a snapshot of the root disk, create an image file in Google Cloud Storage from the snapshot, and create a new virtual machine instance in the US-East region using the image file the root disk.

Explanation:-The instance needs to be deployed in different region and different project, an image needs to be created and used to create a new virtual machine in us-east region.

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

Q2)

Professional Cloud Architect

Sample case study: Mountkirk Games

This is a sample case study that may be used on the Professional Cloud Architect exam. It describes a fictitious business and solution concept to provide additional context to exam questions.

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.

1. 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 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.

2. 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

3. 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

4. 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 and low latency load balancing and frees us up from managing physical servers.

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?

- Verify that the load-testing team is not running their tool against production
- Verify that the new feature code did not introduce any performance bugs
- ✓ Verify that the project quota hasn't been exceeded

Explanation:-The slowness or 503 errors can be experienced due to the quotas being reached and the application has consumed all of its allocated resources. Compute Engine tracks your application's resource usage against system quotas. Compute Engine resets all resource measurements at the beginning of each calendar day (except for Stored Data, which always represents the amount of datastore storage in use). When free applications reach their quota for a resource, they cannot use that resource until the quota is replenished. Paid apps can exceed the free quota until their spending limit is exhausted. When an application consumes all of an allocated resource, the resource becomes unavailable until the quota is replenished. This may mean that your application will not work until the quota is replenished. For resources that are required to initiate a request, when the resource is depleted, Compute Engine by default returns an HTTP 403 or 503 error code for the request instead of calling a request handler.

Q3)

Professional Cloud Architect

Sample case study: Mountkirk Games

This is a sample case study that may be used on the Professional Cloud Architect exam. It describes a fictitious business and solution concept to provide additional context to exam questions.

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.

1. 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 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.

2. 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

3. 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

4. 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 and low latency load balancing and frees us up from managing physical servers.

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:

1. Services are deployed redundantly across multiple regions in the US and Europe
2. Only frontend services are exposed on the public internet
3. They can provide a single frontend IP for their fleet of services
4. Deployment artifacts are immutable.

Which set of products should they use?

● Google Cloud Functions, Google Cloud Pub/Sub, Google Cloud Deployment Manager

✓ Google Container Registry, Google Compute Engine, Google HTTP(S) Load Balancer

Explanation: Google Container Registry to maintain images which are immutable. Global HTTP Load Balancer to allow services across region and expose them as a single frontend IP and Google Compute Engine with managed instance groups to host services.

● Google Cloud Storage, Google App Engine, Google Network Load Balancer

● Google Cloud Storage, Google Cloud Dataflow, Google Compute Engine

Q4)

Professional Cloud Architect

Sample case study: Dress4Win

This is a sample case study that may be used on the Professional Cloud Architect exam. It describes a fictitious business and solution concept to provide additional context to exam questions.

Dress4Win is a web-based company that helps their users organize and manage their personal wardrobe using a web app and mobile application. The company also cultivates an active social network that connects their users with designers and retailers. They monetize their services through advertising, ecommerce, referrals, and a freemium app model. The application has grown from a few servers in the founder's garage to several hundred servers and appliances in a colocated data center. However, the capacity of their infrastructure is now insufficient for the application's rapid growth. Because of this growth and the company's desire to innovate faster, Dress4Win is committing to a full migration to a public cloud.

1. Solution concept

For the first phase of their migration to the cloud, Dress4Win is moving their development and test environments. They are also building a disaster recovery site, because their current infrastructure is at a single location. They are not sure which components of their architecture they can migrate as is and which components they need to change before migrating them.

2. Existing technical environment

Databases:

MySQL. One server for user data, inventory, static data,

MySQL 5.7

8 core CPUs

128 GB of RAM

2x 5 TB HDD (RAID 1)

Compute:

40 web application servers providing micro-services based APIs and static content

Tomcat - Java

Nginx

Four core CPUs

32 GB of RAM

20 Apache Hadoop/Spark servers:

Data analysis

Real-time trending calculations

Eight core CPUs

128 GB of RAM

4x 5 TB HDD (RAID 1)

Three RabbitMQ servers for messaging, social notifications, and events:

Eight core CPUs

32GB of RAM

Miscellaneous servers:

Jenkins, monitoring, bastion hosts, security scanners

Eight core CPUs

32GB of RAM

Storage appliances:

iSCSI for VM hosts

Fibre channel SAN - MySQL databases

1 PB total storage; 400 TB available

NAS - image storage, logs, backups

100 TB total storage; 35 TB available

3. Business requirements

Build a reliable and reproducible environment with scaled parity of production

Improve security by defining and adhering to a set of security and identity and access management (IAM) best practices for cloud

Improve business agility and speed of innovation through rapid provisioning of new resources

Analyze and optimize architecture for performance in the cloud

4. Technical requirements

Easily create non-production environments in the cloud

Implement an automation framework for provisioning resources in cloud

Implement a continuous deployment process for deploying applications to the on-premises data center or cloud

Support failover of the production environment to cloud during an emergency

Encrypt data on the wire and at rest

Support multiple private connections between the production data center and cloud environment.

5. Executive statement

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

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

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

Customers should be able to upload images with minimal latency and also be shown their images quickly on the main application page when they log in. Which configuration should Dress4Win use?

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

Explanation:-The best solution would be to store the images in Cloud Storage and have the metadata with the Cloud Storage link in Datastore to list down the images.

Q5)

Professional Cloud Architect

Sample case study: Dress4Win

This is a sample case study that may be used on the Professional Cloud Architect exam. It describes a fictitious business and solution concept to provide additional context to exam questions.

Dress4Win is a web-based company that helps their users organize and manage their personal wardrobe using a web app and

mobile application. The company also cultivates an active social network that connects their users with designers and retailers. They monetize their services through advertising, ecommerce, referrals, and a freemium app model. The application has grown from a few servers in the founder's garage to several hundred servers and appliances in a colocated data center. However, the capacity of their infrastructure is now insufficient for the application's rapid growth. Because of this growth and the company's desire to innovate faster, Dress4Win is committing to a full migration to a public cloud.

1. Solution concept

For the first phase of their migration to the cloud, Dress4Win is moving their development and test environments. They are also building a disaster recovery site, because their current infrastructure is at a single location. They are not sure which components of their architecture they can migrate as is and which components they need to change before migrating them.

2. Existing technical environment

The Dress4Win application is served out of a single data center location. All servers run Ubuntu LTS v16.04.

Databases:

MySQL. One server for user data, inventory, static data,
MySQL 5.7
8 core CPUs
128 GB of RAM
2x 5 TB HDD (RAID 1)

Compute:

40 web application servers providing micro-services based APIs and static content
Tomcat - Java
Nginx
Four core CPUs
32 GB of RAM
20 Apache Hadoop/Spark servers:
Data analysis
Real-time trending calculations
Eight core CPUs
128 GB of RAM
4x 5 TB HDD (RAID 1)
Three RabbitMQ servers for messaging, social notifications, and events:
Eight core CPUs
32GB of RAM

Miscellaneous servers:

Jenkins, monitoring, bastion hosts, security scanners
Eight core CPUs
32GB of RAM

Storage appliances:

iSCSI for VM hosts
Fibre channel SAN - MySQL databases
1 PB total storage; 400 TB available
NAS - image storage, logs, backups
100 TB total storage; 35 TB available

3. Business requirements

Build a reliable and reproducible environment with scaled parity of production
Improve security by defining and adhering to a set of security and identity and access management (IAM) best practices for cloud
Improve business agility and speed of innovation through rapid provisioning of new resources
Analyze and optimize architecture for performance in the cloud

4. Technical requirements

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

5. Executive statement

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

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

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

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

Explanation:-The servers should allow access from the uptime servers which performs http checks. Stackdriver can verify the availability of your service by accessing it from locations around the world. You can use the results from these uptime checks in your alerting policies, or you can

directly monitor the results in the Monitoring uptime-check dashboards. Your use of uptime checks is affected by any firewalls protecting your service. If the resource you are checking isn't publicly available, you must configure the resource's firewall to permit incoming traffic from the uptime-check servers. See Getting IP addresses to download a list of the IP addresses. If the resource you are checking doesn't have an external IP address, uptime checks are unable to reach it.

Q6)

Professional Cloud Architect

Sample case study: TerramEarth

This is a sample case study that may be used on the Professional Cloud Architect exam. It describes a fictitious business and solution concept to provide additional context to exam questions.

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.

1. 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.

2. 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 three 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 four weeks while they wait for replacement parts.

3. Business requirements

Decrease unplanned vehicle downtime to less than one 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

4. Technical requirements

Expand beyond a single data center to decrease latency to the American Midwest and East Coast

Create a backup strategy

Increase security of data transfer from equipment to the data center

Improve data in the data warehouse

Use customer and equipment data to anticipate customer needs

5. Application 1: Data ingest

A custom Python application reads uploaded data files 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

6. Application 2: Reporting

An off-the-shelf application that business analysts use to run a daily report to see what equipment needs repair. Only two analysts of a team of 10 (five West Coast, five 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

7. Executive statement

Our competitive advantage has always been in our manufacturing process, with our ability to build better vehicles for lower cost than our competitors. However, new products with different approaches are constantly being developed, and I'm concerned that we lack the skills to undergo the next wave of transformations in our industry. My goals are to build our skills while addressing immediate market needs through incremental innovations.

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?

- IoT -> (FTP) -> Google Load Balancing -> (FTP) -> Google Container Engine -> Cloud Pub/Sub -> Cloud Dataflow -> Cloud SQL -> Analysts
 - ✓ IoT -> (FTP) -> App Engine Flexible Environment -> Cloud Pub/Sub -> Cloud Dataflow -> BigQuery -> Analysts
- Explanation:-**App Engine can be used for ingest the data (Java 8 supports all sockets) with Pub/Sub to store, and Dataflow to process and analyze in BigQuery.
- IoT -> (FTP) -> Google Load Balancing -> Google Container Engine -> Cloud Pub/Sub -> Cloud Dataflow -> BigQuery -> Analysts

Q7)

Professional Cloud Architect

Sample case study: TerramEarth

This is a sample case study that may be used on the Professional Cloud Architect exam. It describes a fictitious business and solution concept to provide additional context to exam questions.

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.

1. 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.

2. 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 three 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 four weeks while they wait for replacement parts.

3. Business requirements

Decrease unplanned vehicle downtime to less than one 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

4. Technical requirements

Expand beyond a single data center to decrease latency to the American Midwest and East Coast

Create a backup strategy

Increase security of data transfer from equipment to the data center

Improve data in the data warehouse

Use customer and equipment data to anticipate customer needs

5. Application 1: Data ingest

A custom Python application reads uploaded data files 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

6. Application 2: Reporting

An off-the-shelf application that business analysts use to run a daily report to see what equipment needs repair. Only two analysts of a team of 10 (five West Coast, five 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

7. Executive statement

Our competitive advantage has always been in our manufacturing process, with our ability to build better vehicles for lower cost than our competitors. However, new products with different approaches are constantly being developed, and I'm concerned that we lack the skills to undergo the next wave of transformations in our industry. My goals are to build our skills while addressing immediate market needs through incremental innovations.

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?

- Vehicles continue to write data using the existing system (FTP)
- Vehicles stream data directly to Google BigQuery
- ✓ Vehicles write data directly to Google Cloud Pub/Sub

Explanation:-Vehicles can write data directly to Google Cloud Pub/Sub. Cloud Pub/Sub is fully managed and provides a simple, reliable, scalable foundation for stream analytics and event-driven computing systems.

- Vehicles write data directly to GCS

Q8)

Your company wants to reduce cost on infrequently accessed data by moving it to the cloud.

The data will still be accessed approximately once a month to refresh historical charts. In addition, data older than 5 years needs to be archived for 5 years for compliance reasons.

How should you store and manage the data?

- ✓ In Google Cloud Storage and stored in a Nearline bucket. Set an Object Lifecycle Management policy to change the storage class to Coldline for data older than 5 years.

Explanation:-The access pattern fits Nearline storage class requirements and Nearline is a more cost-effective storage approach than Multi-Regional. The object lifecycle management policy to move data to Coldline is ideal for archival.

- In Google Cloud Storage and stored in a Nearline bucket. Set an Object Lifecycle Management policy to delete data older than 5 years.
- In Google Cloud Storage and stored in a Multi-Regional bucket. Set an Object Lifecycle Management policy to change the storage class to Coldline for data older than 5 years.
- In Google Cloud Storage and stored in a Multi-Regional bucket. Set an Object Lifecycle Management policy to delete data older than 5 years.

Q9)

Your company collects and stores security camera footage in Google Cloud Storage.

Within the first 30 days, footage is processed regularly for threat detection, object detection, trend analysis, and suspicious behavior detection. You want to minimize the cost of storing all the data.

How should you store the videos?

- Use Google Cloud Regional Storage for the first 30 days, and then move to Google Persistent Disk.
- Use Google Cloud Regional Storage for the first 30 days, and then move to Nearline Storage.
- Use Google Cloud Nearline Storage for the first 30 days, and then move to Coldline Storage.
- ✓ Use Google Cloud Regional Storage for the first 30 days, and then move to Coldline Storage.

Explanation:-The data is accessed frequently within the first 30 days, using Google Cloud Regional Storage will enable the most cost-effective solution for storing and accessing the data. For videos older than 30 days, Google Cloud Coldline Storage offers the most cost-effective solution since it won't be accessed.

Q10)

You are running an application in Google App Engine that is serving production traffic.

You want to deploy a risky but necessary change to the application. It could take down your service if not properly coded.

During development of the application, you realized that it can only be properly tested by live user traffic.

How should you test the feature?

- Create a second project with the new app in isolation, and onboard users.
- Set up a second Google App Engine service, and then update a subset of clients to hit the new service.
- ✓ Deploy a new version of the application, and use traffic splitting to send a small percentage of traffic to it.

Explanation:-Deploying a new version without assigning it as the default version will not create downtime for the application. Using traffic splitting allows for easily redirecting a small amount of traffic to the new version and can also be quickly reverted without application downtime. Traffic migration smoothly switches request routing, gradually moving traffic from the versions currently receiving traffic to one or more versions that you specify. Traffic splitting distributes a percentage of traffic to versions of your application. You can split traffic to move 100% of traffic to a single version or to route percentages of traffic to multiple versions. Splitting traffic to two or more versions allows you to conduct A/B testing between your versions and provides control over the pace when rolling out features.

- Deploy the new application version temporarily, and then roll it back.

Q11)

One of the microservices in your application has an intermittent performance problem.

You have not observed the problem when it occurs but when it does, it triggers a particular burst of log lines. You want to debug a machine while the problem is occurring.

What should you do?

- In the Stackdriver Error Reporting dashboard, look for a pattern in the times the problem occurs.
- Configure your microservice to send traces to Stackdriver Trace so you can find what is taking so long.
- ✓ Set up a log metric in Stackdriver Logging, and then set up an alert to notify you when the number of log lines increases past a threshold.

Explanation:-There is a burst of log lines you can set up a metric that identifies those lines. Stackdriver will also allow you to set up a text, email or messaging alert that can notify promptly when the error is detected so you can hop onto the system to debug.

- Log into one of the machines running the microservice and wait for the log storm.

Q12)

Professional Cloud Architect

Sample case study: TerramEarth

This is a sample case study that may be used on the Professional Cloud Architect exam. It describes a fictitious business and solution concept to provide additional context to exam questions.

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.

1. 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.

2. 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 three 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 four weeks while they wait for replacement parts.

3. Business requirements

Decrease unplanned vehicle downtime to less than one 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

4. Technical requirements

Expand beyond a single data center to decrease latency to the American Midwest and East Coast

Create a backup strategy

Increase security of data transfer from equipment to the data center

Improve data in the data warehouse

Use customer and equipment data to anticipate customer needs

5. Application 1: Data ingest

A custom Python application reads uploaded data files 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

6. Application 2: Reporting

An off-the-shelf application that business analysts use to run a daily report to see what equipment needs repair. Only two analysts of a team of 10 (five West Coast, five 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

7. Executive statement

Our competitive advantage has always been in our manufacturing process, with our ability to build better vehicles for lower cost than our competitors. However, new products with different approaches are constantly being developed, and I'm concerned that we lack the skills to undergo the next wave of transformations in our industry. My goals are to build our skills while addressing immediate market needs through incremental innovations.

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?

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

Explanation:-The data needs to be moved from different regions for analytics the best approach to be to preprocess the data in each region so that the data size is reduced to required data only. This will speed up the time to transfer the data and also reduce the cost for both egress and storage. This preprocessed data can be stored in a regional bucket, which reduces cost further, to be processed by Dataproc.

Q13)

You need to allow traffic from specific virtual machines in 'subnet-a' network access to machines in 'subnet-b' without giving the entirety of subnet-a access.

How can you accomplish this?

- ☐ Create a rule to deny all traffic to the entire subnet, then create a second rule with higher priority giving access to tagged VM's in subnet-a.
- ☐ Relocate the subnet-a machines to a different subnet and give the new subnet the needed access.
- ☒ Create a firewall rule to allow traffic from resources with specific network tags, then assign the machines in subnet-a the same tags.

Explanation:-Network tags allow more granular access based on individually tagged instances. Network tags are text attributes you can add to Compute Engine virtual machine (VM) instances. Tags allow you to make firewall rules and routes applicable to specific VM instances. You can only add network tags to VM instances or instance templates. You cannot tag other GCP resources. You can assign network tags to new instances at creation time, or you can edit the set of assigned tags at any time later. Network tags can be edited without stopping an instance. Network tags allow you to apply firewall rules and routes to a specific instance or set of instances:

You make a firewall rule applicable to specific instances by using target tags and source tags.

You make a route applicable to specific instances by using a tag.

- ☐ You can only grant firewall access to an entire subnet and not individual VM's inside.

Q14)

Your company processes high volumes of IoT data that are time-stamped. The total data volume can be several petabytes.

The data needs to be written and changed at a high speed. You want to use the most performant storage option for your data.

Which product should you use?

- ☐ BigQuery
- ☒ Cloud Bigtable

Explanation:-Cloud Bigtable is the most performant storage option to work with IoT and time series data. Google Cloud Bigtable is a fast, fully managed, highly-scalable NoSQL database service. It is designed for the collection and retention of data from 1TB to hundreds of PB.

- ☐ Cloud Storage
- ☐ Cloud Datastore

Q15)

An Organization has setup an IAM policy at the organization level, the folder level, the project level, and on the resource level.

They want to understand what policy takes effect on the entity.

What would be the correct option?

- ☐ Effective policy for a resource is the policy inherited from its ancestors overriding the policy defined on the resource
- ☒ Effective policy for a resource is the union of the policy set on the resource and the policy inherited from its ancestors

Explanation:-The effective policy for a resource is the union of the policy set on the resource and the policy inherited from its ancestors. Link - <https://cloud.google.com/resource-manager/docs/cloud-platform-resource-hierarchy#inheritance>

- ☐ Effective policy for a resource is the policy defined overriding the policy inherited from its ancestors
- ☐ Effective policy for a resource is the Intersection of the policy set on the resource and the policy inherited from its ancestors

Q16)

Using principal of least privilege and allowing for maximum automation.

What steps can you take to store audit logs for long-term access and to allow access for external auditors to view? (Choose two)

- ☐ Export audit logs to BigQuery via an export sink.
- ☒ Export audit logs to Cloud Storage via an export sink.

Explanation:-Stackdriver logging allows export to Cloud Storage which can be used for long term access and exposed to external auditors using signed urls. Stackdriver Logging provides an operational datastore for logs and provides rich export capabilities. You might export your logs for several reasons, such as retaining logs for long-term storage (months or years) to meet compliance requirements or for running data analytics against the metrics extracted from the logs. Stackdriver Logging can export to Cloud Storage, BigQuery, and Cloud Pub/Sub.

- ☐ Create an account for auditors to have view access to Stackdriver Logging.
- ☒ Generate a signed URL to the Stackdriver export destination for auditors to access.

Explanation:-Stackdriver logging allows export to Cloud Storage which can be used for long term access and exposed to external auditors using signed urls. Stackdriver Logging provides an operational datastore for logs and provides rich export capabilities. You might export your logs for several reasons, such as retaining logs for long-term storage (months or years) to meet compliance requirements or for running data analytics against the metrics extracted from the logs. Stackdriver Logging can export to Cloud Storage, BigQuery, and Cloud Pub/Sub.

Q17)

You have a web server that is working properly for http traffic, but you can't connect to its instance VM over SSH.

Which of these troubleshooting methods can you use without disrupting production traffic? (Select 2 answers)

- ☐ Create a startup script to collect information.
- ☐ Check Stackdriver logs for any connectivity errors
- ☒ Access the serial console output

Explanation:-You debug the issue with disrupting production traffic by using Debug the issue in serial console or creating a snapshot and image and associating it with new instance. You can enable read-write access to an instance's serial console so you can log into the console and

troubleshoot problems with the instance. This is particularly useful when you cannot log in with SSH or if the instance has no connection to the network. The serial console remains accessible in both these conditions. You might have an instance you can't connect to that continues to correctly serve production traffic. In this case, you might want to inspect the disk without interrupting the instance's ability to serve users. First, take a snapshot of the instance's boot disk, then create a new disk from that snapshot, create a temporary instance, and finally attach and mount the new persistent disk to your temporary instance to troubleshoot the disk.

- ☐ Connect to the instance using Cloud Shell and check the logs
- ☒ Create a snapshot of the disk and use it to create a new disk; then attach the new disk to a new instance

Explanation:-You debug the issue with disrupting production traffic by using Debug the issue in serial console or creating a snapshot and image and associating it with new instance. You can enable read-write access to an instance's serial console so you can log into the console and troubleshoot problems with the instance. This is particularly useful when you cannot log in with SSH or if the instance has no connection to the network. The serial console remains accessible in both these conditions. You might have an instance you can't connect to that continues to correctly serve production traffic. In this case, you might want to inspect the disk without interrupting the instance's ability to serve users. First, take a snapshot of the instance's boot disk, then create a new disk from that snapshot, create a temporary instance, and finally attach and mount the new persistent disk to your temporary instance to troubleshoot the disk.

Q18)

You have a Kubernetes cluster with 1 node-pool. The cluster receives a lot of traffic and needs to grow.

You decide to add a node. What should you do?

- ☐ Edit the managed instance group of the cluster and increase the number of VMs by 1.
- ☐ Use "kubectl container clusters resize" with the desired number of nodes.
- ☒ Use "gcloud container clusters resize" with the desired number of nodes.

Explanation:-The kubernetes cluster can be resized using the gcloud command. `gcloud container clusters resize [CLUSTER_NAME] --node-pool [POOL_NAME] \ --size [SIZE]`

- ☐ Edit the managed instance group of the cluster and enable autoscaling.

Q19)

Your company is planning the infrastructure for a new large-scale application that will need to store over 100 TB or a petabyte of data in NoSQL format for Low-latency read/write and High-throughput analytics.

Which storage option should you use?

- ☐ Cloud SQL
- ☐ Cloud Spanner
- ☐ Cloud Datastore
- ☒ Cloud Bigtable

Explanation:-Bigtable is an ideal solution to provide low latency, high throughput data processing storage option with analytics Cloud Bigtable

Q20)

You created an update for your application on App Engine. You want to deploy the update without impacting your users.

You want to be able to roll back as quickly as possible if it fails. What should you do?

- ☐ Notify your users of an upcoming maintenance window. Deploy the update in that maintenance window.
- ☐ Deploy the update as the same version that is currently running.
- ☒ Deploy the update as a new version. Migrate traffic from the current version to the new version.

Explanation:-The deployment can be done seamlessly by deploying a new version and migrating the traffic gradually from the old version to the new version. If any issue is encountered, the traffic can be migrated 100% to the old version. Manage how much traffic is received by a version of your application by migrating or splitting traffic. Traffic migration smoothly switches request routing, gradually moving traffic from the versions currently receiving traffic to one or more versions that you specify. Traffic splitting distributes a percentage of traffic to versions of your application. You can split traffic to move 100% of traffic to a single version or to route percentages of traffic to multiple versions. Splitting traffic to two or more versions allows you to conduct A/B testing between your versions and provides control over the pace when rolling out features.

- ☐ Delete the current version of your application. Deploy the update using the same version identifier as the deleted version.

Q21)

Professional Cloud Architect

Sample case study: TerramEarth

This is a sample case study that may be used on the Professional Cloud Architect exam. It describes a fictitious business and solution concept to provide additional context to exam questions.

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.

1. 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.

2. 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 three 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 four weeks while they wait for replacement parts.

3. Business requirements

Decrease unplanned vehicle downtime to less than one 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

4. Technical requirements

Expand beyond a single data center to decrease latency to the American Midwest and East Coast

Create a backup strategy

Increase security of data transfer from equipment to the data center

Improve data in the data warehouse

Use customer and equipment data to anticipate customer needs

5. Application 1: Data ingest

A custom Python application reads uploaded data files 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

6. Application 2: Reporting

An off-the-shelf application that business analysts use to run a daily report to see what equipment needs repair. Only two analysts of a team of 10 (five West Coast, five 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

7. Executive statement

Our competitive advantage has always been in our manufacturing process, with our ability to build better vehicles for lower cost than our competitors. However, new products with different approaches are constantly being developed, and I'm concerned that we lack the skills to undergo the next wave of transformations in our industry. My goals are to build our skills while addressing immediate market needs through incremental innovations.

Based on TerramEarth's current data flow environment (FTP -> ETL -> EDW)

What are the direct GCP services needed to replicate the same structure for batch uploads?

☐ Cloud Dataproc - Cloud Storage - BigQuery

☒ Cloud Storage - Cloud Dataflow - BigQuery

Explanation:-As the current data flow for TerramEarth for batch uploads is FTP -> ETL -> EDW, it can be directly migrated to Cloud Storage for storage, Dataflow for ETL processing and transform and BigQuery as the data warehousing solution.

☐ Cloud Dataflow - Cloud Bigtable - Cloud Dataproc

☐ Cloud Spanner - Cloud SQL - BigQuery

Q22)

A company is planning to setup VPN connection with their on-premises data center.

What information is required to connect to an on-premises network router over VPN using Cloud Router for dynamic routing? (Select 3 answers)

☒ Border Gateway Protocol address

Explanation:-For setting up VPN dynamic routes Remote router (peer) IP address, Border Gateway Protocol address, Shared secret are required. Shared secret — Provide a pre-shared key used for authentication. The shared secret for the Cloud VPN tunnel must match the one used when you configure the counterpart tunnel on the on-premises VPN gateway. You can follow these directions to generate a cryptographically strong shared secret. Remote peer IP address — Supply the public IP address of the on-premises VPN gateway. Cloud Router BGP IP and BGP peer IP — The two BGP interface IP addresses must be link-local IP addresses belonging to a common /30 CIDR from the 169.254.0.0/16 block. Each BGP IP defines the respective link-local IP used to exchange route information. For example, 169.254.1.1 and 169.254.1.2 belong to a common /30 block.

☐ IPsec Protocol address

☒ Shared secret

Explanation:-For setting up VPN dynamic routes Remote router (peer) IP address, Border Gateway Protocol address, Shared secret are required. Shared secret — Provide a pre-shared key used for authentication. The shared secret for the Cloud VPN tunnel must match the one used when you configure the counterpart tunnel on the on-premises VPN gateway. You can follow these directions to generate a cryptographically strong shared secret. Remote peer IP address — Supply the public IP address of the on-premises VPN gateway. Cloud Router BGP IP and BGP peer IP — The

two BGP interface IP addresses must be link-local IP addresses belonging to a common /30 CIDR from the 169.254.0.0/16 block. Each BGP IP defines the respective link-local IP used to exchange route information. For example, 169.254.1.1 and 169.254.1.2 belong to a common /30 block.

● Remote router DNS name

✓ Remote router (peer) IP address

Explanation:—For setting up VPN dynamic routes Remote router (peer) IP address, Border Gateway Protocol address, Shared secret are required. Shared secret — Provide a pre-shared key used for authentication. The shared secret for the Cloud VPN tunnel must match the one used when you configure the counterpart tunnel on the on-premises VPN gateway. You can follow these directions to generate a cryptographically strong shared secret. Remote peer IP address — Supply the public IP address of the on-premises VPN gateway. Cloud Router BGP IP and BGP peer IP — The two BGP interface IP addresses must be link-local IP addresses belonging to a common /30 CIDR from the 169.254.0.0/16 block. Each BGP IP defines the respective link-local IP used to exchange route information. For example, 169.254.1.1 and 169.254.1.2 belong to a common /30 block.

Q23)

Professional Cloud Architect

Sample case study: Dress4Win

This is a sample case study that may be used on the Professional Cloud Architect exam. It describes a fictitious business and solution concept to provide additional context to exam questions.

Dress4Win is a web-based company that helps their users organize and manage their personal wardrobe using a web app and mobile application. The company also cultivates an active social network that connects their users with designers and retailers. They monetize their services through advertising, ecommerce, referrals, and a freemium app model. The application has grown from a few servers in the founder's garage to several hundred servers and appliances in a colocated data center. However, the capacity of their infrastructure is now insufficient for the application's rapid growth. Because of this growth and the company's desire to innovate faster, Dress4Win is committing to a full migration to a public cloud.

1. Solution concept

For the first phase of their migration to the cloud, Dress4Win is moving their development and test environments. They are also building a disaster recovery site, because their current infrastructure is at a single location. They are not sure which components of their architecture they can migrate as is and which components they need to change before migrating them.

2. Existing technical environment

The Dress4Win application is served out of a single data center location. All servers run Ubuntu LTS v16.04.

Databases:

MySQL. One server for user data, inventory, static data,

MySQL 5.7

8 core CPUs

128 GB of RAM

2x 5 TB HDD (RAID 1)

Compute:

40 web application servers providing micro-services based APIs and static content

Tomcat - Java

Nginx

Four core CPUs

32 GB of RAM

20 Apache Hadoop/Spark servers:

Data analysis

Real-time trending calculations

Eight core CPUs

128 GB of RAM

4x 5 TB HDD (RAID 1)

Three RabbitMQ servers for messaging, social notifications, and events:

Eight core CPUs

32GB of RAM

Miscellaneous servers:

Jenkins, monitoring, bastion hosts, security scanners

Eight core CPUs

32GB of RAM

Storage appliances:

iSCSI for VM hosts

Fibre channel SAN - MySQL databases

1 PB total storage; 400 TB available

NAS - image storage, logs, backups

100 TB total storage; 35 TB available

3. Business requirements

Build a reliable and reproducible environment with scaled parity of production

Improve security by defining and adhering to a set of security and identity and access management (IAM) best practices for cloud

Improve business agility and speed of innovation through rapid provisioning of new resources

Analyze and optimize architecture for performance in the cloud

4. Technical requirements

Easily create non-production environments in the cloud

Implement an automation framework for provisioning resources in cloud

Implement a continuous deployment process for deploying applications to the on-premises data center or cloud

Support failover of the production environment to cloud during an emergency

Encrypt data on the wire and at rest

Support multiple private connections between the production data center and cloud environment.

5. Executive statement

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

developing better features. Our traffic patterns are highest in the mornings and weekend evenings; during other times, 80% of our capacity is sitting idle.

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

The Dress4Win developers are evaluating using Google Cloud Platform. They have identified some applications that can easily move to Google App Engine Flexible Environment. The developers will deploy their code using the Google Cloud SDK tools.

Which two of their stated technical requirements does this solution meet?(Choose 2)

- ☐ Support multiple private connections between the production data center and cloud environment.
- ☒ Use managed services whenever possible

Explanation:-App Engine is a fully managed serverless application platform. Build and deploy applications on a fully managed platform. Scale your applications seamlessly from zero to planet scale without having to worry about managing the underlying infrastructure. With zero server management and zero configuration deployments, developers can focus only on building great applications without the management overhead. App Engine enables developers to stay more productive and agile by supporting popular development languages and a wide range of developer tools. Focus just on writing code, without the worry of managing the underlying infrastructure. With capabilities such as automatic scaling-up and scaling-down of your application between zero and planet scale, fully managed patching and management of your servers, you can offload all your infrastructure concerns to Google. Protect your applications from security threats using App Engine firewall capabilities, Identity and Access Management (IAM) rules, and managed SSL/ TLS certificates.

- ☐ Encrypt data on the wire and at rest.
- ☒ Implement an automation framework for provisioning resources in cloud.

Explanation:-App Engine provide automation and lets you focus on the code only. Build and deploy applications on a fully managed platform. Scale your applications seamlessly from zero to planet scale without having to worry about managing the underlying infrastructure. With zero server management and zero configuration deployments, developers can focus only on building great applications without the management overhead. App Engine enables developers to stay more productive and agile by supporting popular development languages and a wide range of developer tools. Focus just on writing code, without the worry of managing the underlying infrastructure. With capabilities such as automatic scaling-up and scaling-down of your application between zero and planet scale, fully managed patching and management of your servers, you can offload all your infrastructure concerns to Google. Protect your applications from security threats using App Engine firewall capabilities, Identity and Access Management (IAM) rules, and managed SSL/ TLS certificates.

Q24) What is the command for creating a storage bucket that has once per month access and is named 'archive_bucket'?

- ☒ gsutil mb -c nearline gs://archive_bucket

Explanation:-The data needs to be accessed on monthly basis Nearline is and ideal storage class. Also gsutil needs -c parameter to pass the class. Nearline - Data you do not expect to access frequently (i.e., no more than once per month). Ideal for back-up and serving long-tail multimedia content.

- ☐ gsutil mb -c coldline gs://archive_bucket
- ☐ gsutil mb gs://archive_bucket
- ☐ gsutil rm -coldline gs://archive_bucket

Q25)

Professional Cloud Architect

Sample case study: Dress4Win

This is a sample case study that may be used on the Professional Cloud Architect exam. It describes a fictitious business and solution concept to provide additional context to exam questions.

Dress4Win is a web-based company that helps their users organize and manage their personal wardrobe using a web app and mobile application. The company also cultivates an active social network that connects their users with designers and retailers. They monetize their services through advertising, ecommerce, referrals, and a freemium app model. The application has grown from a few servers in the founder's garage to several hundred servers and appliances in a colocated data center. However, the capacity of their infrastructure is now insufficient for the application's rapid growth. Because of this growth and the company's desire to innovate faster, Dress4Win is committing to a full migration to a public cloud.

1. Solution concept

For the first phase of their migration to the cloud, Dress4Win is moving their development and test environments. They are also building a disaster recovery site, because their current infrastructure is at a single location. They are not sure which components of their architecture they can migrate as is and which components they need to change before migrating them.

2. Existing technical environment

The Dress4Win application is served out of a single data center location. All servers run Ubuntu LTS v16.04.

Databases:

MySQL. One server for user data, inventory, static data,
MySQL 5.7
8 core CPUs
128 GB of RAM
2x 5 TB HDD (RAID 1)

Compute:

40 web application servers providing micro-services based APIs and static content
Tomcat - Java
Nginx
Four core CPUs
32 GB of RAM
20 Apache Hadoop/Spark servers:
Data analysis
Real-time trending calculations
Eight core CPUs
128 GB of RAM

4x 5 TB HDD (RAID 1)

Three RabbitMQ servers for messaging, social notifications, and events:

Eight core CPUs

32GB of RAM

Miscellaneous servers:

Jenkins, monitoring, bastion hosts, security scanners

Eight core CPUs

32GB of RAM

Storage appliances:

iSCSI for VM hosts

Fibre channel SAN - MySQL databases

1 PB total storage; 400 TB available

NAS - image storage, logs, backups

100 TB total storage; 35 TB available

3. Business requirements

Build a reliable and reproducible environment with scaled parity of production

Improve security by defining and adhering to a set of security and identity and access management (IAM) best practices for cloud

Improve business agility and speed of innovation through rapid provisioning of new resources

Analyze and optimize architecture for performance in the cloud

4. Technical requirements

Easily create non-production environments in the cloud

Implement an automation framework for provisioning resources in cloud

Implement a continuous deployment process for deploying applications to the on-premises data center or cloud

Support failover of the production environment to cloud during an emergency

Encrypt data on the wire and at rest

Support multiple private connections between the production data center and cloud environment.

5. Executive statement

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

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

Dress4Win has a number of internal backends that they do not want to be exposed to the public Internet.

How can they reduce their external exposure while still allowing maintenance access to resources when working remotely?

- ☐ Hide the external IP address behind a load balancer and restrict load balancer access to the internal company network.
- ☐ Remove the external IP address and have remote employees dial into the company VPN connection for maintenance work.
- ☒ Remove the external IP address and use a bastion host to access internal-only resources.

Explanation:-It is a best practice to remove external ip address from the instances so that they are not reachable from the internet and have a Bastion host of Jump server to be able to login into the servers. Bastion hosts provide an external facing point of entry into a network containing private network instances. This host can provide a single point of fortification or audit and can be started and stopped to enable or disable inbound SSH communication from the Internet. By using a bastion host, you can connect to an instance that does not have an external IP address. This approach allows you to connect to a development environment or manage the database instance for your external application, for example, without configuring additional firewall rules. A complete hardening of a bastion host is outside the scope of this article, but some initial steps taken can include: Limit the CIDR range of source IPs that can communicate with the bastion. Configure firewall rules to allow SSH traffic to private instances from only the bastion host. By default, SSH on instances is configured to use private keys for authentication. When using a bastion host, you log into the bastion host first, and then into your target private instance. Because of this two-step login, which is why bastion hosts are sometimes called "jump servers," you should use ssh-agent forwarding instead of storing the target machine's private key on the bastion host as a way of reaching the target machine. You need to do this even if using the same key-pair for both bastion and target instances, as the bastion has direct access to only the public half of the key-pair.

- ☐ Remove the external IP address and use Cloud Shell to access internal-only resources

Q26)

A company uses Cloud Storage for storing their critical data. As a part of compliance, the objects need to be encrypted using customer-supplied encryption keys.

How should the object be handled to support customer-supplied encryption?

- ☐ Use gcloud config to define the encryption
- ☒ Use boto config to define the encryption key

Explanation:-The encryption key can be configured in the boto configuration file.

1. Add the following option to the [GSUtil] section of your boto configuration file:

encryption_key = [YOUR_ENCRYPTION_KEY]

where [YOUR_ENCRYPTION_KEY] is the key for encrypting the uploaded file. Note: You can alternatively include this information in each gsutil command by using the -o top level flag: -o "GSUtil:encryption_key=[YOUR_ENCRYPTION_KEY]"

2. Use the gsutil cp command, replacing [VALUES_IN_BRACKETS] with the appropriate values:

gsutil cp [LOCAL_OBJECT_LOCATION] gs://[DESTINATION_BUCKET_NAME]/

- ☐ Create bucket with —encryption-key and use gsutil to upload files
- ☐ Use gsutil with —encryption-key to pass the encryption key

Q27)

A company is hosting their Echo application on Google Cloud using Google Kubernetes Engine.

The application is deployed with deployment echo-deployment exposed with echo-service. They have a new image that needs to be deployed for the application.

How can the change be deployed with minimal downtime?

- ☐ Update image in instance template and use rolling deployment of instance group with Kubernetes engine.
- ☐ Delete the service and create a new service with the updated image
- ☐ Delete the deployment and create a new deployment with the updated image
- ☒ Update image using kubectl set image deployment

Explanation:-The image can be directly updated using the kubectl command and Kubernetes Engine performs a rolling update. You can perform a rolling update to update the images, configuration, labels, annotations, and resource limits/requests of the workloads in your clusters. Rolling updates incrementally replace your resource's Pods with new ones, which are then scheduled on nodes with available resources. Rolling updates are designed to update your workloads without downtime. You can use kubectl set to make changes to an object's image, resources (compute resource such as CPU and memory), or selector fields. For example, to update a Deployment from nginx version 1.7.9 to 1.9.1, run the following command:

```
kubectl set image deployment nginx nginx=nginx:1.9.1
```

The kubectl set image command updates the nginx image of the Deployment's Pods one at a time.

Q28)

A company is storing their data in BigQuery for analysis. The data science team accessing the data needs only the data of last 45 days.

How would you design to ensure cost optimization and easy of analysis?

- ☐ Configure cron job to run bq command to delete data every 45 days
- ☒ Create time-partitioned table with expiration

Explanation:-BigQuery tables can be configured to create partitioned tables with the ability to specify expiration time to automatically expire the data. Both ingestion-time partitioned and partitioned tables are managed the same way. You can perform the following management tasks for partitioned tables:

Update a time-partitioned table's:

Table expiration time

Partition expiration time

Description

Schema definition

Labels

Rename (copy) a time-partitioned table

Copy a time-partitioned table

Copy partitions

Delete a time-partitioned table

Delete partitions in a time-partitioned table

- ☐ Create Table and configure expiration on the table
- ☐ Create time-partitioned table in Bigtable instead

Q29)

A company wants building an application stores images in a Cloud Storage bucket and want to generate thumbnails as well resize the images.

They want to use managed service which will help them scale automatically from zero to scale and back to zero.

Which GCP service satisfies the requirement?

- ☐ Google Kubernetes Engine
- ☐ Google App Engine
- ☒ Cloud Functions

Explanation:-Cloud Functions can help automatically scale as per the demand, with no invocations if no demand. Google Cloud Functions is a serverless execution environment for building and connecting cloud services. With Cloud Functions you write simple, single-purpose functions that are attached to events emitted from your cloud infrastructure and services. Your function is triggered when an event being watched is fired. Your code executes in a fully managed environment. There is no need to provision any infrastructure or worry about managing any servers. Cloud Functions removes the work of managing servers, configuring software, updating frameworks, and patching operating systems. The software and infrastructure are fully managed by Google so that you just add code. Furthermore, provisioning of resources happens automatically in response to events. This means that a function can scale from a few invocations a day to many millions of invocations without any work from you.

- ☐ Google Compute Engine

Q30)

Using the principle of least privilege, your colleague Bob needs to be able to create new instances on Compute Engine in project 'Project A'.

How should you give him access without giving more permissions than is necessary?

- ☐ Give Bob Project Editor IAM role for Project A.
- ☐ Create a shared VPC that Bob can access Compute resources from.
- ☐ Give Bob Compute Engine Admin Role for Project A.
- ☒ Give Bob Compute Engine Instance Admin Role for Project A.

Explanation:-the access needs to be given only to create instances, the user should be given compute instance admin role, which provides the least privilege.

```
roles/compute.instanceAdmin.v1
```

```
roles/compute.admin
```


Q31)

You need to take streaming data from thousands of Internet of Things (IoT) devices, ingest it, run it through a processing pipeline, and store it for analysis.

You want to run SQL queries against your data for analysis. What services in which order should you use for this task?

- ☐ App Engine, Cloud Dataflow, BigQuery
- ☒ Cloud Pub/Sub, Cloud Dataflow, BigQuery

Explanation:-The need to ingest it, transform and store the Cloud Pub/Sub, Cloud Dataflow, BigQuery is ideal stack to handle the IoT data. Google Cloud Pub/Sub provides a globally durable message ingestion service. By creating topics for streams or channels, you can enable different components of your application to subscribe to specific streams of data without needing to construct subscriber-specific channels on each device. Cloud Pub/Sub also natively connects to other Cloud Platform services, helping you to connect ingestion, data pipelines, and storage systems. Google Cloud Dataflow provides the open Apache Beam programming model as a managed service for processing data in multiple ways, including batch operations, extract-transform-load (ETL) patterns, and continuous, streaming computation. Cloud Dataflow can be particularly useful for managing the high-volume data processing pipelines required for IoT scenarios. Cloud Dataflow is also designed to integrate seamlessly with the other Cloud Platform services you choose for your pipeline. Google BigQuery provides a fully managed data warehouse with a familiar SQL interface, so you can store your IoT data alongside any of your other enterprise analytics and logs. The performance and cost of BigQuery means you might keep your valuable data longer, instead of deleting it just to save disk space. Sample Arch - Mobile Gaming Analysis Telemetry

- ☐ Cloud Pub/Sub, Cloud Dataflow, Cloud Dataproc
- ☐ Cloud Dataflow, Cloud Pub/Sub, BigQuery

Q32)

Your company is planning on deploying a web application to Google Cloud hosted on a custom Linux distribution.

Your website will be accessible globally and needs to scale to meet demand.

Choose all of the components that will be necessary to achieve this goal. (Select TWO)

- ☒ Managed Instance Group on Compute Engine

Explanation:-The requirement is to support custom Linux distribution, only Compute Engine supports the same. HTTP(S) load balancing can balance HTTP and HTTPS traffic across multiple backend instances, across multiple regions. Your entire app is available via a single global IP address, resulting in a simplified DNS setup. HTTP(S) load balancing is scalable, fault-tolerant, requires no pre-warming, and enables content-based load balancing. For HTTPS traffic, it provides SSL termination and load balancing.

- ☒ HTTP Load Balancer

Explanation:-Only HTTP load balancer support global access.

- ☐ Network Load Balancer
- ☐ App Engine Standard environment

Q33)

Your business is connected to a VPN connection to GCP. On GCP, you are hosting an internal website using an HTTP load balancer which serves traffic to a managed instance group.

Your users are unable to access the internal website, though all systems are up. What is a possible cause for this problem?

- ☐ You need to enable Stackdriver debugger to collect more information.
- ☐ Your load balancer is not set for the correct region.
- ☒ Firewall rules are not correctly configured.

Explanation:-The systems are up and the users are not able to login, the instances firewall rules would not be configured to allow ingress traffic from the users network. When you configure Cloud VPN tunnels to connect your Google Cloud Platform network to an on-premises network, you should review and modify firewall rules in both networks to make sure that they meet your needs. This page provides guidance for configuring GCP firewall rules and your on-premises network firewall rules. The implied allow egress rule allows instances and other resources in your GCP network to make outgoing requests and receive established responses, but the implied deny ingress rule blocks all incoming traffic to your GCP resources. At minimum, you need to create firewall rules to allow ingress traffic from your on-premises network to GCP. You may also need to create egress rules if you have created other egress rules to deny certain types of traffic.

- ☐ You've enabled static IP address for the managed instance group servers.

Q34)

Your infrastructure runs on AWS and includes a set of multi-TB enterprise databases that are backed up nightly on the S3.

You need to create a redundant backup to Google Cloud. You are responsible for performing scheduled monthly disaster recovery drills.

You want to create a cost-effective solution. What should you do?

- ☐ Use Storage Transfer Service to transfer the backup files to a Cloud Storage Coldline storage bucket as a final destination.
- ☒ Use Storage Transfer Service to transfer the backup files to a Cloud Storage Nearline storage bucket as a final destination.

Explanation:-The data needs to be backed up nightly and accessed monthly, Storage Transfer service can be used to transfer it from S3 and Nearline storage for cost-effective storage solution. is correct because you will need to access your backup data monthly to test your disaster recovery process, so you should use a Nearline bucket; also because you will be performing ongoing, regular data transfers, so you should use the storage transfer service.

- ☐ Use Transfer Appliance to transfer the backup files to a Cloud Storage Coldline bucket as a final destination.
- ☐ Use Transfer Appliance to transfer the backup files to a Cloud Storage Nearline storage bucket as a final destination.

Q35)

You need to create a new Kubernetes Cluster on Google Cloud Platform that can autoscale the number of worker nodes.

What should you do?

- Create Compute Engine instances for the workers and the master and install Kubernetes. Rely on Kubernetes to create additional Compute Engine instances when needed.
- Create a cluster on Kubernetes Engine and enable autoscaling on the instance group of the cluster.
- Configure a Compute Engine instance as a worker and add it to an unmanaged instance group. Add a load balancer to the instance group and rely on the load balancer to create additional Compute Engine instances when needed.
- ✔ Create a cluster on Kubernetes Engine and enable autoscaling on Kubernetes Engine.

Explanation:-Kubernetes cluster provides auto scaling feature which can be enabled on the cluster engine. GKE's cluster autoscaler automatically resizes clusters based on the demands of the workloads you want to run. With autoscaling enabled, GKE automatically adds a new node to your cluster if you've created new Pods that don't have enough capacity to run; conversely, if a node in your cluster is underutilized and its Pods can be run on other nodes, GKE can delete the node. Cluster autoscaling allows you to pay only for resources that are needed at any given moment, and to automatically get additional resources when demand increases.

Q36)

Your application has a large international audience and runs stateless virtual machines within a managed instance group across multiple locations. One feature of the application lets users upload files and share them with other users.

Files must be available for 30 days; after that, they are removed from the system entirely. Which storage solution should you choose?

- A managed instance group of Filestore servers.
- Persistent SSD on virtual machine instances.
- ✔ A multi-regional Cloud Storage bucket.

Explanation:-The key storage requirements is it being global, allow lifecycle management and sharing capability. Cloud Storage is an ideal choice as it can be configured to be multi-regional, have lifecycle management rules to auto delete the files after 30 days and share them with others.

- A Cloud Datastore database.

Q37)

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?

- ✔ Choose larger instances for your database
- Create snapshots of your database more regularly
- Implement routinely scheduled failovers of your databases

Explanation:-Routinely scheduled failovers would help to test the process regularly. Replicas provide a very good RPO and RTO and the data loss would be minimal. Replicas can be configured to be automatically promoted and it doesn't work there might a chance of misconfiguration.

- Use a different database

Q38)

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?

- 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
- Work with your ISP to diagnose the problem
- Open a support ticket to ask for network capture and flow data to diagnose the problem, then roll back your application
- ✔ 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

Explanation:-The recommended approach would be to roll back the application. The application should then be debugged in the lower environment using Stackdriver.

Q39)

Your customer is moving their corporate applications to Google Cloud Platform. The security team wants detailed visibility of all projects in the organization. You provision the Google Cloud Resource Manager and set up yourself as the org admin.

What Google Cloud Identity and Access Management (Cloud IAM) roles should you give to the security team?

- Project owner, network admin
- Org admin, project browser
- ✔ Org viewer, project viewer

Explanation:-The security team only needs visibility to the projects, project viewer provides the same with the best practice of least privilege.

- Org viewer, project owner

Q40)

Your marketing department wants to send out a promotional email campaign. The development team wants to minimize direct operation management.

They project a wide range of possible customer responses, from 100 to 500,000 click-through per day. The link leads to a simple website that explains the promotion and collects user information and preferences.

Which infrastructure should you recommend? Choose 2 answers.

- Use a single Compute Engine virtual machine (VM) to host a web server, backend by Google Cloud SQL.

✔ Use a managed instance group to serve the website and Google Cloud Bigtable to store user data.

Explanation:-App Engine for the website and Datastore to store responses. App Engine & Datastore are GCP managed service and would provide the scalability.

- Use a Google Container Engine cluster to serve the website and store data to persistent disk.

Explanation:-Managed instance group with Google Compute Engine for the website and Bigtable to store responses. Bigtable provides a NoSQL datastore with better scalability but is more expensive.

✔ Use Google App Engine to serve the website and Google Cloud Datastore to store user data.

Explanation:-App Engine for the website and Datastore to store responses. App Engine & Datastore are GCP managed service and would provide the scalability.

Q41)

One of your primary business objectives is being able to trust the data stored in your application. You want to log all changes to the application data.

How can you design your logging system to verify authenticity of your logs?

✔ Digitally sign each timestamp and log entry and store the signature

Explanation:-Digitally signing the log is the only way to verify authenticity of the logs.

References: KMS - Digital Signatures

A digital signature is a cryptographic output used to verify the authenticity of data. A digital signature algorithm allows for two distinct operations: a signing operation, which uses a signing key to produce a signature over raw data

a verification operation, where the signature can be validated by a party who has no knowledge of the signing key

The main purposes of a digital signature are:

verification of the integrity of the signed data

non-repudiation if the signer claims the signature is not authentic

Digital signatures rely on asymmetric cryptography, also known as public key cryptography. An asymmetric key consists of a public/private key pair.

The private key is used to create a signature, and the corresponding public key is used to verify the signature.

- Use a SQL database and limit who can modify the log table
- Create a JSON dump of each log entry and store it in Google Cloud Storage
- Write the log concurrently in the cloud and on premises

Q42)

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

Which three practices should you recommend? Choose 3 answers.

- Select an automation framework to reliably provision the cloud infrastructure
- ✔ Deploy a continuous integration tool with automated testing in a staging environment

Explanation:-The application can be migrated to App Engine which is a managed Google platform.

Even though Stackdriver Debugger is enabled by default and not explicit instrumentation is required, if you go for custom images for J2EE application it would need to be enabled.

A continuous integration tool with automated testing for deployment

- Integrate Cloud Dataflow into the application to capture real-time metrics
- ✔ Instrument the application with a monitoring tool like Stackdriver Debugger

Explanation:-The application can be migrated to App Engine which is a managed Google platform.

Even though Stackdriver Debugger is enabled by default and not explicit instrumentation is required, if you go for custom images for J2EE application it would need to be enabled.

A continuous integration tool with automated testing for deployment

- Migrate from MySQL to a managed NoSQL database like Google Cloud Datastore or Bigtable
- ✔ Port the application code to run on Google App Engine

Explanation:-The application can be migrated to App Engine which is a managed Google platform.

Even though Stackdriver Debugger is enabled by default and not explicit instrumentation is required, if you go for custom images for J2EE application it would need to be enabled.

A continuous integration tool with automated testing for deployment

Q43)

An application development team believes their current logging tool will not meet their needs for their new cloud-based product on Google Cloud.

They want a better tool to capture errors and help them analyze their historical log data. You want to help them find a solution that meets their needs.

What should you do?

- Help them upgrade their current tool to take advantage of any new features
- Help them define their requirements and assess viable logging tools
- Send them a list of online resources about logging best practices
- ✔ Direct them to download and install the Google StackDriver logging agent

Explanation:-Stackdriver can be used to capture logs and ability to analyze. The logs can be streamed to Stackdriver using the Stackdriver logging agent. The Stackdriver Logging agent streams logs from your VM instances and from selected third-party software packages to Stackdriver Logging. It is a best practice to run the Stackdriver Logging agent on all your VM instances.

Q44)

To reduce costs, the Director of Engineering has required all developers to move their development infrastructure resources from on-premises virtual machines (VMs) to Google Cloud Platform. These resources go through multiple start/stop events

during the day and require state to persist.

You have been asked to design the process of running a development environment in Google Cloud while providing cost visibility to the finance department.

Which two steps should you take? Choose 2 answers.

- ☐ Store all state into local SSD, snapshot the persistent disks, and terminate the VM
- ☒ Use Google BigQuery billing export and labels to associate cost to groups

Explanation:-The disk would be persisted, allowing the developer to get back their VM.

BigQuery would give the cost visibility by labeling the resources.

- ☐ Apply VM CPU utilization label and include it in the BigQuery billing export
- ☐ Use the --auto-delete flag on all persistent disks and terminate the VM
- ☒ Use the --no-auto-delete flag on all persistent disks and stop the VM

Explanation:-The disk would be persisted, allowing the developer to get back their VM.

BigQuery would give the cost visibility by labeling the resources.

- ☐ Store all state in Google Cloud Storage, snapshot the persistent disks, and terminate the VM

Q45)

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?

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

Explanation:-The object lifecycle in Cloud Storage can be automatically controlled using a JSON document defining the rules. Sets the lifecycle configuration on one or more buckets. The config-json-file specified on the command line should be a path to a local file containing the lifecycle configuration JSON document.

- ☐ Write a lifecycle management rule in XML and push it to the bucket with gsutil

Q46)

You set up an autoscaling instance group to serve web traffic for an upcoming launch.

After configuring the instance group as a backend service to an HTTP(S) load balancer, you notice that virtual machine (VM) instances are being terminated and re-launched every minute. The instances do not have a public IP address.

You have verified the appropriate web response is coming from each instance using the curl command. You want to ensure the backend is configured correctly. What should you do?

- ☐ Create a tag on each instance with the name of the load balancer. Configure a firewall rule with the name of the load balancer as the source and the instance tag as the destination.
- ☐ Assign a public IP to each instance and configure a firewall rule to allow the load balancer to reach the instance public IP.
- ☒ Ensure that a firewall rule exists to allow load balancer health checks to reach the instances in the instance group.

Explanation:-The auto scaling group health check failures can result in auto healing leading to VM being terminated and relaunched. The instance should be configured to allow traffic from health check probes which depend on the type of load balancer. You need to create ingress firewall rules applicable to all VMs being load balanced to allow traffic from health check prober IP ranges. The following examples create firewall rules that are applicable to VM instances by target tag. For more information about specifying targets for firewall rules, see the explanation of targets in the Firewall Rules Overview and Configuring Network Tags. Each of these examples allows all TCP traffic from GCP health check systems to your VM instances. (TCP traffic includes SSL, HTTP, HTTPS, and HTTP/2 traffic.) If you prefer, you can specify ports along with the TCP protocol; however, if you specify ports, your firewall rules might become specific to a particular health check. If you use tcp:80 for the protocol and port, this allows TCP traffic on port 80, so GCP could contact your VMs using HTTP on port 80, but it could not contact them using HTTPS on port 443.

- ☐ Ensure that a firewall rules exists to allow source traffic on HTTP/HTTPS to reach the load balancer.

Q47)

Your company wants to track whether someone is present in a meeting room reserved for a scheduled meeting.

There are 1000 meeting rooms across 5 offices on 3 continents. Each room is equipped with a motion sensor that reports its status every second. The data from the motion detector includes only a sensor ID and several different discrete items of information.

Analysts will use this data, together with information about account owners and office locations. Which database type should you use?

- ☐ Relational
- ☐ Blobstore
- ☒ NoSQL

Explanation:-NoSQL like Bigtable and Datastore solution is an ideal solution to store sensor ID and several different discrete items of information. It also provides an ability to join with other data. Datastore can also be configured to store data in multi-region locations.

- ☐ Flat file

Q48)

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?

- Google Compute Engine with Google BigQuery
- Google Container Engine with Bigtable
- ✓ Google Cloud Dataflow

Explanation:- There is no existing scripts or code, Dataflow suits the need perfectly as it supports both batch and streaming jobs. 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. And with its serverless approach to resource provisioning and management, you have access to virtually limitless capacity to solve your biggest data processing challenges, while paying only for what you use. Cloud Dataflow unlocks transformational use cases across industries, including:

- Clickstream, Point-of-Sale, and segmentation analysis in retail
- Fraud detection in financial services
- Personalized user experience in gaming
- IoT analytics in manufacturing, healthcare, and logistics

- Google Cloud Dataproc

Q49)

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

```
1. Uncaught exception from servlet
2. java.lang.SecurityException: SHA1 digest error for com/example/SomeClass.class
3. at com.google.appengine.runtime.Request.process-d36f818a24b8cf1d(Request.java)
4. at sun.security.util.ManifestEntryVerifier.verify(ManifestEntryVerifier.java:210)
5. at java.util.jar.JarVerifier.processEntry(JarVerifier.java:218)
6. at java.util.jar.JarVerifier.update(JarVerifier.java:205)
7. at java.util.jar.JarVerifier$VerifierStream.read(JarVerifier.java:428)
8. at sun.misc.Resource.getBytes(Resource.java:124)
9. at java.net.URLClassLoader.defineClass(URLClassLoader.java:273)
10. at sun.reflect.GeneratedMethodAccessor5.invoke(Unknown Source)
11. at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)
12. at java.lang.reflect.Method.invoke(Method.java:616)
13. at java.lang.ClassLoader.loadClass(ClassLoader.java:266)
```

What should you do?

- Remove the JAR file which is causing issue and redeploy the application
- Upgrade the Jar versions, as the versions are incompatible
- ✓ Turn off precompilation for App Engine and redeploy your application.

Explanation:- App Engine uses a "precompilation" process with the Java bytecode of an app to enhance the performance of the app in the Java runtime environment. Precompiled code functions identically to the original bytecode. However, if the jar is already signed and receive the SHA1 digest error either strip the jar of its signature or turn off the precompilation. App Engine's precompilation isn't compatible with signed JAR files. If your application is precompiled (the default), it can't load signed JAR files. If the application tries to load a signed JAR, at runtime App Engine will generate an exception like `java.lang.SecurityException: SHA1 digest error for com/example/SomeClass.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.JarVerifier$VerifierStream.read(JarVerifier.java:428) at sun.misc.Resource.getBytes(Resource.java:124) at java.net.URLClassLoader.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)`

There are two ways to work around this:

Strip the JAR's signature; or Disable precompilation for the whole application.

- Upload missing JAR files and redeploy your application.

Q50)

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?

- 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
- 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
- ✓ 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

Explanation:- You need to simulate more load to represent the additional users. Also for testing the resiliency, you need to terminate instances and check if it works and provides the same SLA.

- 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

Q51)

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?

- Enable Google Cloud Storage (GCS) log export to audit logs into a GCS bucket and delegate access to the bucket
- Use Cloud Functions to transfer log entries to Google Cloud SQL and use ACLs and views to limit an auditor's view
- ✔ Enable Logging export to Google BigQuery and use ACLs and views to scope the data shared with the auditor

Explanation:-BigQuery is a good storage option with analysis capability. Also, the access to the data can be controlled using ACLs and Views. BigQuery uses access control lists (ACLs) to manage permissions on projects and datasets. BigQuery is a petabyte-scale analytics data warehouse that you can use to run SQL queries over vast amounts of data in near realtime. Giving a view access to a dataset is also known as creating an authorized view in BigQuery. An authorized view allows you to share query results with particular users and groups without giving them access to the underlying tables. You can also use the view's SQL query to restrict the columns (fields) the users are able to query. In this tutorial, you create an authorized view.

- Create custom Google Stackdriver alerts and send them to the auditor

Q52)

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

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

Which combination of technologies will meet all of his requirements?

- Google Kubernetes Engine, Jenkins, and Cloud Load Balancing
- Google Kubernetes Engine and Cloud Deployment Manager
- Google Kubernetes Engine and Cloud Load Balancing
- ✔ Google Kubernetes Engine, Jenkins, and Helm

Explanation:-Google Kubernetes Engine (Point 1,2,4 & 6) is based on open source Kubernetes and can help host the application, autoscale and also provides load balancing capabilities.. Jenkins (Point 1 & 3) can help provide continuous software delivery and Helm (Point 5) provides the ability to deploy application bundles using dynamic templates. Helm helps you manage Kubernetes applications — Helm Charts help you define, install, and upgrade even the most complex Kubernetes application. You can create HTTP(S) load balancers by using an Ingress resource. HTTP(S) load balancers are designed to terminate HTTP(S) requests and can make better context-aware load balancing decisions. They offer features like customizable URL maps and TLS termination. GKE automatically configures health checks for HTTP(S) load balancers.

Q53)

Professional Cloud Architect

Sample case study: Mountkirk Games

This is a sample case study that may be used on the Professional Cloud Architect exam. It describes a fictitious business and solution concept to provide additional context to exam questions.

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.

1. 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 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.

2. 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

3. 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

4. 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 and low latency load balancing and frees us up from managing physical servers.

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?

☐ Cloud SQL, Cloud Storage, Cloud Pub/Sub, and Cloud Dataflow

☒ Cloud Dataflow, Cloud Storage, Cloud Pub/Sub, and BigQuery

Explanation: It meets the technical requirements. Cloud Pub/Sub to capture the real time data. CloudStorage to store files uploaded by users. Dataflow to process data on the fly as well as handle data that arrives late. BigQuery to be able to query historical data. 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 queries to access at least 10 TB of historical data.
5. Process files that are regularly uploaded by users' mobile devices.

☐ Google Kubernetes Engine, Cloud Pub/Sub, and Cloud SQL

☐ Cloud Dataproc, Cloud Pub/Sub, Cloud SQL, and Cloud Dataflow

Q54)

Professional Cloud Architect

Sample case study: Dress4Win

This is a sample case study that may be used on the Professional Cloud Architect exam. It describes a fictitious business and solution concept to provide additional context to exam questions.

Dress4Win is a web-based company that helps their users organize and manage their personal wardrobe using a web app and mobile application. The company also cultivates an active social network that connects their users with designers and retailers. They monetize their services through advertising, ecommerce, referrals, and a freemium app model. The application has grown from a few servers in the founder's garage to several hundred servers and appliances in a colocated data center. However, the capacity of their infrastructure is now insufficient for the application's rapid growth. Because of this growth and the company's desire to innovate faster, Dress4Win is committing to a full migration to a public cloud.

1. Solution concept

For the first phase of their migration to the cloud, Dress4Win is moving their development and test environments. They are also building a disaster recovery site, because their current infrastructure is at a single location. They are not sure which components of their architecture they can migrate as is and which components they need to change before migrating them.

2. Existing technical environment

The Dress4Win application is served out of a single data center location. All servers run Ubuntu LTS v16.04.

Databases:

MySQL. One server for user data, inventory, static data,

MySQL 5.7

8 core CPUs

128 GB of RAM

2x 5 TB HDD (RAID 1)

Compute:

40 web application servers providing micro-services based APIs and static content

Tomcat - Java

Nginx

Four core CPUs

32 GB of RAM

20 Apache Hadoop/Spark servers:

Data analysis

Real-time trending calculations

Eight core CPUs

128 GB of RAM

4x 5 TB HDD (RAID 1)

Three RabbitMQ servers for messaging, social notifications, and events:

Eight core CPUs

32GB of RAM

Miscellaneous servers:

Jenkins, monitoring, bastion hosts, security scanners

Eight core CPUs

32GB of RAM

Storage appliances:

iSCSI for VM hosts

Fibre channel SAN - MySQL databases

1 PB total storage; 400 TB available

NAS - image storage, logs, backups

100 TB total storage; 35 TB available

3. Business requirements

Build a reliable and reproducible environment with scaled parity of production

Improve security by defining and adhering to a set of security and identity and access management (IAM) best practices for cloud

Improve business agility and speed of innovation through rapid provisioning of new resources

Analyze and optimize architecture for performance in the cloud

4. Technical requirements

Easily create non-production environments in the cloud

Implement an automation framework for provisioning resources in cloud

Implement a deployment process for deploying applications to the on-premises data center or cloud
Support failover of the production environment to cloud during an emergency
Encrypt data on the wire and at rest
Support multiple private connections between the production data center and cloud environment.

5. Executive statement

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

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

Dress4Win would like to become familiar with deploying applications to the cloud by successfully deploying some applications quickly, as is.

They have asked for your recommendation. What should you advise?

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

Explanation:-The first application to be moved should be the applications that are easy to move. Self contained applications with no internal dependencies should be the best option.

Assessing the items to migrate

After you have evaluated the cost of the move, then you can start looking at what to migrate. In modern enterprises, there are many different kinds of applications, from customer-facing apps, to back office apps, to developer tools, to experimental applications. Moving all these applications at the same time and the same way doesn't make sense.

We recommend sorting applications into three broad buckets:

Applications that are easy to move. These have fewer dependencies, are newer, are written internally so have no licensing considerations, and are more tolerant to scaling and other cloud patterns.

Applications that are difficult to move. These have more dependencies, are less tolerant to scaling, or have complex license requirements.

Applications that can't be moved. Some applications that might not be good candidates to migrate run on specialized or older hardware, have business or regulatory requirements that make it necessary for them to stay in your data center, or have complex license requirements that don't allow them to move to the cloud.

These are just examples of each of these three buckets, and it is likely your applications have many more deciding factors. Your Cloud Platform sales team can help you.

These considerations all apply whether you're moving from a data center or another cloud provider.

When this work is done, you can pick your first application or applications to migrate. We strongly recommend you migrate a only a few applications at first. The first ones will provide not only the template for future migrations, but also help you define your migration processes.

Q55)

Professional Cloud Architect

Sample case study: TerramEarth

This is a sample case study that may be used on the Professional Cloud Architect exam. It describes a fictitious business and solution concept to provide additional context to exam questions.

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.

1. 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.

2. 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 three 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 four weeks while they wait for replacement parts.

3. Business requirements

Decrease unplanned vehicle downtime to less than one 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

4. Technical requirements

Expand beyond a single data center to decrease latency to the American Midwest and East Coast

Create a backup strategy

Increase security of data transfer from equipment to the data center

Improve data in the data warehouse

Use customer and equipment data to anticipate customer needs

5. Application 1: Data ingest

A custom Python application reads uploaded data files 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

6. Application 2: Reporting

An off-the-shelf application that business analysts use to run a daily report to see what equipment needs repair. Only two analysts of a team of 10 (five West Coast, five 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

7. Executive statement

Our competitive advantage has always been in our manufacturing process, with our ability to build better vehicles for lower cost than our competitors. However, new products with different approaches are constantly being developed, and I'm concerned that we lack the skills to undergo the next wave of transformations in our industry. My goals are to build our skills while addressing immediate market needs through incremental innovations.

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?

- ☒ Have the vehicle's computer compress the data in hourly snapshots, and store it in a GCS Coldline bucket

Explanation:-The analytics need to be done next year and the storage costs needs to be kept low, GCS Coldline storage is the best and lowest cost storage option.

- ☐ Push the telemetry data in real-time to a streaming Dataflow job that compresses the data, and store it in Cloud Bigtable
- ☐ Push the telemetry data in real-time to a streaming Dataflow job that compresses the data, and store it in Google BigQuery
- ☐ Have the vehicle's computer compress the data in hourly snapshots, and store it in a Google Cloud Storage (GCS) Nearline bucket

Q56)

Professional Cloud Architect

Sample case study: TerramEarth

This is a sample case study that may be used on the Professional Cloud Architect exam. It describes a fictitious business and solution concept to provide additional context to exam questions.

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.

1. 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.

2. 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 three 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 four weeks while they wait for replacement parts.

3. Business requirements

Decrease unplanned vehicle downtime to less than one 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

4. Technical requirements

Expand beyond a single data center to decrease latency to the American Midwest and East Coast

Create a backup strategy

Increase security of data transfer from equipment to the data center

Improve data in the data warehouse

Use customer and equipment data to anticipate customer needs

5. Application 1: Data ingest

A custom Python application reads uploaded data files 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

6. Application 2: Reporting

An off-the-shelf application that business analysts use to run a daily report to see what equipment needs repair. Only two analysts of a team of 10 (five West Coast, five 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

7. Executive statement

Our competitive advantage has always been in our manufacturing process, with our ability to build better vehicles for lower cost than our competitors. However, new products with different approaches are constantly being developed, and I'm concerned that we lack the skills to undergo the next wave of transformations in our industry. My goals are to build our skills while addressing immediate market needs through incremental innovations.

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 three)

- ☐ Export a debug VM into an image, and run the image on a local server where kernel log messages will be displayed on the
- ☐ Identify whether a live migration event of the failed server occurred, using the activity log.
- ☒ Adjust the Google Stackdriver timeline to match the failure time, and observe the batch server metrics.

Explanation:-Stackdriver logging can help with the logs which can indicate the errors or exceptions resulting failure of the batch jobs

- ☒ Use gcloud or Cloud Console to connect to the serial console and observe the logs.

Explanation:-Stackdriver logging can help with the logs which can indicate the errors or exceptions resulting failure of the batch jobs

- ☐ Read the debug GCE Activity log using the API or Cloud Console.

Explanation:-gcloud or serial port console can be used to login into the servers to check the logs. Interacting with the Serial Console - The instance's operating system, BIOS, and other system-level entities often write output to the serial ports, and can accept input such as commands or answers to prompts.

- ☒ Use Stackdriver Logging to search for the module log entries.

Explanation:-Stackdriver logging can help with the logs which can indicate the errors or exceptions resulting failure of the batch jobs

Q57)

You are monitoring a service with uptime checks.

The services are reported as unavailable from different GCP regions.

You know that the service is up and running.

How can you solve the monitoring issues?

- ☐ Download the source IPs from the uptime check console and create an egress firewall rule for the service.
- ☐ Use a third-party tool, outside GCP, to create the uptime checks.
- ☐ Install Stackdriver monitoring agents on all instances that are hosting the service.
- ☒ Download the source IPs from the uptime check console and create an ingress firewall rule for the service.

Explanation:-You need to open firewall rules to allow uptime checks. Reference: <https://cloud.google.com/monitoring/uptime-checks/using-uptime-checks#get-ips>: