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.pass4itsure.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

Latest PROFESSIONAL-CLOUD-DEVOPS-ENGINEER Dumps | PROFESSIONAL-CLOUD-DEVOPS-ENGINEER
VCE Dumps | PROFESSIONAL-CLOUD-DEVOPS-ENGINEER Exam Questions

1 / 5

Latest PROFESSIONAL-CLOUD-DEVOPS-ENGINEER Dumps | PROFESSIONAL-CLOUD-DEVOPS-ENGINEER
VCE Dumps | PROFESSIONAL-CLOUD-DEVOPS-ENGINEER Exam Questions

2 / 5

QUESTION 1

Your product is currently deployed in three Google Cloud Platform (GCP) zones with your users divided between the zones. You can fail over from one zone to another, but it causes a 10-minute service disruption for the affected users. You typically experience a database failure once per quarter and can detect it within five minutes. You are cataloging the reliability risks of a new real-time chat feature for your product. You catalog the following information for each risk: Mean Time to Detect (MTTD) in minutes Mean Time to Repair (MTTR) in minutes Mean Time Between Failure (MTBF) in days User Impact Percentage

The chat feature requires a new database system that takes twice as long to successfully fail over between zones. You want to account for the risk of the new database failing in one zone. What would be the values for the risk of database failover with the new system?

A. MTTD: 5 MTTR: 10 MTBF: 90 Impact: 33%

B. MTTD: 5 MTTR: 20 MTBF: 90 Impact: 33%

C. MTTD: 5 MTTR: 10 MTBF: 90 Impact: 50%

D. MTTD: 5 MTTR: 20 MTBF: 90 Impact: 50%

Correct Answer: B

QUESTION 2

You manage several production systems that run on Compute Engine in the same Google Cloud Platform (GCP) project. Each system has its own set of dedicated Compute Engine instances. You want to know how must it costs to run each of the systems. What should you do?

A. In the Google Cloud Platform Console, use the Cost Breakdown section to visualize the costs per system.

B. Assign all instances a label specific to the system they run. Configure BigQuery billing export and query costs per label.

C. Enrich all instances with metadata specific to the system they run. Configure Stackdriver Logging to export to BigQuery, and query costs based on the metadata.

D. Name each virtual machine (VM) after the system it runs. Set up a usage report export to a Cloud Storage bucket. Configure the bucket as a source in BigQuery to query costs based on VM name.

Correct Answer: B

https://cloud.google.com/billing/docs/how-to/export-data-bigquery

QUESTION 3

You need to deploy a new service to production. The service needs to automatically scale using a managed instance group and should be deployed across multiple regions. The service needs a large number of resources for each instance and you need to plan for capacity. What should you do?

A. Monitor results of Cloud Trace to determine the optimal sizing.

B. Use the n2-highcpu-96 machine type in the configuration of the managed instance group.

C. Deploy the service in multiple regions and use an internal load balancer to route traffic.

D. Validate that the resource requirements are within the available project quota limits of each region.

Correct Answer: D

QUESTION 4

Your company has a Google Cloud resource hierarchy with folders for production, test, and development. Your cyber security team needs to review your company\\'s Google Cloud security posture to accelerate security issue identification and resolution. You need to centralize the logs generated by Google Cloud services from all projects only inside your production folder to allow for alerting and near-real time analysis. What should you do?

A. Enable the Workflows API and route all the logs to Cloud Logging.

B. Create a central Cloud Monitoring workspace and attach all related projects.

C. Create an aggregated log sink associated with the production folder that uses a Pub/Sub topic as the destination.

D. Create an aggregated log sink associated with the production folder that uses a Cloud Logging bucket as the destination.

Correct Answer: C

https://cloudplatform.googleblog.com/2015/06/Real-Time-Log-Streaming-and-Analysis-with-Google-Cloud-Platform-Logentries.html

QUESTION 5

You are building the CI/CD pipeline for an application deployed to Google Kubernetes Engine (GKE). The application is deployed by using a Kubernetes Deployment, Service, and Ingress. The application team asked you to deploy the application by using the blue/green deployment methodology. You need to implement the rollback actions. What should you do?

A. Run the kubectl rollout undo command.

B. Delete the new container image, and delete the running Pods.

C. Update the Kubernetes Service to point to the previous Kubernetes Deployment.

D. Scale the new Kubernetes Deployment to zero.

Correct Answer: C

The point of using blue/green deployment is to have both blue and green versions deployed, and we chose the active one at service level, so rollback should only consist of updating the service to serve the other version.

Latest PROFESSIONAL-CLOUD-DEVOPS-ENGINEER Dumps

PROFESSIONAL-CLOUD-DEVOPS-ENGINEER VCE Dumps

PROFESSIONAL-CLOUD-DEVOPS-ENGINEER Exam Questions