PROFESSIONAL-CLOUD-DEVOPS-ENGINEER Q&As

Professional Cloud DevOps Engineer

Pass Google PROFESSIONAL-CLOUD-DEVOPS-ENGINEER Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.pass2lead.com/professional-cloud-devops-engineer.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Google Official Exam Center

QUESTION 1

Your company is using HTTPS requests to trigger a public Cloud Run-hosted service accessible at the https://booking-engine-abcdef.a.run.app URL. You need to give developers the ability to test the latest revisions of the service before the service is exposed to customers. What should you do?

A. Run the gcloud run deploy booking-engine --no-traffic --tag dev command. Use the https://dev--booking-engine-abcdef.a.run.app URL for testing.

B. Run the gcloud run services update-traffic booking-engine --to-revisions LATEST=1 command. Use the https://booking-engine-abcdef.a.run.app URL for testing.

C. Pass the curl –H “Authorization:Bearer $(gcloud auth print-identity-token)” auth token. Use the https://booking-engine-abcdef.a.run.app URL to test privately.

D. Grant the roles/run.invoker role to the developers testing the booking-engine service. Use the https://booking-engine-abcdef.private.run.app URL for testing.

Correct Answer: A

Reference: https://cloud.google.com/sdk/gcloud/reference/run/deploy

QUESTION 2

You support a web application that runs on App Engine and uses CloudSQL and Cloud Storage for data storage. After a short spike in website traffic, you notice a big increase in latency for all user requests, increase in CPU use, and the

number of processes running the application. Initial troubleshooting reveals:

After the initial spike in traffic, load levels returned to normal but users still experience high latency.

Requests for content from the CloudSQL database and images from Cloud Storage show the same high latency.

No changes were made to the website around the time the latency increased.

There is no increase in the number of errors to the users.

You expect another spike in website traffic in the coming days and want to make sure users don\\'t experience latency. What should you do?

A. Upgrade the GCS buckets to Multi-Regional.

B. Enable high availability on the CloudSQL instances.

C. Move the application from App Engine to Compute Engine.

D. Modify the App Engine configuration to have additional idle instances.

Correct Answer: D

Scaling App Engine scales the number of instances automatically in response to processing volume. This scaling factors in the automatic_scaling settings that are provided on a per-version basis in the configuration file. A service with basic scaling is configured by setting the maximum number of instances in the max_instances parameter of the

basic_scaling setting. The number of live instances scales with the processing volume. You configure the number of instances of each version in that service\\'s configuration file. The number of instances usually corresponds to the size of a dataset being held in memory or the desired throughput for offline work. You can adjust the number of instances of a manually-scaled version very quickly, without stopping instances that are currently running, using the Modules API set_num_instances function.

https://cloud.google.com/appengine/docs/standard/python/how-instances-are-managed

QUESTION 3

You are the Site Reliability Engineer responsible for managing your company\\'s data services and products. You regularly navigate operational challenges, such as unpredictable data volume and high cost, with your company\\'s data ingestion processes. You recently learned that a new data ingestion product will be developed in Google Cloud. You need to collaborate with the product development team to provide operational input on the new product. What should you do?

A. Deploy the prototype product in a test environment, run a load test, and share the results with the product development team.

B. When the initial product version passes the quality assurance phase and compliance assessments, deploy the product to a staging environment. Share error logs and performance metrics with the product development team.

C. When the new product is used by at least one internal customer in production, share error logs and monitoring metrics with the product development team.

D. Review the design of the product with the product development team to provide feedback early in the design phase.

Correct Answer: D

QUESTION 4

Your team of Infrastructure DevOps Engineers is growing, and you are starting to use Terraform to manage infrastructure. You need a way to implement code versioning and to share code with other team members. What should you do?

A. Store the Terraform code in a version-control system. Establish procedures for pushing new versions and merging with the master.

B. Store the Terraform code in a network shared folder with child folders for each version release. Ensure that everyone works on different files.

C. Store the Terraform code in a Cloud Storage bucket using object versioning. Give access to the bucket to every team member so they can download the files.

D. Store the Terraform code in a shared Google Drive folder so it syncs automatically to every team member\\'s computer. Organize files with a naming convention that identifies each new version.

Correct Answer: A

Reference: https://www.terraform.io/docs/cloud/guides/recommended-practices/part3.3.html

QUESTION 5

Some of your production services are running in Google Kubernetes Engine (GKE) in the eu-west-1 region. Your build system runs in the us-west-1 region. You want to push the container images from your build system to a scalable registry to maximize the bandwidth for transferring the images to the cluster. What should you do?

A. Push the images to Google Container Registry (GCR) using the gcr.io hostname.

B. Push the images to Google Container Registry (GCR) using the us.gcr.io hostname.

C. Push the images to Google Container Registry (GCR) using the eu.gcr.io hostname.

D. Push the images to a private image registry running on a Compute Engine instance in the eu-west-1 region.

Correct Answer: C

Pushing the images to Google Container Registry (GCR) using the eu.gcr.io hostname will allow the images to be transferred to the GKE cluster in the eu-west-1 region with the best possible network performance. This will minimize the latency when the cluster pulls the images from the registry, maximizing the bandwidth for transferring the images to the cluster.

QUESTION 6

Your company recently migrated to Google Cloud. You need to design a fast, reliable, and repeatable solution for your company to provision new projects and basic resources in Google Cloud. What should you do?

A. Use the Google Cloud console to create projects.

B. Write a script by using the gcloud CLI that passes the appropriate parameters from the request. Save the script in a Git repository.

C. Write a Terraform module and save it in your source control repository. Copy and run the terraform apply command to create the new project.

D. Use the Terraform repositories from the Cloud Foundation Toolkit. Apply the code with appropriate parameters to create the Google Cloud project and related resources.

Correct Answer: D

QUESTION 7

Your application artifacts are being built and deployed via a CI/CD pipeline. You want the CI/CD pipeline to securely access application secrets. You also want to more easily rotate secrets in case of a security breach. What should you do?

A. Prompt developers for secrets at build time. Instruct developers to not store secrets at rest.

B. Store secrets in a separate configuration file on Git. Provide select developers with access to the configuration file.

C. Store secrets in Cloud Storage encrypted with a key from Cloud KMS. Provide the CI/CD pipeline with access to Cloud KMS via IAM.

D. Encrypt the secrets and store them in the source code repository. Store a decryption key in a separate repository and

grant your pipeline access to it.

Correct Answer: C

By storing secrets in Cloud Storage, you can take advantage of the security features provided by the platform and encrypt them using Cloud KMS, a GCP service that allows you to create, manage, and use encryption keys. This way you can control who has access to the secrets, and you can easily rotate the encryption keys in case of a security breach. Additionally, you can use IAM to give the CI/CD pipeline the necessary permissions to access the secrets and use them during the deployment process, without the need to store them in the source code or give access to them to specific developers.

QUESTION 8

Your organization is starting to containerize with Google Cloud. You need a fully managed storage solution for container images and Helm charts. You need to identify a storage solution that has native integration into existing Google Cloud services, including Google Kubernetes Engine (GKE), Cloud Run, VPC Service Controls, and Identity and Access Management (IAM). What should you do?

A. Use Docker to configure a Cloud Storage driver pointed at the bucket owned by your organization.

B. Configure an open source container registry server to run in GKE with a restrictive role-based access control (RBAC) configuration.

C. Configure Artifact Registry as an OCI-based container registry for both Helm charts and container images.

D. Configure Container Registry as an OCI-based container registry for container images.

Correct Answer: C

https://cloud.google.com/artifact-registry/docs/helm

QUESTION 9

Your team is writing a postmortem after an incident on your external facing application. Your team wants to improve the postmortem policy to include triggers that indicate whether an incident requires a postmortem. Based on Site Reliability Engineering (SRE) practices, what triggers should be defined in the postmortem policy? (Choose two.)

A. An external stakeholder asks for a postmortem

B. Data is lost due to an incident.

C. An internal stakeholder requests a postmortem.

D. The monitoring system detects that one of the instances for your application has failed.

E. The CD pipeline detects an issue and rolls back a problematic release.

Correct Answer: BE

QUESTION 10

You have deployed a fleet of Compute Engine instances in Google Cloud. You need to ensure that monitoring metrics and logs for the instances are visible in Cloud Logging and Cloud Monitoring by your company\\'s operations and cyber security teams. You need to grant the required roles for the Compute Engine service account by using Identity and Access Management (IAM) while following the principle of least privilege. What should you do?

A. Grant the logging.logWriter and monitoring.metricWriter roles to the Compute Engine service accounts.

B. Grant the logging.admin and monitoring.editor roles to the Compute Engine service accounts.

C. Grant the logging.editor and monitoring.metricWriter roles to the Compute Engine service accounts.

D. Grant the logging.logWriter and monitoring.editor roles to the Compute Engine service accounts.

Correct Answer: A

Remove admin role from the options and there is no such role as logging.editor

---

QUESTION 11

You are deploying a Cloud Build job that deploys Terraform code when a Git branch is updated. While testing, you noticed that the job fails. You see the following error in the build logs:

Initializing the backend...

Error: Failed to get existing workspaces: querying Cloud Storage failed: googleapi: Error 403

You need to resolve the issue by following Google-recommended practices. What should you do?

A. Change the Terraform code to use local state.

B. Create a storage bucket with the name specified in the Terraform configuration.

C. Grant the roles/owner Identity and Access Management (IAM) role to the Cloud Build service account on the project.

D. Grant the roles/storage.objectAdmin Identity and Access Management (1AM) role to the Cloud Build service account on the state file bucket.

Correct Answer: D

QUESTION 12

A third-party application needs to have a service account key to work properly. When you try to export the key from your cloud project, you receive an error: “The organization policy constraint iam.disableServiceAccounKeyCreation is enforced.” You need to make the third-party application work while following Google-recommended security practices.

What should you do?

A. Enable the default service account key, and download the key.

B. Remove the iam.disableServiceAccountKeyCreation policy at the organization level, and create a key.

C. Disable the service account key creation policy at the project\\'s folder, and download the default key.

D. Add a rule to set the iam.disableServiceAccountKeyCreation policy to off in your project, and create a key.

Correct Answer: B

QUESTION 13

You support a multi-region web service running on Google Kubernetes Engine (GKE) behind a Global HTTP/S Cloud Load Balancer (CLB). For legacy reasons, user requests first go through a third-party Content Delivery Network (CDN), which then routes traffic to the CLB. You have already implemented an availability Service Level Indicator (SLI) at the CLB level. However, you want to increase coverage in case of a potential load balancer misconfiguration, CDN failure, or other global networking catastrophe. Where should you measure this new SLI? (Choose two.)

A. Your application servers\\' logs.

B. Instrumentation coded directly in the client.

C. Metrics exported from the application servers.

D. GKE health checks for your application servers.

E. A synthetic client that periodically sends simulated user requests.

Correct Answer: BE

https://cloud.google.com/architecture/adopting-slos#choosing_a_measurement_method

B > Using client instrumentation.

E > Implementing synthetic testing.

QUESTION 14

Your team is designing a new application for deployment both inside and outside Google Cloud Platform (GCP). You need to collect detailed metrics such as system resource utilization. You want to use centralized GCP services while minimizing the amount of work required to set up this collection system. What should you do?

A. Import the Stackdriver Profiler package, and configure it to relay function timing data to Stackdriver for further analysis.

B. Import the Stackdriver Debugger package, and configure the application to emit debug messages with timing information.

C. Instrument the code using a timing library, and publish the metrics via a health check endpoint that is scraped by Stackdriver.

D. Install an Application Performance Monitoring (APM) tool in both locations, and configure an export to a central data storage location for analysis.

Correct Answer: A

https://cloud.google.com/profiler/docs/about-profiler

Cloud Profiler is a statistical, low-overhead profiler that continuously gathers CPU usage and memory-allocation

information from your production applications.

QUESTION 15

You encounter a large number of outages in the production systems you support. You receive alerts for all the outages, the alerts are due to unhealthy systems that are automatically restarted within a minute. You want to set up a process that would prevent staff burnout while following Site Reliability Engineering (SRE) practices. What should you do?

A. Eliminate alerts that are not actionable

B. Redefine the related SLO so that the error budget is not exhausted

C. Distribute the alerts to engineers in different time zones

D. Create an incident report for each of the alerts

Correct Answer: A