

Latest Version: 6.0

Question: 1

Your application starts on the VM as a systemd service. Your application outputs its log information to stdout. You need to send the application logs to Stackdriver without changing the application. What should you do?

Response:

- A. Review the application logs from the Compute Engine VM Instance activity logs in Stackdriver.
- B. Review the application logs from the Compute Engine VM Instance data access logs in Stackdriver.
- C. Install Stackdriver Logging Agent. Review the application logs from the Compute Engine VM Instance syslog logs in Stackdriver.
- D. Install Stackdriver Logging Agent. Review the application logs from the Compute Engine VM Instance system event logs in Stackdriver.

Answer: C

Question: 2

Your company plans to expand their analytics use cases. One of the new use cases requires your data analysts to analyze events using SQL on a near real-time basis. You expect rapid growth and want to use managed services as much as possible.

What should you do?

Response:

- A. Create a Cloud Pub/Sub topic and a subscription. Stream your events from the source into the Pub/Sub topic. Leverage Cloud Dataflow to ingest these events into BigQuery.
- B. Create a Cloud Pub/Sub topic and a subscription. Stream your events from the source into the Pub/Sub topic. Leverage Cloud Dataflow to ingest these events into Cloud Storage.
- C. Create a Kafka instance on a large Compute Engine instance. Stream your events from the source into a Kafka pipeline. Leverage Cloud Dataflow to ingest these events into Cloud Storage.
- D. Create a Cloud Pub/Sub topic and a subscription. Stream your events from the source into the Pub/Sub topic. Leverage Cloud Dataflow to ingest these events into Cloud Datastore.

Answer: A

Question: 3

Which architecture should HipLocal use for log analysis?

Response:

- A. Use Cloud Spanner to store each event.
- B. Start storing key metrics in Cloud Memorystore.
- C. Use Stackdriver Logging with a BigQuery sink.
- D. Use Stackdriver Logging with a Cloud Storage sink.

Answer: C

Question: 4

As part of their expansion, HipLocal is creating new projects in order to separate resources. They want to build a system to automate enabling of their APIs. What should they do?

Response:

- A. Copy existing persistent disks to the new project.
- B. Use the service management API to define a new service.
- C. Use the service management API to enable the Compute API.
- D. Use the service management API to enable the Cloud Storage API.

Answer: C

Question: 5

You have a service running on Compute Engine virtual machine instances behind a global load balancer. You need to ensure that when the instance fails, it is recovered. What should you do?

Response:

- A. Set up health checks in the load balancer configuration.
- B. Deploy a service to the instances to notify you when they fail.
- C. Use Stackdriver alerting to trigger a workflow to reboot the instance.
- D. Set up health checks in the managed instance group configuration.

Answer: D

Question: 6

You are capturing important audit activity in Stackdriver Logging. You need to read the information from Stackdriver Logging to perform real-time analysis of the logs. You will have multiple processes performing different types of analysis on the logging data.

What should you do?

Response:

- A. Read the logs directly from the Stackdriver Logging API.
- B. Set up a Stackdriver Logging sync to BigQuery, and read the logs from the BigQuery table.
- C. Set up a Stackdriver Logging sync to Cloud Pub/Sub, and read the logs from a Cloud Pub/Sub topic.
- D. Set up a Stackdriver Logging sync to Cloud Storage, and read the logs from a Cloud Storage bucket.

Answer: C

Question: 7

Your organization has grown, and new teams need access to manage network connectivity within and across projects. You are now seeing intermittent timeout errors in your application. You want to find the cause of the problem. What should you do?

Response:

- A. Set up Wireshark on each Google Cloud Virtual Machine instance.
- B. Configure VPC flow logs for each of the subnets in your VPC.
- C. Review the instance admin activity logs in Stackdriver for the application instances.
- D. Configure firewall rules logging for each of the firewalls in your VPC.

Answer: B

Question: 8

Your company has a successful multi-player game that has become popular in the US. Now, it wants to expand to other regions. It is launching a new feature that allows users to trade points. This feature will work for users across the globe.

Your company's current MySQL backend is reaching the limit of the Compute Engine instance that hosts the game. Your company wants to migrate to a different database that will provide global consistency and high availability across the regions.

Which database should they choose?

Response:

- A. BigQuery
- B. Cloud Spanner
- C. Cloud SQL
- D. Cloud Bigtable

Answer: B

Question: 9

You have an application that accepts inputs from users. The application needs to kick off different background tasks based on these inputs. You want to allow for automated asynchronous execution of these tasks as soon as input is submitted by the user.

Which product should you use?

Response:

- A. Cloud Tasks
- B. Cloud Bigtable
- C. Cloud Pub/Sub
- D. Cloud Composer

Answer: A

Question: 10

You are building a storage layer for an analytics Hadoop cluster for your company. This cluster will run multiple jobs on a nightly basis, and you need to access the data frequently. You want to use Cloud Storage for this purpose.

Which storage option should you choose?

Response:

- A. Multi-regional storage
- B. Regional storage
- C. Nearline storage
- D. Coldline storage

Answer: B