Free Questions for Professional-Cloud-DevOps-Engineer

Shared by Blackwell on 29-01-2024

Question 1

Question Type: MultipleChoice

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?

Options:

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 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.

Answer:

D

Explanation:

Terraform is an open-source tool that allows you to define and provision infrastructure as code1.Terraform can be used to create and manage Google Cloud resources, such as projects, networks, and services2.The Cloud Foundation Toolkit is a set of open-source Terraform modules and tools that provide best practices and guidance for deploying Google Cloud infrastructure3.The Cloud Foundation Toolkit includes Terraform repositories for creating Google Cloud projects and related resources, such as IAM policies, APIs, service accounts, and billing4. By using the Terraform repositories from the Cloud Foundation Toolkit, you can design a fast, reliable, and repeatable solution for your company to provision new projects and basic resources in Google Cloud. You can also customize the Terraform code to suit your specific needs and preferences.

Question 2

Question Type: MultipleChoice

You support a user-facing web application. When analyzing the application's error budget over the previous six months, you notice that the application has never consumed more than 5% of its error budget in any given time window. You hold a Service Level Objective (SLO) review with business stakeholders and confirm that the SLO is set appropriately. You want your application's SLO to more closely reflect its observed reliability. What steps can you take to further that goal while balancing velocity, reliability, and business needs? (Choose two.)

Options:

A- Add more serving capacity to all of your application's zones.

B- Have more frequent or potentially risky application releases.

C- Tighten the SLO match the application's observed reliability.

D- Implement and measure additional Service Level Indicators (SLIs) fro the application.

E- Announce planned downtime to consume more error budget, and ensure that users are not depending on a tighter SLO.

Answer:

D, E

Explanation:

https://sre.google/sre-book/service-level-objectives/

You want the application's SLO to more closely reflect it's observed reliability. The key here is error budget never goes over 5%. This means they can have additional downtime and still stay within their budget.

Question 3

Question Type: MultipleChoice

You need to define SLOs for a high-traffic web application. Customers are currently happy with the application performance and availability. Based on current measurement, the 90th percentile Of latency is 160 ms and the 95th percentile of latency is 300 ms over a 28-day window. What latency SLO should you publish?

Options:

A- 90th percentile - 150 ms
95th percentile - 290 ms

B- 90th percentile - 160 ms
95th percentile - 300 ms
C- 90th percentile - 190 ms
95th percentile - 330 ms
D- 90th percentile - 300 ms
95th percentile - 450 ms

Answer:

B

Explanation:

a latency SLO is a service level objective that specifies a target level of responsiveness for a web application1.A latency SLO can be expressed as a percentile of latency over a time window, such as the 90th percentile of latency over 28 days2. A percentile of latency is the maximum amount of time that a given percentage of requests take to complete.For example, the 90th percentile of latency is the maximum amount of time that 90% of requests take to complete3.

To define a latency SLO, you need to consider the following factors24:

The expectations and satisfaction of your customers. You want to set a latency SLO that reflects the level of performance that your customers are happy with and willing to pay for.

The current and historical measurements of your latency. You want to set a latency SLO that is based on data and realistic for your web application.

The trade-offs and costs of improving your latency. You want to set a latency SLO that balances the benefits of faster response times with the costs of engineering work, infrastructure, and complexity.

Based on these factors, the best option for defining a latency SLO for your web application is option B. Option B sets the latency SLO to match the current measurement of your latency, which means that you are meeting the expectations and satisfaction of your customers. Option B also sets a realistic and achievable target for your web application, which means that you do not need to invest extra resources or effort to improve your latency.Option B also aligns with the best practice of setting conservative SLOs, which means that you have some buffer or margin for error in case your latency fluctuates or degrades5.

Question 4

Question Type: MultipleChoice

You are the on-call Site Reliability Engineer for a microservice that is deployed to a Google

Kubernetes Engine (GKE) Autopilot cluster. Your company runs an online store that publishes order messages to Pub/Sub and a microservice receives these messages and updates stock information in the warehousing system. A sales event caused an increase in orders, and the stock information is not being updated quickly enough. This is causing a large number of orders to be accepted for products that are out of stock You check the metrics for the microservice and compare them to typical levels.

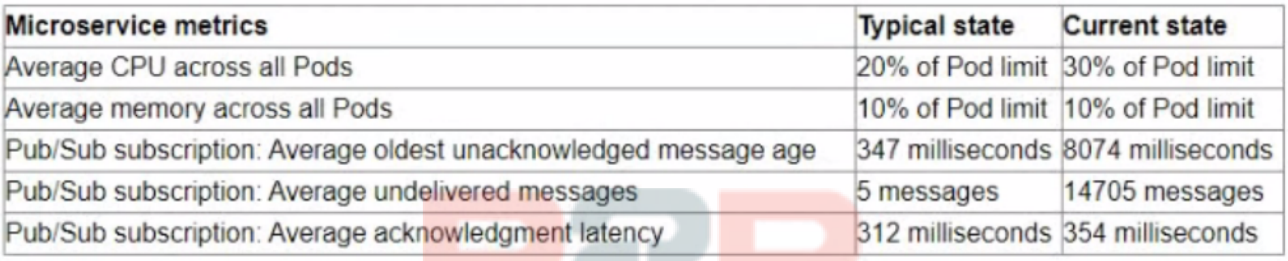

Microservice metrics	Typical state	Current state
Average CPU across all Pods	20% of Pod limit	30% of Pod limit
Average memory across all Pods	10% of Pod limit	10% of Pod limit
Pub/Sub subscription: Average oldest unacknowledged message age	347 milliseconds	8074 milliseconds
Pub/Sub subscription: Average undelivered messages	5 messages	14705 messages
Pub/Sub subscription: Average acknowledgment latency	312 milliseconds	354 milliseconds

You need to ensure that the warehouse system accurately reflects product inventory at the time orders are placed and minimize the impact on customers What should you do?

Options:

A- Decrease the acknowledgment deadline on the subscription

B- Add a virtual queue to the online store that allows typical traffic levels

C- Increase the number of Pod replicas

D- Increase the Pod CPU and memory limits

Answer:

C

Explanation:

The best option for ensuring that the warehouse system accurately reflects product inventory at the time orders are placed and minimizing the impact on customers is to increase the number of Pod replicas. Increasing the number of Pod replicas will increase the scalability and availability of your microservice, which will allow it to handle more Pub/Sub messages and update stock information faster. This way, you can reduce the backlog of undelivered messages and oldest unacknowledged message age, which are causing delays in updating product inventory. You can use Horizontal Pod Autoscaler or Cloud Monitoring metrics-based autoscaling to automatically adjust the number of Pod replicas based on load or custom metrics.

Question 5

Question Type: MultipleChoice

As a Site Reliability Engineer, you support an application written in GO that runs on Google Kubernetes Engine (GKE) in production. After releasing a new version Of the application, you notice the application runs for about 15 minutes and then restarts. You decide to add Cloud Profiler to your application and now notice that the heap usage grows constantly until the application restarts. What should you do?

Options:

A- Add high memory compute nodes to the cluster.

B- Increase the memory limit in the application deployment.

C- Add Cloud Trace to the application, and redeploy.

D- Increase the CPU limit in the application deployment.

Answer:

B

Explanation:

The correct answer is B, Increase the memory limit in the application deployment.

The application is experiencing a memory leak, which means that it is allocating memory that is not freed or reused. This causes the heap usage to grow constantly until it reaches the memory limit of the pod, which triggers a restart by Kubernetes. Increasing the memory limit in the application deployment can help mitigate the problem by allowing the application to run longer before reaching the limit. However, this is not a permanent solution, as the memory leak will still occur and eventually exhaust the available memory. The best solution is to identify and fix the source of the memory leak in the application code, using tools like Cloud Profiler and pprof12.

Using Cloud Profiler with Go, Troubleshooting memory leaks. Profiling Go Programs, Heap profiles.

Question 6

Question Type: MultipleChoice

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?

Options:

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 (IAM) role to the Cloud Build service account on the state file bucket.

Answer:

D

Explanation:

The correct answer is D. Grant the roles/storage.objectAdmin Identity and Access Management (IAM) role to the Cloud Build service account on the state file bucket.

According to the Google Cloud documentation, Cloud Build is a service that executes your builds on Google Cloud Platform infrastructure1. Cloud Build uses a service account to execute your build steps and access resources, such as Cloud Storage buckets2. Terraform is an open-source tool that allows you to define and provision infrastructure as code3. Terraform uses a state file to store and track the state of your infrastructure4. You can configure Terraform to use a Cloud Storage bucket as a backend to store and share the state file across multiple users or environments5.

The error message indicates that Cloud Build failed to access the Cloud Storage bucket that contains the Terraform state file. This is likely because the Cloud Build service account does not have the necessary permissions to read and write objects in the bucket. To resolve this issue, you need to grant the roles/storage.objectAdmin IAM role to the Cloud Build service account on the state file bucket. This role allows the service account to create, delete, and manage objects in the bucket6. You can use the gcloud command-line tool or the Google Cloud Console to grant this role.

The other options are incorrect because they do not follow Google-recommended practices. Option A is incorrect because it changes the Terraform code to use local state, which is not recommended for production or collaborative environments, as it can cause conflicts, data loss, or inconsistency. Option B is incorrect because it creates a new storage bucket with the name specified in the Terraform configuration, but it does not grant any permissions to the Cloud Build service account on the new bucket. Option C is incorrect because it grants the roles/owner IAM role to the Cloud Build service account on the project, which is too broad and violates the principle of least privilege. The roles/owner role grants full access to all resources in the project, which can pose a security risk if misused or compromised.

Cloud Build Documentation, Overview. Service accounts, Service accounts. Terraform by HashiCorp, Terraform by HashiCorp. State, State. Google Cloud Storage Backend, Google Cloud Storage Backend. Predefined roles, Predefined roles. [Granting roles to service accounts for specific resources], Granting roles to service accounts for specific resources. [Local Backend], Local Backend. [Understanding roles], Understanding roles.

Question 7

Question Type: MultipleChoice

Your organization wants to collect system logs that will be used to generate dashboards in Cloud Operations for their Google Cloud project. You need to configure all current and future Compute Engine instances to collect the system logs and you must ensure that the Ops Agent remains up to date. What should you do?

Options:

A- Use the gcloud CLI to install the Ops Agent on each VM listed in the Cloud Asset Inventory

B- Select all VMs with an Agent status of Not detected on the Cloud Operations VMs dashboard Then select Install agents

C- Use the gcloud CLI to create an Agent Policy.

D- Install the Ops Agent on the Compute Engine image by using a startup script

Answer:

C

Explanation:

The best option for configuring all current and future Compute Engine instances to collect system logs and ensure that the Ops Agent remains up to date is to use the gcloud CLI to create an Agent Policy. An Agent Policy is a resource that defines how Ops Agents are installed and configured on VM instances that match certain criteria, such as labels or zones. Ops Agents are software agents that collect metrics and logs from VM instances and send them to Cloud

Operations products, such as Cloud Monitoring and Cloud Logging. By creating an Agent Policy, you can ensure that all current and future VM instances that match the policy criteria will have the Ops Agent installed and updated automatically. This way, you can collect system logs from all VM instances and use them to generate dashboards in Cloud Operations.

Question 8

Question Type: MultipleChoice

Your company runs applications in Google Kubernetes Engine (GKE). Several applications rely on ephemeral volumes. You noticed some applications were unstable due to the DiskPressure node condition on the worker nodes. You need to identify which Pods are causing the issue, but you do not have execute access to workloads and nodes. What should you do?

Options:

A- Check the node/ephemeral_storage/used_bytes metric by using Metrics Explorer.

B- Check the metric by using Metrics Explorer.

C- Locate all the Pods with emptyDir volumes. use the df-h command to measure volume disk usage.

D- Locate all the Pods with emptyDir volumes. Use the du -sh * command to measure volume disk usage.

Answer:

A

Explanation:

The correct answer is A, Check the node/ephemeral_storage/used_bytes metric by using Metrics Explorer.

The node/ephemeral_storage/used_bytes metric reports the total amount of ephemeral storage used by Pods on each node1. You can use Metrics Explorer to query and visualize this metric and filter it by node name, namespace, or Pod name2. This way, you can identify which Pods are consuming the most ephemeral storage and causing disk pressure on the nodes. You do not need to have execute access to the workloads or nodes to use Metrics Explorer.

The other options are incorrect because they require execute access to the workloads or nodes, which you do not have. The df -h and du -sh * commands are Linux commands that can measure disk usage, but you need to run them inside the Pods or on the nodes, which is not possible in your scenario34.

Monitoring metrics for Kubernetes system components, Node metrics, node/ephemeral_storage/used_bytes. Using Metrics Explorer, Querying metrics. How do I find out disk space utilization information using Linux command line?, df command. How to check disk space in Linux from the command line, du command.

Question 9

Question Type: MultipleChoice

You are building and running client applications in Cloud Run and Cloud Functions Your client requires that all logs must be available for one year so that the client can import the logs into their logging service You must minimize required code changes What should you do?

Options:

A- Update all images in Cloud Run and all functions in Cloud Functions to send logs to both Cloud Logging and

the client's logging service Ensure that all the ports required to send logs are open in the VPC firewall

B- Create a Pub/Sub topic subscription and logging sink Configure the logging sink to send all logs into the

topic Give your client access to the topic to retrieve the logs

C- Create a storage bucket and appropriate VPC firewall rules Update all images in Cloud Run and all

functions in Cloud Functions to send logs to a file within the storage bucket

D- Create a logs bucket and logging sink. Set the retention on the logs bucket to 365 days Configure the

logging sink to send logs to the bucket Give your client access to the bucket to retrieve the logs

Answer:

D

Explanation:

The best option for storing all logs for one year and minimizing required code changes is to create a logs bucket and logging sink, set the retention on the logs bucket to 365 days, configure the logging sink to send logs to the bucket, and give your client access to the bucket to retrieve the logs. A logs bucket is a Cloud Storage bucket that is used to store logs from Cloud Logging. A logging sink is a resource that defines where log entries are sent, such as a logs bucket, BigQuery dataset, or Pub/Sub topic. You can create a logs bucket and logging sink in Cloud Logging and set the retention on the logs bucket to 365 days. This way, you can ensure that all logs are stored for one year and protected from deletion. You can also configure the logging sink to send logs from Cloud Run and Cloud Functions to the logs bucket without any code changes. You can then give your client access to the logs bucket by using IAM policies or signed URLs.

Question 10

Question Type: MultipleChoice

You want to share a Cloud Monitoring custom dashboard with a partner team What should you do?

Options:

A- Provide the partner team with the dashboard URL to enable the partner team to create a copy of the dashboard

B- Export the metrics to BigQuery Use Looker Studio to create a dashboard, and share the dashboard with the partner team

C- Copy the Monitoring Query Language (MQL) query from the dashboard; and send the MQL query to the partner team

D- Download the JSON definition of the dashboard, and send the JSON file to the partner team

Answer:

A

Explanation:

The best option for sharing a Cloud Monitoring custom dashboard with a partner team is to provide the partner team with the dashboard URL to enable the partner team to create a copy of the dashboard. A Cloud Monitoring custom dashboard is a dashboard that allows you to create and customize charts and widgets to display metrics, logs, and traces from your Google Cloud resources and applications. You can share a custom dashboard with a partner team by providing them with the dashboard URL, which is a link that allows them to view the dashboard in their browser. The partner team can then create a copy of the dashboard in their own project by using the Copy Dashboard option. This way, they can access and modify the dashboard without affecting the original one.

Question 11

Question Type: MultipleChoice

You are designing a new Google Cloud organization for a client. Your client is concerned with the risks associated with long-lived credentials created in Google Cloud. You need to design a solution to completely eliminate the risks associated with the use of JSON service account keys while minimizing operational overhead. What should you do?

Options:

A- Use custom versions of predefined roles to exclude all iam.serviceAccountKeys. * service account role permissions.

B- Apply the constraints/iam.disableserviceAccountKeycreation constraint to the organization.

C- Apply the constraints/iam. disableServiceAccountKeyUp10ad constraint to the organization.

D- Grant the roles/ iam.serviceAccountKeyAdmin IAM role to organization administrators only.

Answer:

B

Explanation:

The correct answer is B, Apply the constraints/iam.disableServiceAccountKeyCreation constraint to the organization.

According to the Google Cloud documentation, the constraints/iam.disableServiceAccountKeyCreation constraint is an organization policy constraint that prevents the creation of user-managed service account keys1. User-managed service account keys are long-lived credentials that can be downloaded as JSON or P12 files and used to authenticate as a service account2. These keys pose severe security risks if they are leaked, stolen, or misused by unauthorized entities34. By applying this constraint to the organization, you can completely eliminate the risks associated with the use of JSON service account keys and enforce a more secure alternative for authentication, such as Workload Identity or short-lived access tokens12. This also minimizes operational overhead by avoiding the need to manage, rotate, or revoke user-managed service account keys.

The other options are incorrect because they do not completely eliminate the risks associated with the use of JSON service account keys. Option A is incorrect because it only restricts the IAM permissions to create, list, get, delete, or sign service account keys, but it does not prevent existing keys from being used or leaked. Option C is incorrect because it only disables the upload of user-managed service account keys, but it does not prevent the creation or download of such keys. Option D is incorrect because it only limits the IAM role that can create and manage service account keys, but it does not prevent the keys from being distributed or exposed to unauthorized entities.

Disable user-managed service account key creation, Disable user-managed service account key creation. Service accounts, User-managed service accounts. Help keep your Google Cloud service account keys safe, Help keep your Google Cloud service account keys safe. Stop Downloading Google Cloud Service Account Keys!, Stop Downloading Google Cloud Service Account Keys! [Service Account Keys], Service Account Keys. [Disable user-managed service account key upload], Disable user-managed service account key upload. [Granting roles to service accounts], Granting roles to service accounts.