Google.Premium.Professional Cloud Architect.by.VCEplus.35q - DEMO

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

Website: www.vceplus.com

Free Exam: https://vceplus.com/exam-Professional-Cloud-Architect /

Questions & Answers Exam Engine is rigorously checked before being put up for sale. We make sure there is nothing irrelevant in Professional Cloud Architect exam products and you get latest questions. We strive to deliver the best Professional Cloud Architect exam product for top grades in your first attempt.

Website: https://vceplus.com
VCE to PDF Converter: https://vceplus.com/vce-to-pdf/
Facebook: https://www.facebook.com/VCE.For.All.VN/
Twitter : https://twitter.com/VCE_Plus

JencoMart

QUESTION 1
Company Overview

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

Company Background

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

Solution Concept

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

Existing Technical Environment

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

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

Application: Customer loyalty portal

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

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

- Authenticates all users

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

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

Business Requirements

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

Technical Requirements

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

CEO Statement

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

CTO Statement

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

CFO Statement

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

A. **Overview**

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 2
The JencoMart security team requires that all Google Cloud Platform infrastructure is deployed using a least privilege model with separation of duties for administration between production and development resources.
What Google domain and project structure should you recommend?

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

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Note: The principle of least privilege and separation of duties are concepts that, although semantically different, are intrinsically related from the standpoint of security. The intent behind both is to prevent people from having higher privilege levels than they actually need
  Principle of Least Privilege: Users should only have the least amount of privileges required to perform their job and no more. This reduces authorization

exploitation by limiting access to resources such as targets, jobs, or monitoring templates for which they are not authorized.
 Separation of Duties: Beyond limiting user privilege level, you also limit user duties, or the specific jobs they can perform. No user should be given responsibility for more than one related function. This limits the ability of a user to perform a malicious action and then cover up that action.
References: https://cloud.google.com/kms/docs/separation-of-duties

QUESTION 3
A few days after JencoMart migrates the user credentials database to Google Cloud Platform and shuts down the old server, the new database server stops responding to SSH connections. It is still serving database requests to the application servers correctly.
What three steps should you take to diagnose the problem? Choose 3 answers.

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

Correct Answer: CDF
Section: (none)
Explanation

Explanation/Reference:
Explanation:
D: Handling "Unable to connect on port 22" error message Possible causes include:
 There is no firewall rule allowing SSH access on the port. SSH access on port 22 is enabled on all Compute Engine instances by default. If you have disabled access, SSH from the Browser will not work. If you run sshd on a port other than 22, you need to enable the access to that port with a custom firewall rule.
 The firewall rule allowing SSH access is enabled, but is not configured to allow connections from GCP Console services. Source IP addresses for browser-based SSH sessions are dynamically allocated by GCP Console and can vary from session to session.
F: Handling "Could not connect, retrying..." error
You can verify that the daemon is running by navigating to the serial console output page and looking for output lines prefixed with the accounts-from-metadata: string. If you are using a standard image but you do not see these output prefixes in the serial console output, the daemon might be stopped. Reboot the instance to restart the daemon.
References:
https://cloud.google.com/compute/docs/ssh-in-browser https://cloud.google.com/compute/docs/ssh-in-browser

QUESTION 4
JencoMart has decided to migrate user profile storage to Google Cloud Datastore and the application servers to Google Compute Engine (GCE). During the migration, the existing infrastructure will need access to Datastore to upload the data.
What service account key-management strategy should you recommend?

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

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
 Explanation:
Migrating data to Google Cloud Platform
Let's say that you have some data processing that happens on another cloud provider and you want to transfer the processed data to Google Cloud Platform. You can use a service account from the virtual machines on the external cloud to push the data to Google Cloud Platform. To do this, you must create and download a service account key when you create the service account and then use that key from the external process to call the Cloud Platform APIs.
References: https://cloud.google.com/iam/docs/understanding-service-accounts#migrating_data_to_google_cloud_platform

QUESTION 5
JencoMart has built a version of their application on Google Cloud Platform that serves traffic to Asia. You want to measure success against their business and technical goals.
Which metrics should you track?

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

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
 Explanation:
From scenario:
Business Requirements include: Expand services into Asia
Technical Requirements include: Decrease latency in Asia

QUESTION 6

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

What are three potential bottlenecks? Choose 3 answers.
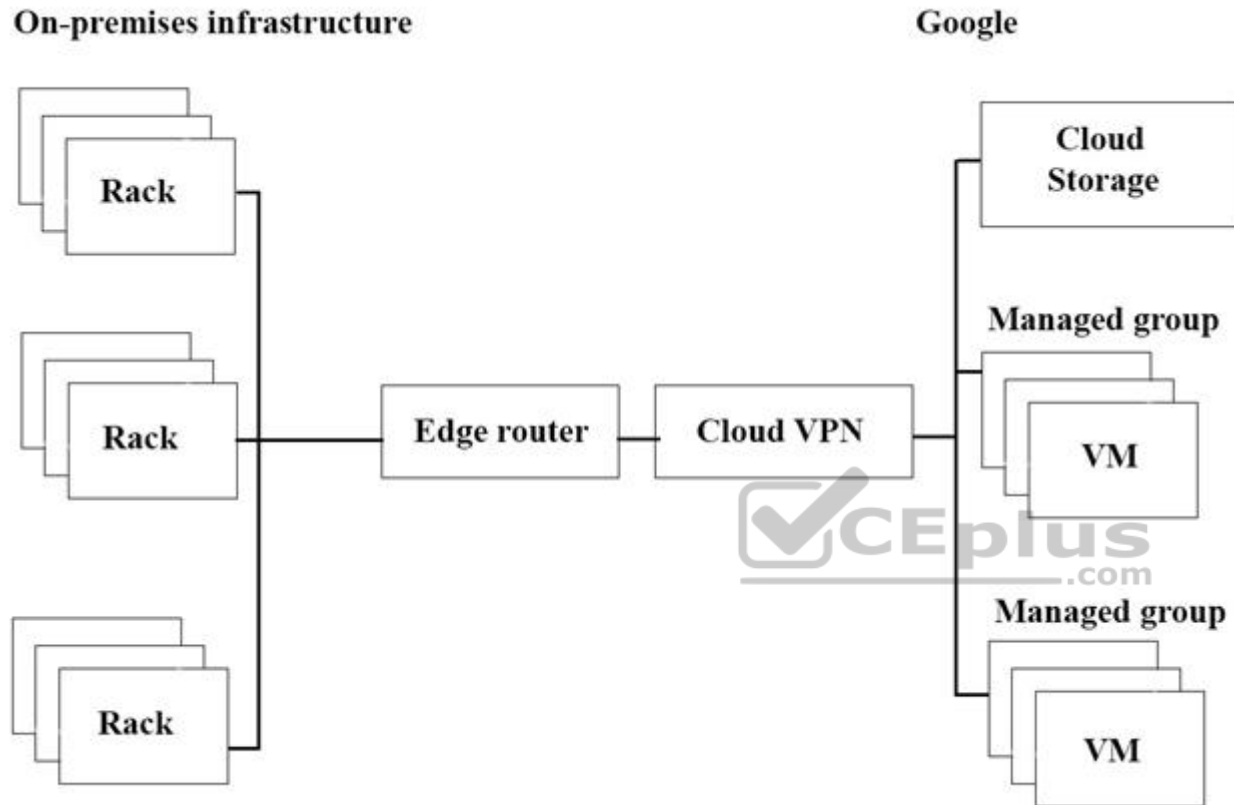
A. A single VPN tunnel, which limits throughput

B. A tier of Google Cloud Storage that is not suited for this task

C. A copy command that is not suited to operate over long distances

D. Fewer virtual machines (VMs) in GCP than on-premises machines

E. A separate storage layer outside the VMs, which is not suited for this task

F. Complicated internet connectivity between the on-premises infrastructure and GCP

Correct Answer: ACE

Section: (none)
Explanation

Explanation/Reference:

QUESTION 7
JencoMart wants to move their User Profiles database to Google Cloud Platform.
Which Google Database should they use?

A. Cloud Spanner

B. Google BigQuery

C. Google Cloud SQL

D. Google Cloud Datastore

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
 Explanation:
Common workloads for Google Cloud Datastore:
  User profiles
  Product catalogs
  Game state
References: https://cloud.google.com/storage-options/ https://cloud.google.com/datastore/docs/concepts/overview"ƒ

Mountkirk Games, A

QUESTION 1
Company Overview

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

Company Background

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

Solution Concept

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

Technical Requirements

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

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

CEO Statement

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

CTO Statement

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

CFO Statement

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

A. **Overview**

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 2
Mountkirk Games wants you to design their new testing strategy. How should the test coverage differ from their existing backends on the other platforms?

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

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
 Explanation:
From Scenario:
A few of their games were more popular than expected, and they had problems scaling their application servers, MySQL databases, and analytics tools.
Requirements for Game Analytics Platform include: Dynamically scale up or down based on game activity

QUESTION 3
Mountkirk Games has deployed their new backend on Google Cloud Platform (GCP). You want to create a through testing process for new versions of the backend before they are released to the public. You want the testing environment to scale in an economical way. How should you design the process?

A. Create a scalable environment in GCP for simulating production load

B. Use the existing infrastructure to test the GCP-based backend at scale

C. Build stress tests into each component of your application using resources internal to GCP to simulate load

D. Create a set of static environments in GCP to test different levels of load - for example, high, medium, and low

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
 Explanation:
From scenario: Requirements for Game Backend Platform
1. Dynamically scale up or down based on game activity
2. Connect to a managed NoSQL database service
3. Run customize Linux distro

QUESTION 4
Mountkirk Games wants to set up a continuous delivery pipeline. Their architecture includes many small services that they want to be able to update and roll back quickly. Mountkirk Games has the following requirements:
  Services are deployed redundantly across multiple regions in the US and Europe
  Only frontend services are exposed on the public internet
  They can provide a single frontend IP for their fleet of services
  Deployment artifacts are immutable
Which set of products should they use?

A. Google Cloud Storage, Google Cloud Dataflow, Google Compute Engine

B. Google Cloud Storage, Google App Engine, Google Network Load Balancer

C. Google Kubernetes Registry, Google Container Engine, Google HTTP(S) Load Balancer

D. Google Cloud Functions, Google Cloud Pub/Sub, Google Cloud Deployment Manager

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 5
Mountkirk Games' gaming servers are not automatically scaling properly. Last month, they rolled out a new feature, which suddenly became very popular. A

record number of users are trying to use the service, but many of them are getting 503 errors and very slow response times. What should they investigate first?

A. Verify that the database is online
B. Verify that the project quota hasn't been exceeded
C. Verify that the new feature code did not introduce any performance bugs
D. Verify that the load-testing team is not running their tool against production

Correct Answer: B
Section: (none)
Explanation

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

QUESTION 6
Mountkirk Games needs to create a repeatable and configurable mechanism for deploying isolated application environments. Developers and testers can access each other's environments and resources, but they cannot access staging or production resources. The staging environment needs access to some services from production.
What should you do to isolate development environments from staging and production?

A. Create a project for development and test and another for staging and production
B. Create a network for development and test and another for staging and production
C. Create one subnetwork for development and another for staging and production
D. Create one project for development, a second for staging and a third for production

Correct Answer: A
Section: (none)
Explanation

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

QUESTION 7
Mountkirk Games wants to set up a real-time analytics platform for their new game. The new platform must meet their technical requirements.
Which combination of Google technologies will meet all of their requirements?

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

C. Cloud SQL, Cloud Storage, Cloud Pub/Sub, and Cloud Dataflow

D. Cloud Dataproc, Cloud Pub/Sub, Cloud SQL, and Cloud Dataflow

E. Cloud Pub/Sub, Compute Engine, Cloud Storage, and Cloud Dataproc

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Ingest millions of streaming events per second from anywhere in the world with Cloud Pub/Sub, powered by Google's unique, high-speed private network.
Process the streams with Cloud Dataflow to ensure reliable, exactly-once, low-latency data transformation. Stream the transformed data into BigQuery, the cloud-native data warehousing service, for immediate analysis via SQL or popular visualization tools.
From scenario: They plan to deploy the game's backend on Google Compute Engine so they can capture streaming metrics, run intensive analytics.
Requirements for Game Analytics Platform
1. Dynamically scale up or down based on game activity
2. Process incoming data on the fly directly from the game servers
3. Process data that arrives late because of slow mobile networks
4. Allow SQL queries to access at least 10 TB of historical data
5. Process files that are regularly uploaded by users' mobile devices
6. Use only fully managed services
References: https://cloud.google.com/solutions/big-data/stream-analytics/"*ƒ*

Mountkirk Games, B

QUESTION 1
Company Overview

Mountkirk Games makes online, session-based, multiplayer games for mobile platforms. They build all of their games using some server-side integration.
Historically, they have used cloud providers to lease physical servers.
Due to the unexpected popularity of some of their games, they have had problems scaling their global audience, application servers, MySQL databases, and analytics tools.
Their current model is to write game statistics to files and send them through an ETL tool that loads them into a centralized MySQL database for reporting.

Solution Concept

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

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

Technical Requirements
Requirements for Game Backend Platform

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

Requirements for Game Analytics Platform

▪ Dynamically scale up or down based on game activity
▪ Process incoming data on the fly directly from the game servers
▪ Process data that arrives late because of slow mobile networks
▪ Allow queries to access at least 10 TB of historical data
▪ Process files that are regularly uploaded by users' mobile devices

Executive Statement

Our last successful game did not scale well with our previous cloud provider, resulting in lower user adoption and affecting the game's reputation. Our investors want more key performance indicators (KPIs) to evaluate the speed and stability of the game, as well as other metrics that provide deeper insight into usage patterns so we can adapt the game to target users. Additionally, our current technology stack cannot provide the scale we need, so we want to replace MySQL and move to an environment that provides autoscaling, low latency load balancing, and frees us up from managing physical servers.

A. **Overview**

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 2
For this question, refer to the Mountkirk Games case study. Mountkirk Games wants to migrate from their current analytics and statistics reporting model to one that meets their technical requirements on Google Cloud Platform.
Which two steps should be part of their migration plan? (Choose two.)

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

Correct Answer: AB
Section: (none)
Explanation

Explanation/Reference:

QUESTION 3
For this question, refer to the Mountkirk Games case study. You need to analyze and define the technical architecture for the compute workloads for your company, Mountkirk Games. Considering the Mountkirk Games business and technical requirements, what should you do?

A. Create network load balancers. Use preemptible Compute Engine instances.
B. Create network load balancers. Use non-preemptible Compute Engine instances.

C. Create a global load balancer with managed instance groups and autoscaling policies. Use preemptible Compute Engine instances.

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

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 4
For this question, refer to the Mountkirk Games case study. Mountkirk Games wants to design their solution for the future in order to take advantage of cloud and technology improvements as they become available. Which two steps should they take? (Choose two.)

A. Store as much analytics and game activity data as financially feasible today so it can be used to train machine learning models to predict user behavior in the future.

B. Begin packaging their game backend artifacts in container images and running them on Kubernetes Engine to improve the availability to scale up or down based on game activity.

C. Set up a CI/CD pipeline using Jenkins and Spinnaker to automate canary deployments and improve development velocity.

D. Adopt a schema versioning tool to reduce downtime when adding new game features that require storing additional player data in the database.

E. Implement a weekly rolling maintenance process for the Linux virtual machines so they can apply critical kernel patches and package updates and reduce the risk of 0-day vulnerabilities.

Correct Answer: CE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 5
For this question, refer to the Mountkirk Games case study. Mountkirk Games wants you to design a way to test the analytics platform's resilience to changes in mobile network latency. What should you do?

A. Deploy failure injection software to the game analytics platform that can inject additional latency to mobile client analytics traffic.

B. Build a test client that can be run from a mobile phone emulator on a Compute Engine virtual machine, and run multiple copies in Google Cloud Platform regions all over the world to generate realistic traffic.

C. Add the ability to introduce a random amount of delay before beginning to process analytics files uploaded from mobile devices.

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

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 6
For this question, refer to the Mountkirk Games case study. You need to analyze and define the technical architecture for the database workloads for your company, Mountkirk Games. Considering the business and technical requirements, what should you do?

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

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 7
For this question, refer to the Mountkirk Games case study. Which managed storage option meets Mountkirk's technical requirement for storing game activity in a time series database service?

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

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 8
For this question, refer to the Mountkirk Games case study. You need to analyze and define the technical architecture for the compute workloads for your company, Mounkirk Games. Considering the Mountkirk games business and technical requirements, what should you do?

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

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 9
For this question, refer to the Mountkirk Games case study. You are in charge of the new Game Backend Platform architecture. The game communicates with the backend over a REST API.
You want to follow Google-recommended practices. How should you design the backend?

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

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

Mix Questions

QUESTION 1
Your web application uses Google Kubernetes Engine to manage several workloads. One workload requires a consistent set of hostnames even after pod scaling and relaunches.
Which feature of Kubernetes should you use to accomplish this?

A. StatefulSets

B. Role-based access control

C. Container environment variables

D. Persistent Volumes

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 2
You are using Cloud CDN to deliver static HTTP(S) website content hosted on a Compute Engine instance group. You want to improve the cache hit ratio.
What should you do?

A. Customize the cache keys to omit the protocol from the key.

B. Shorten the expiration time of the cached objects.

C. Make sure the HTTP(S) header "Cache-Region" points to the closest region of your users.

D. Replicate the static content in a Cloud Storage bucket. Point CloudCDN toward a load balancer on that bucket.

Correct Answer: A
Section: (none)
Explanation

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

QUESTION 3
Your architecture calls for the centralized collection of all admin activity and VM system logs within your project.
How should you collect these logs from both VMs and services?

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

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Reference https://cloud.google.com/logging/docs/agent/

QUESTION 4
You have an App Engine application that needs to be updated. You want to test the update with production traffic before replacing the current application version.
What should you do?

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

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 5
All compute Engine instances in your VPC should be able to connect to an Active Directory server on specific ports. Any other traffic emerging from your instances is not allowed. You want to enforce this using VPC firewall rules.
How should you configure the firewall rules?

A. Create an egress rule with priority 1000 to deny all traffic for all instances. Create another egress rule with priority 100 to allow the Active Directory traffic for all instances.
B. Create an egress rule with priority 100 to deny all traffic for all instances. Create another egress rule with priority 1000 to allow the Active Directory traffic for all instances.

C. Create an egress rule with priority 1000 to allow the Active Directory traffic. Rely on the implied deny egress rule with priority 100 to block all traffic for all instances.
D. Create an egress rule with priority 100 to allow the Active Directory traffic. Rely on the implied deny egress rule with priority 1000 to block all traffic for all instances.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 6
Your customer runs a web service used by e-commerce sites to offer product recommendations to users. the company has begun experimenting with a machine learning model on Google Cloud Platform to improve the quality of results.
What should the customer do to improve their model's results over time?

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

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 7
A development team at your company has created a dockerized HTTPS web application. You need to deploy the application on Google Kubernetes Engine (GKE) and make sure that the application scales automatically.
How should you deploy to GKE?

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

D. Enable autoscaling on the Compute Engine instance group. Use a Service resource of type LoadBalancer to load-balance the HTTPS traffic.

Correct Answer: B
Section: (none)
Explanation

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

QUESTION 8
You need to design a solution for global load balancing based on the URL path being requested. You need to ensure operations reliability and end-to-end in-transit encryption based on Google best practices.
What should you do?

A. Create a cross-region load balancer with URL Maps.

B. Create an HTTPS load balancer with URL maps.

C. Create appropriate instance groups and instances. Configure SSL proxy load balancing.

D. Create a global forwarding rule. Configure SSL proxy balancing.

Correct Answer: B
Section: (none)
Explanation

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

QUESTION 9
You have an application that makes HTTP requests to Cloud Storage. Occasionally the requests fail with HTTP status codes of 5xx and 429.
How should you handle these types of errors?

A. Use gRPC instead of HTTP for better performance.

B. Implement retry logic using a truncated exponential backoff strategy.

C. Make sure the Cloud Storage bucket is multi-regional for geo-redundancy.

D. Monitor https://status.cloud.google.com/feed.atom and only make requests if Cloud Storage is not reporting an incident.

Correct Answer: B
Section: (none)
Explanation

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

QUESTION 10
You need to develop procedures to test a disaster plan for a mission-critical application. You want to use Google-recommended practices and native capabilities within GCP.
What should you do?

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

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 11
Your company creates rendering software which users can download from the company website. Your company has customers all over the world. You want to minimize latency for all your customers. You want to follow Googlerecommended practices.
How should you store the files?

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

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 12

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

A. Store the data in Google Drive and manually delete records as they expire.

B. Anonymize the data using the Cloud Data Loss Prevention API and store it indefinitely.

C. Store the data in Cloud Storage and use lifecycle management to delete files when they expire.

D. Store the data in Cloud Storage and run a nightly batch script that deletes all expired data.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 13
You are deploying a PHP App Engine Standard service with SQL as the backend. You want to minimize the number of queries to the database.
What should you do?

A. Set the memcache service level to dedicated. Create a key from the hash of the query, and return database values from memcache before issuing a query to Cloud SQL.

B. Set the memcache service level to dedicated. Create a cron task that runs every minute to populate the cache with keys containing query results.

C. Set the memcache service level to shared. Create a cron task that runs every minute to save all expected queries to a key called "cached-queries".

D. Set the memcache service level to shared. Create a key called "cached-queries", and return database values from the key before using a query to Cloud SQL.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 14
You need to ensure reliability for your application and operations by supporting reliable task a scheduling for compute on GCP. Leveraging Google best practices, what should you do?

A. Using the Cron service provided by App Engine, publishing messages directly to a message-processing utility service running on Compute Engine instances.

B. Using the Cron service provided by App Engine, publish messages to a Cloud Pub/Sub topic. Subscribe to that topic using a message-processing utility service running on Compute Engine instances.

C. Using the Cron service provided by Google Kubernetes Engine (GKE), publish messages directly to a message-processing utility service running on Compute Engine instances.

D. Using the Cron service provided by GKE, publish messages to a Cloud Pub/Sub topic. Subscribe to that topic using a message-processing utility service running on Compute Engine instances.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 15
Your company is building a new architecture to support its data-centric business focus. You are responsible for setting up the network. Your company's mobile and web-facing applications will be deployed on-premises, and all data analysis will be conducted in GCP. The plan is to process and load 7 years of archived .csv files totaling 900 TB of data and then continue loading 10 TB of data daily. You currently have an existing 100-MB internet connection.
What actions will meet your company's needs?

A. Compress and upload both archieved files and files uploaded daily using the qsutil -m option.

B. Lease a Transfer Appliance, upload archived files to it, and send it, and send it to Google to transfer archived data to Cloud Storage. Establish a connection with Google using a Dedicated Interconnect or Direct Peering connection anduse it to upload files daily.

C. Lease a Transfer Appliance, upload archived files to it, and send it, and send it to Google to transfer archived data to Cloud Storage. Establish one Cloud VPN Tunnel to VPC networks over the public internet, and compares and uploadfiles daily using the gsutil -m option.

D. Lease a Transfer Appliance, upload archived files to it, and send it to Google to transfer archived data to Cloud Storage. Establish a Cloud VPN Tunnel to VPC networks over the public internet, and compress and upload files daily.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference: