

## Professional-Cloud-Developer Dumps

### Google Certified Professional - Cloud Developer

<https://www.certleader.com/Professional-Cloud-Developer-dumps.html>



**NEW QUESTION 1**

- (Exam Topic 1)

In order to meet their business requirements, how should HipLocal store their application state?

- A. Use local SSDs to store state.
- B. Put a memcache layer in front of MySQL.
- C. Move the state storage to Cloud Spanner.
- D. Replace the MySQL instance with Cloud SQL.

**Answer: B**

**NEW QUESTION 2**

- (Exam Topic 1)

For this question, refer to the HipLocal case study.

HipLocal is expanding into new locations. They must capture additional data each time the application is launched in a new European country. This is causing delays in the development process due to constant schema changes and a lack of environments for conducting testing on the application changes. How should they resolve the issue while meeting the business requirements?

- A. Create new Cloud SQL instances in Europe and North America for testing and deployment.
- B. Provide developers with local MySQL instances to conduct testing on the application changes.
- C. Migrate data to Bigtable.
- D. Instruct the development teams to use the Cloud SDK to emulate a local Bigtable development environment.
- E. Move from Cloud SQL to MySQL hosted on Compute Engine.
- F. Replicate hosts across regions in the Americas and Europe.
- G. Provide developers with local MySQL instances to conduct testing on the application changes.
- H. Migrate data to Firestore in Native mode and set up instances.

**Answer: B**

**NEW QUESTION 3**

- (Exam Topic 1)

HipLocal wants to improve the resilience of their MySQL deployment, while also meeting their business and technical requirements. Which configuration should they choose?

- A. Use the current single instance MySQL on Compute Engine and several read-only MySQL servers on Compute Engine.
- B. Use the current single instance MySQL on Compute Engine, and replicate the data to Cloud SQL in an external master configuration.
- C. Replace the current single instance MySQL instance with Cloud SQL, and configure high availability.
- D. Replace the current single instance MySQL instance with Cloud SQL, and Google provides redundancy without further configuration.

**Answer: B**

**NEW QUESTION 4**

- (Exam Topic 1)

HipLocal is configuring their access controls.

Which firewall configuration should they implement?

- A. Block all traffic on port 443.
- B. Allow all traffic into the network.
- C. Allow traffic on port 443 for a specific tag.
- D. Allow all traffic on port 443 into the network.

**Answer: D**

**NEW QUESTION 5**

- (Exam Topic 1)

HipLocal's APIs are showing occasional failures, but they cannot find a pattern. They want to collect some metrics to help them troubleshoot. What should they do?

- A. Take frequent snapshots of all of the VMs.
- B. Install the Stackdriver Logging agent on the VMs.
- C. Install the Stackdriver Monitoring agent on the VMs.
- D. Use Stackdriver Trace to look for performance bottlenecks.

**Answer: C**

**NEW QUESTION 6**

- (Exam Topic 2)

You support an application that uses the Cloud Storage API. You review the logs and discover multiple HTTP 503 Service Unavailable error responses from the API. Your application logs the error and does not take any further action. You want to implement Google-recommended retry logic to improve success rates. Which approach should you take?

- A. Retry the failures in batch after a set number of failures is logged.
- B. Retry each failure at a set time interval up to a maximum number of times.
- C. Retry each failure at increasing time intervals up to a maximum number of tries.
- D. Retry each failure at decreasing time intervals up to a maximum number of tries.

**Answer:** C

**Explanation:**

<https://cloud.google.com/storage/docs/retry-strategy>

#### NEW QUESTION 7

- (Exam Topic 2)

You work for an organization that manages an ecommerce site. Your application is deployed behind a global HTTP(S) load balancer. You need to test a new product recommendation algorithm. You plan to use A/B testing to determine the new algorithm's effect on sales in a randomized way. How should you test this feature?

- A. Split traffic between versions using weights.
- B. Enable the new recommendation feature flag on a single instance.
- C. Mirror traffic to the new version of your application.
- D. Use HTTP header-based routing.

**Answer:** A

**Explanation:**

[https://cloud.google.com/load-balancing/docs/https/traffic-management-global#traffic\\_actions\\_weight-based\\_tra](https://cloud.google.com/load-balancing/docs/https/traffic-management-global#traffic_actions_weight-based_tra) Deploying a new version of an existing production service generally incurs some risk. Even if your tests pass in staging, you probably don't want to subject 100% of your users to the new version immediately. With traffic management, you can define percentage-based traffic splits across multiple backend services.

For example, you can send 95% of the traffic to the previous version of your service and 5% to the new version of your service. After you've validated that the new production version works as expected, you can gradually shift the percentages until 100% of the traffic reaches the new version of your service. Traffic splitting is typically used for deploying new versions, A/B testing, service migration, and similar processes.

[https://cloud.google.com/traffic-director/docs/advanced-traffic-management#weight-based\\_traffic\\_splitting\\_for\\_](https://cloud.google.com/traffic-director/docs/advanced-traffic-management#weight-based_traffic_splitting_for_) [https://cloud.google.com/architecture/implementing-deployment-and-testing-strategies-on-gke#split\\_the\\_traffic\\_](https://cloud.google.com/architecture/implementing-deployment-and-testing-strategies-on-gke#split_the_traffic_) [https://cloud.google.com/load-balancing/docs/https/traffic-management-global#traffic\\_actions\\_weight-based\\_tra](https://cloud.google.com/load-balancing/docs/https/traffic-management-global#traffic_actions_weight-based_tra)

#### NEW QUESTION 8

- (Exam Topic 2)

Your company needs a database solution that stores customer purchase history and meets the following requirements:

Customers can query their purchase immediately after submission. Purchases can be sorted on a variety of fields. Distinct record formats can be stored at the same time. Which storage option satisfies these requirements?

- A. Firestore in Native mode
- B. Cloud Storage using an object read
- C. Cloud SQL using a SQL SELECT statement
- D. Firestore in Datastore mode using a global query

**Answer:** A

#### NEW QUESTION 9

- (Exam Topic 2)

You have written a Cloud Function that accesses other Google Cloud resources. You want to secure the environment using the principle of least privilege. What should you do?

- A. Create a new service account that has Editor authority to access the resource
- B. The deployer is given permission to get the access token.
- C. Create a new service account that has a custom IAM role to access the resource
- D. The deployer is given permission to get the access token.
- E. Create a new service account that has Editor authority to access the resource
- F. The deployer is given permission to act as the new service account.
- G. Create a new service account that has a custom IAM role to access the resource
- H. The deployer is given permission to act as the new service account.

**Answer:** D

**Explanation:**

Reference:

<https://cloud.google.com/blog/products/application-development/least-privilege-for-cloud-functions-using-cloud>

#### NEW QUESTION 10

- (Exam Topic 2)

You are developing a microservice-based application that will be deployed on a Google Kubernetes Engine cluster. The application needs to read and write to a Spanner database. You want to follow security best practices while minimizing code changes. How should you configure your application to retrieve Spanner credentials?

- A. Configure the appropriate service accounts, and use Workload Identity to run the pods.
- B. Store the application credentials as Kubernetes Secrets, and expose them as environment variables.
- C. Configure the appropriate routing rules, and use a VPC-native cluster to directly connect to the database.
- D. Store the application credentials using Cloud Key Management Service, and retrieve them whenever a database connection is made.

**Answer:** A

**Explanation:**

<https://cloud.google.com/kubernetes-engine/docs/concepts/workload-identity>

**NEW QUESTION 10**

- (Exam Topic 2)

You want to create “fully baked” or “golden” Compute Engine images for your application. You need to bootstrap your application to connect to the appropriate database according to the environment the application is running on (test, staging, production). What should you do?

- A. Embed the appropriate database connection string in the image.
- B. Create a different image for each environment.
- C. When creating the Compute Engine instance, add a tag with the name of the database to be connected. In your application, query the Compute Engine API to pull the tags for the current instance, and use the tag to construct the appropriate database connection string.
- D. When creating the Compute Engine instance, create a metadata item with a key of “DATABASE” and a value for the appropriate database connection string.
- E. In your application, read the “DATABASE” environment variable, and use the value to connect to the appropriate database.
- F. When creating the Compute Engine instance, create a metadata item with a key of “DATABASE” and a value for the appropriate database connection string.
- G. In your application, query the metadata server for the “DATABASE” value, and use the value to connect to the appropriate database.

**Answer: C**

**NEW QUESTION 14**

- (Exam Topic 2)

You are developing a marquee stateless web application that will run on Google Cloud. The rate of the incoming user traffic is expected to be unpredictable, with no traffic on some days and large spikes on other days. You need the application to automatically scale up and down, and you need to minimize the cost associated with running the application. What should you do?

- A. Build the application in Python with Firestore as the database.
- B. Deploy the application to Cloud Run.
- C. Build the application in C# with Firestore as the database.
- D. Deploy the application to App Engine flexible environment.
- E. Build the application in Python with CloudSQL as the database.
- F. Deploy the application to App Engine standard environment.
- G. Build the application in Python with Firestore as the database.
- H. Deploy the application to a Compute Engine managed instance group with autoscaling.

**Answer: A**

**NEW QUESTION 19**

- (Exam Topic 2)

You have a mixture of packaged and internally developed applications hosted on a Compute Engine instance that is running Linux. These applications write log records as text in local files. You want the logs to be written to Cloud Logging. What should you do?

- A. Pipe the content of the files to the Linux Syslog daemon.
- B. Install a Google version of fluentd on the Compute Engine instance.
- C. Install a Google version of collectd on the Compute Engine instance.
- D. Using cron, schedule a job to copy the log files to Cloud Storage once a day.

**Answer: B**

**Explanation:**

Reference: <https://cloud.google.com/logging/docs/agent/logging/configuration>

**NEW QUESTION 22**

- (Exam Topic 2)

You want to re-architect a monolithic application so that it follows a microservices model. You want to accomplish this efficiently while minimizing the impact of this change to the business.

Which approach should you take?

- A. Deploy the application to Compute Engine and turn on autoscaling.
- B. Replace the application's features with appropriate microservices in phases.
- C. Refactor the monolithic application with appropriate microservices in a single effort and deploy it.
- D. Build a new application with the appropriate microservices separate from the monolith and replace it when it is complete.

**Answer: C**

**Explanation:**

Reference: <https://cloud.google.com/solutions/migrating-a-monolithic-app-to-microservices-gke>

**NEW QUESTION 27**

- (Exam Topic 2)

You work for a web development team at a small startup. Your team is developing a Node.js application using Google Cloud services, including Cloud Storage and Cloud Build. The team uses a Git repository for version control. Your manager calls you over the weekend and instructs you to make an emergency update to one of the company's websites, and you're the only developer available. You need to access Google Cloud to make the update, but you don't have your work laptop. You are not allowed to store source code locally on a non-corporate computer. How should you set up your developer environment?

- A. Use a text editor and the Git command line to send your source code updates as pull requests from a public computer.
- B. Use a text editor and the Git command line to send your source code updates as pull requests from a virtual machine running on a public computer.
- C. Use Cloud Shell and the built-in code editor for development.
- D. Send your source code updates as pull requests.
- E. Use a Cloud Storage bucket to store the source code that you need to edit.
- F. Mount the bucket to a public computer as a drive, and use a code editor to update the code.
- G. Turn on versioning for the bucket, and point it to the team's Git repository.



**Answer:** C

**Explanation:**

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

### NEW QUESTION 31

- (Exam Topic 2)

You need to migrate a standalone Java application running in an on-premises Linux virtual machine (VM) to Google Cloud in a cost-effective manner. You decide not to take the lift-and-shift approach, and instead you plan to modernize the application by converting it to a container. How should you accomplish this task?

- A. Use Migrate for Anthos to migrate the VM to your Google Kubernetes Engine (GKE) cluster as a container.
- B. Export the VM as a raw disk and import it as an image.
- C. Create a Compute Engine instance from the Imported image.
- D. Use Migrate for Compute Engine to migrate the VM to a Compute Engine instance, and use Cloud Build to convert it to a container.
- E. Use Jib to build a Docker image from your source code, and upload it to Artifact Registry.
- F. Deploy the application in a GKE cluster, and test the application.

**Answer:** D

**Explanation:**

<https://cloud.google.com/blog/products/application-development/introducing-jib-build-java-docker-images-better>

### NEW QUESTION 32

- (Exam Topic 2)

Your security team is auditing all deployed applications running in Google Kubernetes Engine. After completing the audit, your team discovers that some of the applications send traffic within the cluster in clear text. You need to ensure that all application traffic is encrypted as quickly as possible while minimizing changes to your applications and maintaining support from Google. What should you do?

- A. Use Network Policies to block traffic between applications.
- B. Install Istio, enable proxy injection on your application namespace, and then enable mTLS.
- C. Define Trusted Network ranges within the application, and configure the applications to allow traffic only from those networks.
- D. Use an automated process to request SSL Certificates for your applications from Let's Encrypt and add them to your applications.

**Answer:** D

### NEW QUESTION 33

- (Exam Topic 2)

You are developing an application that will be launched on Compute Engine instances into multiple distinct projects, each corresponding to the environments in your software development process (development, QA, staging, and production). The instances in each project have the same application code but a different configuration. During deployment, each instance should receive the application's configuration based on the environment it serves. You want to minimize the number of steps to configure this flow.

What should you do?

- A. When creating your instances, configure a startup script using the `gcloud` command to determine the project name that indicates the correct environment.
- B. In each project, configure a metadata key "environment" whose value is the environment it serves.
- C. Use your deployment tool to query the instance metadata and configure the application based on the "environment" value.
- D. Deploy your chosen deployment tool on an instance in each project.
- E. Use a deployment job to retrieve the appropriate configuration file from your version control system, and apply the configuration when deploying the application on each instance.
- F. During each instance launch, configure an instance custom-metadata key named "environment" whose value is the environment the instance serves.
- G. Use your deployment tool to query the instance metadata, and configure the application based on the "environment" value.

**Answer:** B

**Explanation:**

Reference: <https://cloud.google.com/compute/docs/metadata/overview>

### NEW QUESTION 36

- (Exam Topic 2)

Your team is writing a backend application to implement the business logic for an interactive voice response (IVR) system that will support a payroll application. The IVR system has the following technical characteristics:

- Each customer phone call is associated with a unique IVR session.
- The IVR system creates a separate persistent gRPC connection to the backend for each session.
- If the connection is interrupted, the IVR system establishes a new connection, causing a slight latency for that call.

You need to determine which compute environment should be used to deploy the backend application. Using current call data, you determine that:

- Call duration ranges from 1 to 30 minutes.
- Calls are typically made during business hours.
- There are significant spikes of calls around certain known dates (e.g., pay days), or when large payroll changes occur.

You want to minimize cost, effort, and operational overhead. Where should you deploy the backend application?

- A. Compute Engine
- B. Google Kubernetes Engine cluster in Standard mode
- C. Cloud Functions
- D. Cloud Run

**Answer:** D

**Explanation:**

This page shows Cloud Run-specific details for developers who want to use gRPC to connect a Cloud Run service with other services, for example, to provide

simple, high performance communication between internal microservices. You can use all gRPC types, streaming or unary, with Cloud Run.

Possible use cases include:

Communication between internal microservices.

High loads of data (gRPC uses protocol buffers, which are up to seven times faster than REST calls). Only a simple service definition is needed, you don't want to write a full client library.

Use streaming gRPCs in your gRPC server to build more responsive applications and APIs. <https://cloud.google.com/run/docs/tutorials/secure-services#:~:text=The%20backend%20service%20is%20priva>

#### NEW QUESTION 38

- (Exam Topic 2)

Before promoting your new application code to production, you want to conduct testing across a variety of different users. Although this plan is risky, you want to test the new version of the application with production users and you want to control which users are forwarded to the new version of the application based on their operating system. If bugs are discovered in the new version, you want to roll back the newly deployed version of the application as quickly as possible.

What should you do?

- A. Deploy your application on Cloud Ru
- B. Use traffic splitting to direct a subset of user traffic to the new version based on the revision tag.
- C. Deploy your application on Google Kubernetes Engine with Anthos Service Mes
- D. Use traffic splitting to direct a subset of user traffic to the new version based on the user-agent header.
- E. Deploy your application on App Engin
- F. Use traffic splitting to direct a subset of user traffic to the new version based on the IP address.
- G. Deploy your application on Compute Engin
- H. Use Traffic Director to direct a subset of user traffic to the new version based on predefined weights.

**Answer:** B

#### NEW QUESTION 40

- (Exam Topic 2)

You have a container deployed on Google Kubernetes Engine. The container can sometimes be slow to launch, so you have implemented a liveness probe. You notice that the liveness probe occasionally fails on launch. What should you do?

- A. Add a startup probe.
- B. Increase the initial delay for the liveness probe.
- C. Increase the CPU limit for the container.
- D. Add a readiness probe.

**Answer:** B

#### Explanation:

<https://kubernetes.io/docs/tasks/configure-pod-container/configure-liveness-readiness-startup-probes/#configure>

#### NEW QUESTION 44

- (Exam Topic 2)

You are designing an application that consists of several microservices. Each microservice has its own RESTful API and will be deployed as a separate Kubernetes Service. You want to ensure that the consumers of these APIs aren't impacted when there is a change to your API, and also ensure that third-party systems aren't interrupted when new versions of the API are released. How should you configure the connection to the application following Google-recommended best practices?

- A. Use an Ingress that uses the API's URL to route requests to the appropriate backend.
- B. Leverage a Service Discovery system, and connect to the backend specified by the request.
- C. Use multiple clusters, and use DNS entries to route requests to separate versioned backends.
- D. Combine multiple versions in the same service, and then specify the API version in the POST request.

**Answer:** A

#### NEW QUESTION 45

- (Exam Topic 2)

You are deploying your application on a Compute Engine instance that communicates with Cloud SQL. You will use Cloud SQL Proxy to allow your application to communicate to the database using the service account associated with the application's instance. You want to follow the Google-recommended best practice of providing minimum access for the role assigned to the service account. What should you do?

- A. Assign the Project Editor role.
- B. Assign the Project Owner role.
- C. Assign the Cloud SQL Client role.
- D. Assign the Cloud SQL Editor role.

**Answer:** C

#### Explanation:

Reference: <https://cloud.google.com/sql/docs/mysql/sql-proxy>

#### NEW QUESTION 46

- (Exam Topic 2)

Your team is building an application for a financial institution. The application's frontend runs on Compute Engine, and the data resides in Cloud SQL and one Cloud Storage bucket. The application will collect data containing PII, which will be stored in the Cloud SQL database and the Cloud Storage bucket. You need to secure the PII data. What should you do?

- A. 1) Create the relevant firewall rules to allow only the frontend to communicate with the Cloud SQL database2) Using IAM, allow only the frontend service account to access the Cloud Storage bucket
- B. 1) Create the relevant firewall rules to allow only the frontend to communicate with the Cloud SQL database2) Enable private access to allow the frontend to access the Cloud Storage bucket privately
- C. 1) Configure a private IP address for Cloud SQL2) Use VPC-SC to create a service perimeter3) Add the Cloud SQL database and the Cloud Storage bucket to the same service perimeter
- D. 1) Configure a private IP address for Cloud SQL2) Use VPC-SC to create a service perimeter3) Add the Cloud SQL database and the Cloud Storage bucket to different service perimeters

**Answer: C**

#### NEW QUESTION 48

- (Exam Topic 2)

You need to deploy resources from your laptop to Google Cloud using Terraform. Resources in your Google Cloud environment must be created using a service account. Your Cloud Identity has the roles/iam.serviceAccountTokenCreator Identity and Access Management (IAM) role and the necessary permissions to deploy the resources using Terraform. You want to set up your development environment to deploy the desired resources following Google-recommended best practices. What should you do?

- A. 1) Download the service account's key file in JSON format, and store it locally on your laptop.2) Set the GOOGLE\_APPLICATION\_CREDENTIALS environment variable to the path of your downloaded key file.
- B. 1) Run the following command from a command line: gcloud config set auth/impersonate\_service\_account service-account-name@project.iam.gserviceaccount.com.2) Set the GOOGLE\_OAUTH\_ACCESS\_TOKEN environment variable to the value that is returned by the gcloud auth print-access-token command.
- C. 1) Run the following command from a command line: gcloud auth application-default login.2) In the browser window that opens, authenticate using your personal credentials.
- D. 1) Store the service account's key file in JSON format in Hashicorp Vault.2) Integrate Terraform with Vault to retrieve the key file dynamically, and authenticate to Vault using a short-lived access token.

**Answer: D**

#### Explanation:

<https://cloud.google.com/iam/docs/best-practices-for-managing-service-account-keys#file-system> Whenever possible, avoid storing service account keys on a file system. If you can't avoid storing keys on disk, make sure to restrict access to the key file, configure file access auditing, and encrypt the underlying disk.

<https://cloud.google.com/iam/docs/best-practices-for-managing-service-account-keys#software-keystore> In situations where using a hardware-based key store isn't viable, use a software-based key store to manage service account keys. Similar to hardware-based options, a software-based key store lets users or applications use service account keys without revealing the private key. Software-based key store solutions can help you control key access in a fine-grained manner and can also ensure that each key access is logged.

#### NEW QUESTION 52

- (Exam Topic 2)

Your teammate has asked you to review the code below, which is adding a credit to an account balance in Cloud Datastore. Which improvement should you suggest your teammate make?

```
public Entity creditAccount(long accountId, long
creditAmount) {
    Entity account = datastore.get
(keyFactory.newKey(accountId));
    account = Entity.builder(account).set(
        "balance", account.getLong("balance")
+ creditAmount).build()
    datastore.put(account);
    return account;
}
```

- A. Get the entity with an ancestor query.
- B. Get and put the entity in a transaction.
- C. Use a strongly consistent transactional database.
- D. Don't return the account entity from the function.

**Answer: A**

#### NEW QUESTION 53

- (Exam Topic 2)

You are monitoring a web application that is written in Go and deployed in Google Kubernetes Engine. You notice an increase in CPU and memory utilization. You need to determine which source code is consuming the most CPU and memory resources. What should you do?

- A. Download, install, and start the Snapshot Debugger agent in your V
- B. Take debug snapshots of the functions that take the longest time
- C. Review the call stack frame, and identify the local variables at that level in the stack.
- D. Import the Cloud Profiler package into your application, and initialize the Profiler agent
- E. Review the generated flame graph in the Google Cloud console to identify time-intensive functions.
- F. Import OpenTelemetry and Trace export packages into your application, and create the trace provider. Review the latency data for your application on the Trace overview page, and identify where bottlenecks are occurring.
- G. Create a Cloud Logging query that gathers the web application's log
- H. Write a Python script that calculates the difference between the timestamps from the beginning and the end of the application's longest functions to identify time-intensive functions.

**Answer:** B

#### NEW QUESTION 54

- (Exam Topic 2)

You are developing an ecommerce application that stores customer, order, and inventory data as relational tables inside Cloud Spanner. During a recent load test, you discover that Spanner performance is not scaling linearly as expected. Which of the following is the cause?

- A. The use of 64-bit numeric types for 32-bit numbers.
- B. The use of the STRING data type for arbitrary-precision values.
- C. The use of Version 1 UUIDs as primary keys that increase monotonically.
- D. The use of LIKE instead of STARTS\_WITH keyword for parameterized SQL queries.

**Answer:** C

#### NEW QUESTION 56

- (Exam Topic 2)

You are developing a new web application using Cloud Run and committing code to Cloud Source Repositories. You want to deploy new code in the most efficient way possible. You have already created a Cloud Build YAML file that builds a container and runs the following command: `gcloud run deploy`. What should you do next?

- A. Create a Pub/Sub topic to be notified when code is pushed to the repositor
- B. Create a Pub/Sub trigger that runs the build file when an event is published to the topic.
- C. Create a build trigger that runs the build file in response to a repository code being pushed to the development branch.
- D. Create a webhook build trigger that runs the build file in response to HTTP POST calls to the webhook URL.
- E. Create a Cron job that runs the following command every 24 hours: `gcloud builds submit`.

**Answer:** B

#### Explanation:

<https://cloud.google.com/build/docs/triggers>

Cloud Build uses build triggers to enable CI/CD automation. You can configure triggers to listen for incoming events, such as when a new commit is pushed to a repository or when a pull request is initiated, and then automatically execute a build when new events come in. You can also configure triggers to build code on any changes to your source repository or only on changes that match certain criteria.

#### NEW QUESTION 61

- (Exam Topic 2)

Your development team is using Cloud Build to promote a Node.js application built on App Engine from your staging environment to production. The application relies on several directories of photos stored in a Cloud Storage bucket named `webphotos-staging` in the staging environment. After the promotion, these photos must be available in a Cloud Storage bucket named `webphotos-prod` in the production environment. You want to automate the process where possible. What should you do?

- A) Manually copy the photos to `webphotos-prod`.
- B) Add a startup script in the application's `app.yaml` file to move the photos from `webphotos-staging` to `webphotos-prod`.
- C) Add a build step in the `cloudbuild.yaml` file before the promotion step with the arguments:

```
- name: gcr.io/cloud-builders/gsutil
  args: ['cp', '-r', 'gs://webphotos-staging',
'gs://webphotos-prod']
  waitFor: ['-']
```

- D) Add a build step in the `cloudbuild.yaml` file before the promotion step with the arguments:

```
- name: gcr.io/cloud-builders/gcloud
  args: ['cp', '-A', 'gs://webphotos-staging',
'gs://webphotos-prod']
  waitFor: ['-']
```

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

**Answer:** C

#### Explanation:

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

#### NEW QUESTION 65

- (Exam Topic 2)

Your application is running on Compute Engine and is showing sustained failures for a small number of requests. You have narrowed the cause down to a single Compute Engine instance, but the instance is unresponsive to SSH. What should you do next?

- A. Reboot the machine.
- B. Enable and check the serial port output.
- C. Delete the machine and create a new one.
- D. Take a snapshot of the disk and attach it to a new machine.

**Answer:** A



**NEW QUESTION 69**

- (Exam Topic 2)

You manage an application that runs in a Compute Engine instance. You also have multiple backend services executing in stand-alone Docker containers running in Compute Engine instances. The Compute Engine instances supporting the backend services are scaled by managed instance groups in multiple regions. You want your calling application to be loosely coupled. You need to be able to invoke distinct service implementations that are chosen based on the value of an HTTP header found in the request. Which Google Cloud feature should you use to invoke the backend services?

- A. Traffic Director
- B. Service Directory
- C. Anthos Service Mesh
- D. Internal HTTP(S) Load Balancing

**Answer:** D

**NEW QUESTION 70**

- (Exam Topic 2)

You are designing a deployment technique for your new applications on Google Cloud. As part of your deployment planning, you want to use live traffic to gather performance metrics for both new and existing applications. You need to test against the full production load prior to launch. What should you do?

- A. Use canary deployment
- B. Use blue/green deployment
- C. Use rolling updates deployment
- D. Use A/B testing with traffic mirroring during deployment

**Answer:** A

**Explanation:**

Reference: <https://cloud.google.com/architecture/application-deployment-and-testing-strategies>

**NEW QUESTION 75**

- (Exam Topic 2)

You want to migrate an on-premises container running in Knative to Google Cloud. You need to make sure that the migration doesn't affect your application's deployment strategy, and you want to use a fully managed service. Which Google Cloud service should you use to deploy your container?

- A. Cloud Run
- B. Compute Engine
- C. Google Kubernetes Engine
- D. App Engine flexible environment

**Answer:** A

**Explanation:**

<https://cloud.google.com/blog/products/serverless/knative-based-cloud-run-services-are-ga>

**NEW QUESTION 76**

- (Exam Topic 2)

You are deploying a single website on App Engine that needs to be accessible via the URL <http://www.altostrat.com/>. What should you do?

- A. Verify domain ownership with Webmaster Centra
- B. Create a DNS CNAME record to point to the App Engine canonical name [ghs.googlehosted.com](https://ghs.googlehosted.com).
- C. Verify domain ownership with Webmaster Centra
- D. Define an A record pointing to the single global App Engine IP address.
- E. Define a mapping in `dispatch.yaml` to point the domain [www.altostrat.com](http://www.altostrat.com) to your App Engine service. Create a DNS CNAME record to point to the App Engine canonical name [ghs.googlehosted.com](https://ghs.googlehosted.com).
- F. Define a mapping in `dispatch.yaml` to point the domain [www.altostrat.com](http://www.altostrat.com) to your App Engine service. Define an A record pointing to the single global App Engine IP address.

**Answer:** A

**Explanation:**

Reference: <https://cloud.google.com/appengine/docs/flexible/dotnet/mapping-custom-domains?hl=fa>

**NEW QUESTION 77**

- (Exam Topic 2)

You are in the final stage of migrating an on-premises data center to Google Cloud. You are quickly approaching your deadline, and discover that a web API is running on a server slated for decommissioning. You need to recommend a solution to modernize this API while migrating to Google Cloud. The modernized web API must meet the following requirements:

- Autoscales during high traffic periods at the end of each month
  - Written in Python 3.x
  - Developers must be able to rapidly deploy new versions in response to frequent code changes
- You want to minimize cost, effort, and operational overhead of this migration. What should you do?

- A. Modernize and deploy the code on App Engine flexible environment.
- B. Modernize and deploy the code on App Engine standard environment.
- C. Deploy the modernized application to an n1-standard-1 Compute Engine instance.
- D. Ask the development team to re-write the application to run as a Docker container on Google Kubernetes Engine.

**Answer:** B

**Explanation:**

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

**NEW QUESTION 82**

- (Exam Topic 2)

Your team detected a spike of errors in an application running on Cloud Run in your production project. The application is configured to read messages from Pub/Sub topic A, process the messages, and write the messages to topic B. You want to conduct tests to identify the cause of the errors. You can use a set of mock messages for testing. What should you do?

- A. Deploy the Pub/Sub and Cloud Run emulators on your local machine
- B. Deploy the application locally, and change the logging level in the application to DEBUG or INFO
- C. Write mock messages to topic A, and then analyze the logs.
- D. Use the gcloud CLI to write mock messages to topic
- E. Change the logging level in the application to DEBUG or INFO, and then analyze the logs.
- F. Deploy the Pub/Sub emulator on your local machine
- G. Point the production application to your local Pub/Sub topic
- H. Write mock messages to topic A, and then analyze the logs.
- I. Use the Google Cloud console to write mock messages to topic
- J. Change the logging level in the application to DEBUG or INFO, and then analyze the logs.

**Answer:** A

**NEW QUESTION 83**

- (Exam Topic 2)

You are a developer working on an internal application for payroll processing. You are building a component of the application that allows an employee to submit a timesheet, which then initiates several steps:

- An email is sent to the employee and manager, notifying them that the timesheet was submitted.
- A timesheet is sent to payroll processing for the vendor's API.
- A timesheet is sent to the data warehouse for headcount planning.

These steps are not dependent on each other and can be completed in any order. New steps are being considered and will be implemented by different development teams. Each development team will implement the error handling specific to their step. What should you do?

- A. Deploy a Cloud Function for each step that calls the corresponding downstream system to complete the required action.
- B. Create a Pub/Sub topic for each step
- C. Create a subscription for each downstream development team to subscribe to their step's topic.
- D. Create a Pub/Sub topic for timesheet submission
- E. Create a subscription for each downstream development team to subscribe to the topic.
- F. Create a timesheet microservice deployed to Google Kubernetes Engine
- G. The microservice calls each downstream step and waits for a successful response before calling the next step.

**Answer:** C

**NEW QUESTION 88**

- (Exam Topic 2)

Your company has a data warehouse that keeps your application information in BigQuery. The BigQuery data warehouse keeps 2 PBs of user data. Recently, your company expanded your user base to include EU users and needs to comply with these requirements:

Your company must be able to delete all user account information upon user request. All EU user data must be stored in a single region specifically for EU users. Which two actions should you take? (Choose two.)

- A. Use BigQuery federated queries to query data from Cloud Storage.
- B. Create a dataset in the EU region that will keep information about EU users only.
- C. Create a Cloud Storage bucket in the EU region to store information for EU users only.
- D. Re-upload your data using a Cloud Dataflow pipeline by filtering your user records out.
- E. Use DML statements in BigQuery to update/delete user records based on their requests.

**Answer:** CE

**Explanation:**

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

**NEW QUESTION 91**

- (Exam Topic 2)

Your development team has been tasked with maintaining a .NET legacy application. The application incurs occasional changes and was recently updated. Your goal is to ensure that the application provides consistent results while moving through the CI/CD pipeline from environment to environment. You want to minimize the cost of deployment while making sure that external factors and dependencies between hosting environments are not problematic. Containers are not yet approved in your organization. What should you do?

- A. Rewrite the application using .NET Core, and deploy to Cloud Run
- B. Use revisions to separate the environments.
- C. Use Cloud Build to deploy the application as a new Compute Engine image for each build
- D. Use this image in each environment.
- E. Deploy the application using MS Web Deploy, and make sure to always use the latest, patched MS Windows Server base image in Compute Engine.
- F. Use Cloud Build to package the application, and deploy to a Google Kubernetes Engine cluster
- G. Use namespaces to separate the environments.

**Answer:** B

**Explanation:**

[https://cloud.google.com/architecture/modernization-path-dotnet-applications-google-cloud#phase\\_1\\_rehost\\_in\\_](https://cloud.google.com/architecture/modernization-path-dotnet-applications-google-cloud#phase_1_rehost_in_)

<https://cloud.google.com/architecture/modernization-path-dotnet-applications-google-cloud>

**NEW QUESTION 95**

- (Exam Topic 2)

You are developing an application that reads credit card data from a Pub/Sub subscription. You have written code and completed unit testing. You need to test the Pub/Sub integration before deploying to Google Cloud. What should you do?

- A. Create a service to publish messages, and deploy the Pub/Sub emulator
- B. Generate random content in the publishing service, and publish to the emulator.
- C. Create a service to publish messages to your applicatio
- D. Collect the messages from Pub/Sub in production, and replay them through the publishing service.
- E. Create a service to publish messages, and deploy the Pub/Sub emulator
- F. Collect the messages from Pub/Sub in production, and publish them to the emulator.
- G. Create a service to publish messages, and deploy the Pub/Sub emulator
- H. Publish a standard set of testing messages from the publishing service to the emulator.

**Answer:** D

**NEW QUESTION 97**

- (Exam Topic 2)

You recently migrated an on-premises monolithic application to a microservices application on Google Kubernetes Engine (GKE). The application has dependencies on backend services on-premises, including a CRM system and a MySQL database that contains personally identifiable information (PII). The backend services must remain on-premises to meet regulatory requirements.

You established a Cloud VPN connection between your on-premises data center and Google Cloud. You notice that some requests from your microservices application on GKE to the backend services are failing due to latency issues caused by fluctuating bandwidth, which is causing the application to crash. How should you address the latency issues?

- A. Use Memorystore to cache frequently accessed PII data from the on-premises MySQL database
- B. Use Istio to create a service mesh that includes the microservices on GKE and the on-premises services
- C. Increase the number of Cloud VPN tunnels for the connection between Google Cloud and the on-premises services
- D. Decrease the network layer packet size by decreasing the Maximum Transmission Unit (MTU) value from its default value on Cloud VPN

**Answer:** C

**Explanation:**

<https://cloud.google.com/network-connectivity/docs/vpn/concepts/choosing-networks-routing#route-alignment>

**NEW QUESTION 99**

- (Exam Topic 2)

You have an application deployed in production. When a new version is deployed, some issues don't arise until the application receives traffic from users in production. You want to reduce both the impact and the number of users affected.

Which deployment strategy should you use?

- A. Blue/green deployment
- B. Canary deployment
- C. Rolling deployment
- D. Recreate deployment

**Answer:** A

**Explanation:**

Reference: <https://thenewstack.io/deployment-strategies/>

**NEW QUESTION 104**

- (Exam Topic 2)

You have an application deployed in production. When a new version is deployed, you want to ensure that all production traffic is routed to the new version of your application. You also want to keep the previous version deployed so that you can revert to it if there is an issue with the new version.

Which deployment strategy should you use?

- A. Blue/green deployment
- B. Canary deployment
- C. Rolling deployment
- D. Recreate deployment

**Answer:** A

**NEW QUESTION 106**

- (Exam Topic 2)

You are developing an application that consists of several microservices running in a Google Kubernetes Engine cluster. One microservice needs to connect to a third-party database running on-premises. You need to store credentials to the database and ensure that these credentials can be rotated while following security best practices. What should you do?

- A. Store the credentials in a sidecar container proxy, and use it to connect to the third-party database.
- B. Configure a service mesh to allow or restrict traffic from the Pods in your microservice to the database.
- C. Store the credentials in an encrypted volume mount, and associate a Persistent Volume Claim with the client Pod.
- D. Store the credentials as a Kubernetes Secret, and use the Cloud Key Management Service plugin to handle encryption and decryption.

**Answer:** D

**Explanation:**

<https://cloud.google.com/kubernetes-engine/docs/how-to/encrypting-secrets>

By default, Google Kubernetes Engine (GKE) encrypts customer content stored at rest, including Secrets. GKE handles and manages this default encryption for you without any additional action on your part.

Application-layer secrets encryption provides an additional layer of security for sensitive data, such as Secrets, stored in etcd. Using this functionality, you can use a key managed with Cloud KMS to encrypt data at the application layer. This encryption protects against attackers who gain access to an offline copy of etcd.

**NEW QUESTION 108**

- (Exam Topic 2)

You are developing an HTTP API hosted on a Compute Engine virtual machine instance that needs to be invoked by multiple clients within the same Virtual Private Cloud (VPC). You want clients to be able to get the IP address of the service. What should you do?

- A. Reserve a static external IP address and assign it to an HTTP(S) load balancing service's forwarding rule. Clients should use this IP address to connect to the service.
- B. Reserve a static external IP address and assign it to an HTTP(S) load balancing service's forwarding rule. Then, define an A record in Cloud DN
- C. Clients should use the name of the A record to connect to the service.
- D. Ensure that clients use Compute Engine internal DNS by connecting to the instance name with the url [https://\[INSTANCE\\_NAME\].\[ZONE\].c.\[PROJECT\\_ID\].internal/](https://[INSTANCE_NAME].[ZONE].c.[PROJECT_ID].internal/).
- E. Ensure that clients use Compute Engine internal DNS by connecting to the instance name with the url [https://\[API\\_NAME\]/\[API\\_VERSION\]/](https://[API_NAME]/[API_VERSION]/).

**Answer:** D

**NEW QUESTION 112**

- (Exam Topic 2)

Your company has created an application that uploads a report to a Cloud Storage bucket. When the report is uploaded to the bucket, you want to publish a message to a Cloud Pub/Sub topic. You want to implement a solution that will take a small amount of effort to implement. What should you do?

- A. Configure the Cloud Storage bucket to trigger Cloud Pub/Sub notifications when objects are modified.
- B. Create an App Engine application to receive the file; when it is received, publish a message to the Cloud Pub/Sub topic.
- C. Create a Cloud Function that is triggered by the Cloud Storage bucket.
- D. In the Cloud Function, publish a message to the Cloud Pub/Sub topic.
- E. Create an application deployed in a Google Kubernetes Engine cluster to receive the file; when it is received, publish a message to the Cloud Pub/Sub topic.

**Answer:** C

**Explanation:**

<https://cloud.google.com/storage/docs/pubsub-notifications>

**NEW QUESTION 115**

- (Exam Topic 2)

Your organization has recently begun an initiative to replatform their legacy applications onto Google Kubernetes Engine. You need to decompose a monolithic application into microservices. Multiple instances have read and write access to a configuration file, which is stored on a shared file system. You want to minimize the effort required to manage this transition, and you want to avoid rewriting the application code. What should you do?

- A. Create a new Cloud Storage bucket, and mount it via FUSE in the container.
- B. Create a new persistent disk, and mount the volume as a shared PersistentVolume.
- C. Create a new Filestore instance, and mount the volume as an NFS PersistentVolume.
- D. Create a new ConfigMap and volumeMount to store the contents of the configuration file.

**Answer:** D

**Explanation:**

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

ConfigMaps bind non-sensitive configuration artifacts such as configuration files, command-line arguments, and environment variables to your Pod containers and system components at runtime.

A ConfigMap separates your configurations from your Pod and components, which helps keep your workloads portable. This makes their configurations easier to change and manage, and prevents hardcoding configuration data to Pod specifications.

**NEW QUESTION 118**

- (Exam Topic 2)

You are designing a schema for a Cloud Spanner customer database. You want to store a phone number array field in a customer table. You also want to allow users to search customers by phone number. How should you design this schema?

- A. Create a table named Customer
- B. Add an Array field in a table that will hold phone numbers for the customer.
- C. Create a table named Customer
- D. Create a table named Phone
- E. Add a CustomerId field in the Phones table to find the CustomerId from a phone number.
- F. Create a table named Customer
- G. Add an Array field in a table that will hold phone numbers for the customer
- H. Create a secondary index on the Array field.
- I. Create a table named Customers as a parent table
- J. Create a table named Phones, and interleave this table into the Customer table
- K. Create an index on the phone number field in the Phones table.

**Answer:** C



**NEW QUESTION 120**

- (Exam Topic 2)

Your application is controlled by a managed instance group. You want to share a large read-only data set between all the instances in the managed instance group. You want to ensure that each instance can start quickly and can access the data set via its filesystem with very low latency. You also want to minimize the Total cost of the solution. What should you do?

- A. Move the data to a Cloud Storage bucket, and mount the bucket on the filesystem using Cloud Storage FUSE.
- B. Move the data to a Cloud Storage bucket, and copy the data to the boot disk of the instance via a startup script.
- C. Move the data to a Compute Engine persistent disk, and attach the disk in read-only mode to multiple Compute Engine virtual machine instances.
- D. Move the data to a Compute Engine persistent disk, take a snapshot, create multiple disks from the snapshot, and attach each disk to its own instance.

**Answer: C**

**NEW QUESTION 122**

- (Exam Topic 2)

You have an application running in App Engine. Your application is instrumented with Stackdriver Trace. The /product-details request reports details about four known unique products at /sku-details as shown below. You want to reduce the time it takes for the request to complete. What should you do?

**Timeline**

- A. Increase the size of the instance class.
- B. Change the Persistent Disk type to SSD.
- C. Change /product-details to perform the requests in parallel.
- D. Store the /sku-details information in a database, and replace the webservice call with a database query.

**Answer: C**

**NEW QUESTION 125**

- (Exam Topic 2)

You are building a CI/CD pipeline that consists of a version control system, Cloud Build, and Container Registry. Each time a new tag is pushed to the repository, a Cloud Build job is triggered, which runs unit tests on the new code builds a new Docker container image, and pushes it into Container Registry. The last step of your pipeline should deploy the new container to your production Google Kubernetes Engine (GKE) cluster. You need to select a tool and deployment strategy that meets the following requirements:

- Zero downtime is incurred
  - Testing is fully automated
  - Allows for testing before being rolled out to users
  - Can quickly rollback if needed
- What should you do?

- A. Trigger a Spinnaker pipeline configured as an A/B test of your new code and, if it is successful, deploy the container to production.
- B. Trigger a Spinnaker pipeline configured as a canary test of your new code and, if it is successful, deploy the container to production.
- C. Trigger another Cloud Build job that uses the Kubernetes CLI tools to deploy your new container to your GKE cluster, where you can perform a canary test.
- D. Trigger another Cloud Build job that uses the Kubernetes CLI tools to deploy your new container to your GKE cluster, where you can perform a shadow test.

**Answer: D**

**Explanation:**

[https://cloud.google.com/architecture/implementing-deployment-and-testing-strategies-on-gke#perform\\_a\\_shadow\\_test](https://cloud.google.com/architecture/implementing-deployment-and-testing-strategies-on-gke#perform_a_shadow_test) With a shadow test, you test the new version of your application by mirroring user traffic from the current application version without impacting the user requests.

**NEW QUESTION 128**

- (Exam Topic 2)

You are building an API that will be used by Android and iOS apps. The API must:

- Support HTTPs
  - Minimize bandwidth cost
  - Integrate easily with mobile apps
- Which API architecture should you use?

- A. RESTful APIs
- B. MQTT for APIs
- C. gRPC-based APIs
- D. SOAP-based APIs

**Answer: A**

**Explanation:**

Reference: <https://www.devteam.space/blog/how-to-build-restful-api-for-your-mobile-app/>

**NEW QUESTION 130**

- (Exam Topic 2)

You want to use the Stackdriver Logging Agent to send an application's log file to Stackdriver from a Compute Engine virtual machine instance. After installing the Stackdriver Logging Agent, what should you do first?

- A. Enable the Error Reporting API on the project.
- B. Grant the instance full access to all Cloud APIs.
- C. Configure the application log file as a custom source.
- D. Create a Stackdriver Logs Export Sink with a filter that matches the application's log entries.

**Answer: B**

**NEW QUESTION 133**

- (Exam Topic 2)

Your operations team has asked you to create a script that lists the Cloud Bigtable, Memorystore, and Cloud SQL databases running within a project. The script should allow users to submit a filter expression to limit the results presented. How should you retrieve the data?

- A. Use the HBase API, Redis API, and MySQL connection to retrieve database list
- B. Combine the results, and then apply the filter to display the results
- C. Use the HBase API, Redis API, and MySQL connection to retrieve database list
- D. Filter the results individually, and then combine them to display the results
- E. Run `gcloud bigtable instances list`, `gcloud redis instances list`, and `gcloud sql databases list`
- F. Use a filter within the application, and then display the results
- G. Run `gcloud bigtable instances list`, `gcloud redis instances list`, and `gcloud sql databases list`
- H. Use `--filter` flag with each command, and then display the results

**Answer: D**

**Explanation:**

<https://cloud.google.com/sdk/gcloud/reference/topic/filters>

Most gcloud commands return a list of resources on success. By default they are pretty-printed on the standard output. The `--format=NAME[ATTRIBUTES](PROJECTION)` and `--filter=EXPRESSION` flags along with projections can be used to format and change the default output to a more meaningful result. Use the `--format` flag to change the default output format of a command. For details run `$ gcloud topic formats`.

**NEW QUESTION 135**

- (Exam Topic 2)

You are planning to migrate a MySQL database to the managed Cloud SQL database for Google Cloud. You have Compute Engine virtual machine instances that will connect with this Cloud SQL instance. You do not want to whitelist IPs for the Compute Engine instances to be able to access Cloud SQL. What should you do?

- A. Enable private IP for the Cloud SQL instance.
- B. Whitelist a project to access Cloud SQL, and add Compute Engine instances in the whitelisted project.
- C. Create a role in Cloud SQL that allows access to the database from external instances, and assign the Compute Engine instances to that role.
- D. Create a CloudSQL instance on one project
- E. Create Compute engine instances in a different project. Create a VPN between these two projects to allow internal access to CloudSQL.

**Answer: C**

**Explanation:**

Reference: <https://cloud.google.com/sql/docs/mysql/connect-external-app>

**NEW QUESTION 139**

- (Exam Topic 2)

You have an on-premises application that authenticates to the Cloud Storage API using a user-managed service account with a user-managed key. The application connects to Cloud Storage using Private Google Access over a Dedicated Interconnect link. You discover that requests from the application to access objects in the Cloud Storage bucket are failing with a 403 Permission Denied error code. What is the likely cause of this issue?

- A. The folder structure inside the bucket and object paths have changed.
- B. The permissions of the service account's predefined role have changed.
- C. The service account key has been rotated but not updated on the application server.
- D. The Interconnect link from the on-premises data center to Google Cloud is experiencing a temporary outage.

**Answer: C**

**NEW QUESTION 142**

- (Exam Topic 2)

Your company stores their source code in a Cloud Source Repositories repository. Your company wants to build and test their code on each source code commit to the repository and requires a solution that is managed and has minimal operations overhead. Which method should they use?

- A. Use Cloud Build with a trigger configured for each source code commit.
- B. Use Jenkins deployed via the Google Cloud Platform Marketplace, configured to watch for source code commits.
- C. Use a Compute Engine virtual machine instance with an open source continuous integration tool, configured to watch for source code commits.
- D. Use a source code commit trigger to push a message to a Cloud Pub/Sub topic that triggers an App Engine service to build the source code.

**Answer: A**

**Explanation:**

[https://cloud.google.com/build/docs/automating-builds/create-manage-triggers#:~:text=A%20Cloud%20Build%](https://cloud.google.com/build/docs/automating-builds/create-manage-triggers#:~:text=A%20Cloud%20Build%20)

**NEW QUESTION 144**

- (Exam Topic 2)

You manage your company's ecommerce platform's payment system, which runs on Google Cloud. Your company must retain user logs for 1 year for internal auditing purposes and for 3 years to meet compliance requirements. You need to store new user logs on Google Cloud to minimize on-premises storage usage and ensure that they are easily searchable. You want to minimize effort while ensuring that the logs are stored correctly. What should you do?

- A. Store the logs in a Cloud Storage bucket with bucket lock turned on.
- B. Store the logs in a Cloud Storage bucket with a 3-year retention period.
- C. Store the logs in Cloud Logging as custom logs with a custom retention period.
- D. Store the logs in a Cloud Storage bucket with a 1-year retention period.
- E. After 1 year, move the logs to another bucket with a 2-year retention period.

**Answer: C**

**Explanation:**

<https://cloud.google.com/logging/docs/buckets#custom-retention>

**NEW QUESTION 148**

.....

## Thank You for Trying Our Product

\* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

\* One year free update

You can enjoy free update one year. 24x7 online support.

\* Trusted by Millions

We currently serve more than 30,000,000 customers.

\* Shop Securely

All transactions are protected by VeriSign!

**100% Pass Your Professional-Cloud-Developer Exam with Our Prep Materials Via below:**

<https://www.certleader.com/Professional-Cloud-Developer-dumps.html>