Free Questions for Professional-Cloud-Developer

Shared by Caldwell on 12-12-2023

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Case Study: Mix Questions

Mix Questions

Professional-Cloud-Developer Mix Questions IN THIS CASE STUDY

A governmental regulation was recently passed that affects your application. For compliance purposes, you are now required to send a duplicate of specific application logs from your application's project to a project that is restricted to the security team. What should you do?

Options:

A- Modify the _Default tog bucket sink rules to reroute the logs into the security team's log bucket.
B- Create user-defined log buckets in the security team's project. Configure a Cloud Logging sink to route your application s logs to log buckets in the security team's project.
C- Create a job that copies the System Event logs from the _Required log bucket into the security team's log bucket in their project.
D- Create a job that copies the togs from the _Required log bucket into the security team's log bucket in their project.

Answer:

B

Question 2

Question Type: MultipleChoice

Case Study: Mix Questions

Mix Questions

Professional-Cloud-Developer Mix Questions IN THIS CASE STUDY

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

Options:

A- Import the Cloud Profiler package into your application, and initialize the Profiler agent. Review the generated flame graph in the Google Cloud console to identify time-intensive functions.
B- Create a Cloud Logging query that gathers the web application's logs. Write a Python script that calculates the difference between the
timestamps from the beginning and the end of the application's longest functions to identity time-intensive functions.
C- Import OpenTelemetry and Trace export packages into your application, and create the trace provider. Review the latency data for your
application on the Trace overview page, and identify which functions cause the most latercy.
D- Add print commands to the application source code to log when each function is called, and redeploy the application.

Answer:

A

Explanation:

Google Cloud Profiler efficiently analyzes cloud-based applications, like those on GKE, using minimal code changes and flame graphs for in-depth insights into CPU and memory usage, thus identifying performance bottlenecks with ease. For more information, visit:
https://dev.to/billcchung/profiling-go-grpc-service-with-google-cloud-profiler-1k8l and
https://www.cncf.io/blog/2020/03/27/continuous-profiling-go-application-running-in-kubernetes/

Question 3

Question Type: MultipleChoice

Case Study: Mix Questions

Mix Questions

Professional-Cloud-Developer Mix Questions IN THIS CASE STUDY

You plan to deploy a new application revision with a Deployment resource to Google Kubernetes Engine (GKE) in production. The container might not work correctly. You want to minimize risk in case there are issues after deploying the revision. You want to follow Google-recommended best practices. What should you do?

Options:

A- Perform a rolling update with a PodDisruptionBudget of 80%.
B- Perform a rolling update with a HorizontalPodAutoscaler scale-down policy value of 0.
C- Convert the Deployment to a StatefulSet, and perform a rolling update with a PodDisruptionBudget of 80%.
D- Convert the Deployment to a StatefulSet, and perform a rolling update with a HorizontalPodAutoscaler scale-down policy value of 0.

Answer:

A

Explanation:

https://cloud.google.com/blog/products/containers-kubernetes/ensuring-reliability-and-uptime-for-your-gke-cluster

Setting PodDisruptionBudget ensures that your workloads have a sufficient number of replicas, even during maintenance. Using the PDB, you can define a number (or percentage) of pods that can be terminated, even if terminating them brings the current replica count below the desired value. With PDB configured, Kubernetes will drain a node following the configured disruption schedule. New pods will be deployed on other available nodes. This approach ensures Kubernetes schedules workloads in an optimal way while controlling the disruption based on the PDB configuration.

https://blog.knoldus.com/how-to-avoid-outages-in-your-kubernetes-cluster-using-pdb/

Question 4

Question Type: MultipleChoice

Case Study: Mix Questions

Mix Questions

Professional-Cloud-Developer Mix Questions IN THIS CASE STUDY

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

Options:

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

Answer:

D

Question 5

Question Type: MultipleChoice

Case Study: Mix Questions

Mix Questions

Professional-Cloud-Developer Mix Questions IN THIS CASE STUDY

You recently developed an application that monitors a large number of stock prices. You need to configure Pub/Sub to receive a high volume messages and update the current stock price in a single large in-memory database The downstream service needs only the most up-to-date prices in the in-memory database to perform stock trading transactions Each message contains three pieces of information

* Stock symbol

* Stock price

* Timestamp for the update

.

How should you set up your Pub/Sub subscription?

Options:

A- Create a pull subscription with both ordering and exactly-once delivery turned off
B- Create a pull subscription with exactly-once delivery enabled
C- Create a push subscription with exactly-once delivery enabled
D- Create a push subscription with both ordering and exactly-once delivery turned off

Answer:

C

Question 6

Question Type: MultipleChoice

Case Study: Mix Questions

Mix Questions

Professional-Cloud-Developer Mix Questions IN THIS CASE STUDY

You are deploying a microservices application to Google Kubernetes Engine (GKE) that will broadcast livestreams. You expect unpredictable traffic patterns and large variations in the number of concurrent users. Your application must meet the following requirements:

* Scales automatically during popular events and maintains high availability

* Is resilient in the event of hardware failures

How should you configure the deployment parameters? (Choose two.)

Options:

A- Distribute your workload evenly using a multi-zonal node pool.
B- Distribute your workload evenly using multiple zonal node pools.
C- Use cluster autoscaler to resize the number of nodes in the node pool, and use a Horizontal Pod Autoscaler to scale the workload.
D- Create a managed instance group for Compute Engine with the cluster nodes. Configure autoscaling rules for the managed instance group.
E- Create alerting policies in Cloud Monitoring based on GKE CPU and memory utilization. Ask an on-duty engineer to scale the workload by executing a script when CPU and memory usage exceed predefined thresholds.

Answer:

A, C

Question 7

Question Type: MultipleChoice

Case Study: Mix Questions

Mix Questions

Professional-Cloud-Developer Mix Questions IN THIS CASE STUDY

You have a web application that publishes messages to Pub/Sub. You plan to build new versions of the application locally and need to quickly test Pub/Sub integration tor each new build. How should you configure local testing?

Options:

A- Run the gclcud config set api_endpoint_overrides/pubsub https: / 'pubsubemulator.googleapi3.com. coin/ command to change the Pub/Sub endpoint prior to starting the application
B- In the Google Cloud console, navigate to the API Library and enable the Pub/Sub API When developing locally, configure your application to call pubsub.googleapis com
C- Install Cloud Code on the integrated development environment (IDE) Navigate to Cloud APIs, and enable Pub/Sub against a valid Google Project ID. When developing locally, configure your application to call pubsub.googleapis com
D- Install the Pub/Sub emulator using gcloud and start the emulator with a valid Google Project ID. When developing locally, configure your application to use the local emulator by exporting the fuhsub emulator Host variable

Answer:

C

Question 8

Question Type: MultipleChoice

Case Study: Mix Questions

Mix Questions

Professional-Cloud-Developer Mix Questions IN THIS CASE STUDY

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

Options:

A- Download, install, and start the Snapshot Debugger agent in your VM. Take debug snapshots

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

Answer:

B

Question 9

Question Type: MultipleChoice

Case Study: Mix Questions

Mix Questions

Professional-Cloud-Developer Mix Questions IN THIS CASE STUDY

You are reviewing and updating your Cloud Build steps to adhere to Google-recommended practices. Currently, your build steps include:

1. Pull the source code from a source repository.

2. Build a container image

3. Upload the built image to Artifact Registry.

You need to add a step to perform a vulnerability scan of the built container image, and you want the results of the scan to be available to your deployment pipeline running in Google Cloud. You want to minimize changes that could disrupt other teams' processes What should you do?

Options:

A- Enable Binary Authorization, and configure it to attest that no vulnerabilities exist in a container image.
B- Enable the Container Scanning API in Artifact Registry, and scan the built container images for vulnerabilities.

C- Upload the built container images to your Docker Hub instance, and scan them for vulnerabilities.
D- Add Artifact Registry to your Aqua Security instance, and scan the built container images for vulnerabilities

Answer:

B

Question 10

Question Type: MultipleChoice

Case Study: Mix Questions

Mix Questions

Professional-Cloud-Developer Mix Questions IN THIS CASE STUDY

Your company's corporate policy states that there must be a copyright comment at the very beginning of all source files. You want to write a custom step in Cloud Build that is triggered by each source commit. You need the trigger to validate that the source contains a copyright and add one for subsequent steps if not there. What should you do?

Options:

A- Build a new Docker container that examines the files in /workspace and then checks and adds a copyright for each source file. Changed files are explicitly committed back to the source repository.
B- Build a new Docker container that examines the files in /workspace and then checks and adds a copyright for each source file. Changed files do not need to be committed back to the source repository.
C- Build a new Docker container that examines the files in a Cloud Storage bucket and then checks and adds a copyright for each source file. Changed files are written back to the Cloud Storage bucket.
D- Build a new Docker container that examines the files in a Cloud Storage bucket and then checks and adds a copyright for each source file. Changed files are explicitly committed back to the source repository.

Answer:

A

Explanation:

https://cloud.google.com/build/docs/configuring-builds/pass-data-between-steps#passing_data_using_workspaces

To pass data between build steps, store the assets produced by the build step in /workspace and these assets will be available to any subsequent build steps.

P2P exams

P2P exams

To Get Premium Files for Professional-Cloud-Developer Visit
https://www.p2pexams.com/products/professional-cloud-developer

For More Free Questions Visit
https://www.p2pexams.com/google/pdf/professional-cloud-developer

