

Google

Exam Questions Professional-Cloud-Architect

Google Certified Professional - Cloud Architect (GCP)



NEW QUESTION 1

- (Topic 1)

For this question, refer to the Mountkirk Games case study.

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

Answer: C

NEW QUESTION 2

- (Topic 1)

For this question, refer to the Mountkirk Games case study

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.

Answer: D

NEW QUESTION 3

- (Topic 1)

For this question, refer to the Mountkirk Games case study.

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

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

Answer: A

Explanation:

From scenario: Requirements for Game Backend Platform

? Dynamically scale up or down based on game activity

? Connect to a managed NoSQL database service

? Run customize Linux distro

NEW QUESTION 4

- (Topic 2)

For this question, refer to the TerramEarth case study

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

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

Answer: A

Explanation:

<https://cloud.google.com/appengine/docs/flexible/go/authorizing-apps>

https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#delegate_application_authorization_with_oauth2

Delegate application authorization with OAuth2

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

References: https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#delegate_application_authorization_with_oauth2

<https://cloud.google.com/appengine/docs/flexible/go/authorizing-apps>

NEW QUESTION 5

- (Topic 2)

For this question, refer to the TerramEarth case study.

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

How should you design the data ingestion?

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

Answer: B

Explanation:

<https://cloud.google.com/solutions/data-lifecycle-cloud-platform> <https://cloud.google.com/solutions/designing-connected-vehicle-platform>

NEW QUESTION 6

- (Topic 2)

For this question, refer to the TerramEarth case study.

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

- A. Use one Google Container Engine cluster of FTP server
- B. Save the data to a Multi-Regional bucket
- C. Run the ETL process using data in the bucket.
- D. Use multiple Google Container Engine clusters running FTP servers located in different region
- E. Save the data to Multi-Regional buckets in us, eu, and asi
- F. Run the ETL process using the data in the bucket.
- G. Directly transfer the files to different Google Cloud Multi-Regional Storage bucket locations in us, eu, and asia using Google APIs over HTTP(S). Run the ETL process using the data in the bucket.
- H. Directly transfer the files to a different Google Cloud Regional Storage bucket location in us, eu, and asia using Google APIs over HTTP(S). Run the ETL process to retrieve the data from each Regional bucket.

Answer: D

Explanation:

<https://cloud.google.com/storage/docs/locations>

NEW QUESTION 7

- (Topic 2)

Your agricultural division is experimenting with fully autonomous vehicles.

You want your architecture to promote strong security during vehicle operation. Which two architecture should you consider?

Choose 2 answers:

- A. Treat every micro service call between modules on the vehicle as untrusted.
- B. Require IPv6 for connectivity to ensure a secure address space.
- C. Use a trusted platform module (TPM) and verify firmware and binaries on boot.
- D. Use a functional programming language to isolate code execution cycles.
- E. Use multiple connectivity subsystems for redundancy.
- F. Enclose the vehicle's drive electronics in a Faraday cage to isolate chips.

Answer: AC

NEW QUESTION 8

- (Topic 2)

For this question, refer to the TerramEarth case study.

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

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

Answer: A

Explanation:

https://cloud.google.com/endpoints/docs/openapi/about-cloud-endpoints?hl=en_US&_ga=2.21787131.-1712523161.1522785064

<https://cloud.google.com/endpoints/docs/openapi/architecture-overview> <https://cloud.google.com/storage/docs/gsutil/commands/test>

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

From scenario: Business Requirements

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

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

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

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

NEW QUESTION 9

- (Topic 2)

For this question, refer to the TerramEarth case study

You analyzed TerramEarth's business requirement to reduce downtime, and found that they can achieve a majority of time saving by reducing customers' wait time for parts. You decided to focus on reduction of the 3 weeks aggregate reporting time. Which modifications to the company's processes should you recommend?

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

Answer: C

Explanation:

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

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

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

NEW QUESTION 10

- (Topic 3)

For this question, refer to the JencoMart case study.

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

Answer: D

Explanation:

<https://cloud.google.com/datastore/docs/concepts/overview>

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>

NEW QUESTION 10

- (Topic 3)

For this question, refer to the JencoMart case study.

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

Answer: ADF

NEW QUESTION 11

- (Topic 3)

For this question, refer to the JencoMart case study.

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
- B. Each account should contain one project for every application.
- C. 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.
- D. Create a single G Suite account to manage users with each stage of each application in its own project.
- E. 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.

Answer: D

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>

NEW QUESTION 16

- (Topic 4)

For this question, refer to the Dress4Win case study.

As part of their new application experience, Dress4Win allows customers to upload images of themselves. The customer has exclusive control over who may view these images. Customers should be able to upload images with minimal latency and also be shown their images quickly on the main application page when they log in. Which configuration should Dress4Win use?

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

Answer: A

NEW QUESTION 19

- (Topic 4)

For this question, refer to the Dress4Win case study.

Dress4Win has end-to-end tests covering 100% of their endpoints. They want to ensure that the move to the cloud does not introduce any new bugs. Which additional testing methods should the developers employ to prevent an outage?

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

Answer: B

NEW QUESTION 21

- (Topic 4)

For this question, refer to the Dress4Win case study.

You want to ensure Dress4Win's sales and tax records remain available for infrequent viewing by auditors for at least 10 years. Cost optimization is your top priority. Which cloud services should you choose?

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

Answer: A

Explanation:

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

NEW QUESTION 26

- (Topic 4)

For this question, refer to the Dress4Win case study.

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

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

Answer: B

NEW QUESTION 28

- (Topic 5)

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

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

Answer: A

Explanation:

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

- ? <https://cloud.google.com/architecture/framework/system-design>
- ? <https://cloud.google.com/architecture/identity/best-practices-for-planning>
- ? <https://cloud.google.com/resource-manager/docs/creating-managing-folders>

NEW QUESTION 31

- (Topic 5)

Your company has an application running on Compute Engine that allows users to play their favorite music. There are a fixed number of instances. Files are stored in Cloud Storage and data is streamed directly to users. Users are reporting that they sometimes need to attempt to play popular songs multiple times before they are successful. You need to improve the performance of the application. What should you do?

- A.
 - * 1. Copy popular songs into CloudSQL as a blob
 - * 2. Update application code to retrieve data from CloudSQL when Cloud Storage is overloaded
- B.
 - * 1. Create a managed instance group with Compute Engine instances
 - * 2. Create a global load balancer and configure it with two backends
 - * Managed instance group
 - * Cloud Storage bucket
 - * 3. Enable Cloud CDN on the bucket backend
- C.
 - * 1. Mount the Cloud Storage bucket using gcsfuse on all backend Compute Engine instances
 - * 2. Serve music files directly from the backend Compute Engine instance
- D.
 - * 1. Create a Cloud Filestore NFS volume and attach it to the backend Compute Engine instances
 - * 2. Download popular songs in Cloud Filestore
 - * 3. Serve music files directly from the backend Compute Engine instance

A.

Answer: B

NEW QUESTION 36

- (Topic 5)

Your company has a networking team and a development team. The development team runs applications on Compute Engine instances that contain sensitive data. The development team requires administrative permissions for Compute Engine. Your company requires all network resources to be managed by the networking team. The development team does not want the networking team to have access to the sensitive data on the instances. What should you do?

- A. * 1. Create a project with a standalone VPC and assign the Network Admin role to the networking team. * 2. Create a second project with a standalone VPC and assign the Compute Admin role to the development team. * 3. Use Cloud VPN to join the two VPCs.
- B. * 1. Create a project with a standalone Virtual Private Cloud (VPC), assign the Network Admin role to the networking team, and assign the Compute Admin role to the development team.
- C. * 1. Create a project with a Shared VPC and assign the Network Admin role to the networking team. * 2. Create a second project without a VPC, configure it as a

Shared VPC service project, and assign the Compute Admin role to the development team.

D. * 1. Create a project with a standalone VPC and assign the Network Admin role to the networking team.* 2. Create a second project with a standalone VPC and assign the Compute Admin role to the development team.* 3. Use VPC Peering to join the two VPCs.

Answer: C

Explanation:

In this scenario, a large organization has a central team that manages security and networking controls for the entire organization. Developers do not have permissions to make changes to any network or security settings defined by the security and networking team but they are granted permission to create resources such as virtual machines in shared subnets. To facilitate this the organization makes use of a shared VPC (Virtual Private Cloud). A shared VPC allows creation of a VPC network of RFC 1918 IP spaces that associated projects (service projects) can then use. Developers using the associated projects can create VM instances in the shared VPC network spaces. The organization's network and security admins can create subnets, VPNs, and firewall rules usable by all the projects in the VPC network. https://cloud.google.com/iam/docs/job-functions/networking#single_team_manages_security_network_for_organization

Reference: <https://cloud.google.com/vpc/docs/shared-vpc>

NEW QUESTION 38

- (Topic 5)

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

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

Answer: B

NEW QUESTION 39

- (Topic 5)

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

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

Answer: B

NEW QUESTION 43

- (Topic 5)

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

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

Answer: D

Explanation:

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

NEW QUESTION 48

- (Topic 5)

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

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

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

D. Google Compute Engine, Jenkins, and Cloud Load Balancing

Answer: A

Explanation:

Helm for managing Kubernetes

Kubernetes can base on the URL to route traffic to different location (path)

<https://cloud.google.com/kubernetes-engine/docs/tutorials/http-balancer> eg.apiVersion: networking.k8s.io/v1beta1

kind: Ingress metadata:

name: fanout-ingress spec:

rules:

- http: paths:

- path: /* backend:

serviceName: web servicePort: 8080

- path: /v2/* backend: serviceName: web2 servicePort: 8080

NEW QUESTION 51

- (Topic 5)

Your company is developing a web-based application. You need to make sure that production deployments are linked to source code commits and are fully auditable. What should you do?

- A. Make sure a developer is tagging the code commit with the date and time of commit
- B. Make sure a developer is adding a comment to the commit that links to the deployment.
- C. Make the container tag match the source code commit hash.
- D. Make sure the developer is tagging the commits with :latest

Answer: C

Explanation:

From: <https://cloud.google.com/architecture/best-practices-for-building-containers>

Under: Tagging using the Git commit hash (bottom of page almost)

"In this case, a common way of handling version numbers is to use the Git commit SHA-1 hash (or a short version of it) as the version number. By design, the Git commit hash is immutable and references a specific version of your software.

You can use this commit hash as a version number for your software, but also as a tag for the Docker image built from this specific version of your software. Doing so makes Docker images traceable: because in this case the image tag is immutable, you instantly know which specific version of your software is running inside a given container."

NEW QUESTION 54

- (Topic 5)

You are helping the QA team to roll out a new load-testing tool to test the scalability of your primary cloud services that run on Google Compute Engine with Cloud Bigtable. Which three requirements should they include? Choose 3 answers

- A. Ensure that the load tests validate the performance of Cloud Bigtable.
- B. Create a separate Google Cloud project to use for the load-testing environment.
- C. Schedule the load-testing tool to regularly run against the production environment.
- D. Ensure all third-party systems your services use are capable of handling high load.
- E. Instrument the production services to record every transaction for replay by the load-testing tool.
- F. Instrument the load-testing tool and the target services with detailed logging and metrics collection.

Answer: ABF

NEW QUESTION 58

- (Topic 5)

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

- A. * 1. Define a migration plan based on the list of the applications and their dependencies.* 2. Migrate all virtual machines into Compute Engine individually with Migrate for Compute Engine.
- B. * 1. Perform an assessment of virtual machines running in the current VMware environment.* 2. Create images of all disk
- C. Import disks on Compute Engine.* 3. Create standard virtual machines where the boot disks are the ones you have imported.
- D. * 1. Perform an assessment of virtual machines running in the current VMware environment.* 2. Define a migration plan, prepare a Migrate for Compute Engine migration RunBook, and execute the migration.
- E. * 1. Perform an assessment of virtual machines running in the current VMware environment.* 2.Install a third-party agent on all selected virtual machine
- F. 3.Migrate all virtual machines into Compute Engine.

Answer: C

Explanation:

The framework illustrated in the preceding diagram has four phases:

•Assess. In this phase, you assess your source environment, assess the workloads that you want to migrate to Google Cloud, and assess which VMs support each workload.

•Plan. In this phase, you create the basic infrastructure for Migrate for Compute Engine, such as provisioning the resource hierarchy and setting up network access.

•Deploy. In this phase, you migrate the VMs from the source environment to Compute Engine.

•Optimize. In this phase, you begin to take advantage of the cloud technologies and capabilities.

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

NEW QUESTION 62

- (Topic 5)

You are developing your microservices application on Google Kubernetes Engine. During testing, you want to validate the behavior of your application in case a specific microservice should suddenly crash. What should you do?

- A. Add a taint to one of the nodes of the Kubernetes cluster
- B. For the specific microservice, configure a pod anti-affinity label that has the name of the tainted node as a value.
- C. Use Istio's fault injection on the particular microservice whose faulty behavior you want to simulate.
- D. Destroy one of the nodes of the Kubernetes cluster to observe the behavior.
- E. Configure Istio's traffic management features to steer the traffic away from a crashing microservice.

Answer: B

Explanation:

Microservice runs on all nodes. The Micro service runs on Pod, Pod runs on Nodes. Nodes is nothing but Virtual machines. Once deployed the application microservices will get deployed across all Nodes. Destroying one node may not mimic the behaviour of microservice crashing as it may be running in other nodes.
link: <https://istio.io/latest/docs/tasks/traffic-management/fault-injection/>

NEW QUESTION 65

- (Topic 5)

You are developing a globally scaled frontend for a legacy streaming backend data API.

This API expects

events in strict chronological order with no repeat data for proper processing.

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

- A. Cloud Pub/Sub alone
- B. Cloud Pub/Sub to Cloud DataFlow
- C. Cloud Pub/Sub to Stackdriver
- D. Cloud Pub/Sub to Cloud SQL

Answer: B

Explanation:

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

NEW QUESTION 69

- (Topic 5)

Your application needs to process credit card transactions. You want the smallest scope of

Payment Card Industry (PCI) compliance without compromising the ability to analyze transactional data and trends relating to which payment methods are used.

How should you design your architecture?

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

Answer: A

Explanation:

<https://cloud.google.com/solutions/pci-dss-compliance-in-gcp>

NEW QUESTION 72

- (Topic 5)

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

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

Answer: D

Explanation:

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

NEW QUESTION 74

- (Topic 5)

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

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

Answer: C

NEW QUESTION 79

- (Topic 5)

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.

Answer: D

Explanation:

<https://cloud.google.com/solutions/building-a-serverless-ml-model>

NEW QUESTION 84

- (Topic 5)

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

What should you do?

- A. Engage with a security company to run web scrapes that look your users' authentication data on malicious websites and notify you if any is found.
- B. Deploy intrusion detection software to your virtual machines to detect and log unauthorized access.
- C. Schedule a disaster simulation exercise during which you can shut off all VMs in a zone to see how your application behaves.
- D. Configure a read replica for your Cloud SQL instance in a different zone than the master, and then manually trigger a failover while monitoring KPIs for our REST API.

Answer: C

NEW QUESTION 87

- (Topic 5)

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

```
java.lang.SecurityException: SHA1 digest error for
com/Altostrat/CloakedServlet.class
    at com.google.appengine.runtime.Request.process
-d36f818a24b8cf1d (Request.java)
    at
sun.security.util.ManifestEntryVerifier.verify
(ManifestEntryVerifier.java:210)
    at java.util.jar.JarVerifier.processEntry
(JarVerifier.java:218)
    at java.util.jar.JarVerifier.update
(JarVerifier.java:205)
    at
java.util.jar.JarVerifiersVerifierStream.read
(JarVerifier.java:428)
    at sun.misc.Resource.getBytes
(Resource.java:124)
    at java.net.URL.ClassLoader.defineClass
(URLClassLoader.java:273)
    at sun.reflect.GeneratedMethodAccessor5.invoke
(Unknown Source)
    at
sun.reflect.DelegatingMethodAccessorImpl.invoke
(DelegatingMethodAccessorImpl.java:43)
    at java.lang.reflect.Method.invoke
(Method.java:616)
    at java.lang.ClassLoader.loadClass
(ClassLoader.java:266)
```

What should you do?

- A. Upload missing JAR files and redeploy your application.

- B. Digitally sign all of your JAR files and redeploy your application
- C. Recompile the CLoakedServlet class using and MD5 hash instead of SHA1

Answer: B

NEW QUESTION 91

- (Topic 5)

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

```
FROM ubuntu:16.04

COPY . /src

RUN apt-get update && apt-get install -y python python-pip

RUN pip install -r requirements.txt
```

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

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

Answer: CE

Explanation:

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

Note: Alpine Linux is built around musl libc and busybox. This makes it smaller and more resource efficient than traditional GNU/Linux distributions. A container requires no more

than 8 MB and a minimal installation to disk requires around 130 MB of storage. Not only do you get a fully-fledged Linux environment but a large selection of packages from the repository.

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

NEW QUESTION 93

- (Topic 5)

Your operations team has asked you to help diagnose a performance issue in a production application that runs on Compute Engine. The application is dropping requests that reach it when under heavy load. The process list for affected instances shows a single application process that is consuming all available CPU, and autoscaling has reached the upper limit of instances. There is no abnormal load on any other related systems, including the database. You want to allow production traffic to be served again as quickly as possible. Which action should you recommend?

- A. Change the autoscaling metric to `agent.googleapis.com/memory/percent_used`.
- B. Restart the affected instances on a staggered schedule.
- C. SSH to each instance and restart the application process.
- D. Increase the maximum number of instances in the autoscaling group.

Answer: D

Explanation:

Reference: <https://cloud.google.com/blog/products/sap-google-cloud/best-practices-for-sap-app-server-autoscaling-on-google-cloud>

NEW QUESTION 94

- (Topic 5)

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

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

Answer: D

Explanation:

Overview of storage classes, price, and use cases <https://cloud.google.com/storage/docs/storage-classes>

Why export logs? <https://cloud.google.com/logging/docs/export/>

StackDriver Quotas and Limits for Monitoring <https://cloud.google.com/monitoring/quotas> The BigQuery pricing. <https://cloud.google.com/bigquery/pricing>

NEW QUESTION 99

- (Topic 5)

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

- A. Configure a cron job to use the `gcloud` tool to take regular backups using persistent disk snapshots.
- B. Mount a Local SSD volume as the backup location.
- C. After the backup is complete, use `gsutil` to move the backup to Google Cloud Storage.

- D. Use gcsfuse to mount a Google Cloud Storage bucket as a volume directly on the instance and write backups to the mounted location using mysqldump
- E. Mount additional persistent disk volumes onto each virtual machine (VM) instance in a RAID10 array and use LVM to create snapshots to send to Cloud Storage.

Answer: B

Explanation:

<https://cloud.google.com/compute/docs/instances/sql-server/best-practices>

NEW QUESTION 102

- (Topic 5)

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

- A. Use the --no-auto-delete flag on all persistent disks and stop the VM.
- B. Use the -auto-delete flag on all persistent disks and terminate the VM.
- C. Apply VM CPU utilization label and include it in the BigQuery billing export.
- D. Use Google BigQuery billing export and labels to associate cost to groups.
- E. Store all state into local SSD, snapshot the persistent disks, and terminate the VM.
- F. Store all state in Google Cloud Storage, snapshot the persistent disks, and terminate the VM.

Answer: AD

Explanation:

<https://cloud.google.com/billing/docs/how-to/export-data-bigquery>

NEW QUESTION 103

- (Topic 5)

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

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

Answer: A

Explanation:

<https://cloud.google.com/storage/docs/encryption/customer-supplied-keys#gsutil>

NEW QUESTION 105

- (Topic 5)

You want to enable your running Google Container Engine cluster to scale as demand for your application changes. What should you do?

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

Answer: B

Explanation:

<https://cloud.google.com/kubernetes-engine/docs/concepts/cluster-autoscaler>

Cluster autoscaling

--enable-autoscaling

Enables autoscaling for a node pool.

Enables autoscaling in the node pool specified by --node-pool or the default node pool if --node-pool is not provided.

Where:

--max-nodes=MAX_NODES

Maximum number of nodes in the node pool.

Maximum number of nodes to which the node pool specified by --node-pool (or default node pool if unspecified) can scale.

NEW QUESTION 109

- (Topic 5)

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

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

- C. Use gcloud to create a Kubernetes cluster.
- D. Use kubectl to create the deployment.
- E. Use kubectl to create a Kubernetes cluster.
- F. Use Deployment Manager to create the deployment.
- G. Use kubectl to create a Kubernetes cluster.
- H. Use kubectl to create the deployment.

Answer: B

Explanation:

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

NEW QUESTION 110

- (Topic 5)

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

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

Answer: D

Explanation:

<https://cloud.google.com/security/compliance/pci-dss>

NEW QUESTION 114

- (Topic 5)

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

```

Java.lang.SecurityException : SHA1 digest

    At com.google.appengine.runtime.Request.pro
    At
    Sun.security.util.ManifestEntryVerifier.ver
    At java.net.URLClassLoader.defineClass
    At sun.reflect.GeneratedMethodAccessor
    At
    Sun.reflect.DelegatingMethodAccessorImpl.
    At java.lang.reflect.Method.invoke
  
```

- A. Recompile the CLoakedServlet class using and MD5 hash instead of SHA1
- B. Digitally sign all of your JAR files and redeploy your application.
- C. Upload missing JAR files and redeploy your application

Answer: B

NEW QUESTION 117

- (Topic 5)

You are deploying an application to Google Cloud. The application is part of a system. The application in Google Cloud must communicate over a private network with applications in a non-Google Cloud environment. The expected average throughput is 200 kbps. The business requires:

- 99.99% system availability
- cost optimization

You need to design the connectivity between the locations to meet the business requirements. What should you provision?

- A. A Classic Cloud VPN gateway connected with one tunnel to an on-premises VPN gateway.
- B. A Classic Cloud VPN gateway connected with two tunnels to an on-premises VPN gateway.
- C. An HA Cloud VPN gateway connected with two tunnels to an on-premises VPN gateway.
- D. Two HA Cloud VPN gateways connected to two on-premises VPN gateway
- E. Configure each HA CloudVPN gateway to have two tunnels, each connected to different on-premises VPN gateways.

Answer: C

Explanation:

https://cloud.google.com/network-connectivity/docs/vpn/concepts/topologies#configurations_that_support_9999_availability

NEW QUESTION 118

- (Topic 5)

You are running a cluster on Kubernetes Engine to serve a web application. Users are reporting that a specific part of the application is not responding anymore. You notice that all pods of your deployment keep restarting after 2 seconds. The application writes logs to standard output. You want to inspect the logs to find the cause of the issue. Which approach can you take?

- A. Review the Stackdriver logs for each Compute Engine instance that is serving as a node in the cluster.
- B. Review the Stackdriver logs for the specific Kubernetes Engine container that is serving the unresponsive part of the application.
- C. Connect to the cluster using gcloud credentials and connect to a container in one of the pods to read the logs.
- D. Review the Serial Port logs for each Compute Engine instance that is serving as a node in the cluster.

Answer: B

NEW QUESTION 123

- (Topic 5)

Your company uses Google Kubernetes Engine (GKE) as a platform for all workloads. Your company has a single large GKE cluster that contains batch, stateful, and stateless workloads. The GKE cluster is configured with a single node pool with 200 nodes. Your company needs to reduce the cost of this cluster but does not want to compromise availability. What should you do?

- A. Create a second GKE cluster for the batch workloads onl
- B. Allocate the 200 original nodes across both clusters.
- C. Configure a HorizontalPodAutoscaler for all stateless workloads and for all compatible stateful workload
- D. Configure the cluster to use node auto scaling.
- E. Configure CPU and memory limits on the namespaces in the cluste
- F. Configure all Pods to have a CPU and memory limits.
- G. Change the node pool to use spot VMs.

Answer: C

Explanation:

One way to reduce the cost of a Google Kubernetes Engine (GKE) cluster without compromising availability is to use horizontal pod autoscalers (HPA) and node auto scaling. HPA allows you to automatically scale the number of Pods in a deployment based on the resource usage of the Pods. By configuring HPA for stateless workloads and for compatible stateful workloads, you can ensure that the number of Pods is automatically adjusted based on the actual resource usage, which can help to reduce costs. Node auto scaling allows you to automatically add or remove nodes from the node pool based on the resource usage of the cluster. By configuring node auto scaling, you can ensure that the cluster has the minimum number of nodes needed to meet the resource requirements of the workloads, which can also help to reduce costs.

NEW QUESTION 127

- (Topic 5)

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

- A. Use gcloud sql instances restart.
- B. Validate that the Service Account used by the Cloud SQL proxy container still has the Cloud Build Editor role.
- C. In the GCP Console, navigate to Stackdriver Loggin
- D. Consult logs for Kubernetes Engine and Cloud SQL.
- E. In the GCP Console, navigate to Cloud SQ
- F. Restore the latest backu
- G. Use kubect1 to restart all pods.

Answer: C

NEW QUESTION 130

- (Topic 5)

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

- A. The effective policy is determined only by the policy set at the node
- B. The effective policy is the policy set at the node and restricted by the policies of its ancestors
- C. The effective policy is the union of the policy set at the node and policies inherited from its ancestors
- D. The effective policy is the intersection of the policy set at the node and policies inherited from its ancestors

Answer: B

Explanation:

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

NEW QUESTION 131

- (Topic 5)

Your company and one of its partners each have a Google Cloud project in separate organizations. Your company's project (prj-a) runs in Virtual Private Cloud (vpc-a). The partner's project (prj-b) runs in vpc-b. There are two instances running on vpc-a and one instance running on vpc-b. Subnets defined in both VPCs are not overlapping. You need to ensure that all instances communicate with each other via internal IPs minimizing latency and maximizing throughput. What should you do?

- A. Set up a network peering between vpc-a and vpc-b
- B. Set up a VPN between vpc-a and vpc-b using Cloud VPN
- C. Configure IAP TCP forwarding on the instance in vpc-b and then launch the following gcloud command from one of the instances in vpc-a: `gcloud compute instances tcp-forward --remote-addresses 0.0.0.0/0 --remote-ports 0-65535 --instance prj-a-vm --target-instance prj-b-vm`
- * 1. Create an additional instance in vpc-a* 2. Create an additional instance in vpc-b* 3. Install OpenVPN in newly created instances* 4. Configure a VPN tunnel between vpc-a and vpc-b with the help of OpenVPN

Answer: C

NEW QUESTION 135

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 139

- (Topic 5)

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

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

Answer: B

NEW QUESTION 140

- (Topic 5)

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

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

Answer: C

Explanation:

<https://cloud.google.com/compute/docs/instance-groups/autohealing-instances-in-migs> Health checks used for autohealing should be conservative so they don't preemptively delete and recreate your instances. When an autohealer health check is too aggressive, the autohealer might mistake busy instances for failed instances and unnecessarily restart them, reducing availability

NEW QUESTION 143

- (Topic 5)

An application development team has come to you for advice. They are planning to write and deploy an HTTP(S) API using Go 1.12. The API will have a very unpredictable workload and must remain reliable during peaks in traffic. They want to minimize operational overhead for this application. What approach should you recommend?

- A. Use a Managed Instance Group when deploying to Compute Engine
- B. Develop an application with containers, and deploy to Google Kubernetes Engine (GKE)
- C. Develop the application for App Engine standard environment
- D. Develop the application for App Engine Flexible environment using a custom runtime

Answer: C

Explanation:

<https://cloud.google.com/appengine/docs/the-appengine-environments>

NEW QUESTION 144

- (Topic 5)

You have created several preemptible Linux virtual machine instances using Google Compute Engine. You want to properly shut down your application before the virtual machines are preempted. What should you do?

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

Answer: C

NEW QUESTION 147

- (Topic 5)

Your company has a stateless web API that performs scientific calculations. The web API runs on a single Google Kubernetes Engine (GKE) cluster. The cluster is currently deployed in us-central1. Your company has expanded to offer your API to customers in Asia. You want to reduce the latency for the users in Asia. What should you do?

- A. Use a global HTTP(s) load balancer with Cloud CDN enabled
- B. Create a second GKE cluster in asia-southeast1, and expose both API's using a Service of type Load Balance
- C. Add the public IPs to the Cloud DNS zone
- D. Increase the memory and CPU allocated to the application in the cluster
- E. Create a second GKE cluster in asia-southeast1, and use kubemci to create a global HTTP(s) load balancer

Answer: D

Explanation:

https://cloud.google.com/kubernetes-engine/docs/concepts/multi-cluster-ingress#how_works

<https://github.com/GoogleCloudPlatform/k8s-multicloud-ingress> <https://cloud.google.com/blog/products/gcp/how-to-deploy-geographically-distributed-services-on-kubernetes-engine-with-kubemci>

NEW QUESTION 150

- (Topic 5)

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

What should you do?

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

Answer: D

Explanation:

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

NEW QUESTION 155

- (Topic 5)

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

Which primary compute resource should you choose?

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

Answer: A

Explanation:

<https://cloud.google.com/serverless-options>

NEW QUESTION 156

- (Topic 5)

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

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

Answer: D

Explanation:

<https://cloud.google.com/security/gdpr/?tab=tab4>

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

NEW QUESTION 159

- (Topic 5)

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

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

Answer: C

Explanation:

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

NEW QUESTION 161

- (Topic 5)

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

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

Answer: B

Explanation:

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

NEW QUESTION 165

- (Topic 5)

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

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

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

Answer: D

Explanation:

<https://cloud.google.com/compute/docs/disks>

NEW QUESTION 169

- (Topic 5)

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.

Answer: B

Explanation:

<https://cloud.google.com/appengine/docs/standard/python/splitting-traffic>

NEW QUESTION 174

- (Topic 5)

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

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

Answer: B

Explanation:

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

NEW QUESTION 179

- (Topic 5)

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

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

Answer: C

NEW QUESTION 184

- (Topic 5)

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

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

Answer: B

Explanation:

https://services.google.com/fh/files/misc/cloud_center_of_excellence.pdf

NEW QUESTION 189

- (Topic 5)

You write a Python script to connect to Google BigQuery from a Google Compute Engine virtual machine. The script is printing errors that it cannot connect to BigQuery. What should you do to fix the script?

- A. Install the latest BigQuery API client library for Python
- B. Run your script on a new virtual machine with the BigQuery access scope enabled
- C. Create a new service account with BigQuery access and execute your script with that user
- D. Install the bq component for gcloud with the command `gcloud components install bq`.

Answer: B

Explanation:

The error is most likely caused by the access scope issue. When creating a new instance, you have the default Compute Engine default service account but most services access including BigQuery is not enabled. Create an instance with Most access are not enabled by default. You have a default service account but don't have the

permission (scope) you can stop the instance, edit, change scope and restart it to enable the scope access. Of course, if you Run your script on a new virtual machine with the BigQuery access scope enabled, it also works
<https://cloud.google.com/compute/docs/access/service-accounts>

NEW QUESTION 194

- (Topic 5)

One of your primary business objectives is being able to trust the data stored in your application. You want to log all changes to the application data. How can you design your logging system to verify authenticity of your logs?

- A. Write the log concurrently in the cloud and on premises.
- B. Use a SQL database and limit who can modify the log table.
- C. Digitally sign each timestamp and log entry and store the signature.
- D. Create a JSON dump of each log entry and store it in Google Cloud Storage.

Answer: C

Explanation:

<https://cloud.google.com/storage/docs/access-logs>

References: <https://cloud.google.com/logging/docs/reference/tools/gcloud-logging>

NEW QUESTION 196

- (Topic 5)

You are designing a Data Warehouse on Google Cloud and want to store sensitive data in BigQuery. Your company requires you to generate encryption keys outside of Google Cloud. You need to implement a solution. What should you do?

- A. Generate a new key in Cloud Key Management Service (Cloud KMS). Store all data in Cloud Storage using the customer-managed key option and select the created key
- B. Set up a Dataflow pipeline to decrypt the data and to store it in a BigQuery dataset.
- C. Generate a new key in Cloud Key Management Service (Cloud KMS). Create a dataset in BigQuery using the customer-managed key option and select the created key
- D. Import a key in Cloud KM
- E. Store all data in Cloud Storage using the customer- managed key option and select the created key
- F. Set up a Dataflow pipeline to decrypt the data and to store it in a new BigQuery dataset.
- G. Import a key in Cloud KM
- H. Create a dataset in BigQuery using the customer-supplied key option and select the created key.

Answer: D

Explanation:

<https://cloud.google.com/bigquery/docs/customer-managed-encryption>

NEW QUESTION 198

- (Topic 5)

Your team needs to create a Google Kubernetes Engine (GKE) cluster to host a newly built application that requires access to third-party services on the internet. Your company does not allow any Compute Engine instance to have a public IP address on Google Cloud. You need to create a deployment strategy that adheres to these guidelines. What should you do?

- A. Create a Compute Engine instance, and install a NAT Proxy on the instance
- B. Configure all workloads on GKE to pass through this proxy to access third-party services on the Internet
- C. Configure the GKE cluster as a private cluster, and configure Cloud NAT Gateway for the cluster subnet
- D. Configure the GKE cluster as a route-based cluster
- E. Configure Private Google Access on the Virtual Private Cloud (VPC)
- F. Configure the GKE cluster as a private cluster
- G. Configure Private Google Access on the Virtual Private Cloud (VPC)

Answer: B

Explanation:

A Cloud NAT gateway can perform NAT for nodes and Pods in a private cluster, which is a type of VPC-native cluster. The Cloud NAT gateway must be configured to apply to at least the following subnet IP address ranges for the subnet that your cluster uses:

Subnet primary IP address range (used by nodes)

Subnet secondary IP address range used for Pods in the cluster Subnet secondary IP address range used for Services in the cluster

The simplest way to provide NAT for an entire private cluster is to configure a Cloud NAT gateway to apply to all of the cluster's subnet's IP address ranges.

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

NEW QUESTION 199

- (Topic 5)

A lead engineer wrote a custom tool that deploys virtual machines in the legacy data center. He wants to migrate the custom tool to the new cloud environment. You want to advocate for the adoption of Google Cloud Deployment Manager. What are two business risks of migrating to Cloud Deployment Manager? Choose 2 answers.

- A. Cloud Deployment Manager uses Python.
- B. Cloud Deployment Manager APIs could be deprecated in the future.
- C. Cloud Deployment Manager is unfamiliar to the company's engineers.
- D. Cloud Deployment Manager requires a Google APIs service account to run.
- E. Cloud Deployment Manager can be used to permanently delete cloud resources.
- F. Cloud Deployment Manager only supports automation of Google Cloud resources.

Answer: CF

Explanation:

<https://cloud.google.com/deployment-manager/docs/deployments/deleting-deployments>

NEW QUESTION 202

- (Topic 5)

Auditors visit your teams every 12 months and ask to review all the Google Cloud Identity and Access Management (Cloud IAM) policy changes in the previous 12 months. You want to streamline and expedite the analysis and audit process. What should you do?

- A. Create custom Google Stackdriver alerts and send them to the auditor.
- B. Enable Logging export to Google BigQuery and use ACLs and views to scope the data shared with the auditor.
- C. Use cloud functions to transfer log entries to Google Cloud SQL and use ACLS and views to limit an auditor's view.
- D. Enable Google Cloud Storage (GCS) log export to audit logs into a GCS bucket and delegate access to the bucket.

Answer: D

Explanation:

Export the logs to Google Cloud Storage bucket - Archive Storage, as it will not be used for 1 year, price for which is \$0.004 per GB per Month. The price for long term storage in BigQuery is \$0.01 per GB per Month (250% more). Also for analysis purpose, whenever Auditors are there (once per year), you can use BigQuery and use GCS bucket as external data source. BigQuery supports querying Cloud Storage data from these storage classes:
Standard Nearline Coldline Archive

NEW QUESTION 205

- (Topic 5)

Your company pushes batches of sensitive transaction data from its application server VMs to Cloud Pub/Sub for processing and storage. What is the Google-recommended way for your application to authenticate to the required Google Cloud services?

- A. Ensure that VM service accounts are granted the appropriate Cloud Pub/Sub IAM roles.
- B. Ensure that VM service accounts do not have access to Cloud Pub/Sub, and use VM access scopes to grant the appropriate Cloud Pub/Sub IAM roles.
- C. Generate an OAuth2 access token for accessing Cloud Pub/Sub, encrypt it, and store it in Cloud Storage for access from each VM.
- D. Create a gateway to Cloud Pub/Sub using a Cloud Function, and grant the Cloud Function service account the appropriate Cloud Pub/Sub IAM roles.

Answer: A

NEW QUESTION 210

- (Topic 5)

You are using a single Cloud SQL instance to serve your application from a specific zone.

You want to

introduce high availability. What should you do?

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

Answer: B

Explanation:

<https://cloud.google.com/sql/docs/mysql/high-availability>

NEW QUESTION 213

- (Topic 5)

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

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

Answer: B

Explanation:

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

NEW QUESTION 218

- (Topic 5)

The operations manager asks you for a list of recommended practices that she should consider when migrating a J2EE application to the cloud. Which three practices should you recommend? Choose 3 answers

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

Answer: AEF

Explanation:

References: <https://cloud.google.com/appengine/docs/standard/java/tools/uploadinganapp> <https://cloud.google.com/appengine/docs/standard/java/building-app/cloud-sql>

NEW QUESTION 220

- (Topic 5)

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

What should you do?

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

Answer: A

NEW QUESTION 225

- (Topic 5)

You are migrating third-party applications from optimized on-premises virtual machines to Google Cloud. You are unsure about the optimum CPU and memory options. The application has a consistent usage pattern across multiple weeks. You want to optimize resource usage for the lowest cost. What should you do?

- A. Create a Compute engine instance with CPU and Memory options similar to your application's current on-premises virtual machine
- B. Install the cloud monitoring agent, and deploy the third party application
- C. Run a load with normal traffic levels on third party application and follow the Rightsizing Recommendations in the Cloud Console
- D. Create an App Engine flexible environment, and deploy the third party application using a Docker file and a custom runtime
- E. Set CPU and memory options similar to your application's current on-premises virtual machine in the app.yaml file.
- F. Create an instance template with the smallest available machine type, and use an image of the third party application taken from the current on-premises virtual machine
- G. Create a managed instance group that uses average CPU to autoscale the number of instances in the group
- H. Modify the average CPU utilization threshold to optimize the number of instances running.
- I. Create multiple Compute Engine instances with varying CPU and memory options
- J. Install the cloud monitoring agent and deploy the third-party application on each of the instances
- K. Run a load test with high traffic levels on the application and use the results to determine the optimal settings.

Answer: A

Explanation:

Create a Compute engine instance with CPU and Memory options similar to your application's current on-premises virtual machine. Install the cloud monitoring agent, and deploy the third party application. Run a load with normal traffic levels on third party application and follow the Rightsizing Recommendations in the Cloud Console <https://cloud.google.com/migrate/compute-engine/docs/4.9/concepts/planning-a-migration/cloud-instance-rightsizing?hl=en>

NEW QUESTION 228

- (Topic 5)

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

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

Answer: D

Explanation:

Reference: <https://cloud.google.com/compute/docs/ip-addresses/reserve-static-external-ip-address>
<https://cloud.google.com/compute/docs/ip-addresses/reserve-static-external-ip-address#disableexternalip>
You might want to restrict external IP addresses so that only specific VM instances can use them. This option can help to prevent data exfiltration or maintain network isolation. Using an Organization Policy, you can restrict external IP addresses to specific VM instances with constraints to control use of external IP addresses for your VM instances within an organization or a project.

NEW QUESTION 229

- (Topic 5)

You are working at an institution that processes medical data. You are migrating several workloads onto Google Cloud. Company policies require all workloads to run on physically separated hardware, and workloads from different clients must also be separated. You created a sole-tenant node group and added a node for each client. You need to deploy the workloads on these dedicated hosts. What should you do?

- A. Add the node group name as a network tag when creating Compute Engine instances in order to host each workload on the correct node group.
- B. Add the node name as a network tag when creating Compute Engine instances in order to host each workload on the correct node.
- C. Use node affinity labels based on the node group name when creating Compute Engine instances in order to host each workload on the correct node group.
- D. Use node affinity labels based on the node name when creating Compute Engine instances in order to host each workload on the correct node.

Answer: C

Explanation:

https://cloud.google.com/compute/docs/nodes/provisioning-sole-tenant-vms#provision_a_sole-tenant_vm
https://cloud.google.com/compute/docs/nodes/provisioning-sole-tenant-vms#gcloud_2 When you create a VM, you request sole-tenancy by specifying node affinity or anti-affinity, referencing one or more node affinity labels. You specify custom node affinity labels when you create a node template, and Compute Engine automatically includes some default affinity labels on each node. By specifying affinity when you create a VM, you can schedule VMs together on a specific node or nodes in a node group. By specifying anti-affinity when you create a VM, you can ensure that certain VMs are not scheduled together on the same node or nodes in a node group.

NEW QUESTION 233

- (Topic 5)

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

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

Answer: B

Explanation:

<https://cloud.google.com/solutions/authenticating-corporate-users-in-a-hybrid-environment>

NEW QUESTION 238

- (Topic 5)

You want to establish a Compute Engine application in a single VPC across two regions. The application must communicate over VPN to an on-premises network. How should you deploy the VPN?

- A. Use VPC Network Peering between the VPC and the on-premises network.
- B. Expose the VPC to the on-premises network using IAM and VPC Sharing.
- C. Create a global Cloud VPN Gateway with VPN tunnels from each region to the on-premises peer gateway.
- D. Deploy Cloud VPN Gateway in each region.
- E. Ensure that each region has at least one VPN tunnel to the on-premises peer gateway.

Answer: C

Explanation:

<https://cloud.google.com/vpn/docs/how-to/creating-static-vpns>

NEW QUESTION 241

- (Topic 5)

Your company wants to try out the cloud with low risk. They want to archive approximately 100 TB of their log data to the cloud and test the analytics features available to them there, while also retaining that data as a long-term disaster recovery backup. Which two steps should they take? Choose 2 answers

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

Answer: AE

NEW QUESTION 244

- (Topic 5)

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

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

Answer: A

Explanation:

<https://cloud.google.com/monitoring/api/v3/metrics-details#metric-kinds>

NEW QUESTION 249

- (Topic 5)

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

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

Answer: A

Explanation:

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


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

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

NEW QUESTION 252

- (Topic 5)

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



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

- A. Create a cloud router to advertise subnet #2 and subnet #3 to subnet #1.
- B. Add two additional NICs to Instance #1 with the following configuration:
 •NIC1VPC: VPC #2SUBNETWORK: subnet #2•NIC2VPC: VPC #3SUBNETWORK: subnet #3Update firewall rules to enable traffic between instances.
- C. Create two VPN tunnels via CloudVPN:•1 between VPC #1 and VPC #2.•1 between VPC #2 and VPC #3.Update firewall rules to enable traffic between the instances.
- D. Peer all three VPCs:•Peer VPC #1 with VPC #2.•Peer VPC #2 with VPC #3.Update firewall rules to enable traffic between the instances.

Answer: B

Explanation:

As per GCP documentation: "By default, every instance in a VPC network has a single network interface. Use these instructions to create additional network interfaces. Each interface is attached to a different VPC network, giving that instance access to different VPC networks in Google Cloud. You cannot attach multiple network interfaces to the same VPC network." Refer to: <https://cloud.google.com/vpc/docs/create-use-multiple-interfaces>
https://cloud.google.com/vpc/docs/create-use-multiple-interfaces#i_am_not_able_to_connect_to_secondary_interfaces_internal_ip

NEW QUESTION 256

- (Topic 5)

A production database virtual machine on Google Compute Engine has an ext4-formatted persistent disk for data files The database is about to run out of storage space How can you remediate the problem with the least amount of downtime?

- A. In the Cloud Platform Console, increase the size of the persistent disk and use the resize2fs command in Linux.
- B. Shut down the virtual machine, use the Cloud Platform Console to increase the persistent disk size, then restart the virtual machine.
- C. In the Cloud Platform Console, increase the size of the persistent disk and verify the new space is ready to use with the fdisk command in Linux.
- D. In the Cloud Platform Console, create a new persistent disk attached to the virtual machine, format and mount it, and configure the database service to move the files to the new disk.
- E. In the Cloud Platform Console, create a snapshot of the persistent disk, restore the snapshot to a new larger disk, unmount the old disk, mount the new disk, and restart the database service.

Answer: A

Explanation:

On Linux instances, connect to your instance and manually resize your partitions and file systems to use the additional disk space that you added. Extend the file system on the disk or the partition to use the added space. If you grew a partition on your disk, specify the partition. If your disk does not have a partition table, specify only the disk ID.

```
sudo resize2fs /dev/[DISK_ID][PARTITION_NUMBER]
```

where [DISK_ID] is the device name and [PARTITION_NUMBER] is the partition number for the device where you are resizing the file system.

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

NEW QUESTION 260

- (Topic 5)

Your company provides a recommendation engine for retail customers. You are providing retail customers with an API where they can submit a user ID and the API returns a list of recommendations for that user. You are responsible for the API lifecycle and want to ensure stability for your customers in case the API makes backward-incompatible changes. You want to follow Google-recommended practices. What should you do?

- A. Create a distribution list of all customers to inform them of an upcoming backward- incompatible change at least one month before replacing the old API with the new API.
- B. Create an automated process to generate API documentation, and update the public API documentation as part of the CI/CD process when deploying an update to the API.
- C. Use a versioning strategy for the APIs that increases the version number on every backward-incompatible change.
- D. Use a versioning strategy for the APIs that adds the suffix "DEPRECATED" to the current API version number on every backward-incompatible change.
- E. Use the current version number for the new API.

Answer: C

Explanation:

<https://cloud.google.com/apis/design/versioning>

All Google API interfaces must provide a major version number, which is encoded at the end of the protobuf package, and included as the first part of the URI path for REST APIs. If an API introduces a breaking change, such as removing or renaming a field, it must increment its API version number to ensure that existing user code does not suddenly break.

NEW QUESTION 261

- (Topic 5)

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

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

Answer: A

Explanation:

<https://cloud.google.com/transfer-appliance/docs/2.0/faq>

NEW QUESTION 262

- (Topic 5)

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.

Answer: A

Explanation:

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

NEW QUESTION 267

- (Topic 5)

Your company's user-feedback portal comprises a standard LAMP stack replicated across two zones. It is deployed in the us-central1 region and uses autoscaled managed instance groups on all layers, except the database. Currently, only a small group of select customers have access to the portal. The portal meets a 99.99% availability SLA under these conditions. However, next quarter, your company will be making the portal available to all users, including unauthenticated users. You need to develop a resiliency testing strategy to ensure the system maintains the SLA once they introduce additional user load. What should you do?

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

Answer: A

NEW QUESTION 272

- (Topic 5)

You have deployed several instances on Compute Engine. As a security requirement, instances cannot have a public IP address. There is no VPN connection between Google

Cloud and your office, and you need to connect via SSH into a specific machine without violating the security requirements. What should you do?

- A. Configure Cloud NAT on the subnet where the instance is hosted
- B. Create an SSH connection to the Cloud NAT IP address to reach the instance.
- C. Add all instances to an unmanaged instance group
- D. Configure TCP Proxy Load Balancing with the instance group as a backend
- E. Connect to the instance using the TCP Proxy IP.
- F. Configure Identity-Aware Proxy (IAP) for the instance and ensure that you have the role of IAP-secured Tunnel User
- G. Use the gcloud command line tool to ssh into the instance.
- H. Create a bastion host in the network to SSH into the bastion host from your office location
- I. From the bastion host, SSH into the desired instance.

Answer: C

Explanation:

https://cloud.google.com/iap/docs/using-tcp-forwarding#tunneling_with_ssh

Leveraging the BeyondCorp security model. "This January, we enhanced context-aware access capabilities in Cloud Identity-Aware Proxy (IAP) to help you protect SSH and RDP access to your virtual machines (VMs)—without needing to provide your VMs with public IP addresses, and without having to set up bastion hosts. "

<https://cloud.google.com/blog/products/identity-security/cloud-iap-enables-context-aware-access-to-vm-via-ssh-and-rdp-without-bastion-hosts>

Reference: <https://cloud.google.com/solutions/connecting-securely>

NEW QUESTION 274

- (Topic 5)

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

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

Answer: C

NEW QUESTION 276

- (Topic 5)

Your company is using BigQuery as its enterprise data warehouse. Data is distributed over several Google Cloud projects. All queries on BigQuery need to be billed on a single

project. You want to make sure that no query costs are incurred on the projects that contain the data. Users should be able to query the datasets, but not edit them.

How should you configure users' access roles?

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

Answer: A

Explanation:

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

NEW QUESTION 278

- (Topic 5)

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

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

Answer: B

NEW QUESTION 279

- (Topic 5)

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

- A. Use kubectl set image deployment/echo-deployment <new-image>

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

Answer: A

Explanation:

https://cloud.google.com/kubernetes-engine/docs/how-to/updating-apps#updating_an_application

NEW QUESTION 283

- (Topic 6)

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

What should you do?

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

Answer: A

Explanation:

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

NEW QUESTION 285

- (Topic 7)

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

Which two actions should you take?

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

Answer: A

NEW QUESTION 287

- (Topic 7)

For this question, refer to the TerramEarth case study.

You start to build a new application that uses a few Cloud Functions for the backend. One use case requires a Cloud Function `func_display` to invoke another Cloud Function `func_query`. You want `func_query` only to accept invocations from `func_display`. You also want to follow Google's recommended best practices. What should you do?

- A. Create a token and pass it in as an environment variable to `func_display`
- B. When invoking `func_query`, include the token in the request. Pass the same token to `func_query` and reject the invocation if the tokens are different.
- C. Make `func_query` 'Require authentication.' Create a unique service account and associate it to `func_display`
- D. Grant the service account invoker role for `func_query`
- E. Create an id token in `func_display` and include the token to the request when invoking `func_query`.
- F. Make `func_query` 'Require authentication' and only accept internal traffic
- G. Create those two functions in the same VPC
- H. Create an ingress firewall rule for `func_query` to only allow traffic from `func_display`.
- I. Create those two functions in the same project and VPC
- J. Make `func_query` only accept internal traffic
- K. Create an ingress firewall for `func_query` to only allow traffic from `func_display`
- L. Also, make sure both functions use the same service account.

Answer: B

Explanation:

https://cloud.google.com/functions/docs/securing/authenticating#authenticating_function_to_function_calls

NEW QUESTION 291

- (Topic 7)

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

- A. Use BigQuery as the data warehouse
- B. Connect all vehicles to the network and stream data into BigQuery using Cloud Pub/Sub and Cloud Dataflow
- C. Use Google Data Studio for analysis and reporting.
- D. Use BigQuery as the data warehouse
- E. Connect all vehicles to the network and upload gzip files to a Multi-Regional Cloud Storage bucket using `gcloud`
- F. Use Google Data Studio for analysis and reporting.

- G. Use Cloud Dataproc Hive as the data warehouse
- H. Upload gzip files to a MultiRegional Cloud Storage bucket
- I. Upload this data into BigQuery using gcloud
- J. Use Google data Studio for analysis and reporting.
- K. Use Cloud Dataproc Hive as the data warehouse
- L. Directly stream data into partitioned Hive table
- M. Use Pig scripts to analyze data.

Answer: A

NEW QUESTION 292

- (Topic 7)

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

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

Answer: C

Explanation:

<https://cloud.google.com/bigquery/docs/managing-partitioned-tables#partition-expiration>
<https://cloud.google.com/storage/docs/lifecycle>

NEW QUESTION 295

- (Topic 7)

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

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

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

Answer: A

Explanation:

<https://cloud.google.com/solutions/iot-overview> <https://cloud.google.com/iot/quotas>

NEW QUESTION 296

- (Topic 7)

For this question, refer to the TerramEarth case study. You need to implement a reliable, scalable GCP solution for the data warehouse for your company, TerramEarth. Considering the TerramEarth business and technical requirements, what should you do?

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

Answer: C

Explanation:

https://cloud.google.com/solutions/bigquery-data-warehouse#external_sources <https://cloud.google.com/solutions/bigquery-data-warehouse>

NEW QUESTION 301

- (Topic 8)

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

Answer: A

Explanation:

<https://cloud.google.com/blog/products/databases/getting-started-with-time-series-trend-predictions-using-gcp>

NEW QUESTION 304

- (Topic 8)

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

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

Answer: A

NEW QUESTION 308

- (Topic 8)

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

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

Answer: A

NEW QUESTION 310

- (Topic 8)

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

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

Answer: B

Explanation:

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

NEW QUESTION 314

- (Topic 8)

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.

Answer: D

Explanation:

<https://cloud.google.com/bigtable/docs/schema-design-time-series>

NEW QUESTION 315

- (Topic 8)

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

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

Answer: A

NEW QUESTION 316

- (Topic 8)

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

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

Answer: C

NEW QUESTION 317

- (Topic 9)

For this question, refer to the Helicopter Racing League (HRL) case study. The HRL development team releases a new version of their predictive capability application every Tuesday evening at 3

- A. a.
- B. UTC to a repository
- C. The security team at HRL has developed an in-house penetration test Cloud Function called Airwolf. The security team wants to run Airwolf against the predictive capability application as soon as it is released every Tuesday
- D. You need to set up Airwolf to run at the recurring weekly cadence
- E. What should you do?
- F. Set up Cloud Tasks and a Cloud Storage bucket that triggers a Cloud Function.
- G. Set up a Cloud Logging sink and a Cloud Storage bucket that triggers a Cloud Function.
- H. Configure the deployment job to notify a Pub/Sub queue that triggers a Cloud Function.
- I. Set up Identity and Access Management (IAM) and Confidential Computing to trigger a Cloud Function.

Answer: A

NEW QUESTION 321

- (Topic 9)

For this question, refer to the Helicopter Racing League (HRL) case study. HRL wants better prediction accuracy from their ML prediction models. They want you to use Google's AI Platform so HRL can understand and interpret the predictions. What should you do?

- A. Use Explainable AI.
- B. Use Vision AI.
- C. Use Google Cloud's operations suite.
- D. Use Jupyter Notebooks.

Answer: A

Explanation:

Reference: <https://cloud.google.com/ai-platform/prediction/docs/ai-explanations/preparing-metadata>

NEW QUESTION 324

- (Topic 9)

For this question, refer to the Helicopter Racing League (HRL) case study. Your team is in charge of creating a payment card data vault for card numbers used to bill tens of thousands of viewers, merchandise consumers, and season ticket holders. You need to implement a custom card tokenization service that meets the following requirements:

- It must provide low latency at minimal cost.
- It must be able to identify duplicate credit cards and must not store plaintext card numbers.
- It should support annual key rotation.

Which storage approach should you adopt for your tokenization service?

- A. Store the card data in Secret Manager after running a query to identify duplicates.
- B. Encrypt the card data with a deterministic algorithm stored in Firestore using Datastore mode.
- C. Encrypt the card data with a deterministic algorithm and shard it across multiple Memorystore instances.
- D. Use column-level encryption to store the data in Cloud SQL.

Answer: B

NEW QUESTION 326

- (Topic 9)

For this question, refer to the Helicopter Racing League (HRL) case study. A recent finance audit of cloud infrastructure noted an exceptionally high number of Compute Engine instances are allocated to do video encoding and transcoding. You suspect that these Virtual Machines are zombie machines that were not deleted after their workloads completed. You need to quickly get a list of which VM instances are idle. What should you do?

- A. Log into each Compute Engine instance and collect disk, CPU, memory, and network usage statistics for analysis.
- B. Use the `gcloud compute instances list` to list the virtual machine instances that have the `idle: true` label set.
- C. Use the `gcloud recomender` command to list the idle virtual machine instances.
- D. From the Google Console, identify which Compute Engine instances in the managed instance groups are no longer responding to health check probes.

Answer: C

Explanation:

Reference: <https://cloud.google.com/compute/docs/instances/viewing-and-applying-idle-vm-recommendations>

NEW QUESTION 329

- (Topic 10)

For this question, refer to the EHR Healthcare case study. You are a developer on the EHR customer portal team. Your team recently migrated the customer portal application to Google Cloud. The load has increased on the application servers, and now the application is logging many timeout errors. You recently incorporated Pub/Sub into the application architecture, and the application is not logging any Pub/Sub publishing errors. You want to improve publishing latency. What should you do?

- A. Increase the Pub/Sub Total Timeout retry value.
- B. Move from a Pub/Sub subscriber pull model to a push model.
- C. Turn off Pub/Sub message batching.
- D. Create a backup Pub/Sub message queue.

Answer: C

Explanation:

<https://cloud.google.com/pubsub/docs/publisher?hl=en#batching>

NEW QUESTION 333

- (Topic 10)

For this question, refer to the EHR Healthcare case study. EHR has single Dedicated Interconnect connection between their primary data center and Googles network. This connection satisfies EHR's network and security policies:

- On-premises servers without public IP addresses need to connect to cloud resources without public IP addresses
- Traffic flows from production network mgmt. servers to Compute Engine virtual machines should never traverse the public internet.

You need to upgrade the EHR connection to comply with their requirements. The new connection design must support business critical needs and meet the same network and security policy requirements. What should you do?

- A. Add a new Dedicated Interconnect connection
- B. Upgrade the bandwidth on the Dedicated Interconnect connection to 100 G
- C. Add three new Cloud VPN connections
- D. Add a new Carrier Peering connection

Answer: A

Explanation:

The case does not call out the throughput being an issue. However, to achieve 99.99%, you need to have 4 connections as per Google recommendations. However, in the options only A has the option to add an additional Interconnect connection. <https://cloud.google.com/network-connectivity/docs/interconnect/concepts/dedicated-overview#availability>

NEW QUESTION 336

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

Professional-Cloud-Architect Practice Exam Features:

- * Professional-Cloud-Architect Questions and Answers Updated Frequently
- * Professional-Cloud-Architect Practice Questions Verified by Expert Senior Certified Staff
- * Professional-Cloud-Architect Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * Professional-Cloud-Architect Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The Professional-Cloud-Architect Practice Test Here](#)