

# 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 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 2

- (Topic 1)

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

Mountkirk Games' gaming servers are not automatically scaling properly. Last month, they rolled out a new feature, which suddenly became very popular. A record number of users are trying to use the service, but many of them are getting 503 errors and very slow response times. What should they investigate first?

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

**Answer:** B

#### Explanation:

503 is service unavailable error. If the database was online everyone would get the 503 error.

[https://cloud.google.com/docs/quota#capping\\_usage](https://cloud.google.com/docs/quota#capping_usage)

### NEW QUESTION 3

- (Topic 1)

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

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

- A. Container Engine, Cloud Pub/Sub, and Cloud SQL
- B. Cloud Dataflow, Cloud Storage, Cloud Pub/Sub, and BigQuery
- C. Cloud SQL, Cloud Storage, Cloud Pub/Sub, and Cloud Dataflow
- D. Cloud Dataproc, Cloud Pub/Sub, Cloud SQL, and Cloud Dataflow
- E. Cloud Pub/Sub, Compute Engine, Cloud Storage, and Cloud Dataproc

**Answer:** B

#### Explanation:

A real time requires Stream / Messaging so Pub/Sub, Analytics by Big Query.

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

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

Requirements for Game Analytics Platform

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

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

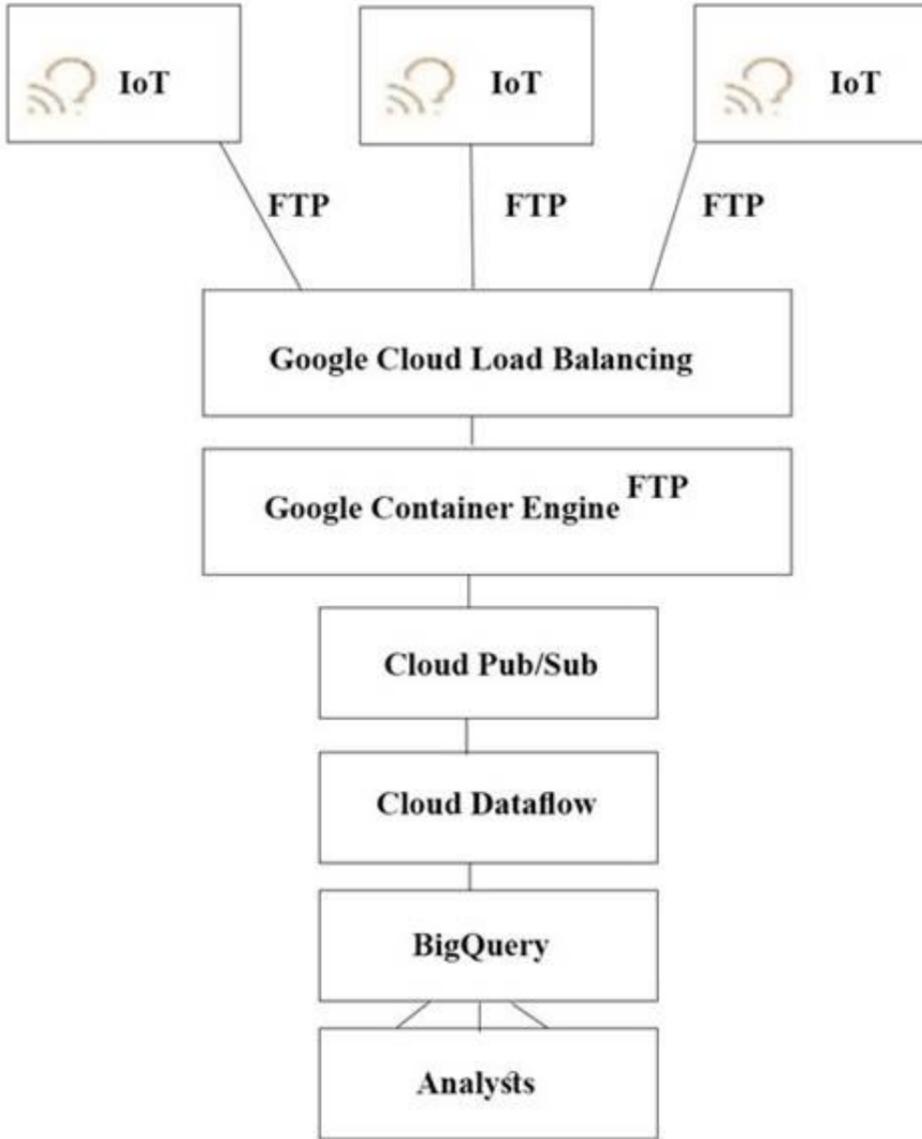
### NEW QUESTION 4

- (Topic 2)

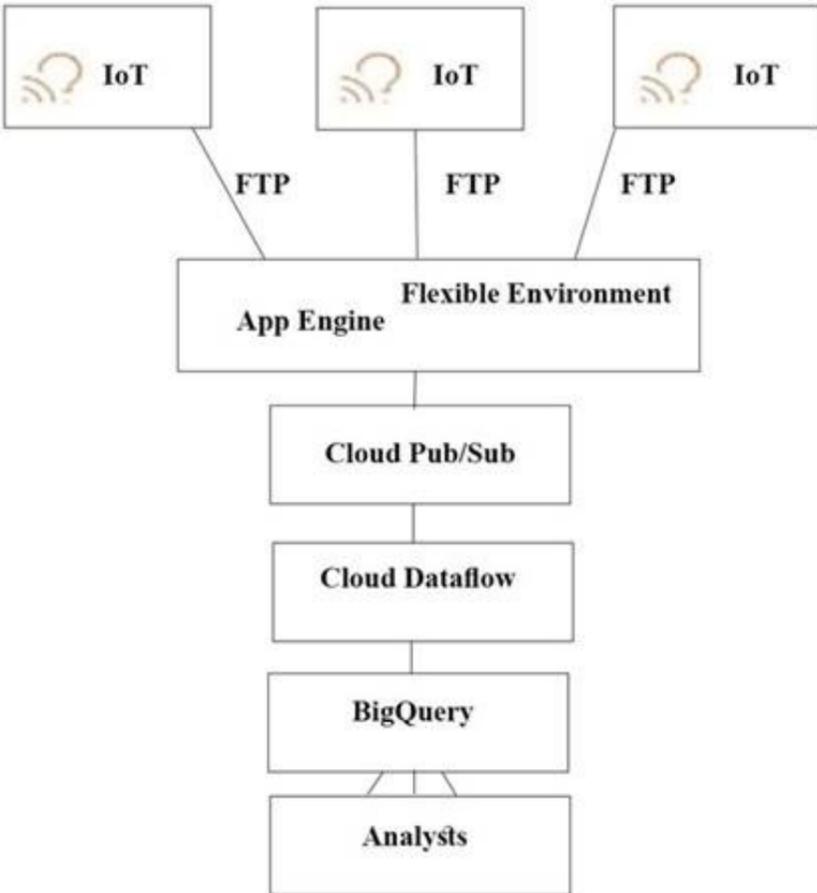
For this question, refer to the TerramEarth case study.

TerramEarth's CTO wants to use the raw data from connected vehicles to help identify approximately when a vehicle in the development team to focus their failure. You want to allow analysts to centrally query the vehicle data. Which architecture should you recommend?

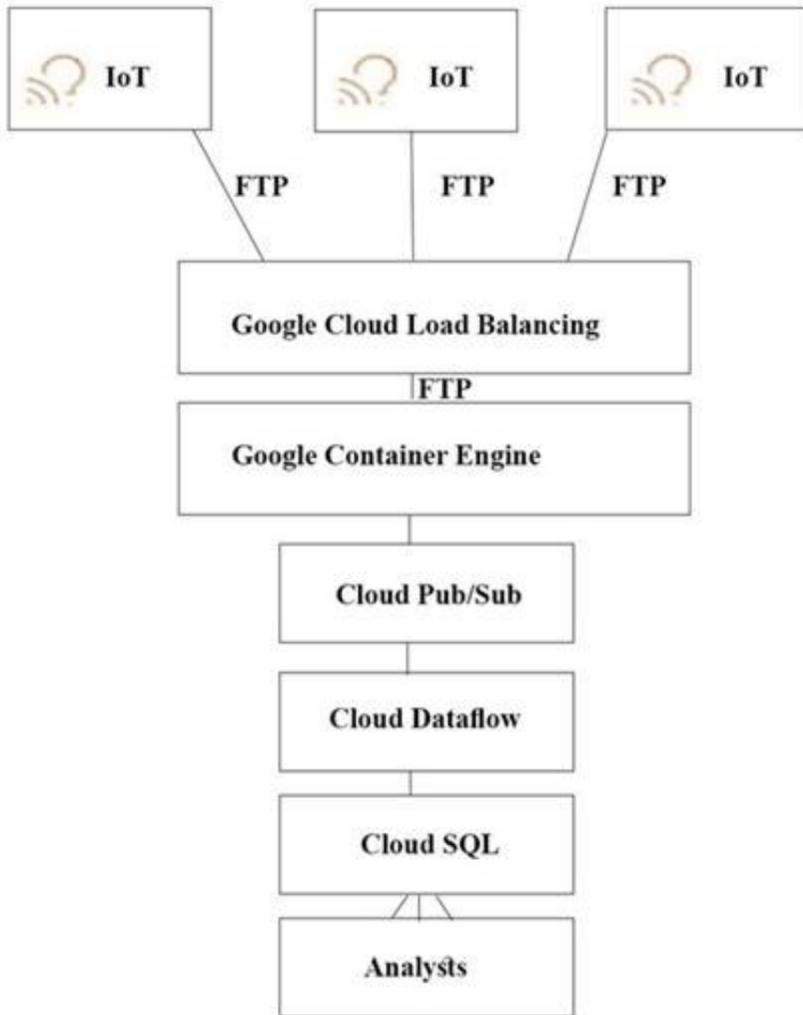
A)



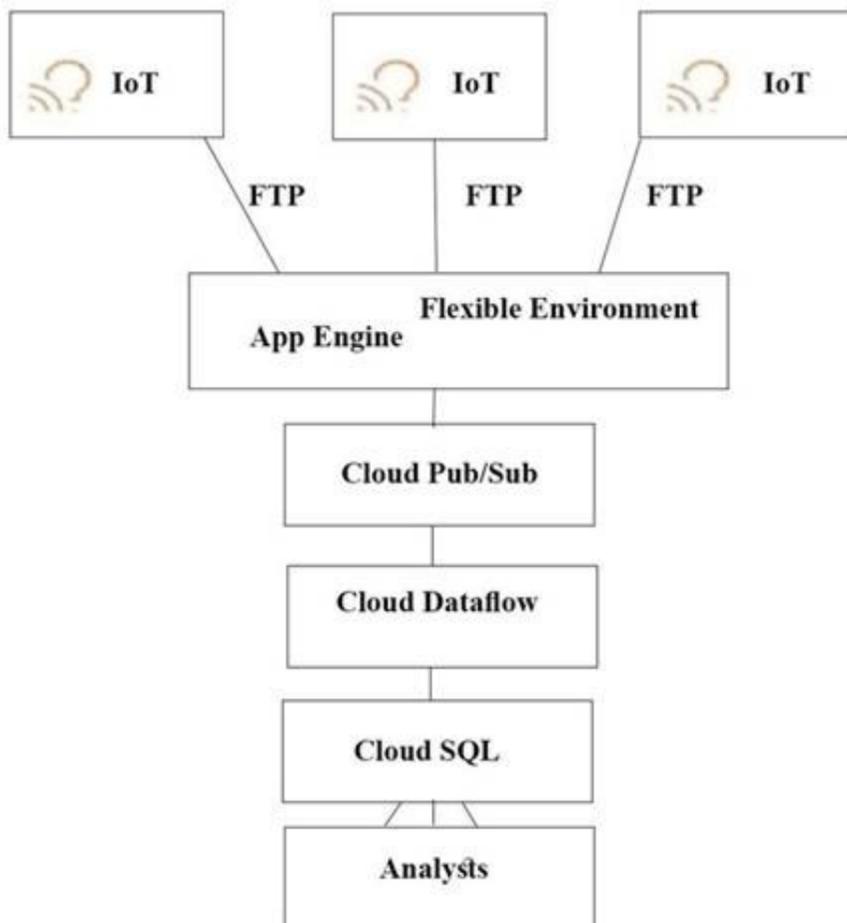
B)



C)



D)

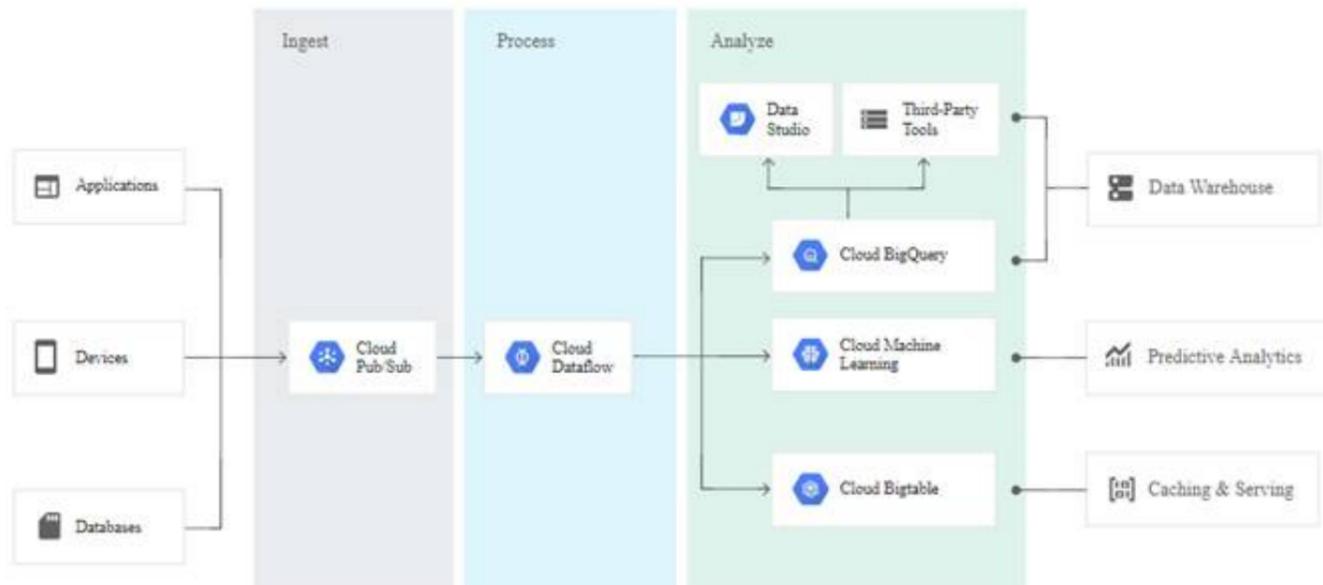


- A. Option A
- B. Option B
- C. Option C
- D. Option D

**Answer:** A

**Explanation:**

<https://cloud.google.com/solutions/iot/> <https://cloud.google.com/solutions/designing-connected-vehicle-platform> [https://cloud.google.com/solutions/designing-connected-vehicle-platform#data\\_ingestion](https://cloud.google.com/solutions/designing-connected-vehicle-platform#data_ingestion)  
<http://www.eweek.com/big-data-and-analytics/google-touts-value-of-cloud-iot-core-for-analyzing-connected-car-data>  
<https://cloud.google.com/solutions/iot/>  
 The push endpoint can be a load balancer.  
 A container cluster can be used. Cloud Pub/Sub for Stream Analytics



References: <https://cloud.google.com/pubsub/> <https://cloud.google.com/solutions/iot/> <https://cloud.google.com/solutions/designing-connected-vehicle-platform>  
[https://cloud.google.com/solutions/designing-connected-vehicle-platform#data\\_ingestion](https://cloud.google.com/solutions/designing-connected-vehicle-platform#data_ingestion) <http://www.eweek.com/big-data-and-analytics/google-touts-value-of-cloud-iot-core-for-analyzing-connected-car-data>  
<https://cloud.google.com/solutions/iot/>

**NEW QUESTION 5**

- (Topic 2)

For this question, refer to the TerramEarth case study.

TerramEarth has equipped unconnected trucks with servers and sensors to collect telemetry data. Next year they want to use the data to train machine learning models. They want to store this data in the cloud while reducing costs. What should they do?

- A. Have the vehicle's computer compress the data in hourly snapshots, and store it in a Google Cloud storage (GCS) Nearline bucket.
- B. Push the telemetry data in Real-time to a streaming dataflow job that compresses the data, and store it in Google BigQuery.
- C. Push the telemetry data in real-time to a streaming dataflow job that compresses the data, and store it in Cloud Bigtable.
- D. Have the vehicle's computer compress the data in hourly snapshots, and store it in a GCS Coldline bucket.

**Answer: D**

**Explanation:**

Coldline Storage is the best choice for data that you plan to access at most once a year, due to its slightly lower availability, 90-day minimum storage duration, costs for data access, and higher per-operation costs. For example:

Cold Data Storage - Infrequently accessed data, such as data stored for legal or regulatory reasons, can be stored at low cost as Coldline Storage, and be available when you need it. Disaster recovery - In the event of a disaster recovery event, recovery time is key. Cloud Storage provides low latency access to data stored as Coldline Storage.

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

**NEW QUESTION 6**

- (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 7**

- (Topic 2)

For this question refer to the TerramEarth case study.

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

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

**Answer: B**

**Explanation:**

Capacity planning, TCO calculations, opex/capex allocation From the case study, it can conclude that Management (CXO) all concern rapid provision of resources (infrastructure)

for growing as well as cost management, such as Cost optimization in Infrastructure, trade up front capital expenditures (Capex) for ongoing operating expenditures (Opex), and Total cost of ownership (TCO)

### NEW QUESTION 8

- (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 9

- (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/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 10

- (Topic 3)

For this question, refer to the JencoMart case study.

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

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

**Answer:** A

#### Explanation:

<https://cloud.google.com/iam/docs/understanding-service-accounts>

Migrating data to Google Cloud Platform

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

References: [https://cloud.google.com/iam/docs/understanding-service-accounts#migrating\\_data\\_to\\_google\\_cloud\\_platform](https://cloud.google.com/iam/docs/understanding-service-accounts#migrating_data_to_google_cloud_platform)

### NEW QUESTION 10

- (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 12**

- (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' images
- 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' images
- 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 16**

- (Topic 4)

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 run the end-to-end tests in the cloud staging environment to determine if the code is working as intended.
- B. They should enable Google Stackdriver Debugger on the application code to show errors in the code
- C. They should add additional unit tests and production scale load tests on their cloud staging environment.
- 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 17**

- (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 20**

- (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 25

- (Topic 4)

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

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

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

**Answer: A**

#### NEW QUESTION 27

- (Topic 5)

Your company captures all web traffic data in Google Analytics 260 and stores it in BigQuery. Each country has its own dataset. Each dataset has multiple tables. You want analysts from each country

to be able to see and query only the data for their respective countries. How should you configure the access rights?

- A. Create a group per country
- B. Add analysts to their respective country-group
- C. Create a single group 'all\_analysts', and add all country-groups as member
- D. Grant the 'all-analysis' group the IAM role of BigQuery jobUser
- E. Share the appropriate dataset with view access with each respective analyst country-group.
- F. Create a group per country
- G. Add analysts to their respective country-group
- H. Create a single group 'all\_analysts', and add all country-groups as member
- I. Grant the 'all-analysis' group the IAM role of BigQuery jobUser
- J. Share the appropriate tables with view access with each respective analyst countrygroup.
- K. Create a group per country
- L. Add analysts to their respective country-group
- M. Create a single group 'all\_analysts', and add all country-groups as member
- N. Grant the 'all-analysis' group the IAM role of BigQuery dataView
- O. Share the appropriate dataset with view access with each respective analystcountry-group.
- P. Create a group per country
- Q. Add analysts to their respective country-group
- R. Create a single group 'all\_analysts', and add all country-groups as member
- S. Grant the 'all-analysis' group the IAM role of BigQuery dataView
- T. Share the appropriate table with view access with each respective analyst countrygroup.

**Answer: A**

#### NEW QUESTION 30

- (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 32

- (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](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 36**

- (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 i the different Shared VPCs
- C. Migrate the protects from the acquired company into your company's Google Cloud organization Re launch the instances in your companies Shared VPC
- D. Set up a Cloud VPN gateway in each Shared VPC and peer Cloud VPNs

**Answer: B**

**NEW QUESTION 38**

- (Topic 5)

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

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

**Answer: C**

**Explanation:**

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

**NEW QUESTION 40**

- (Topic 5)

The application reliability team at your company has added a debug feature to their backend service to send all server events to Google Cloud Storage for eventual analysis. The event records are at least 50 KB and at most 15 MB and are expected to peak at 3,000 events per second. You want to minimize data loss. Which process should you implement?

- A. • Append metadata to file body. • Compress individual files. • Name files with serverName-Timestamp. • Create a new bucket if bucket is older than 1 hour and save individual files to the new bucke
- B. Otherwise, save files to existing bucket
- C. • Batch every 10,000 events with a single manifest file for metadata. • Compress event files and manifest file into a single archive file. • Name files using serverName-EventSequence. • Create a new bucket if bucket is older than 1 day and save the single archive file to the new bucke
- D. Otherwise, save the single archive file to existing bucket.
- E. • Compress individual files. • Name files with serverName-EventSequence. • Save files to one bucket • Set custom metadata headers for each object after saving.
- F. • Append metadata to file body. • Compress individual files. • Name files with a random prefix pattern. • Save files to one bucket

**Answer: D**

**Explanation:**

In order to maintain a high request rate, avoid using sequential names. Using completely random object names will give you the best load distribution. Randomness after a common prefix is effective under the prefix <https://cloud.google.com/storage/docs/request-rate>

#### NEW QUESTION 41

- (Topic 5)

Your company has a Google Cloud project that uses BigQuery for data warehousing. They have a VPN tunnel between the on-premises environment and Google Cloud that is configured with Cloud VPN. The security team wants to avoid data exfiltration by malicious insiders, compromised code, and accidental oversharing. What should they do?

- A. Configure Private Google Access for on-premises only.
- B. Perform the following tasks: 1) Create a service account. 2) Give the BigQuery JobUser role and Storage Reader role to the service account. 3) Remove all other IAM access from the project.
- C. Configure VPC Service Controls and configure Private Google Access.
- D. Configure Private Google Access.

**Answer: C**

#### Explanation:

<https://cloud.google.com/vpc-service-controls/docs/overview>

VPC Service Controls improves your ability to mitigate the risk of data exfiltration from Google Cloud services such as Cloud Storage and BigQuery.

#### NEW QUESTION 45

- (Topic 5)

Your BigQuery project has several users. For audit purposes, you need to see how many queries each user ran in the last month.

- A. Connect Google Data Studio to BigQuery.
- B. Create a dimension for the users and a metric for the amount of queries per user.
- C. In the BigQuery interface, execute a query on the JOBS table to get the required information.
- D. Use 'bq show' to list all jobs.
- E. Per job, use 'bq ls' to list job information and get the required information.
- F. Use Cloud Audit Logging to view Cloud Audit Logs, and create a filter on the query operation to get the required information.

**Answer: C**

#### Explanation:

<https://cloud.google.com/bigquery/docs/managing-jobs>

#### NEW QUESTION 48

- (Topic 5)

Your company has a Google Cloud project that uses BigQuery for data warehousing on a pay-per-use basis. You want to monitor queries in real time to discover the most costly queries and which users spend the most. What should you do?

- A.
  - \* 1. Create a Cloud Logging sink to export BigQuery data access logs to Cloud Storage.
  - \* 2. Develop a Dataflow pipeline to compute the cost of queries split by users.
- B.
  - \* 1. Create a Cloud Logging sink to export BigQuery data access logs to BigQuery.
  - \* 2. Perform a BigQuery query on the generated table to extract the information you need.
- C.
  - \* 1. Activate billing export into BigQuery.
  - \* 2. Perform a BigQuery query on the billing table to extract the information you need.
- D.
  - \* 1. In the BigQuery dataset that contains all the tables to be queried, add a label for each user that can launch a query.
  - \* 2. Open the Billing page of the project.
  - \* 3. Select Reports.
  - \* 4. Select BigQuery as the product and filter by the user you want to check.

A.

**Answer: C**

#### Explanation:

<https://cloud.google.com/blog/products/data-analytics/taking-a-practical-approach-to-bigquery-cost-monitoring>

#### NEW QUESTION 50

- (Topic 5)

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

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

**Answer: A**

#### Explanation:

<https://cloud.google.com/resource-manager/docs/organization-policy/restricting-domains>

#### NEW QUESTION 53

- (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-vms-migrate-for-compute-engine-getting-started>

**NEW QUESTION 55**

- (Topic 5)

You are tasked with building an online analytical processing (OLAP) marketing analytics and reporting tool.

This requires a relational database that can operate on hundreds of terabytes of data. What is the Google recommended tool for such applications?

- A. Cloud Spanner, because it is globally distributed
- B. Cloud SQL, because it is a fully managed relational database
- C. Cloud Firestore, because it offers real-time synchronization across devices
- D. BigQuery, because it is designed for large-scale processing of tabular data

**Answer: A**

**Explanation:**

Reference: <https://cloud.google.com/files/BigQueryTechnicalWP.pdf>

**NEW QUESTION 59**

- (Topic 5)

You are deploying a PHP App Engine Standard service with SQL as the backend. You want to minimize the number of queries to the database.

What should you do?

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

**Answer: A**

**Explanation:**

<https://cloud.google.com/appengine/docs/standard/php/memcache/using>

**NEW QUESTION 64**

- (Topic 5)

Your company is designing its data lake on Google Cloud and wants to develop different ingestion pipelines to collect unstructured data from different sources. After the data is stored in Google Cloud, it will be processed in several data pipelines to build a recommendation engine for end users on the website. The structure of the data retrieved from the source systems can change at any time. The data must be stored exactly as it was retrieved for reprocessing purposes in case the data structure is incompatible with the current processing pipelines. You need to design an architecture to support the use case after you retrieve the data. What should you do?

- A. Send the data through the processing pipeline, and then store the processed data in a BigQuery table for reprocessing.
- B. Store the data in a BigQuery tabl
- C. Design the processing pipelines to retrieve the data from the table.
- D. Send the data through the processing pipeline, and then store the processed data in a Cloud Storage bucket for reprocessing.
- E. Store the data in a Cloud Storage bucke
- F. Design the processing pipelines to retrieve the data from the bucket

**Answer: D**

**NEW QUESTION 68**

- (Topic 5)

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

spend. What should you do?

- A. Write a lifecycle management rule in XML and push it to the bucket with gsutil.
- B. Write a lifecycle management rule in JSON and push it to the bucket with gsutil.
- C. Schedule a cron script using gsutil ls -lr gs://backups/\*\* to find and remove items older than 90 days.
- D. Schedule a cron script using gsutil ls -1 gs://backups/\*\* to find and remove items older than 90 days and schedule it with cron.

**Answer: B**

**Explanation:**

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

#### NEW QUESTION 70

- (Topic 5)

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

- A. Upload the required installation files to Cloud Storage
- B. Configure the VM on a subnet with a Private Google Access subne
- C. Assign only an internal IP address to the V
- D. Download the installation files to the VM using gsutil.
- E. Upload the required installation files to Cloud Storage and use firewall rules to block all traffic except the IP address range for Cloud Storage
- F. Download the files to the VM using gsutil.
- G. Upload the required installation files to Cloud Source Repositorie
- H. Configure the VM on a subnet with a Private Google Access subne
- I. Assign only an internal IP address to the V
- J. Download the installation files to the VM using gcloud.
- K. Upload the required installation files to Cloud Source Repositories and use firewall rules to block all traffic except the IP address range for Cloud Source Repositorie
- L. Download the files to the VM using gsutil.

**Answer: A**

**Explanation:**

<https://cloud.google.com/vpc/docs/private-access-options#pga-supported>

#### NEW QUESTION 74

- (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 75

- (Topic 5)

You want to optimize the performance of an accurate, real-time, weather-charting application. The data comes from 50,000 sensors sending 10 readings a second, in the format of a timestamp and sensor reading. Where should you store the data?

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

**Answer: C**

**Explanation:**

It is time-series data, So Big Table. <https://cloud.google.com/bigtable/docs/schema-design-time-series>

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

Good for:

- ? Low-latency read/write access
- ? High-throughput analytics
- ? Native time series support
- ? Common workloads:
- ? IoT, finance, adtech
- ? Personalization, recommendations
- ? Monitoring
- ? Geospatial datasets
- ? Graphs

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

#### NEW QUESTION 76

- (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 78

- (Topic 5)

Your company has decided to make a major revision of their API in order to create better experiences for their developers. They need to keep the old version of the API available and deployable, while allowing new customers and testers to try out the new API. They want to keep the same SSL and DNS records in place to serve both APIs. What should they do?

- A. Configure a new load balancer for the new version of the API.
- B. Reconfigure old clients to use a new endpoint for the new API.
- C. Have the old API forward traffic to the new API based on the path.
- D. Use separate backend pools for each API path behind the load balancer.

**Answer: D**

#### Explanation:

<https://cloud.google.com/endpoints/docs/openapi/lifecycle-management>

#### NEW QUESTION 80

- (Topic 5)

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

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

**Answer: B**

#### Explanation:

<https://cloud.google.com/architecture/migration-to-google-cloud-transferring-your-large-datasets#transfer-options>  
<https://cloud.google.com/storage-transfer-service>

#### NEW QUESTION 85

- (Topic 5)

Your organization has stored sensitive data in a Cloud Storage bucket. For regulatory reasons, your company must be able to rotate the encryption key used to encrypt the data in the bucket. The data will be processed in Dataproc. You want to follow Google-recommended practices for security. What should you do?

- A. Create a key with Cloud Key Management Service (KMS). Encrypt the data using the encrypt method of Cloud KMS.
- B. Create a key with Cloud Key Management Service (KMS). Set the encryption key on the bucket to the Cloud KMS key.
- C. Generate a GPG key pair
- D. Encrypt the data using the GPG key
- E. Upload the encrypted data to the bucket.
- F. Generate an AES-256 encryption key
- G. Encrypt the data in the bucket using the customer-supplied encryption keys feature.

**Answer: AD**

#### Explanation:

<https://cloud.google.com/storage/docs/encryption/using-customer-managed-keys#add-object-key>  
<https://cloud.google.com/storage/docs/encryption/using-customer-managed-keys>

#### NEW QUESTION 90

- (Topic 5)

You are implementing the infrastructure for a web service on Google Cloud. The web service needs to receive and store the data from 500,000 requests per second. The data will be queried later in real time, based on exact matches of a known set of attributes. There will be periods where the web service will not receive any requests. The business wants to keep costs low. Which web service platform and database should you use for the application?

- A. Cloud Run and BigQuery
- B. Cloud Run and Cloud Bigtable
- C. A Compute Engine autoscaling managed instance group and BigQuery
- D. A Compute Engine autoscaling managed instance group and Cloud Bigtable

**Answer: B**

#### Explanation:

<https://cloud.google.com/run/docs/about-instance-autoscaling> <https://cloud.google.com/blog/topics/developers-practitioners/bigtable-vs-bigquery-whats-difference>

**NEW QUESTION 93**

- (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 if 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 red 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 97**

- (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 101**

- (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 106**

- (Topic 5)

Your company has an application that is running on multiple instances of Compute Engine. It generates 1 TB per day of logs. For compliance reasons, the logs need to be kept for at least two years. The logs need to be available for active query for 30 days. After that, they just need to be retained for audit purposes. You want to implement a storage solution that is compliant, minimizes costs, and follows Google-recommended practices. What should you do?

- A.
  - \* 1. Install the Cloud Ops agent on all instances.
  - \* 2. Create a sink to export logs into a partitioned BigQuery table.
  - \* 3. Set a time\_partitioning\_expiration of 30 days.
- B.
  - \* 1. Install the Cloud Ops agent on all instances.
  - \* 2. Create a sink to export logs into a regional Cloud Storage bucket.
  - \* 3. Create an Object Lifecycle rule to move files into a Coldline Cloud Storage bucket after one month.
  - \* 4. Configure a retention policy at the bucket level to create a lock.
- C.
  - \* 1. Create a daily cron job, running on all instances, that uploads logs into a partitioned BigQuery table.
  - \* 2. Set a time\_partitioning\_expiration of 30 days.
- D.
  - \* 1. Write a daily cron job, running on all instances, that uploads logs into a Cloud Storage bucket.
  - \* 2. Create a sink to export logs into a regional Cloud Storage bucket.
  - \* 3. Create an Object Lifecycle rule to move files into a Coldline Cloud Storage bucket after one month.

A.

**Answer:** B

**Explanation:**

The practice for managing logs generated on Compute Engine on Google Cloud is to install the Cloud Logging agent and send them to Cloud Logging. The sent logs will be aggregated into a Cloud Logging sink and exported to Cloud Storage. The reason for using Cloud Storage as the destination for the logs is that the requirement in question requires setting up a lifecycle based on the storage period. In this case, the log will be used for active queries for 30 days after it is saved, but after that, it needs to be stored for a longer period of time for auditing purposes. If the data is to be used for active queries, we can use BigQuery's Cloud Storage data query feature and move the data past 30 days to Coldline to build a cost-optimal solution.

Therefore, the correct answer is as follows

- \* 1. Install the Cloud Logging agent on all instances.
- Create a sync that exports the logs to the region's Cloud Storage bucket.
- \* 3. Create an Object Lifecycle rule to move the files to the Coldline Cloud Storage bucket after one month. \* 4.
  - \* 4. set up a bucket-level retention policy using bucket locking."

**NEW QUESTION 111**

- (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 114**

- (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 116**

- (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 119**

- (Topic 5)

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

- A. Create an egress rule with priority 1000 to deny all traffic for all instances.
- B. Create another egress rule with priority 100 to allow the Active Directory traffic for all instances.
- C. Create an egress rule with priority 100 to deny all traffic for all instances.
- D. Create another egress rule with priority 1000 to allow the Active Directory traffic for all instances.
- E. Create an egress rule with priority 1000 to allow the Active Directory traffic.
- F. Rely on the implied deny egress rule with priority 100 to block all traffic for all instances.
- G. Create an egress rule with priority 100 to allow the Active Directory traffic.
- H. Rely on the implied deny egress rule with priority 1000 to block all traffic for all instances.

**Answer:** B

**Explanation:**

<https://cloud.google.com/vpc/docs/firewalls>

**NEW QUESTION 122**

- (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 124**

- (Topic 5)

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

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

**Answer: C**

**Explanation:**

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

Cloud Filestore: Cloud Filestore is a scalable and highly available shared file service fully managed by Google. Cloud Filestore provides persistent storage ideal for shared workloads. It is best suited for enterprise applications requiring persistent, durable, shared storage which is accessed by NFS or requires a POSIX compliant file system.

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

**NEW QUESTION 128**

- (Topic 5)

You are creating an App Engine application that uses Cloud Datastore as its persistence layer. You need to retrieve several root entities for which you have the identifiers. You want to minimize the overhead in operations performed by Cloud Datastore. What should you do?

- A. Create the Key object for each Entity and run a batch get operation
- B. Create the Key object for each Entity and run multiple get operations, one operation for each entity
- C. Use the identifiers to create a query filter and run a batch query operation
- D. Use the identifiers to create a query filter and run multiple query operations, one operation for each entity

**Answer: C**

**Explanation:**

<https://cloud.google.com/datastore/docs/concepts/entities#datastore-datastore-batch-upsert-nodejs>

**NEW QUESTION 132**

- (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 136**

- (Topic 5)

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

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

**Answer: CD**

**Explanation:**

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

Enabling automated backups, along with binary logging, is also required for some operations, such as clone and replica creation.

Reference: <https://cloud.google.com/sql/docs/mysql/backup-recovery/backups>

**NEW QUESTION 138**

- (Topic 5)

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

- A. Create a Dataflow pipeline to retrieve the data from the external source
- B. As part of the pipeline use the Cloud Data Loss Prevention (Cloud DLP) API to remove any PII data Store the result in BigQuery
- C. Create a Dataflow pipeline to retrieve the data from the external source
- D. As part of the pipeline store all non-PII data in BigQuery and store all PII data in a Cloud Storage bucket that has a retention policy set.

E. Ask the external partners to upload an data on Cloud Storage Configure Bucket Lock for the bucket Create a Dataflow pipeline to read the data from the bucket As part of the pipeline, use the Cloud Data Loss Prevention (Cloud DIP) API to remove any PII data Store the result in BigQuery  
F. Ask the external partners to import ail data in your BigQuery dataset Create a dataflow pipeline to copy the data into a new table As part of the Dataflow bucket skip all data in columns that have PII data

**Answer:** A

**Explanation:**

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

**NEW QUESTION 142**

- (Topic 5)

Your company sends all Google Cloud logs to Cloud Logging. Your security team wants to monitor the logs. You want to ensure that the security team can react quickly if an anomaly such as an unwanted firewall change or server breach is detected. You want to follow Google-recommended practices. What should you do?

- A. Schedule a cron job with Cloud Schedule
- B. The scheduled job queries the logs every minute for the relevant events.
- C. Export logs to BigQuery, and trigger a query in BigQuery to process the log data for the relevant events.
- D. Export logs to a Pub/Sub topic, and trigger Cloud Function with the relevant log events.
- E. Export logs to a Cloud Storage bucket, and trigger Cloud Run with the relevant log events.

**Answer:** C

**Explanation:**

<https://cloud.google.com/blog/products/management-tools/automate-your-response-to-a-cloud-logging-event>

**NEW QUESTION 143**

- (Topic 5)

Your company operates nationally and plans to use GCP for multiple batch workloads, including some that are not time-critical. You also need to use GCP services that are HIPAA-certified and manage service costs. How should you design to meet Google best practices?

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

**Answer:** B

**Explanation:**

<https://cloud.google.com/security/compliance/hipaa/>

**NEW QUESTION 145**

- (Topic 5)

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

- A. \* 1. Update your GKE cluster to use Cloud Operations for GKE.\* 2. Use the GKE Monitoring dashboard to investigate logs from affected Pods.
- B. \* 1. Create a new GKE cluster with Cloud Operations for GKE enabled.\* 2. Migrate the affected Pods to the new cluster, and redirect traffic for those Pods to the new cluster.\* 3. Use the GKE Monitoring dashboard to investigate logs from affected Pods.
- C. \* 1. Update your GKE cluster to use Cloud Operations for GKE, and deploy Prometheus.\* 2. Set an alert to trigger whenever the application returns an error.
- D. \* 1. Create a new GKE cluster with Cloud Operations for GKE enabled, and deploy Prometheus.\* 2. Migrate the affected Pods to the new cluster, and redirect traffic for those Pods to the new cluster.\* 3. Set an alert to trigger whenever the application returns an error.

**Answer:** A

**Explanation:**

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

**NEW QUESTION 149**

- (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 DSScompliant. 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 150**

- (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.pr

At

Sun.securityutil.manifestEntryVerifier.ver

At java . net . URLClassLoader . defineCla

At sun . reflect . GeneratedMethodAccessors

At

Sun.reflect . DelegatingMethodAccesorImpl.

At java . lang . reflect . MThod . 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 151**

- (Topic 5)

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

- A. Create an aggregated export on the Production folde
- B. Set the log sink to be a Cloud Storage bucket in an operations project
- C. Create an aggregated export on the Organization resourc
- D. Set the tog sink to be a Cloud Storage bucket in an operations project.
- E. Create log exports in the production project
- F. Set the log sinks to be a Cloud Storagebucket in an operations project.
- G. Create tog exports in the production project
- H. Set the tog sinks to be BigQuery datasets in the production projects and grant IAM access to the operations team to run queries on the datasets

**Answer: A**

**Explanation:**

? An aggregated export is a type of sink that combines and routes log entries from the Google Cloud resources contained by an organization or folder1. By creating an aggregated export on the Production folder, you can capture all the logs from the existing and new production projects in that folder automatically1.  
 ? A log sink is a destination for log entries that match a filter1. By setting the log sink to be a Cloud Storage bucket in an operations project, you can store the log entries in Cloud Storage and allow your operations team to access them1.

**NEW QUESTION 154**

- (Topic 5)

Your team will start developing a new application using microservices architecture on Kubernetes Engine. As part of the development lifecycle, any code change that has been pushed to the remote develop branch on your GitHub repository should be built and tested automatically. When the build and test are successful, the relevant microservice will be deployed automatically in the development environment. You want to ensure that all code deployed in the development environment

follows this process. What should you do?

- A. Have each developer install a pre-commit hook on their workstation that tests the code and builds the container when committing on the development branch
- B. After a successful commit, have the developer deploy the newly built container image on the development cluster.
- C. Install a post-commit hook on the remote git repository that tests the code and builds the container when code is pushed to the development branch
- D. After a successful commit, have the developer deploy the newly built container image on the development cluster.
- E. Create a Cloud Build trigger based on the development branch that tests the code, builds the container, and stores it in Container Registry
- F. Create a deployment pipeline that watches for new images and deploys the new image on the development cluster
- G. Ensure only the deployment tool has access to deploy new versions.
- H. Create a Cloud Build trigger based on the development branch to build a new container image and store it in Container Registry
- I. Rely on Vulnerability Scanning to ensure the code tests succeed
- J. As the final step of the Cloud Build process, deploy the new container image on the development cluster
- K. Ensure only Cloud Build has access to deploy new versions.

**Answer: C**

**Explanation:**

<https://cloud.google.com/container-registry/docs/overview>

Create a Cloud Build trigger based on the development branch that tests the code, builds the container, and stores it in Container Registry. Create a deployment pipeline that watches for new images and deploys the new image on the development cluster. Ensure only the deployment tool has access to deploy new versions.

**NEW QUESTION 157**

- (Topic 5)

Your development team has installed a new Linux kernel module on the batch servers in Google Compute Engine (GCE) virtual machines (VMs) to speed up the nightly batch process. Two days after the installation, 50% of web application deployed in the same nightly batch run. You want to collect details on the failure to pass back to the development team. Which three actions should you take? Choose 3 answers

- A. Use Stackdriver Logging to search for the module log entries.
- B. Read the debug GCE Activity log using the API or Cloud Console.
- C. Use gcloud or Cloud Console to connect to the serial console and observe the logs.
- D. Identify whether a live migration event of the failed server occurred, using in the activity log.
- E. Adjust the Google Stackdriver timeline to match the failure time, and observe the batch server metrics.
- F. Export a debug VM into an image, and run the image on a local server where kernel log messages will be displayed on the native screen.

**Answer: ACE**

**Explanation:**

<https://www.flexera.com/blog/cloud/2013/12/google-compute-engine-live-migration-passes-the-test/>

"With live migration, the virtual machines are moved without any downtime or noticeable service degradation"

**NEW QUESTION 160**

- (Topic 5)

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

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

**Answer: D**

**Explanation:**

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

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

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

**NEW QUESTION 164**

- (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 166**

- (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 --local-host-port=localhost:22 --remote-host-port=remotehost:22`
- \* 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 167

- (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 172

- (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 176

- (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 180

- (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 183

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

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

### NEW QUESTION 188

- (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 190

- (Topic 5)

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

Which approach should you take?

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

**Answer: C**

#### Explanation:

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

### NEW QUESTION 192

- (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 196

- (Topic 5)

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

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

**Answer:**

D

**Explanation:**

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

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

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

**NEW QUESTION 200**

- (Topic 5)

For this question, refer to the TerramEarth case study. You are building a microservice-based application for TerramEarth. The application is based on Docker containers. You want to follow Google-recommended practices to build the application continuously and store the build artifacts. What should you do?

A.

- \* 1. Configure a trigger in Cloud Build for new source changes.
- \* 2. Invoke Cloud Build to build one container image, and tag the image with the label 'latest.'
- \* 3. Push the image to the Artifact Registry.

B.

- \* 1. Configure a trigger in Cloud Build for new source changes.
- \* 2. Invoke Cloud Build to build container images for each microservice, and tag them using the code commit hash.
- \* 3. Push the images to the Artifact Registry.

C.

- \* 1 Create a Scheduler job to check the repo every minute.
- \* 2. For any new change, invoke Cloud Build to build container images for the microservices.
- \* 3. Tag the images using the current timestamp, and push them to the Artifact Registry.

D.

- \* 1. Configure a trigger in Cloud Build for new source changes.
- \* 2. The trigger invokes build jobs and build container images for the microservices.
- \* 3. Tag the images with a version number, and push them to Cloud Storage.

A.

**Answer: C**

**NEW QUESTION 201**

- (Topic 5)

Your company plans to migrate a multi-petabyte data set to the cloud. The data set must be available 24hrs a day. Your business analysts have experience only with using a SQL interface. How should you store the data to optimize it for ease of analysis?

- A. Load data into Google BigQuery.
- B. Insert data into Google Cloud SQL.
- C. Put flat files into Google Cloud Storage.
- D. Stream data into Google Cloud Datastore.

**Answer: A**

**Explanation:**

Google Big Query is for multi peta byte storage , HA(High availability) which means 24 hours, SQL interface .

<https://medium.com/google-cloud/the-12-components-of-google-bigquery-c2b49829a7c7> <https://cloud.google.com/solutions/bigquery-data-warehouse>

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

BigQuery is Google's serverless, highly scalable, low cost enterprise data warehouse designed to make all your data analysts productive. Because there is no infrastructure to manage, you can focus on analyzing data to find meaningful insights using familiar SQL and you don't need a database administrator.

BigQuery enables you to analyze all your data by creating a logical data warehouse over managed, columnar storage as well as data from object storage, and spreadsheets.

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

**NEW QUESTION 202**

- (Topic 5)

You want to automate the creation of a managed instance group and a startup script to install the OS package dependencies. You want to minimize the startup time for VMs in the instance group.

What should you do?

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

**Answer: B**

**Explanation:**

"Custom images are more deterministic and start more quickly than instances with startup scripts. However, startup scripts are more flexible and let you update the apps and settings in your instances more easily." [https://cloud.google.com/compute/docs/instance-templates/create-instance-templates#using\\_custom\\_or\\_public\\_images\\_in\\_your\\_instance\\_templates](https://cloud.google.com/compute/docs/instance-templates/create-instance-templates#using_custom_or_public_images_in_your_instance_templates)

**NEW QUESTION 207**

- (Topic 5)

A news feed web service has the following code running on Google App Engine. During peak load, users report that they can see news articles they already viewed. What is the most likely cause of this problem?

```
import news
from flask import Flask, redirect, request
from flask.ext.api import status
from google.appengine.api import users

app = Flask(__name__)
sessions = {}

@app.route("/")
def homepage():
    user = users.get_current_user()
    if not user:
        return "Invalid login",
        status.HTTP_401_UNAUTHORIZED

    if user not in sessions:
        sessions[user] = {"viewed": []}

    news_articles = news.get_new_news (user, sessions [user]
["viewed"])
    sessions [user] ["viewed"] += [n["id"] for n
in news_articles]

    return news.render(news_articles)

if __name__ == "__main__":
    app.run()
```

- A. The session variable is local to just a single instance.
- B. The session variable is being overwritten in Cloud Datastore.
- C. The URL of the API needs to be modified to prevent caching.
- D. The HTTP Expires header needs to be set to -1 to stop caching.

**Answer:** A

**Explanation:**

<https://stackoverflow.com/questions/3164280/google-app-engine-cache-list-in-session-variable?rq=1>

**NEW QUESTION 210**

- (Topic 5)

An application development team believes their current logging tool will not meet their needs for their new cloud-based product. They want a better tool to capture errors and help them analyze their historical log data. You want to help them find a solution that meets their needs, what should you do?

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

**Answer:** C

**Explanation:**

Help them define their requirements and assess viable logging tools. They know the requirements and the existing tools' problems. While it's true StackDriver Logging and Error Reporting possibly meet all their requirements, there might be other tools also meet their need. They need you to provide expertise to make assessment for new tools, specifically, logging tools that can "capture errors and help them analyze their historical log data".

References: <https://cloud.google.com/logging/docs/agent/installation>

**NEW QUESTION 212**

- (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 216

- (Topic 5)

Your company wants to start using Google Cloud resources but wants to retain their on-premises Active Directory domain controller for identity management. What should you do?

- A. Use the Admin Directory API to authenticate against the Active Directory domain controller.
- B. Use Google Cloud Directory Sync to synchronize Active Directory usernames with cloud identities and configure SAML SSO.
- C. Use Cloud Identity-Aware Proxy configured to use the on-premises Active Directory domain controller as an identity provider.
- D. Use Compute Engine to create an Active Directory (AD) domain controller that is a replica of the on-premises AD domain controller using Google Cloud Directory Sync.

**Answer:** B

**Explanation:**

[https://cloud.google.com/solutions/federating-gcp-with-active-directory-introduction#implementing\\_federation](https://cloud.google.com/solutions/federating-gcp-with-active-directory-introduction#implementing_federation)

#### NEW QUESTION 218

- (Topic 5)

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

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

**Answer:** A

**Explanation:**

The Anthos Service Mesh pages in the Google Cloud Console provide both summary and in-depth metrics, charts, and graphs that enable you to observe service behavior. You can monitor the overall health of your services, or drill down on a specific service to set a service level objective (SLO) or troubleshoot an issue.  
<https://cloud.google.com/service-mesh/docs/observability/explore-dashboard> <https://cloud.google.com/anthos/service-mesh>

#### NEW QUESTION 223

- (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 Insert 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 226

- (Topic 5)

You need to evaluate your team's 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](https://services.google.com/fh/files/misc/cloud_center_of_excellence.pdf)

#### NEW QUESTION 228

- (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 you create 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. Most access is 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 232

- (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 KMS.
- 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 KMS.
- 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 236

- (Topic 5)

Your team is developing a web application that will be deployed on Google Kubernetes Engine (GKE). Your CTO expects a successful launch and you need to ensure your application can handle the expected load of tens of thousands of users. You want to test the current deployment to ensure the latency of your application stays below a certain threshold. What should you do?

- A. Use a load testing tool to simulate the expected number of concurrent users and total requests to your application, and inspect the results.
- B. Enable autoscaling on the GKE cluster and enable horizontal pod autoscaling on your application deployment.
- C. Send curl requests to your application, and validate if the auto scaling works.
- D. Replicate the application over multiple GKE clusters in every Google Cloud region. Configure a global HTTP(S) load balancer to expose the different clusters over a single global IP address.
- E. Use Cloud Debugger in the development environment to understand the latency between the different microservices.

**Answer: B**

#### NEW QUESTION 241

- (Topic 5)

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

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

**Answer: A**

#### Explanation:

<https://cloud.google.com/run/docs/rollouts-rollbacks-traffic-migration>

#### NEW QUESTION 246

- (Topic 5)

Your company has a Kubernetes application that pulls messages from Pub/Sub and stores them in Firestore. Because the application is simple, it was deployed as a single pod. The infrastructure team has analyzed Pub/Sub metrics and discovered that the application cannot process the messages in real time. Most of them wait for minutes before being processed. You need to scale the elaboration process that is I/O-intensive. What should you do?

- A. Configure a Kubernetes autoscaling based on the subscription/push\_request metric.
- B. Use the `--enable-` autoscaling flag when you create the Kubernetes cluster
- C. Configure a Kubernetes autoscaling based on the subscription/num\_undelivered message metric.
- D. Use `kubectl autoscale deployment APP_NAME --max 6 --min 2 --cpu-percent 50` to configure Kubernetes autoscaling deployment

**Answer:** A

**Explanation:**

[https://cloud.google.com/kubernetes-engine/docs/concepts/custom-and-external-metrics#external\\_metrics](https://cloud.google.com/kubernetes-engine/docs/concepts/custom-and-external-metrics#external_metrics)

**NEW QUESTION 251**

- (Topic 5)

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

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

**Answer:** B

**NEW QUESTION 252**

- (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 254**

- (Topic 5)

Your company has an application running on App Engine that allows users to upload music files and share them with other people. You want to allow users to upload files directly into Cloud Storage from their browser session. The payload should not be passed through the backend. What should you do?

- A.
  - \* 1. Set a CORS configuration in the target Cloud Storage bucket where the base URL of the App Engine application is an allowed origin.
  - \* 2. Use the Cloud Storage Signed URL feature to generate a POST URL.
- B.
  - \* 1. Set a CORS configuration in the target Cloud Storage bucket where the base URL of the App Engine application is an allowed origin.
  - \* 2. Assign the Cloud Storage WRITER role to users who upload files.
- C.
  - \* 1. Use the Cloud Storage Signed URL feature to generate a POST URL.
  - \* 2. Use App Engine default credentials to sign requests against Cloud Storage.
- D.
  - \* 1. Assign the Cloud Storage WRITER role to users who upload files.
  - \* 2. Use App Engine default credentials to sign requests against Cloud Storage.

A.

**Answer:** B

**NEW QUESTION 255**

- (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 256

- (Topic 5)

Your company has an application running on a deployment in a GKE cluster. You have a separate cluster for development, staging and production. You have discovered that the team is able to deploy a Docker image to the production cluster without first testing the deployment in development and then staging. You want to allow the team to have autonomy but want to prevent this from happening. You want a Google Cloud solution that can be implemented quickly with minimal effort. What should you do?

- A. Create a Kubernetes admission controller to prevent the container from starting if it is not approved for usage in the given environment
- B. Configure a Kubernetes lifecycle hook to prevent the container from starting if it is not approved for usage in the given environment
- C. Implement a corporate policy to prevent teams from deploying Docker image to an environment unless the Docker image was tested in an earlier environment
- D. Configure the binary authorization policies for the development, staging and production cluster
- E. Create attestations as part of the continuous integration pipeline”

**Answer:** D

#### Explanation:

<https://cloud.google.com/architecture/prep-kubernetes-engine-for-prod#binary-authorization>

The most common Binary Authorization use cases involve attestations. An attestation certifies that a specific image has completed a previous stage, as described previously. You configure the Binary Authorization policy to verify the attestation before allowing the image to be deployed. At deploy time, instead of redoing activities that were completed in earlier stages, Binary Authorization only needs to verify the attestation. <https://cloud.google.com/binary-authorization/docs/overview>

#### NEW QUESTION 258

- (Topic 5)

Your company is forecasting a sharp increase in the number and size of Apache Spark and Hadoop jobs being run on your local datacenter. You want to utilize the cloud to help you scale this upcoming demand with the least amount of operations work and code change. Which product should you use?

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

**Answer:** B

#### Explanation:

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

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

#### NEW QUESTION 262

- (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 264

- (Topic 5)

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

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

**Answer:** C

#### NEW QUESTION 266

- (Topic 5)

Your company is developing a new application that will allow globally distributed users to upload pictures and share them with other selected users. The application will support millions of concurrent users. You want to allow developers to focus on just building code without having to create and maintain the

underlying infrastructure. Which service should you use to deploy the application?

- A. App Engine
- B. Cloud Endpoints
- C. Compute Engine
- D. Google Kubernetes Engine

**Answer:** A

**Explanation:**

Reference: <https://cloud.google.com/terms/services> <https://cloud.google.com/appengine/docs/standard/go/how-requests-are-handled>

#### NEW QUESTION 267

- (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 mod
- 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.vmExternallpAcces
- 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 address 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 269

- (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#provision_a_sole-tenant_vm)

[https://cloud.google.com/compute/docs/nodes/provisioning-sole-tenant-vms#gcloud\\_2](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 273

- (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 274

- (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 regio

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 276**

- (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 277**

- (Topic 5)

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

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

**Answer:** A

**Explanation:**

<https://cloud.google.com/stackdriver/docs/solutions/gke/legacy-stackdriver/monitoring>

#### **NEW QUESTION 279**

- (Topic 5)

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

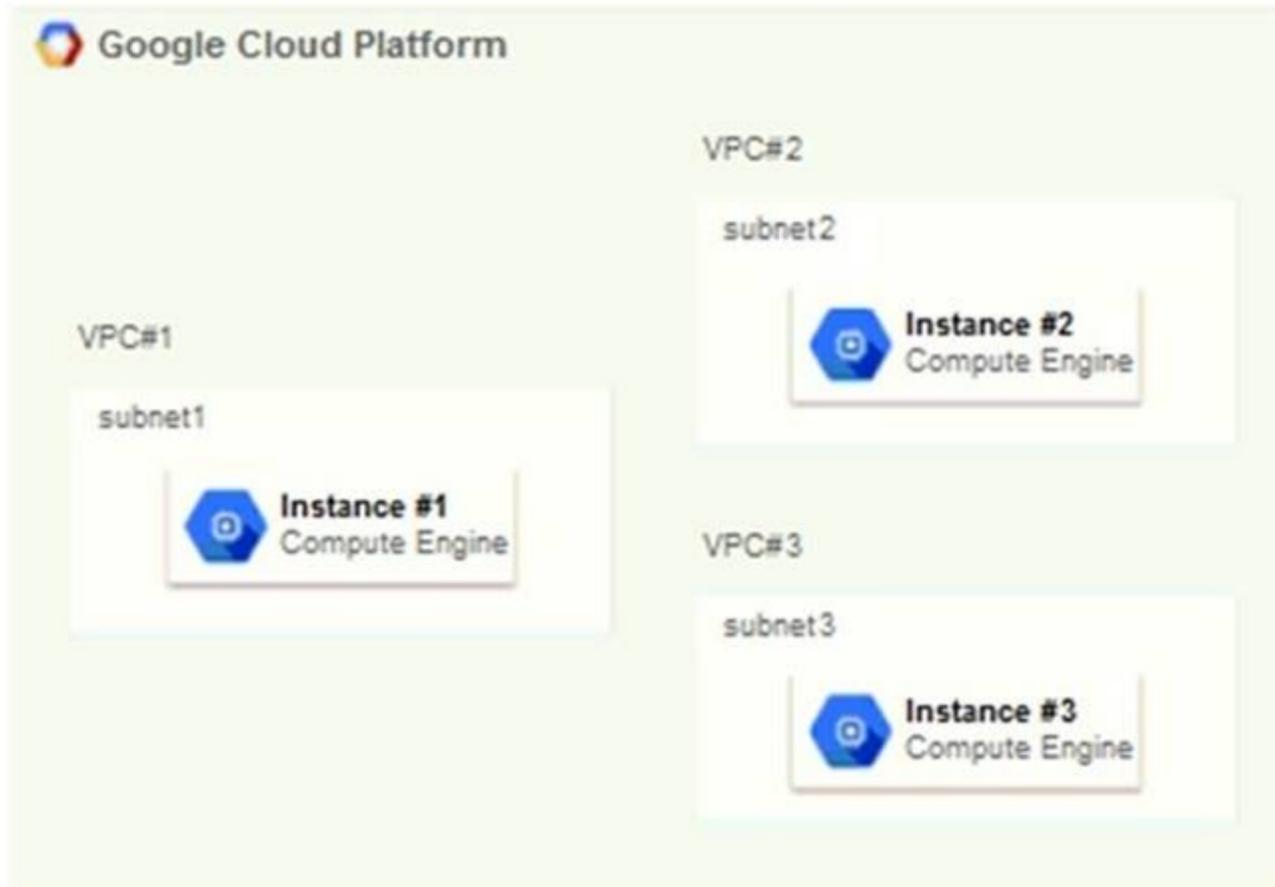
- A. 1) Enable automatic storage increase for the instance.2) Create a Stackdriver alert when CPU usage exceeds 75%, and change the instance type to reduce CPU usage.3) Create a Stackdriver alert for replication lag, and shard the database to reduce replication time.
- B. 1) Enable automatic storage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

**Answer:** A

#### **NEW QUESTION 282**

- (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 #3
 Update 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](https://cloud.google.com/vpc/docs/create-use-multiple-interfaces#i_am_not_able_to_connect_to_secondary_interfaces_internal_ip)

**NEW QUESTION 284**

- (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 289**

- (Topic 5)

Your company has developed a monolithic, 3-tier application to allow external users to upload and share files. The solution cannot be easily enhanced and lacks reliability. The development team would like to re-architect the application to adopt microservices and a fully managed service approach, but they need to convince their leadership that the effort is worthwhile. Which advantage(s) should they highlight to leadership?

- A. The new approach will be significantly less costly, make it easier to manage the underlying infrastructure, and automatically manage the CI/CD pipelines.
- B. The monolithic solution can be converted to a container with Docker.
- C. The generated container can then be deployed into a Kubernetes cluster.
- D. The new approach will make it easier to decouple infrastructure from application, develop and release new features, manage the underlying infrastructure, manage CI/CD pipelines and perform A/B testing, and scale the solution if necessary.
- E. The process can be automated with Migrate for Compute Engine.

**Answer:** C

**Explanation:**

The new approach will make it easier to decouple infrastructure from an application, develop and release new features, manage the underlying infrastructure, manage CI/CD pipelines and perform A/B testing, and scale the solution if necessary.

**NEW QUESTION 293**

.....

## **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](#)**