Google

Exam Questions Professional-Cloud-Developer

Google Certified Professional - Cloud Developer

About Exambible

Your Partner of IT Exam

Found in 1998

Exambible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, Exambible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

   All examinations will be up to date.

* 24/7 Quality Support

   We will provide service round the clock.

* 100% Pass Rate

   Our guarantee that you will pass the exam.

* Unique Gurantee

   If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

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)
Which database should HipLocal use for storing user activity?

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

Answer: A

NEW QUESTION 3
- (Exam Topic 1)
For this question, refer to the HipLocal case study.
HipLocal's application uses Cloud Client Libraries to interact with Google Cloud. HipLocal needs to configure authentication and authorization in the Cloud Client Libraries to implement least privileged access for the application. What should they do?

A. Create an API ke
B. Use the API key to interact with Google Cloud.
C. Use the default compute service account to interact with Google Cloud.
D. Create a service account for the applicatio
E. Export and deploy the private key for the applicatio
F. Use the service account to interact with Google Cloud.
G. Create a service account for the application and for each Google Cloud API used by the application.Export and deploy the private keys used by the applicatio
H. Use the service account with one Google Cloud API to interact with Google Cloud.

Answer: A

NEW QUESTION 4
- (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 5
- (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 6
- (Exam Topic 2)
You are developing a single-player mobile game backend that has unpredictable traffic patterns as users interact with the game throughout the day and night. You want to optimize costs by ensuring that you have enough resources to handle requests, but minimize over-provisioning. You also want the system to handle traffic spikes efficiently. Which compute platform should you use?

A. Cloud Run
B. Compute Engine with managed instance groups
C. Compute Engine with unmanaged instance groups
D. Google Kubernetes Engine using cluster autoscaling

Answer: A

NEW QUESTION 7
- (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 8
- (Exam Topic 2)
Users are complaining that your Cloud Run-hosted website responds too slowly during traffic spikes. You want to provide a better user experience during traffic peaks. What should you do?

A. Read application configuration and static data from the database on application startup.
B. Package application configuration and static data into the application image during build time.
C. Perform as much work as possible in the background after the response has been returned to the user.
D. Ensure that timeout exceptions and errors cause the Cloud Run instance to exit quickly so a replacement instance can be started.

Answer: C

NEW QUESTION 9
- (Exam Topic 2)
You are responsible for deploying a new API. That API will have three different URL paths:
• https://yourcompany.com/students
• https://yourcompany.com/teachers
• https://yourcompany.com/classes
You need to configure each API URL path to invoke a different function in your code. What should you do?

A. Create one Cloud Function as a backend service exposed using an HTTPS load balancer.
B. Create three Cloud Functions exposed directly.
C. Create one Cloud Function exposed directly.
D. Create three Cloud Functions as three backend services exposed using an HTTPS load balancer.

Answer: D

Explanation:
https://cloud.google.com/load-balancing/docs/https/setup-global-ext-https-serverless

NEW QUESTION 10
- (Exam Topic 2)
You made a typo in a low-level Linux configuration file that prevents your Compute Engine instance from booting to a normal run level. You just created the Compute Engine instance today and have done no other maintenance on it, other than tweaking files. How should you correct this error?

A. Download the file using scp, change the file, and then upload the modified version
B. Configure and log in to the Compute Engine instance through SSH, and change the file
C. Configure and log in to the Compute Engine instance through the serial port, and change the file
D. Configure and log in to the Compute Engine instance using a remote desktop client, and change the file

Answer: C

Explanation:
https://cloud.google.com/compute/docs/troubleshooting/troubleshooting-using-serial-console

NEW QUESTION 10
- (Exam Topic 2)
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?

A. 1) Enable Cloud SQL and Cloud Run in the same project.2) Configure a private IP address for Cloud SQ
B. Enable private services access.3) Create a Serverless VPC Access connector.4) Configure Cloud Run to use the connector to connect to Cloud SQL.
C. 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 V
D. 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.
E. 1) Use Cloud SQL and Cloud Run in different projects.2) Configure a private IP address for Cloud SQ
F. Enable private services access.3) Create a Serverless VPC Access connector.4) Set up a VPN connection between the two project
G. Configure Cloud Run to use the connector to connect to Cloud SQL.
H. 1) Install PostgreSQL on a Compute Engine VM, and enable Cloud Run in different projects.2) Configure a private IP address for the V
I. Enable private services access.3) Create a Serverless VPC Access connector.4) Set up a VPN connection between the two project
J. Configure Cloud Run to use the connector to access the VM hosting PostgreSQL

Answer: A

Explanation:
https://cloud.google.com/sql/docs/postgres/connect-run#private-ip

NEW QUESTION 15
- (Exam Topic 2)
You have an application in production. It is deployed on Compute Engine virtual machine instances controlled by a managed instance group. Traffic is routed to the instances via a HTTP(s) load balancer. Your users are unable to access your application. You want to implement a monitoring technique to alert you when the application is unavailable.
Which technique should you choose?

A. Smoke tests
B. Stackdriver uptime checks
C. Cloud Load Balancing - heath checks
D. Managed instance group - heath checks

Answer: B

Explanation:
Reference: https://medium.com/google-cloud/stackdriver-monitoring-automation-part-3-uptime-checks- 476b8507f59c

NEW QUESTION 20
- (Exam Topic 2)
You have an application controlled by a managed instance group. When you deploy a new version of the application, costs should be minimized and the number of instances should not increase. You want to ensure that, when each new instance is created, the deployment only continues if the new instance is healthy. What should you do?

A. Perform a rolling-action with maxSurge set to 1, maxUnavailable set to 0.
B. Perform a rolling-action with maxSurge set to 0, maxUnavailable set to 1
C. Perform a rolling-action with maxHealthy set to 1, maxUnhealthy set to 0.
D. Perform a rolling-action with maxHealthy set to 0, maxUnhealthy set to 1.

Answer: A

Explanation:
Reference:
https://cloud.google.com/compute/docs/instance-groups/rolling-out-updates-to-managed-instance-groups

NEW QUESTION 25
- (Exam Topic 2)
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?

A. Deploy the website on App Engine and use traffic splitting.
B. Deploy the website on App Engine as three separate services.
C. Deploy the website on Cloud Functions and use traffic splitting.
D. Deploy the website on Cloud Functions as three separate functions.

Answer: A

Explanation:
Reference: https://cloud.google.com/appengine/docs/standard/python/splitting-traffic

NEW QUESTION 30
- (Exam Topic 2)
Your data is stored in Cloud Storage buckets. Fellow developers have reported that data downloaded from Cloud Storage is resulting in slow API performance. You want to research the issue to provide details to the GCP support team. Which command should you run?

A. gsutil test –o output.json gs://my-bucket
B. gsutil perfdiag –o output.json gs://my-bucket
C. gcloud compute scp example-instance:~/test-data –o output.json gs://my-bucket
D. gcloud services test –o output.json gs://my-bucket

Answer: B

Explanation:
Reference: https://groups.google.com/forum/#!topic/gce-discussion/xBl9Jq5HDsY

NEW QUESTION 31
- (Exam Topic 2)
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?

A. Replace your entire monitoring platform with Stackdriver.
B. Install the Stackdriver agents on your Compute Engine instances.
C. Use Stackdriver to capture and alert on logs, then ship them to your existing platform.
D. Migrate some traffic back to your old platform and perform AB testing on the two platforms concurrently.

Answer: B

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

NEW QUESTION 35
- (Exam Topic 2)
You are parsing a log file that contains three columns: a timestamp, an account number (a string), and a transaction amount (a number). You want to calculate the sum of all transaction amounts for each unique account number efficiently.
Which data structure should you use?

A. A linked list
B. A hash table
C. A two-dimensional array
D. A comma-delimited string

Answer: B

NEW QUESTION 36
- (Exam Topic 2)
You are creating an App Engine application that writes a file to any user's Google Drive. How should the application authenticate to the Google Drive API?

A. With an OAuth Client ID that uses the https://www.googleapis.com/auth/drive.file scope to obtain an access token for each user.
B. With an OAuth Client ID with delegated domain-wide authority.
C. With the App Engine service account and https://www.googleapis.com/auth/drive.file scope that generates a signed JWT.
D. With the App Engine service account with delegated domain-wide authority.

Answer: B

Explanation:
Reference: https://developers.google.com/drive/api/v3/about-auth

NEW QUESTION 39
- (Exam Topic 2)
You are developing an application that needs to store files belonging to users in Cloud Storage. You want each user to have their own subdirectory in Cloud Storage. When a new user is created, the corresponding empty subdirectory should also be created. What should you do?

A. Create an object with the name of the subdirectory ending with a trailing slash ('/') that is zero bytes in length.
B. Create an object with the name of the subdirectory, and then immediately delete the object within that subdirectory.
C. Create an object with the name of the subdirectory that is zero bytes in length and has WRITER access control list permission.
D. Create an object with the name of the subdirectory that is zero bytes in lengt
E. Set the Content-Type metadata to CLOUDSTORAGE_FOLDER.

Answer: A

Explanation:
https://cloud.google.com/storage/docs/folders
If you create an empty folder using the Google Cloud console, Cloud Storage creates a zero-byte object as a placeholder. For example, if you create a folder called folder in a bucket called my-bucket, a zero- byte object called gs://my-bucket/folder/ is created. This placeholder is discoverable by other tools when listing the objects in the bucket, for example when using the gsutil ls command.

NEW QUESTION 40
- (Exam Topic 2)
You configured your Compute Engine instance group to scale automatically according to overall CPU usage. However, your application's response latency increases sharply before the cluster has finished adding up instances. You want to provide a more consistent latency experience for your end users by changing the configuration ot the instance group autoscaler. Which two configuration changes should you make? (Choose two.)

A. Add the label "AUTOSCALE" to the instance group template.
B. Decrease the cool-down period for instances added to the group.
C. Increase the target CPU usage for the instance group autoscaler.
D. Decrease the target CPU usage for the instance group autoscaler.
E. Remove the health-check for individual VMs in the instance group.

Answer: AC

NEW QUESTION 42
- (Exam Topic 2)
You are designing an application that uses a microservices architecture. You are planning to deploy the application in the cloud and on-premises. You want to make sure the application can scale up on demand and also use managed services as much as possible. What should you do?

A. Deploy open source Istio in a multi-cluster deployment on multiple Google Kubernetes Engine (GKE) clusters managed by Anthos.
B. Create a GKE cluster in each environment with Anthos, and use Cloud Run for Anthos to deploy your application to each cluster.
C. Install a GKE cluster in each environment with Anthos, and use Cloud Build to create a Deployment for your application in each cluster.
D. Create a GKE cluster in the cloud and install open-source Kubernetes on-premise
E. Use an external load balancer service to distribute traffic across the two environments.

Answer: B

Explanation:
https://cloud.google.com/anthos/run
Integrated with Anthos, Cloud Run for Anthos provides a flexible serverless development platform for hybrid and multicloud environments. Cloud Run for Anthos is Google's managed and fully supported Knative offering, an open source project that enables serverless workloads on Kubernetes.

NEW QUESTION 47
- (Exam Topic 2)
Your company has deployed a new API to a Compute Engine instance. During testing, the API is not behaving as expected. You want to monitor the application over 12 hours to diagnose the problem within the application code without redeploying the application. Which tool should you use?

A. Cloud Trace
B. Cloud Monitoring
C. Cloud Debugger logpoints
D. Cloud Debugger snapshots

Answer: C

Explanation:
https://cloud.google.com/debugger/docs/using/logpoints
Logpoints allow you to inject logging into running services without restarting or interfering with the normal function of the service

NEW QUESTION 49
- (Exam Topic 2)
You recently developed an application. You need to call the Cloud Storage API from a Compute Engine instance that doesn't have a public IP address. What should you do?

A. Use Carrier Peering
B. Use VPC Network Peering
C. Use Shared VPC networks
D. Use Private Google Access

Answer: D

Explanation:
https://cloud.google.com/vpc/docs/private-google-access

NEW QUESTION 53
- (Exam Topic 2)
Your team develops stateless services that run on Google Kubernetes Engine (GKE). You need to deploy a new service that will only be accessed by other services running in the GKE cluster. The service will need to scale as quickly as possible to respond to changing load. What should you do?

A. Use a Vertical Pod Autoscaler to scale the containers, and expose them via a ClusterIP Service.
B. Use a Vertical Pod Autoscaler to scale the containers, and expose them via a NodePort Service.
C. Use a Horizontal Pod Autoscaler to scale the containers, and expose them via a ClusterIP Service.
D. Use a Horizontal Pod Autoscaler to scale the containers, and expose them via a NodePort Service.

Answer: C

Explanation:
https://cloud.google.com/kubernetes-engine/docs/concepts/service

NEW QUESTION 55
- (Exam Topic 2)
You want to notify on-call engineers about a service degradation in production while minimizing development time.
What should you do?

A. Use Cloud Function to monitor resources and raise alerts.
B. Use Cloud Pub/Sub to monitor resources and raise alerts.
C. Use Stackdriver Error Reporting to capture errors and raise alerts.
D. Use Stackdriver Monitoring to monitor resources and raise alerts.

Answer: A

NEW QUESTION 56
- (Exam Topic 2)
You are planning to deploy your application in a Google Kubernetes Engine (GKE) cluster The application exposes an HTTP-based health check at /healthz. You want to use this health check endpoint to determine
whether traffic should be routed to the pod by the load balancer.
Which code snippet should you include in your Pod configuration?

A.
```
    livenessProbe:
      httpGet:
        path: /healthz
        port: 80
```

B.
```
    readinessProbe:
      httpGet:
        path: /healthz
        port: 80
```

C.
```
    loadbalancerHealthCheck:
      httpGet:
        path: /healthz
        port: 80
```

D.
```
    healthCheck:
      httpGet:
        path: /healthz
        port: 80
```

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

Answer: B

Explanation:
For the GKE ingress controller to use your readinessProbes as health checks, the Pods for an Ingress must exist at the time of Ingress creation. If your replicas are scaled to 0, the default health check will apply.

NEW QUESTION 57
- (Exam Topic 2)
You are developing a Java Web Server that needs to interact with Google Cloud services via the Google Cloud API on the user's behalf. Users should be able to authenticate to the Google Cloud API using their Google Cloud identities. Which workflow should you implement in your web application?

A. 1) When a user arrives at your application, prompt them for their Google username and password.2) Store an SHA password hash in your application's database along with the user's username.3) The application authenticates to the Google Cloud API using HTTPs requests with the user's username and password hash in the Authorization request header.
B. 1) When a user arrives at your application, prompt them for their Google username and password.2) Forward the user's username and password in an HTTPS request to the Google Cloud authorization server, and request an access token.3) The Google server validates the user's credentials and returns an access token to the application.4) The application uses the access token to call the Google Cloud API.
C. 1) When a user arrives at your application, route them to a Google Cloud consent screen with a list of requested permissions that prompts the user to sign in with SSO to their Google Account.2) After the user signs in and provides consent, your application receives an authorization code from a Google server.3) The Google server returns the authorization code to the user, which is stored in the browser's cookies.4) The user authenticates to the Google Cloud API using the authorization code in the cookie.
D. 1) When a user arrives at your application, route them to a Google Cloud consent screen with a list of requested permissions that prompts the user to sign in with SSO to their Google Account.2) After the user signs in and provides consent, your application receives an authorization code from a Google server.3) The application requests a Google Server to exchange the authorization code with an access token.4) The Google server responds with the access token that is used by the application to call the Google Cloud API.

Answer: D

Explanation:
https://developers.google.com/identity/protocols/oauth2#webserver
The Google OAuth 2.0 endpoint supports web server applications that use languages and frameworks such as PHP, Java, Python, Ruby, and ASP.NET. The authorization sequence begins when your application redirects a browser to a Google URL; the URL includes query parameters that indicate the type of access being requested. Google handles the user authentication, session selection, and user consent. The result is an authorization code, which the application can exchange for an access token and a refresh token.

NEW QUESTION 59
- (Exam Topic 2)
You recently developed a new service on Cloud Run. The new service authenticates using a custom service and then writes transactional information to a Cloud Spanner database. You need to verify that your application can support up to 5,000 read and 1,000 write transactions per second while identifying any bottlenecks that occur. Your test infrastructure must be able to autoscale. What should you do?

A. Build a test harness to generate requests and deploy it to Cloud Ru
B. Analyze the VPC Flow Logs using Cloud Logging.
C. Create a Google Kubernetes Engine cluster running the Locust or JMeter images to dynamically generate load test

D. Analyze the results using Cloud Trace.
E. Create a Cloud Task to generate a test loa
F. Use Cloud Scheduler to run 60,000 Cloud Task transactions per minute for 10 minute
G. Analyze the results using Cloud Monitoring.
H. Create a Compute Engine instance that uses a LAMP stack image from the Marketplace, and use Apache Bench to generate load tests against the servic
I. Analyze the results using Cloud Trace.

Answer: B

Explanation:
https://cloud.google.com/architecture/distributed-load-testing-using-gke

NEW QUESTION 64
- (Exam Topic 2)
You work for an organization that manages an online ecommerce website. Your company plans to expand across the world; however, the estore currently serves one specific region. You need to select a SQL database and configure a schema that will scale as your organization grows. You want to create a table that stores all customer transactions and ensure that the customer (CustomerId) and the transaction (TransactionId) are unique. What should you do?

A. Create a Cloud SQL table that has TransactionId and CustomerId configured as primary key
B. Use anincremental number for the TransactionId.
C. Create a Cloud SQL table that has TransactionId and CustomerId configured as primary key
D. Use a random string (UUID) for the Transactionid.
E. Create a Cloud Spanner table that has TransactionId and CustomerId configured as primary key
F. Use a random string (UUID) for the TransactionId.
G. Create a Cloud Spanner table that has TransactionId and CustomerId configured as primary key
H. Use an incremental number for the TransactionId.

Answer: C

NEW QUESTION 65
- (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 70
- (Exam Topic 2)
You have an application that uses an HTTP Cloud Function to process user activity from both desktop browser and mobile application clients. This function will serve as the endpoint for all metric submissions using HTTP POST.
Due to legacy restrictions, the function must be mapped to a domain that is separate from the domain requested by users on web or mobile sessions. The domain for the Cloud Function is https://fn.example.com. Desktop and mobile clients use the domain https://www.example.com. You need to add a header to the function's HTTP response so that only those browser and mobile sessions can submit metrics to the Cloud Function. Which response header should you add?

A. Access-Control-Allow-Origin: *
B. Access-Control-Allow-Origin: https://*.example.com
C. Access-Control-Allow-Origin: https://fn.example.com
D. Access-Control-Allow-origin: https://www.example.com

Answer: D

NEW QUESTION 71
- (Exam Topic 2)
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?

A. Use the JOIN operator in SQL to combine the tables.
B. Use nested WITH statements to combine the tables.
C. Use the UNION operator in SQL to combine the tables.
D. Use the UNION ALL operator in SQL to combine the tables.

Answer: C

Explanation:
Reference: https://www.techonthenet.com/sql/union_all.php

NEW QUESTION 75
- (Exam Topic 2)
You are deploying a microservices application to Google Kubernetes Engine (GKE). The application will receive daily updates. You expect to deploy a large number of distinct containers that will run on the Linux operating system (OS). You want to be alerted to any known OS vulnerabilities in the new containers. You want to follow Google-recommended best practices. What should you do?

A. Use the gcloud CLI to call Container Analysis to scan new container image
B. Review the vulnerability results before each deployment.
C. Enable Container Analysis, and upload new container images to Artifact Registr
D. Review the vulnerability results before each deployment.
E. Enable Container Analysis, and upload new container images to Artifact Registr
F. Review the critical vulnerability results before each deployment.
G. Use the Container Analysis REST API to call Container Analysis to scan new container image
H. Review the vulnerability results before each deployment.

Answer: B

Explanation:
https://cloud.google.com/container-analysis/docs/automated-scanning-howto https://cloud.google.com/container-analysis/docs/os-overview says: The Container Scanning API allows you to automate OS vulnerability detection, scanning each time you push an image to Container Registry or Artifact Registry. Enabling this API also triggers language package scans for Go and Java vulnerabilities (Preview).

NEW QUESTION 77
- (Exam Topic 2)
You migrated some of your applications to Google Cloud. You are using a legacy monitoring platform deployed on-premises for both on-premises and cloud-deployed applications. You discover that your notification system is responding slowly to time-critical problems in the cloud applications. What should you do?

A. Replace your monitoring platform with Cloud Monitoring.
B. Install the Cloud Monitoring agent on your Compute Engine instances.
C. Migrate some traffic back to your old platfor
D. Perform A/B testing on the two platforms concurrently.
E. Use Cloud Logging and Cloud Monitoring to capture logs, monitor, and send alert
F. Send them to your existing platform.

Answer: D

NEW QUESTION 80
- (Exam Topic 2)
You are using Cloud Build for your CI/CD pipeline to complete several tasks, including copying certain files to Compute Engine virtual machines. Your pipeline requires a flat file that is generated in one builder in the pipeline to be accessible by subsequent builders in the same pipeline. How should you store the file so that all the builders in the pipeline can access it?

A. Store and retrieve the file contents using Compute Engine instance metadata.
B. Output the file contents to a file in /workspac
C. Read from the same /workspace file in the subsequent build step.
D. Use gsutil to output the file contents to a Cloud Storage objec
E. Read from the same object in the subsequent build step.
F. Add a build argument that runs an HTTP POST via curl to a separate web server to persist the value in one builde
G. Use an HTTP GET via curl from the subsequent build step to read the value.

Answer: B

Explanation:
https://cloud.google.com/build/docs/build-config-file-schema

NEW QUESTION 85
- (Exam Topic 2)
One of your deployed applications in Google Kubernetes Engine (GKE) is having intermittent performance issues. Your team uses a third-party logging solution. You want to install this solution on each node in your GKE cluster so you can view the logs. What should you do?

A. Deploy the third-party solution as a DaemonSet
B. Modify your container image to include the monitoring software
C. Use SSH to connect to the GKE node, and install the software manually
D. Deploy the third-party solution using Terraform and deploy the logging Pod as a Kubernetes Deployment

Answer: A

Explanation:

https://cloud.google.com/kubernetes-engine/docs/concepts/daemonset#usage_patterns DaemonSets are useful for deploying ongoing background tasks that you need to run on all or certain nodes, and which do not require user intervention. Examples of such tasks include storage daemons like ceph, log collection daemons like fluent-bit, and node monitoring daemons like collectd.

NEW QUESTION 89
- (Exam Topic 2)
Your team is developing an ecommerce platform for your company. Users will log in to the website and add items to their shopping cart. Users will be automatically logged out after 30 minutes of inactivity. When users log back in, their shopping cart should be saved. How should you store users' session and shopping cart information while following Google-recommended best practices?

A. Store the session information in Pub/Sub, and store the shopping cart information in Cloud SQL.
B. Store the shopping cart information in a file on Cloud Storage where the filename is the SESSION ID.
C. Store the session and shopping cart information in a MySQL database running on multiple Compute Engine instances.
D. Store the session information in Memorystore for Redis or Memorystore for Memcached, and store the shopping cart information in Firestore.

Answer: D

NEW QUESTION 93
- (Exam Topic 2)
You are evaluating developer tools to help drive Google Kubernetes Engine adoption and integration with your development environment, which includes VS Code and IntelliJ. What should you do?

A. Use Cloud Code to develop applications.
B. Use the Cloud Shell integrated Code Editor to edit code and configuration files.
C. Use a Cloud Notebook instance to ingest and process data and deploy models.
D. Use Cloud Shell to manage your infrastructure and applications from the command line.

Answer: A

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

NEW QUESTION 94
- (Exam Topic 2)
You are developing a microservice-based application that will run on Google Kubernetes Engine (GKE). Some of the services need to access different Google Cloud APIs. How should you set up authentication of these services in the cluster following Google-recommended best practices? (Choose two.)

A. Use the service account attached to the GKE node.
B. Enable Workload Identity in the cluster via the gcloud command-line tool.
C. Access the Google service account keys from a secret management service.
D. Store the Google service account keys in a central secret management service.
E. Use gcloud to bind the Kubernetes service account and the Google service account using roles/iam.workloadIdentity.

Answer: BE

Explanation:
https://cloud.google.com/kubernetes-engine/docs/how-to/workload-identity

NEW QUESTION 96
- (Exam Topic 2)
You are writing a Compute Engine hosted application in project A that needs to securely authenticate to a Cloud Pub/Sub topic in project B.
What should you do?

A. Configure the instances with a service account owned by project
B. Add the service account as a Cloud Pub/Sub publisher to project A.
C. Configure the instances with a service account owned by project
D. Add the service account as a publisher on the topic.
E. Configure Application Default Credentials to use the private key of a service account owned by project
F. Add the service account as a Cloud Pub/Sub publisher to project A.
G. Configure Application Default Credentials to use the private key of a service account owned by project
H. Add the service account as a publisher on the topic

Answer: B

Explanation:
https://cloud.google.com/pubsub/docs/access-control
"For example, suppose a service account in Cloud Project A wants to publish messages to a topic in Cloud Project B. You could accomplish this by granting the service account Edit permission in Cloud Project B"

NEW QUESTION 100
- (Exam Topic 2)
Your application is deployed in a Google Kubernetes Engine (GKE) cluster. You want to expose this application publicly behind a Cloud Load Balancing HTTP(S) load balancer. What should you do?

A. Configure a GKE Ingress resource.
B. Configure a GKE Service resource.
C. Configure a GKE Ingress resource with type: LoadBalancer.
D. Configure a GKE Service resource with type: LoadBalancer.

Answer: A

Explanation:
Reference: https://cloud.google.com/kubernetes-engine/docs/concepts/ingress

NEW QUESTION 102
- (Exam Topic 2)
Your company wants to expand their users outside the United States for their popular application. The company wants to ensure 99.999% availability of the database for their application and also wants to minimize the read latency for their users across the globe.
Which two actions should they take? (Choose two.)

A. Create a multi-regional Cloud Spanner instance with "nam-asia-eur1" configuration.
B. Create a multi-regional Cloud Spanner instance with "nam3" configuration.
C. Create a cluster with at least 3 Spanner nodes.
D. Create a cluster with at least 1 Spanner node.
E. Create a minimum of two Cloud Spanner instances in separate regions with at least one node.
F. Create a Cloud Dataflow pipeline to replicate data across different databases.

Answer: BF

NEW QUESTION 103
- (Exam Topic 2)
Your team develops services that run on Google Cloud. You want to process messages sent to a Pub/Sub topic, and then store them. Each message must be processed exactly once to avoid duplication of data and any data conflicts. You need to use the cheapest and most simple solution. What should you do?

A. Process the messages with a Dataproc job, and write the output to storage.
B. Process the messages with a Dataflow streaming pipeline using Apache Beam's PubSubIO package, and write the output to storage.
C. Process the messages with a Cloud Function, and write the results to a BigQuery location where you can run a job to deduplicate the data.
D. Retrieve the messages with a Dataflow streaming pipeline, store them in Cloud Bigtable, and use another Dataflow streaming pipeline to deduplicate messages.

Answer: B

Explanation:
https://cloud.google.com/dataflow/docs/concepts/streaming-with-cloud-pubsub

NEW QUESTION 104
- (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 105
- (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.
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 to your App Engine service.Create a DNS CNAME record to point to the App Engine canonical name ghs.googlehosted.com.
F. Define a mapping in dispatch.yaml to point the domain 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 106
- (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 109
- (Exam Topic 2)
Your company's development teams want to use Cloud Build in their projects to build and push Docker images to Container Registry. The operations team requires all Docker images to be published to a centralized, securely managed Docker registry that the operations team manages.
What should you do?

A. Use Container Registry to create a registry in each development team's projec
B. Configure the Cloud Build build to push the Docker image to the project's registr
C. Grant the operations team access to each development team's registry.
D. Create a separate project for the operations team that has Container Registry configure
E. Assign appropriate permissions to the Cloud Build service account in each developer team's project to allow access to the operation team's registry.
F. Create a separate project for the operations team that has Container Registry configure
G. Create a Service Account for each development team and assign the appropriate permissions to allow it access to the operations team's registr
H. Store the service account key file in the source code repository and use it to authenticate against the operations team's registry.
I. Create a separate project for the operations team that has the open source Docker Registry deployed on a Compute Engine virtual machine instanc
J. Create a username and password for each development tea
K. Store the username and password in the source code repository and use it to authenticate against theoperations team's Docker registry.

Answer: A

Explanation:
Reference: https://cloud.google.com/container-registry/

NEW QUESTION 110
- (Exam Topic 2)
You are running a containerized application on Google Kubernetes Engine. Your container images are stored in Container Registry. Your team uses CI/CD practices. You need to prevent the deployment of containers with known critical vulnerabilities. What should you do?

A. • Use Web Security Scanner to automatically crawl your application• Review your application logs for scan results, and provide an attestation that the container is free of known critical vulnerabilities• Use Binary Authorization to implement a policy that forces the attestation to be provided before the container is deployed
B. • Use Web Security Scanner to automatically crawl your application• Review the scan results in the scan details page in the Cloud Console, and provide an attestation that the container is free of known critical vulnerabilities• Use Binary Authorization to implement a policy that forces the attestation to be provided before the container is deployed
C. • Enable the Container Scanning API to perform vulnerability scanning• Review vulnerability reporting in Container Registry in the Cloud Console, and provide an attestation that the container is free of known critical vulnerabilities• Use Binary Authorization to implement a policy that forces the attestation to be provided before the container is deployed
D. • Enable the Container Scanning API to perform vulnerability scanning• Programmatically review vulnerability reporting through the Container Scanning API, and provide an attestation that the container is free of known critical vulnerabilities• Use Binary Authorization to implement a policy that forces the attestation to be provided before the container is deployed

Answer: D

Explanation:
https://cloud.google.com/binary-authorization/docs/creating-attestations-kritis
https://cloud.google.com/container-analysis/docs/os-overview

NEW QUESTION 115
- (Exam Topic 2)
You are developing an internal application that will allow employees to organize community events within your company. You deployed your application on a single Compute Engine instance. Your company uses Google Workspace (formerly G Suite), and you need to ensure that the company employees can authenticate to the application from anywhere. What should you do?

A. Add a public IP address to your instance, and restrict access to the instance using firewall rule
B. Allow your company's proxy as the only source IP address.
C. Add an HTTP(S) load balancer in front of the instance, and set up Identity-Aware Proxy (IAP).Configure the IAP settings to allow your company domain to access the website.
D. Set up a VPN tunnel between your company network and your instance's VPC location on Google Clou
E. Configure the required firewall rules and routing information to both the on-premises and Google Cloud networks.
F. Add a public IP address to your instance, and allow traffic from the interne
G. Generate a random hash, and create a subdomain that includes this hash and points to your instanc
H. Distribute this DNS address to your company's employees.

Answer: B

Explanation:
https://cloud.google.com/blog/topics/developers-practitioners/control-access-your-web-sites-identity-aware-prox

NEW QUESTION 116
- (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 ste
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 Engin
G. The microservice calls each downstream step and waits for a successful response before calling the next step.

Answer: C

NEW QUESTION 120
- (Exam Topic 2)
You are developing a web application that contains private images and videos stored in a Cloud Storage bucket. Your users are anonymous and do not have Google Accounts. You want to use your
application-specific logic to control access to the images and videos. How should you configure access?

A. Cache each web application user's IP address to create a named IP table using Google Cloud Armor.Create a Google Cloud Armor security policy that allows users to access the backend bucket.
B. Grant the Storage Object Viewer IAM role to allUser
C. Allow users to access the bucket after authenticating through your web application.
D. Configure Identity-Aware Proxy (IAP) to authenticate users into the web applicatio
E. Allow users to access the bucket after authenticating through IAP.
F. Generate a signed URL that grants read access to the bucke
G. Allow users to access the URL after authenticating through your web application.

Answer: D

Explanation:
https://cloud.google.com/storage/docs/access-control/signed-urls#should-you-use
In some scenarios, you might not want to require your users to have a Google account in order to access Cloud Storage, but you still want to control access using your application-specific logic. The typical way to address this use case is to provide a signed URL to a user, which gives the user read, write, or delete access to that resource for a limited time. You specify an expiration time when you create the signed URL. Anyone who knows the URL can access the resource until the expiration time for the URL is reached or the key used to sign the URL is rotated.

NEW QUESTION 124
- (Exam Topic 2)
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?

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

Answer: A

Explanation:
Reference: https://cloud.google.com/storage/docs/request-rate

NEW QUESTION 126
- (Exam Topic 2)
You need to copy directory local-scripts and all of its contents from your local workstation to a Compute Engine virtual machine instance.
Which command should you use?

A. gsutil cp --project "my-gcp-project" -r ~/local-scripts/ gcp-instance-name:~/ server-scripts/ --zone "us-east1-b"
B. gsutil cp --project "my-gcp-project" -R ~/local-scripts/ gcp-instance-name:~/ server-scripts/ --zone "us-east1-b"
C. gcloud compute scp --project "my-gcp-project" --recurse ~/local-scripts/ gcpinstance- name:~/server-scripts/ --zone "us-east1-b"
D. gcloud compute mv --project "my-gcp-project" --recurse ~/local-scripts/ gcpinstance- name:~/server-scripts/ --zone "us-east1-b"

Answer: C

Explanation:
Reference: https://cloud.google.com/sdk/gcloud/reference/compute/copy-files

NEW QUESTION 131
- (Exam Topic 2)
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?

A. 1) Create a managed instance grou

B. 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.
C. 1) Create an unmanaged instance grou
D. Replicate the static content across the VMs.2) Create an external HTTP(S) load balancer3) Enable Cloud CDN, and send traffic to the unmanaged instance group.
E. 1) Create a Standard storage class, regional Cloud Storage bucke
F. Put the static content in the bucket2) Reserve an external IP address, and create an external HTTP(S) load balancer3) Enable Cloud CDN, and send traffic to your backend bucket
G. 1) Create a Standard storage class, multi-regional Cloud Storage bucke
H. 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.

Answer: D

NEW QUESTION 132
- (Exam Topic 2)
Your company has a BigQuery data mart that provides analytics information to hundreds of employees. One user of wants to run jobs without interrupting important workloads. This user isn't concerned about the time it takes to run these jobs. You want to fulfill this request while minimizing cost to the company and the effort required on your part.
What should you do?

A. Ask the user to run the jobs as batch jobs.
B. Create a separate project for the user to run jobs.
C. Add the user as a job.user role in the existing project.
D. Allow the user to run jobs when important workloads are not running.

Answer: B

NEW QUESTION 133
- (Exam Topic 2)
You have an application written in Python running in production on Cloud Run. Your application needs to read/write data stored in a Cloud Storage bucket in the same project. You want to grant access to your
application following the principle of least privilege. What should you do?

A. Create a user-managed service account with a custom Identity and Access Management (IAM) role.
B. Create a user-managed service account with the Storage Admin Identity and Access Management (IAM) role.
C. Create a user-managed service account with the Project Editor Identity and Access Management (IAM) role.
D. Use the default service account linked to the Cloud Run revision in production.

Answer: A

Explanation:
https://cloud.google.com/iam/docs/understanding-roles#storage.admin

NEW QUESTION 136
- (Exam Topic 2)
You are creating a Google Kubernetes Engine (GKE) cluster and run this command:

```
> gcloud container clusters create large-cluster --num-nodes 200
```

The command fails with the error:

```
insufficient regional quota to satisfy request: resource "CPUS": request
requires '200.0' and is short '176.0'. project has a quota of '24.0' with
'24.0' available
```

You want to resolve the issue. What should you do?

A. Request additional GKE quota is the GCP Console.
B. Request additional Compute Engine quota in the GCP Console.
C. Open a support case to request additional GKE quotA.
D. Decouple services in the cluster, and rewrite new clusters to function with fewer cores.

Answer: A

NEW QUESTION 140
- (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 141
- (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/.
E. Ensure that clients use Compute Engine internal DNS by connecting to the instance name with the url https://[API_NAME]/[API_VERSION]/.

Answer: D

NEW QUESTION 146
- (Exam Topic 2)
You are planning to deploy hundreds of microservices in your Google Kubernetes Engine (GKE) cluster. How should you secure communication between the microservices on GKE using a managed service?

A. Use global HTTP(S) Load Balancing with managed SSL certificates to protect your services
B. Deploy open source Istio in your GKE cluster, and enable mTLS in your Service Mesh
C. Install cert-manager on GKE to automatically renew the SSL certificates.
D. Install Anthos Service Mesh, and enable mTLS in your Service Mesh.

Answer: D

Explanation:
https://cloud.google.com/service-mesh/docs/overview#security_benefits
- Ensures encryption in transit. Using mTLS for authentication also ensures that all TCP communications are encrypted in transit.

NEW QUESTION 149
- (Exam Topic 2)
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 --port-name http
--global
gcloud compute backend-services add-backend ${NAME}-http-bes --instance-group ${NAME}-i --balancing-mode RATE --max-rate
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 ${NAME}-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?

A. gcloud compute instances add-access-config ${NAME}-backend-instance-1

B. gcloud compute instances add-tags ${NAME}-backend-instance-1 --tags http-server
C. 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
D. 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

Answer: C

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

NEW QUESTION 153
- (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 156
- (Exam Topic 2)
Your service adds text to images that it reads from Cloud Storage. During busy times of the year, requests to Cloud Storage fail with an HTTP 429 "Too Many Requests" status code.
How should you handle this error?

A. Add a cache-control header to the objects.
B. Request a quota increase from the GCP Console.
C. Retry the request with a truncated exponential backoff strategy.
D. Change the storage class of the Cloud Storage bucket to Multi-regional.

Answer: C

Explanation:
Reference: https://developers.google.com/gmail/api/v1/reference/quota

NEW QUESTION 157
- (Exam Topic 2)
The development teams in your company want to manage resources from their local environments. You have been asked to enable developer access to each team's Google Cloud projects. You want to maximize efficiency while following Google-recommended best practices. What should you do?

A. Add the users to their projects, assign the relevant roles to the users, and then provide the users with each relevant Project ID.
B. Add the users to their projects, assign the relevant roles to the users, and then provide the users with each relevant Project Number.
C. Create groups, add the users to their groups, assign the relevant roles to the groups, and then provide the users with each relevant Project ID.
D. Create groups, add the users to their groups, assign the relevant roles to the groups, and then provide the users with each relevant Project Number.

Answer: C

NEW QUESTION 161
- (Exam Topic 2)
Your application is running in multiple Google Kubernetes Engine clusters. It is managed by a Deployment in each cluster. The Deployment has created multiple replicas of your Pod in each cluster. You want to view the logs sent to stdout for all of the replicas in your Deployment in all clusters. Which command should you use?

A. kubectl logs [PARAM]
B. gcloud logging read [PARAM]
C. kubectl exec –it [PARAM] journalctl
D. gcloud compute ssh [PARAM] –-command= "sudo journalctl"

Answer: D

NEW QUESTION 162
- (Exam Topic 2)
Your development team has built several Cloud Functions using Java along with corresponding integration and service tests. You are building and deploying the functions and launching the tests using Cloud Build. Your Cloud Build job is reporting deployment failures immediately after successfully validating the code. What should you do?

A. Check the maximum number of Cloud Function instances.
B. Verify that your Cloud Build trigger has the correct build parameters.

C. Retry the tests using the truncated exponential backoff polling strategy.
D. Verify that the Cloud Build service account is assigned the Cloud Functions Developer role.

Answer: D

Explanation:
https://cloud.google.com/build/docs/securing-builds/configure-access-for-cloud-build-service-account

NEW QUESTION 166
- (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 167
- (Exam Topic 2)
You deployed a new application to Google Kubernetes Engine and are experiencing some performance degradation. Your logs are being written to Cloud Logging, and you are using a Prometheus sidecar model for capturing metrics. You need to correlate the metrics and data from the logs to troubleshoot the performance issue and send real-time alerts while minimizing costs. What should you do?

A. Create custom metrics from the Cloud Logging logs, and use Prometheus to import the results using the Cloud Monitoring REST API.
B. Export the Cloud Logging logs and the Prometheus metrics to Cloud Bigtabl
C. Run a query to join the results, and analyze in Google Data Studio.
D. Export the Cloud Logging logs and stream the Prometheus metrics to BigQuer
E. Run a recurring query to join the results, and send notifications using Cloud Tasks.
F. Export the Prometheus metrics and use Cloud Monitoring to view them as external metric
G. Configure Cloud Monitoring to create log-based metrics from the logs, and correlate them with the Prometheus data.

Answer: D

Explanation:
Reference:
https://cloud.google.com/blog/products/operations/troubleshoot-gke-faster-with-monitoring-data-in-your-logs

NEW QUESTION 168
- (Exam Topic 2)
You are using Cloud Build to build and test application source code stored in Cloud Source Repositories. The build process requires a build tool not available in the Cloud Build environment.
What should you do?

A. Download the binary from the internet during the build process.
B. Build a custom cloud builder image and reference the image in your build steps.
C. Include the binary in your Cloud Source Repositories repository and reference it in your build scripts.
D. Ask to have the binary added to the Cloud Build environment by filing a feature request against the Cloud Build public Issue Tracker.

Answer: B

NEW QUESTION 172
- (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 175
- (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 179
- (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 184
- (Exam Topic 2)
You are developing a JPEG image-resizing API hosted on Google Kubernetes Engine (GKE). Callers of the service will exist within the same GKE cluster. You want clients to be able to get the IP address of the service.
What should you do?

A. Define a GKE Servic
B. Clients should use the name of the A record in Cloud DNS to find the service's cluster IP address.
C. Define a GKE Servic
D. Clients should use the service name in the URL to connect to the service.
E. Define a GKE Endpoin
F. Clients should get the endpoint name from the appropriate environment variable in the client container.
G. Define a GKE Endpoin
H. Clients should get the endpoint name from Cloud DNS.

Answer: C

NEW QUESTION 189
- (Exam Topic 2)
You have containerized a legacy application that stores its configuration on an NFS share. You need to deploy this application to Google Kubernetes Engine (GKE) and do not want the application serving traffic until after the configuration has been retrieved. What should you do?

A. Use the gsutil utility to copy files from within the Docker container at startup, and start the service using an ENTRYPOINT script.
B. Create a PersistentVolumeClaim on the GKE cluste
C. Access the configuration files from the volume, and start the service using an ENTRYPOINT script.
D. Use the COPY statement in the Dockerfile to load the configuration into the container imag
E. Verify that the configuration is available, and start the service using an ENTRYPOINT script.
F. Add a startup script to the GKE instance group to mount the NFS share at node startu
G. Copy the configuration files into the container, and start the service using an ENTRYPOINT script.

Answer: D

Explanation:
Reference: https://cloud.google.com/compute/docs/instances/startup-scripts/linux

NEW QUESTION 190
- (Exam Topic 2)
Your company's development teams want to use various open source operating systems in their Docker builds. When images are created in published containers in your company's environment, you need to scan them for Common Vulnerabilities and Exposures (CVEs). The scanning process must not impact software development agility. You want to use managed services where possible. What should you do?

A. Enable the Vulnerability scanning setting in the Container Registry.
B. Create a Cloud Function that is triggered on a code check-in and scan the code for CVEs.
C. Disallow the use of non-commercially supported base images in your development environment.
D. Use Cloud Monitoring to review the output of Cloud Build to determine whether a vulnerable version has been used.

Answer: A

Explanation:
https://cloud.google.com/container-analysis/docs/os-overview

NEW QUESTION 194
- (Exam Topic 2)
You want to upload files from an on-premises virtual machine to Google Cloud Storage as part of a data migration. These files will be consumed by Cloud DataProc Hadoop cluster in a GCP environment. Which command should you use?

A. gsutil cp [LOCAL_OBJECT] gs://[DESTINATION_BUCKET_NAME]/
B. gcloud cp [LOCAL_OBJECT] gs://[DESTINATION_BUCKET_NAME]/
C. hadoop fs cp [LOCAL_OBJECT] gs://[DESTINATION_BUCKET_NAME]/
D. gcloud dataproc cp [LOCAL_OBJECT] gs://[DESTINATION_BUCKET_NAME]/

Answer: A

Explanation:
The gsutil cp command allows you to copy data between your local file. storage. boto files generated by running "gsutil config"

NEW QUESTION 196
- (Exam Topic 2)
Your API backend is running on multiple cloud providers. You want to generate reports for the network latency of your API.
Which two steps should you take? (Choose two.)

A. Use Zipkin collector to gather data.
B. Use Fluentd agent to gather data.
C. Use Stackdriver Trace to generate reports.
D. Use Stackdriver Debugger to generate report.
E. Use Stackdriver Profiler to generate report.

Answer: AC

Explanation:
https://cloud.google.com/trace/docs/zipkin
"receive traces from Zipkin clients and forward those traces to Cloud Trace for analysis." https://cloud.google.com/trace/docs/quickstart#analysis_reports_window

NEW QUESTION 199
- (Exam Topic 2)
Your development team has been asked to refactor an existing monolithic application into a set of composable microservices. Which design aspects should you implement for the new application? (Choose two.)

A. Develop the microservice code in the same programming language used by the microservice caller.
B. Create an API contract agreement between the microservice implementation and microservice caller.
C. Require asynchronous communications between all microservice implementations and microservice callers.
D. Ensure that sufficient instances of the microservice are running to accommodate the performance requirements.
E. Implement a versioning scheme to permit future changes that could be incompatible with the current interface.

Answer: BE

NEW QUESTION 204
......

Relate Links

100% Pass Your Professional-Cloud-Developer Exam with Exambible Prep Materials

https://www.exambible.com/Professional-Cloud-Developer-exam/

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - https://www.exambible.com/