Cloud Developer Certified Official Practice Test Professional-Cloud-Developer - Nov-2023 [Q97-Q113

Cloud Developer Certified Official Practice Test Professional-Cloud-Developer - Nov-2023 Ace Google Professional-Cloud-Developer Certification with Actual Questions Nov 27, 2023 Updated

Google Professional-Cloud-Developer certification is recognized globally and is highly valued in the cloud computing industry. Google Certified Professional - Cloud Developer certification provides a competitive edge to software developers and helps them to stand out in the job market. In addition, this certification also helps organizations to identify and hire skilled professionals who can develop and deploy applications on GCP. By becoming a Google Certified Professional - Cloud Developer, individuals can demonstrate their commitment to staying current with the latest cloud technologies and their ability to design, develop, and manage cloud applications on GCP.

Google Professional-Cloud-Developer certification exam is a credential offered by Google Cloud that validates a professional's expertise in designing, developing, and deploying applications on the Google Cloud Platform. Google Certified Professional - Cloud

Developer certification is intended for developers who work with Google Cloud technologies and are responsible for creating cloud-native applications. Professional-Cloud-Developer exam is designed to test one's ability to use Google Cloud's tools and services to build scalable and fault-tolerant applications that meet business requirements.

NEW QUESTION 97

You are developing a corporate tool on Compute Engine for the finance department, which needs to authenticate users and verify that they are in the finance department. All company employees use G Suite.

What should you do?
* Enable Cloud Identity-Aware Proxy on the HTTP(s) load balancer and restrict access to a Google Group containing users in the finance department. Verify the provided JSON Web Token within the application.
* Enable Cloud Identity-Aware Proxy on the HTTP(s) load balancer and restrict access to a Google Group containing users in the finance department. Issue client-side certificates to everybody in the finance team and verify the certificates in the application.
* Configure Cloud Armor Security Policies to restrict access to only corporate IP address ranges. Verify the provided JSON Web Token within the application.
* Configure Cloud Armor Security Policies to restrict access to only corporate IP address ranges. Issue client side certificates to everybody in the finance team and verify the certificates in the application.

NEW QUESTION 98

Your company’s product team has a new requirement based on customer demand to autoscale your stateless and distributed service running in a Google Kubernetes Engine (GKE) duster. You want to find a solution that minimizes changes because this feature will go live in two weeks. What should you do?
* Deploy a Vertical Pod Autoscaler, and scale based on the CPU load.
* Deploy a Vertical Pod Autoscaler, and scale based on a custom metric.
* Deploy a Horizontal Pod Autoscaler, and scale based on the CPU toad.
* Deploy a Horizontal Pod Autoscaler, and scale based on a custom metric.

Explanation

https://cloud.google.com/kubernetes-engine/docs/concepts/horizontalpodautoscaler The Horizontal Pod Autoscaler changes the shape of your Kubernetes workload by automatically increasing or decreasing the number of Pods in response to the workload’s CPU or memory consumption, or in response to custom metrics reported from within Kubernetes or external metrics from sources outside of your cluster.

NEW QUESTION 99

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

Explanation

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

NEW QUESTION 100

You have deployed an HTTP(s) Load Balancer with the gcloud commands shown below.

```
export NAME-load-balancer

# create network
gcloud compute networks create ${NAME}

# add instance
gcloud compute instances create ${NAME}-backend-instance-1  --subnet ${NAME} --no address

# create the instance group
gcloud compute instance-groups unmanaged create ${NAME}-i
gcloud compute instance-groups unmanaged set-named-ports ${NAME}-i --named-ports http:80
gcloud compute instance-groups unmanaged add-instances ${NAME}-i --instances ${NAME}-instance-1

# configure health checks
gcloud compute health-checks create http ${NAME}-http-hc --port 80

# create backend service
gcloud compute backend-services create ${NAME}-http-bes --health-checks ${NAME}-http-hc --protocol HTTP
--global
gcloud compute backend-services add-backend ${NAME}-http-bes --instance-group ${NAME}-i --balancing-mod
100000 --capacity-scaler 1.0 --global --instance-group-zone us-east1-d

# create urls maps and forwarding rule
gcloud compute url-maps create ${NAME}-http-urlmap --default-service ${NAME}-http-bes
gcloud compute target-http-proxies create ${NAME}-http-proxy --url-map ${NAME}-http-urlmap
gcloud compute forwarding-rules create ${NAME}-http-fw --global --ip-protocol ICP --target-http-proxy
--ports 80
```


Health checks to port 80 on the Compute Engine virtual machine instance are failing and no traffic is sent to your instances. You want to resolve the problem.

Which commands should you run?
* gcloud compute instances add-access-config ${NAME}-backend-instance-1
* gcloud compute instances add-tags ${NAME}-backend-instance-1 –tags http-server
* gcloud compute firewall-rules create allow-lb –network load-balancer –allow tcp –source-ranges 130.211.0.0/22,35.191.0.0/16 –direction INGRESS
* gcloud compute firewall-rules create allow-lb –network load-balancer –allow tcp –destination-ranges 130.211.0.0/22,35.191.0.0/16 –direction EGRESS

Reference: https://cloud.google.com/vpc/docs/special-configurations

NEW QUESTION 101

You have recently instrumented a new application with OpenTelemetry, and you want to check the latency of your application requests in Trace. You want to ensure that a specific request is always traced. What should you do?
* Wait 10 minutes, then verify that Trace captures those types of requests automatically.
* Write a custom script that sends this type of request repeatedly from your dev project.

* Use the Trace API to apply custom attributes to the trace.
* Add the X-Cloud-Trace-Context header to the request with the appropriate parameters.

Explanation

https://cloud.google.com/trace/docs/setup#force-trace

Cloud Trace doesn’t sample every request. To force a specific request to be traced, add an X-Cloud-Trace-Context header to the request.

NEW QUESTION 102

You migrated your applications to Google Cloud Platform and kept your existing monitoring platform. You now find that your notification system is too slow for time critical problems.

What should you do?
* Replace your entire monitoring platform with Stackdriver.
* Install the Stackdriver agents on your Compute Engine instances.
* Use Stackdriver to capture and alert on logs, then ship them to your existing platform.
* Migrate some traffic back to your old platform and perform AB testing on the two platforms concurrently.

Explanation/Reference: https://cloud.google.com/monitoring/

NEW QUESTION 103

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?
* Retry the failures in batch after a set number of failures is logged.
* Retry each failure at a set time interval up to a maximum number of times.
* Retry each failure at increasing time intervals up to a maximum number of tries.
* Retry each failure at decreasing time intervals up to a maximum number of tries.

Explanation

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

NEW QUESTION 104

Your application requires service accounts to be authenticated to GCP products via credentials stored on its host Compute Engine virtual machine instances. You want to distribute these credentials to the host instances as securely as possible. What should you do?
* Use HTTP signed URLs to securely provide access to the required resources.
* Use the instance’s service account Application Default Credentials to authenticate to the required resources.
* Generate a P12 file from the GCP Console after the instance is deployed, and copy the credentials to the host instance before starting the application.
* Commit the credential JSON file into your application’s source repository, and have your CI/CD process package it with the software that is deployed to the instance.

NEW QUESTION 105

You are building a highly available and globally accessible application that will serve static content to users.

You need to configure the storage and serving components. You want to minimize management overhead and latency while maximizing reliability for users. What should you do?
* 1) Create a managed instance group. Replicate the static content across the virtual machines (VMs)

2) Create an external HTTP(S) load balancer.

3) Enable Cloud CDN, and send traffic to the managed instance group.
* 1) Create an unmanaged instance group. Replicate the static content across the VMs.

2) Create an external HTTP(S) load balancer

3) Enable Cloud CDN, and send traffic to the unmanaged instance group.
* 1) Create a Standard storage class, regional Cloud Storage bucket. Put the static content in the bucket

2) Reserve an external IP address, and create an external HTTP(S) load balancer

3) Enable Cloud CDN, and send traffic to your backend bucket
* 1) Create a Standard storage class, multi-regional Cloud Storage bucket. Put the static content in the bucket.

2) Reserve an external IP address, and create an external HTTP(S) load balancer.

3) Enable Cloud CDN, and send traffic to your backend bucket.

NEW QUESTION 106

Your team is developing a new application using a PostgreSQL database and Cloud Run. You are responsible for ensuring that all traffic is kept private on Google Cloud. You want to use managed services and follow Google-recommended best practices. What should you do?
* 1. Enable Cloud SQL and Cloud Run in the same project.

2. Configure a private IP address for Cloud SQL. Enable private services access.

3. Create a Serverless VPC Access connector.

4. Configure Cloud Run to use the connector to connect to Cloud SQL.
* 1. Install PostgreSQL on a Compute Engine virtual machine (VM), and enable Cloud Run in the same project.

2. Configure a private IP address for the VM. Enable private services access.

3. Create a Serverless VPC Access connector.

4. Configure Cloud Run to use the connector to connect to the VM hosting PostgreSQL.
* 1. Use Cloud SQL and Cloud Run in different projects.

2. Configure a private IP address for Cloud SQL. Enable private services access.

3. Create a Serverless VPC Access connector.

4. Set up a VPN connection between the two projects. Configure Cloud Run to use the connector to connect to Cloud SQL.
* 1. Install PostgreSQL on a Compute Engine VM, and enable Cloud Run in different projects.

2. Configure a private IP address for the VM. Enable private services access.

3. Create a Serverless VPC Access connector.

4. Set up a VPN connection between the two projects. Configure Cloud Run to use the connector to access the VM hosting PostgreSQL
https://cloud.google.com/sql/docs/postgres/connect-run#private-ip

NEW QUESTION 107

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.yami 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: ['-']
```

* Option A
* Option B
* Option C
* Option D

Explanation

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

NEW QUESTION 108

HipLocal wants to reduce the number of on-call engineers and eliminate manual scaling.

Which two services should they choose? (Choose two.)
* Use Google App Engine services.
* Use serverless Google Cloud Functions.
* Use Knative to build and deploy serverless applications.
* Use Google Kubernetes Engine for automated deployments.
* Use a large Google Compute Engine cluster for deployments.

NEW QUESTION 109

You are developing a corporate tool on Compute Engine for the finance department, which needs to authenticate users and verify that they are in the finance department. All company employees use G Suite.

What should you do?
* Enable Cloud Identity-Aware Proxy on the HTTP(s) load balancer and restrict access to a Google Group containing users in the finance department. Verify the provided JSON Web Token within the application.
* Enable Cloud Identity-Aware Proxy on the HTTP(s) load balancer and restrict access to a Google Group containing users in the finance department. Issue client-side certificates to everybody in the finance team and verify the certificates in the application.
* Configure Cloud Armor Security Policies to restrict access to only corporate IP address ranges. Verify the provided JSON Web Token within the application.
* Configure Cloud Armor Security Policies to restrict access to only corporate IP address ranges. Issue client side certificates to everybody in the finance team and verify the certificates in the application.
Explanation

https://cloud.google.com/iap/docs/signed-headers-howto#securing_iap_headers
(https://cloud.google.com/endpoints/docs/openapi/authenticating-users-google-id).

https://cloud.google.com/armor/docs/security-policy-overview#:~:text=Google%20Cloud%20Armor%20security

“Google Cloud Armor security policies protect your application by providing Layer 7 filtering and by scrubbing incoming requests for common web attacks or other Layer 7 attributes to potentially block traffic before it reaches your load balanced backend services or backend buckets”

NEW QUESTION 110

You have two tables in an ANSI-SQL compliant database with identical columns that you need to quickly combine into a single table, removing duplicate rows from the result set.

What should you do?
* Use the JOIN operator in SQL to combine the tables.
* Use nested WITH statements to combine the tables.
* Use the UNION operator in SQL to combine the tables.
* Use the UNION ALL operator in SQL to combine the tables.
Reference: https://www.techonthenet.com/sql/union_all.php

NEW QUESTION 111

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?
* 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.
* 1) Run the following command from a command line: gcloud config set

auth/impersonate_service_account service-account-name@project.iam.gserviceacccount.com.

2) Set the GOOGLE_OAUTH_ACCESS_TOKEN environment variable to the value that is returned by the gcloud auth print-access-token command.
* 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.
* 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.
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 112

You are load testing your server application. During the first 30 seconds, you observe that a previously inactive Cloud Storage bucket is now servicing 2000 write requests per second and 7500 read requests per second.

Your application is now receiving intermittent 5xx and 429 HTTP responses from the Cloud Storage JSON API as the demand escalates. You want to decrease the failed responses from the Cloud Storage API.

What should you do?
* Distribute the uploads across a large number of individual storage buckets.
* Use the XML API instead of the JSON API for interfacing with Cloud Storage.
* Pass the HTTP response codes back to clients that are invoking the uploads from your application.
* Limit the upload rate from your application clients so that the dormant bucket’s peak request rate is reached more gradually.

NEW QUESTION 113

Your website is deployed on Compute Engine. Your marketing team wants to test conversion rates between 3 different website designs.

Which approach should you use?

* Deploy the website on App Engine and use traffic splitting.
* Deploy the website on App Engine as three separate services.
* Deploy the website on Cloud Functions and use traffic splitting.
* Deploy the website on Cloud Functions as three separate functions.

Who should take the Google Professional Cloud Developer exam

Individuals should pursue the **Google Professional Cloud Developer Exam** if they want to demonstrate their expertise and ability to design highly scalable, available, and reliable cloud-native applications and deploy applications. It's perfect for solutions and/or enterprise architects, systems administrators or operations team members or simply any professional who wants in on this specific area of IT and cloud. A Professional Cloud Developer should have skills at producing meaningful metrics and logs to debug and trace code and proficiency with at least one general-purpose programming language.

Try Free and Start Using Realistic Verified Professional-Cloud-Developer Dumps Instantly.:
https://www.validexam.com/Professional-Cloud-Developer-latest-dumps.html]