

Google.Professional-Cloud-DevOps-Engineer.v2024-03-30.q107

Exam Code:	Professional-Cloud-DevOps-Engineer
Exam Name:	Google Cloud Certified - Professional Cloud DevOps Engineer Exam
Certification Provider:	Google
Free Question Number:	107
Version:	v2024-03-30
# of views:	539
# of Questions views:	1070
https://www.freeqas.com/qa/Google/Professional-Cloud-DevOps-Engineer/Google.Professional-Cloud-DevOps-Engineer.v2024-03-30.q107.html	

NEW QUESTION: 1

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?

- 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 ([LEAVE A REPLY](#))

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.

NEW QUESTION: 2

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?

- 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 (LEAVE A REPLY)

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 entities^{3,4}. 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.

NEW QUESTION: 3

You support an e-commerce application that runs on a large Google Kubernetes Engine (GKE) cluster deployed on-premises and on Google Cloud Platform. The application consists of microservices that run in containers. You want to identify containers that are using the most CPU and memory. What should you do?

- A.** Use Stackdriver Kubernetes Engine Monitoring.
- B.** Use Prometheus to collect and aggregate logs per container, and then analyze the results in Grafana.
- C.** Use the Stackdriver Monitoring API to create custom metrics, and then organize your containers using groups.
- D.** Use Stackdriver Logging to export application logs to BigQuery, aggregate logs per container, and then analyze CPU and memory consumption.

Answer: A ([LEAVE A REPLY](#))

<https://cloud.google.com/anthos/clusters/docs/on-prem/1.7/concepts/logging-and-monitoring>

NEW QUESTION: 4

Your company follows Site Reliability Engineering practices. You are the person in charge of Communications for a large, ongoing incident affecting your customer-facing applications. There is still no estimated time for a resolution of the outage. You are receiving emails from internal stakeholders who want updates on the outage, as well as emails from customers who want to know what is happening. You want to efficiently provide updates to everyone affected by the outage. What should you do?

- A.** Provide all internal stakeholder emails to the Incident Commander, and allow them to manage internal communications. Focus on providing responses directly to customers.
- B.** Delegate the responding to internal stakeholder emails to another member of the Incident Response Team.

Focus on providing responses directly to customers.

- C.** Focus on responding to internal stakeholders at least every 30 minutes. Commit to "next update" times.
- D.** Provide periodic updates to all stakeholders in a timely manner. Commit to a "next update" time in all communications.

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 5

Your organization wants to implement Site Reliability Engineering (SRE) culture and principles. Recently, a service that you support had a limited outage. A manager on

another team asks you to provide a formal explanation of what happened so they can action remediations. What should you do?

- A.** Develop a postmortem that includes the root causes, resolution, lessons learned, and a prioritized list of action items. Share it with the manager only.
- B.** Develop a postmortem that includes the root causes, resolution, lessons learned, the list of people responsible, and a list of action items for each person. Share it on the engineering organization's document portal.
- C.** Develop a postmortem that includes the root causes, resolution, lessons learned, the list of people responsible, and a list of action items for each person. Share it with the manager only.
- D.** Develop a postmortem that includes the root causes, resolution, lessons learned, and a prioritized list of action items. Share it on the engineering organization's document portal.

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 6

You use Terraform to manage an application deployed to a Google Cloud environment. The application runs on instances deployed by a managed instance group. The Terraform code is deployed by using a CI/CD pipeline. When you change the machine type on the instance template used by the managed instance group, the pipeline fails at the terraform apply stage with the following error message:

You need to update the instance template and minimize disruption to the application and the number of pipeline runs. What should you do?

- A.** Delete the managed instance group and recreate it after updating the instance template.
- B.** Add a new instance template, update the managed instance group to use the new instance template, and delete the old instance template.
- C.** Remove the managed instance group from the Terraform state file, update the instance template, and reimport the managed instance group.
- D.** Set the `create_before_destroy` meta-argument to `true` in the lifecycle block on the instance template.

Answer: D ([LEAVE A REPLY](#))

Explanation:

The best option for updating the instance template and minimizing disruption to the application and the number of pipeline runs is to set the `create_before_destroy` meta-argument to `true` in the lifecycle block on the instance template. The `create_before_destroy` meta-argument is a Terraform feature that specifies that a new resource should be created before destroying an existing one during an update. This way, you can avoid downtime and errors when updating a resource that is in use by another resource, such as an instance template that is used by a managed instance group. By setting the `create_before_destroy` meta-argument to `true` in the lifecycle block on the instance template, you can ensure that Terraform creates a new instance template with the updated machine type, updates the

managed instance group to use the new instance template, and then deletes the old instance template.

NEW QUESTION: 7

You manage an application that is writing logs to Stackdriver Logging. You need to give some team members the ability to export logs. What should you do?

- A. Grant the team members the IAM role of logging.configWriter on Cloud IAM.
- B. Configure Access Context Manager to allow only these members to export logs.
- C. Create and grant a custom IAM role with the permissions logging.sinks.list and logging.sink.get.
- D. Create an Organizational Policy in Cloud IAM to allow only these members to create log exports.

Answer: A (LEAVE A REPLY)

<https://cloud.google.com/logging/docs/access-control>

The logging.configWriter role grants permissions to create, update, and delete log exports. This is the correct role to give team members who need to export logs2.

NEW QUESTION: 8

Your team uses Cloud Build for all CI/CO pipelines. You want to use the kubectl builder for Cloud Build to deploy new images to Google Kubernetes Engine (GKE). You need to authenticate to GKE while minimizing development effort. What should you do?

- A. Assign the Container Developer role to the Cloud Build service account.
- B. Specify the Container Developer role for Cloud Build in the cloudbuild.yaml file.
- C. Create a new service account with the Container Developer role and use it to run Cloud Build.
- D. Create a separate step in Cloud Build to retrieve service account credentials and pass these to kubectl.

Answer: A (LEAVE A REPLY)

<https://cloud.google.com/build/docs/deploying-builds/deploy-gke>

<https://cloud.google.com/build/docs/securing-builds/configure-user-specified-service-accounts>

NEW QUESTION: 9

Your organization recently adopted a container-based workflow for application development. Your team develops numerous applications that are deployed continuously through an automated build pipeline to the production environment. A recent security audit alerted your team that the code pushed to production could contain vulnerabilities and that the existing tooling around virtual machine (VM) vulnerabilities no longer applies to the containerized environment. You need to ensure the security and patch level of all code running through the pipeline. What should you do?

- A. Set up Container Analysis to scan and report Common Vulnerabilities and Exposures.

- B.** Configure the containers in the build pipeline to always update themselves before release.
- C.** Reconfigure the existing operating system vulnerability software to exist inside the container.
- D.** Implement static code analysis tooling against the Docker files used to create the containers.

Answer: ([SHOW ANSWER](#))

<https://cloud.google.com/binary-authorization>

Binary Authorization is a deploy-time security control that ensures only trusted container images are deployed on Google Kubernetes Engine (GKE) or Cloud Run. With Binary Authorization, you can require images to be signed by trusted authorities during the development process and then enforce signature validation when deploying. By enforcing validation, you can gain tighter control over your container environment by ensuring only verified images are integrated into the build-and-release process.

NEW QUESTION: 10

You support a high-traffic web application and want to ensure that the home page loads in a timely manner. As a first step, you decide to implement a Service Level Indicator (SLI) to represent home page request latency with an acceptable page load time set to 100 ms. What is the Google-recommended way of calculating this SLI?

- A.** Bucketize the request latencies into ranges, and then compute the median and 90th percentiles.
- B.** Buckelize the request latencies into ranges, and then compute the percentile at 100 ms.
- C.** Count the number of home page requests that load in under 100 ms. and then divide by the total number of all web application requests.
- D.** Count the number of home page requests that load in under 100 ms, and then divide by the total number of home page requests.

Answer: **B** ([LEAVE A REPLY](#))

NEW QUESTION: 11

You support an application that stores product information in cached memory. For every cache miss, an entry is logged in Stackdriver Logging. You want to visualize how often a cache miss happens over time. What should you do?

- A.** Link Stackdriver Logging as a source in Google Data Studio. Filter (he logs on the cache misses.
- B.** Configure Stackdriver Profiler to identify and visualize when the cache misses occur based on the logs.
- C.** Create a logs-based metric in Stackdriver Logging and a dashboard for that metric in Stackdriver Monitoring.
- D.** Configure BigQuery as a sink for Stackdriver Logging. Create a scheduled query to filter the cache miss logs and write them to a separate table

Answer: C (LEAVE A REPLY)

Explanation

<https://cloud.google.com/logging/docs/logs-based-metrics#counter-metric>

NEW QUESTION: 12

You are developing reusable infrastructure as code modules. Each module contains integration tests that launch the module in a test project. You are using GitHub for source control. You need to Continuously test your feature branch and ensure that all code is tested before changes are accepted. You need to implement a solution to automate the integration tests. What should you do?

- A.** Use a Jenkins server for CI/CD pipelines. Periodically run all tests in the feature branch.
- B.** Use Cloud Build to run the tests. Trigger all tests to run after a pull request is merged.
- C.** Ask the pull request reviewers to run the integration tests before approving the code.
- D.** Use Cloud Build to run tests in a specific folder. Trigger Cloud Build for every GitHub pull request.

Answer: D (LEAVE A REPLY)

Cloud Build is a service that executes your builds on Google Cloud Platform infrastructure. Cloud Build can import source code from Google Cloud Storage, Cloud Source Repositories, GitHub, or Bitbucket, execute a build to your specifications, and produce artifacts such as Docker containers or Java archives¹. Cloud Build can also run integration tests as part of your build steps².

You can use Cloud Build to run tests in a specific folder by specifying the path to the folder in the `dir` field of your build step³. For example, if you have a folder named `tests` that contains your integration tests, you can use the following build step to run them:

steps:

```
- name: 'gcr.io/cloud-builders/go'
```

```
  args: ['test', '-v']
```

```
  dir: 'tests'
```

Copy

You can use Cloud Build to trigger builds for every GitHub pull request by using the Cloud Build GitHub app. The app allows you to automatically build on Git pushes and pull requests and view your build results on GitHub and Google Cloud console⁴. You can configure the app to run builds on specific branches, tags, or paths⁵. For example, if you want to run builds on pull requests that target the master branch, you can use the following trigger configuration:

includedFiles:

```
- '**'
```

```
name: 'pull-request-trigger'
```

```
github:
```

```
  name: 'my-repo'
```

```
  owner: 'my-org'
```

pullRequest:

branch: '^master\$'

Using Cloud Build to run tests in a specific folder and trigger builds for every GitHub pull request is a good way to continuously test your feature branch and ensure that all code is tested before changes are accepted. This way, you can catch any errors or bugs early and prevent them from affecting the main branch.

Using a Jenkins server for CI/CD pipelines is not a bad option, but it would require more setup and maintenance than using Cloud Build, which is fully managed by Google Cloud. Periodically running all tests in the feature branch is not as efficient as running tests for every pull request, as it may delay the feedback loop and increase the risk of conflicts or failures.

Using Cloud Build to run the tests after a pull request is merged is not a good practice, as it may introduce errors or bugs into the main branch that could have been prevented by testing before merging.

Asking the pull request reviewers to run the integration tests before approving the code is not a reliable way of ensuring code quality, as it depends on human intervention and may be prone to errors or oversights.

Reference:

- 1: Overview | Cloud Build Documentation | Google Cloud
- 2: Running integration tests | Cloud Build Documentation | Google Cloud
- 3: Build configuration overview | Cloud Build Documentation | Google Cloud
- 4: Building repositories from GitHub | Cloud Build Documentation | Google Cloud
- 5: Creating GitHub app triggers | Cloud Build Documentation | Google Cloud

NEW QUESTION: 13

You are part of an organization that follows SRE practices and principles. You are taking over the management of a new service from the Development Team, and you conduct a Production Readiness Review (PRR). After the PRR analysis phase, you determine that the service cannot currently meet its Service Level Objectives (SLOs). You want to ensure that the service can meet its SLOs in production. What should you do next?

- A.** Bring the service into production with no SLOs and build them when you have collected operational data.
- B.** Identify recommended reliability improvements to the service to be completed before handover.
- C.** Notify the development team that they will have to provide production support for the service.
- D.** Adjust the SLO targets to be achievable by the service so you can bring it into production.

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 14

Your team has recently deployed an NGINX-based application into Google Kubernetes Engine (GKE) and has exposed it to the public via an HTTP Google Cloud Load Balancer (GCLB) ingress. You want to scale the deployment of the application's frontend using an appropriate Service Level Indicator (SLI). What should you do?

- A.** Expose the NGINX stats endpoint and configure the horizontal pod autoscaler to use the request metrics exposed by the NGINX deployment.
- B.** Configure the horizontal pod autoscaler to use the average response time from the Liveness and Readiness probes.
- C.** Configure the vertical pod autoscaler in GKE and enable the cluster autoscaler to scale the cluster as pods expand.
- D.** Install the Stackdriver custom metrics adapter and configure a horizontal pod autoscaler to use the number of requests provided by the GCLB.

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 15

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.** Configure an open source container registry server to run in GKE with a restrictive role-based access control (RBAC) configuration.
- B.** Configure Container Registry as an OCI-based container registry for container images.
- C.** Use Docker to configure a Cloud Storage driver pointed at the bucket owned by your organization.
- D.** Configure Artifact Registry as an OCI-based container registry for both Helm charts and container images.

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 16

You have a pool of application servers running on Compute Engine. You need to provide a secure solution that requires the least amount of configuration and allows developers to easily access application logs for troubleshooting. How would you implement the solution on GCP?

- A.** * Deploy the Stackdriver logging agent to the application servers.
* Give the developers the IAM Logs Viewer role to access Stackdriver and view logs.
- B.** * Deploy the Stackdriver monitoring agent to the application servers.
* Give the developers the IAM Monitoring Viewer role to access Stackdriver and view metrics.
- C.** * Install the gsutil command line tool on your application servers.

* Write a script using gsutil to upload your application log to a Cloud Storage bucket, and then schedule it to run via cron every 5 minutes.

* Give the developers the IAM Object Viewer access to view the logs in the specified bucket.

D. * Deploy the Stackdriver logging agent to the application servers.

* Give the developers the IAM Logs Private Logs Viewer role to access Stackdriver and view logs.

Answer: D (LEAVE A REPLY)

Valid Professional-Cloud-DevOps-Engineer Dumps shared by PrepPdf.com for Helping Passing Professional-Cloud-DevOps-Engineer Exam! PrepPdf.com now offer the **newest Professional-Cloud-DevOps-Engineer exam dumps**, the PrepPdf.com Professional-Cloud-DevOps-Engineer exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com Professional-Cloud-DevOps-Engineer dumps with Test Engine here:

<https://www.preppdf.com/Google/Professional-Cloud-DevOps-Engineer-prepaway-exam-dumps.html> (200 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 17

Your company follows Site Reliability Engineering practices. You are the Incident Commander for a new, customer-impacting incident. You need to immediately assign two incident management roles to assist you in an effective incident response. What roles should you assign?

Choose 2 answers

A. Communications Lead

B. Operations Lead

C. Customer Impact Assessor

D. External Customer Communications Lead

E. Engineering Lead

Answer: B,C (LEAVE A REPLY)

NEW QUESTION: 18

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 2 answers

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.

Answer: A,C ([LEAVE A REPLY](#))

Explanation

The best options for defining triggers that indicate whether an incident requires a postmortem based on Site Reliability Engineering (SRE) practices are an external stakeholder asks for a postmortem and data is lost due to an incident. An external stakeholder is someone who is affected by or has an interest in the service, such as a customer or a partner. If an external stakeholder asks for a postmortem, it means that they are concerned about the impact or root cause of the incident, and they expect an explanation and remediation from the service provider. Therefore, this should trigger a postmortem to address their concerns and improve their satisfaction. Data loss is a serious consequence of an incident that can affect the integrity and reliability of the service. If data is lost due to an incident, it means that there was a failure in the backup or recovery mechanisms, or that there was a corruption or deletion of data. Therefore, this should trigger a postmortem to investigate the cause and impact of the data loss, and to prevent it from happening again.

NEW QUESTION: 19

Your company's security team needs to have read-only access to Data Access audit logs in the _Required bucket. You want to provide your security team with the necessary permissions following the principle of least privilege and Google-recommended practices. What should you do?

- A. Assign the roles/logging.viewer role to each member of the security team
- B. Assign the roles/logging.viewer role to a group with all the security team members
- C. Assign the roles/logging.privateLogViewer role to each member of the security team
- D. Assign the roles/logging.privateLogviewer role to a group with all the security team members

Answer: ([SHOW ANSWER](#))

Explanation

The best option for providing your security team with the necessary permissions following the principle of least privilege and Google-recommended practices is to assign the roles/logging.privateLogViewer role to a group with all the security team members. The roles/logging.privateLogViewer role is a predefined role that grants read-only access to Data Access audit logs and other private logs in Cloud Logging. A group is a collection of users that can be assigned roles and permissions as a single unit. You can assign the roles/logging.privateLogViewer role to a group with all the security team members by using IAM policies.

This way, you can provide your security team with the minimum level of access they need to view Data Access audit logs in the _Required bucket.

NEW QUESTION: 20

You work for a global organization and are running a monolithic application on Compute Engine. You need to select the machine type for the application to use that optimizes CPU utilization by using the fewest number of steps. You want to use historical system metrics to identify the machine type for the application to use. You want to follow Google-recommended practices. What should you do?

- A.** Use the Recommender API and apply the suggested recommendations
- B.** Create an Agent Policy to automatically install Ops Agent in all VMs
- C.** Install the Ops Agent in a fleet of VMs by using the gcloud CLI
- D.** Review the Cloud Monitoring dashboard for the VM and choose the machine type with the lowest CPU utilization

Answer: A ([LEAVE A REPLY](#))

The best option for selecting the machine type for the application to use that optimizes CPU utilization by using the fewest number of steps is to use the Recommender API and apply the suggested recommendations. The Recommender API is a service that provides recommendations for optimizing your Google Cloud resources, such as Compute Engine instances, disks, and firewalls. You can use the Recommender API to get recommendations for changing the machine type of your Compute Engine instances based on historical system metrics, such as CPU utilization. You can also apply the suggested recommendations by using the Recommender API or Cloud Console. This way, you can optimize CPU utilization by using the most suitable machine type for your application with minimal effort.

NEW QUESTION: 21

You are managing an application that exposes an HTTP endpoint without using a load balancer. The latency of the HTTP responses is important for the user experience. You want to understand what HTTP latencies all of your users are experiencing. You use Stackdriver Monitoring. What should you do?

- A.** * In your application, create a metric with a metricKind set to DELTA and a valueType set to DOUBLE.
* In Stackdriver's Metrics Explorer, use a Stacked Bar graph to visualize the metric.
- B.** * In your application, create a metric with a metricKind set to CUMULATIVE and a valueType set to DOUBLE.
* In Stackdriver's Metrics Explorer, use a Line graph to visualize the metric.
- C.** * In your application, create a metric with a metricKind set to gauge and a valueType set to distribution.
* In Stackdriver's Metrics Explorer, use a Heatmap graph to visualize the metric.
- D.** * In your application, create a metric with a metricKind set to METRIC_KIND_UNSPECIFIED and a valueType set to INT64.
* In Stackdriver's Metrics Explorer, use a Stacked Area graph to visualize the metric.

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 22

You are creating a CI/CD pipeline in Cloud Build to build an application container image. The application code is stored in GitHub. Your company requires that production image builds are only run against the main branch and that the change control team approves all pushes to the main branch. You want the image build to be as automated as possible. What should you do?

Choose 2 answers

- A.** Create a trigger on the Cloud Build job. Set the repository event setting to 'Pull request'.
- B.** Add the owners file to the Included files filter on the trigger.
- C.** Create a trigger on the Cloud Build job. Set the repository event setting to 'Push to a branch'.
- D.** Configure a branch protection rule for the main branch on the repository.
- E.** Enable the Approval option on the trigger.

Answer: C,D ([LEAVE A REPLY](#))

The best options for creating a CI/CD pipeline in Cloud Build to build an application container image and ensuring that production image builds are only run against the main branch and that the change control team approves all pushes to the main branch are to create a trigger on the Cloud Build job, set the repository event setting to 'Push to a branch', and configure a branch protection rule for the main branch on the repository. A trigger is a resource that starts a build when an event occurs, such as a code change. By creating a trigger on the Cloud Build job and setting the repository event setting to 'Push to a branch', you can ensure that the image build is only run when code is pushed to a specific branch, such as the main branch. A branch protection rule is a rule that enforces certain policies on a branch, such as requiring reviews, status checks, or approvals before merging code. By configuring a branch protection rule for the main branch on the repository, you can ensure that the change control team approves all pushes to the main branch.

NEW QUESTION: 23

Your application images are built and pushed to Google Container Registry (GCR). You want to build an automated pipeline that deploys the application when the image is updated while minimizing the development effort. What should you do?

- A.** Use Cloud Pub/Sub to trigger a Spinnaker pipeline.
- B.** Use Cloud Build to trigger a Spinnaker pipeline.
- C.** Use Cloud Pub/Sub to trigger a custom deployment service running in Google Kubernetes Engine (GKE).
- D.** Use a custom builder in Cloud Build to trigger a Jenkins pipeline.

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 24

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

Choose 2 answers.

- A.** Add more serving capacity to all of your application's zones.
- B.** Implement and measure all other available SLIs for the application.
- C.** Announce planned downtime to consume more error budget and ensure that users are not depending on a tighter SLO.
- D.** Have more frequent or potentially risky application releases.
- E.** Tighten the SLO to match the application's observed reliability.

Answer: ([SHOW ANSWER](#))

Explanation

The best options for furthering your application's reliability goal while balancing velocity, reliability, and business needs are to have more frequent or potentially risky application releases and to tighten the SLO to match the application's observed reliability. Having more frequent or potentially risky application releases can help you increase the change velocity and deliver new features faster. However, this also increases the likelihood of consuming more error budget and reducing the reliability of your service. Therefore, you should monitor your error budget consumption and adjust your release policies accordingly. For example, you can freeze or slow down releases when the error budget is low, or accelerate releases when the error budget is high. Tightening the SLO to match the application's observed reliability can help you align your service quality with your users' expectations and business needs. However, this also means that you have less room for error and need to maintain a higher level of reliability. Therefore, you should ensure that your SLO is realistic and achievable, and that you have sufficient engineering resources and processes to meet it.

NEW QUESTION: 25

Your development team has created a new version of their service's API. You need to deploy the new versions of the API with the least disruption to third-party developers and end users of third-party installed applications.

What should you do?

- A.** Announce deprecation of the old version of the API.
Contact remaining users on the old API.
Introduce the new version of the API.
Deprecate the old version of the API.
Provide best effort support to users of the old API.
Turn down the old version of the API.

B. Announce deprecation of the old version of the API.

Introduce the new version of the API.

Contact remaining users on the old API.

Deprecate the old version of the API.

Turn down the old version of the API.

Provide best effort support to users of the old API.

C. Introduce the new version of the API.

Announce deprecation of the old version of the API.

Deprecate the old version of the API.

Contact remaining users of the old API.

Provide best effort support to users of the old API.

Turn down the old version of the API.

D. Introduce the new version of the API.

Contact remaining users of the old API.

Announce deprecation of the old version of the API.

Deprecate the old version of the API.

Turn down the old version of the API.

Provide best effort support to users of the old API.

Answer: C (LEAVE A REPLY)

NEW QUESTION: 26

Your company follows Site Reliability Engineering practices. You are the person in charge of Communications for a large, ongoing incident affecting your customer-facing applications. There is still no estimated time for a resolution of the outage. You are receiving emails from internal stakeholders who want updates on the outage, as well as emails from customers who want to know what is happening. You want to efficiently provide updates to everyone affected by the outage. What should you do?

A. Focus on responding to internal stakeholders at least every 30 minutes. Commit to "next update" times.

B. Provide periodic updates to all stakeholders in a timely manner. Commit to a "next update" time in all communications.

C. Delegate the responding to internal stakeholder emails to another member of the Incident Response Team. Focus on providing responses directly to customers.

D. Provide all internal stakeholder emails to the Incident Commander, and allow them to manage internal communications. Focus on providing responses directly to customers.

Answer: B (LEAVE A REPLY)

Explanation

When disaster strikes, the person who declares the incident typically steps into the IC role and directs the high-level state of the incident. The IC concentrates on the 3Cs and does the following: Commands and coordinates the incident response, delegating roles as needed. By default, the IC assumes all roles that have not been delegated yet.

Communicates effectively. Stays in control of the incident response. Works with other responders to resolve the incident. <https://sre.google/workbook/incident-response/>

NEW QUESTION: 27

You are configuring Cloud Logging for a new application that runs on a Compute Engine instance with a public IP address.

A user-managed service account is attached to the instance.

You confirmed that the necessary agents are running on the instance but you cannot see any log entries from the instance in Cloud Logging. You want to resolve the issue by following Google-recommended practices.

What should you do?

- A. Add the Logs Writer role to the service account.
- B. Enable Private Google Access on the subnet that the instance is in.
- C. Update the instance to use the default Compute Engine service account.
- D. Export the service account key and configure the agents to use the key.

Answer: A (LEAVE A REPLY)

To use Cloud Logging, the service account attached to the Compute Engine instance must have the necessary permissions to write log entries. The Logs Writer role (roles/logging.logWriter) provides this permission. You can grant this role to the user-managed service account at the project, folder, or organization level¹.

Private Google Access is not required for Cloud Logging, as it allows instances without external IP addresses to access Google APIs and services². The default Compute Engine service account already has the Logs Writer role, but it is not a recommended practice to use it for user applications³. Exporting the service account key and configuring the agents to use the key is not a secure way of authenticating the service account, as it exposes the key to potential compromise⁴.

Explanation:

The correct answer is

Reference:

1: Access control with IAM | Cloud Logging | Google Cloud

2: Private Google Access overview | VPC | Google Cloud

3: Service accounts | Compute Engine Documentation | Google Cloud

4: Best practices for securing service accounts | IAM Documentation | Google Cloud

NEW QUESTION: 28

Your application services run in Google Kubernetes Engine (GKE). You want to make sure that only images from your centrally-managed Google Container Registry (GCR) image registry in the altostrat-images project can be deployed to the cluster while minimizing development time. What should you do?

- A. Use a Binary Authorization policy that includes the whitelist name pattern `gcr.io/attostrat-images/`.

B. Add logic to the deployment pipeline to check that all manifests contain only images from gcr.io/altostrat-images.

C. Create a custom builder for Cloud Build that will only push images to gcr.io/altostrat-images.

D. Add a tag to each image in gcr.io/altostrat-images and check that this tag is present when the image is deployed.

Answer: D (LEAVE A REPLY)

NEW QUESTION: 29

You are ready to deploy a new feature of a web-based application to production. You want to use Google Kubernetes Engine (GKE) to perform a phased rollout to half of the web server pods.

What should you do?

A. Use a partitioned rolling update.

B. Use Node taints with NoExecute.

C. Use a replica set in the deployment specification.

D. Use a stateful set with parallel pod management policy.

Answer: A (LEAVE A REPLY)

<https://medium.com/velotio-perspectives/exploring-upgrade-strategies-for-stateful-sets-in-kubernetes-c02b8286f251>

NEW QUESTION: 30

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

Choose 2 answers

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

B. Implement and measure all other available SLIs for the application

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

D. Have more frequent or potentially risky application releases

E. Tighten the SLO to match the application's observed reliability

Answer: (SHOW ANSWER)

The best options for furthering your application's reliability goal while balancing velocity, reliability, and business needs are to have more frequent or potentially risky application releases and to tighten the SLO to match the application's observed reliability. Having more frequent or potentially risky application releases can help you increase the change velocity and deliver new features faster. However, this also increases the likelihood of

consuming more error budget and reducing the reliability of your service. Therefore, you should monitor your error budget consumption and adjust your release policies accordingly. For example, you can freeze or slow down releases when the error budget is low, or accelerate releases when the error budget is high. Tightening the SLO to match the application's observed reliability can help you align your service quality with your users' expectations and business needs. However, this also means that you have less room for error and need to maintain a higher level of reliability. Therefore, you should ensure that your SLO is realistic and achievable, and that you have sufficient engineering resources and processes to meet it.

NEW QUESTION: 31

You use Cloud Build to build your application. You want to reduce the build time while minimizing cost and development effort. What should you do?

- A. Use larger Cloud Build virtual machines (VMs) by using the machine-type option.
- B. Run multiple Jenkins agents to parallelize the build.
- C. Use multiple smaller build steps to minimize execution time.
- D. Use Cloud Storage to cache intermediate artifacts.

Answer: C (LEAVE A REPLY)

Valid Professional-Cloud-DevOps-Engineer Dumps shared by PrepPdf.com for Helping Passing Professional-Cloud-DevOps-Engineer Exam! PrepPdf.com now offer the **newest Professional-Cloud-DevOps-Engineer exam dumps**, the PrepPdf.com Professional-Cloud-DevOps-Engineer exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com Professional-Cloud-DevOps-Engineer dumps with Test Engine here:

<https://www.preppdf.com/Google/Professional-Cloud-DevOps-Engineer-prepaway-exam-dumps.html> (200 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 32

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.

Answer: D (LEAVE A REPLY)

Explanation

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>

<https://cloud.google.com/appengine/docs/standard/python/config/appref>

`max_idle_instances` Optional. The maximum number of idle instances that App Engine should maintain for this version. Specify a value from 1 to 1000. If not specified, the default value is automatic, which means App Engine will manage the number of idle instances. Keep the following in mind: A high maximum reduces the number of idle instances more gradually when load levels return to normal after a spike. This helps your application maintain steady performance through fluctuations in request load, but also raises the number of idle instances (and consequent running costs) during such periods of heavy load.

NEW QUESTION: 33

You are reviewing your deployment pipeline in Google Cloud Deploy. You must reduce toil in the pipeline and you want to minimize the amount of time it takes to complete an end-to-end deployment. What should you do?

Choose 2 answers

- A.** Create a trigger to notify the required team to complete the next step when manual intervention is required
- B.** Divide the automation steps into smaller tasks
- C.** Use a script to automate the creation of the deployment pipeline in Google Cloud Deploy
- D.** Add more engineers to finish the manual steps.

E. Automate promotion approvals from the development environment to the test environment

Answer: A,E (LEAVE A REPLY)

Explanation

The best options for reducing toil in the pipeline and minimizing the amount of time it takes to complete an end-to-end deployment are to create a trigger to notify the required team to complete the next step when manual intervention is required and to automate promotion approvals from the development environment to the test environment. A trigger is a resource that initiates a deployment when an event occurs, such as a code change, a schedule, or a manual request. You can create a trigger to notify the required team to complete the next step when manual intervention is required by using Cloud Build or Cloud Functions. This way, you can reduce the waiting time and human errors in the pipeline. A promotion approval is a process that allows you to approve or reject a deployment from one environment to another, such as from development to test. You can automate promotion approvals from the development environment to the test environment by using Google Cloud Deploy or Cloud Build. This way, you can speed up the deployment process and avoid manual steps.

NEW QUESTION: 34

You support a high-traffic web application that runs on Google Cloud Platform (GCP). You need to measure application reliability from a user perspective without making any engineering changes to it. What should you do?

Choose 2 answers

- A.** Analyze the web proxy logs only and capture response time of each request.
- B.** Modify the code to capture additional information for user interaction.
- C.** Use current and historic Request Logs to trace customer interaction with the application.
- D.** Review current application metrics and add new ones as needed.
- E.** Create new synthetic clients to simulate a user journey using the application.

Answer: B,E (LEAVE A REPLY)

NEW QUESTION: 35

Your company experiences bugs, outages, and slowness in its production systems. Developers use the production environment for new feature development and bug fixes. Configuration and experiments are done in the production environment, causing outages for users. Testers use the production environment for load testing, which often slows the production systems. You need to redesign the environment to reduce the number of bugs and outages in production and to enable testers to load test new features. What should you do?

- A.** Create a development environment for writing code and a test environment for configurations, experiments, and load testing.

- B.** Create a development environment with smaller server capacity and give access only to developers and testers.
- C.** Secure the production environment to ensure that developers can't change it and set up one controlled update per year.
- D.** Create an automated testing script in production to detect failures as soon as they occur.

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 36

You are responsible for creating and modifying the Terraform templates that define your Infrastructure. Because two new engineers will also be working on the same code, you need to define a process and adopt a tool that will prevent you from overwriting each other's code. You also want to ensure that you capture all updates in the latest version. What should you do?

- A.** Store your code as text files in Google Drive in a defined folder structure that organizes the files.
 - * At the end of each day, confirm that all changes have been captured in the files within the folder structure and create a new .zip archive with a predefined naming convention.
 - * Upload the .zip archive to a versioned Cloud Storage bucket and accept it as the latest version.
- B.** Store your code in a Git-based version control system.
 - * Establish a process that allows developers to merge their own changes at the end of each day.
 - * Package and upload code to a versioned Cloud Storage bucket as the latest master version.
- C.** Store your code as text files in Google Drive in a defined folder structure that organizes the files.
 - * At the end of each day, confirm that all changes have been captured in the files within the folder structure.
 - * Rename the folder structure with a predefined naming convention that increments the version.
- D.** Store your code in a Git-based version control system.
 - * Establish a process that includes code reviews by peers and unit testing to ensure integrity and functionality before integration of code.
 - * Establish a process where the fully integrated code in the repository becomes the latest master version.

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 37

Your company experiences bugs, outages, and slowness in its production systems. Developers use the production environment for new feature development and bug fixes.

Configuration and experiments are done in the production environment, causing outages for users. Testers use the production environment for load testing, which often slows the production systems. You need to redesign the environment to reduce the number of bugs and outages in production and to enable testers to load test new features. What should you do?

- A.** Create a development environment with smaller server capacity and give access only to developers and testers.
- B.** Secure the production environment to ensure that developers can't change it and set up one controlled update per year.
- C.** Create an automated testing script in production to detect failures as soon as they occur.
- D.** Create a development environment for writing code and a test environment for configurations, experiments, and load testing.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 38

Your organization recently adopted a container-based workflow for application development. Your team develops numerous applications that are deployed continuously through an automated build pipeline to the production environment. A recent security audit alerted your team that the code pushed to production could contain vulnerabilities and that the existing tooling around virtual machine (VM) vulnerabilities no longer applies to the containerized environment. You need to ensure the security and patch level of all code running through the pipeline. What should you do?

- A.** Set up Container Analysis to scan and report Common Vulnerabilities and Exposures.
- B.** Configure the containers in the build pipeline to always update themselves before release.
- C.** Reconfigure the existing operating system vulnerability software to exist inside the container.
- D.** Implement static code analysis tooling against the Docker files used to create the containers.

Answer: **D** ([LEAVE A REPLY](#))

Explanation

<https://cloud.google.com/binary-authorization>

Binary Authorization is a deploy-time security control that ensures only trusted container images are deployed on Google Kubernetes Engine (GKE) or Cloud Run. With Binary Authorization, you can require images to be signed by trusted authorities during the development process and then enforce signature validation when deploying. By enforcing validation, you can gain tighter control over your container environment by ensuring only verified images are integrated into the build-and-release process.

NEW QUESTION: 39

You deploy a new release of an internal application during a weekend maintenance window when there is minimal user traffic. After the window ends, you learn that one of the new features isn't working as expected in the production environment. After an extended outage, you roll back the new release and deploy a fix. You want to modify your release process to reduce the mean time to recovery so you can avoid extended outages in the future. What should you do?

Choose 2 answers

- A.** Integrate a code linting tool to validate coding standards before any code is accepted into the repository.
- B.** Configure a CI server. Add a suite of unit tests to your code and have your CI server run them on commit and verify any changes.
- C.** Before merging new code, require 2 different peers to review the code changes.
- D.** Require developers to run automated integration tests on their local development environments before release.
- E.** Adopt the blue/green deployment strategy when releasing new code via a CD server.

Answer: B,E ([LEAVE A REPLY](#))

NEW QUESTION: 40

You are designing a deployment technique for your applications on Google Cloud. As part of your deployment planning, you want to use live traffic to gather performance metrics for new versions of your applications. You need to test against the full production load before your applications are launched. What should you do?

- A.** Use A/B testing with blue/green deployment.
- B.** Use shadow testing with continuous deployment.
- C.** Use canary testing with continuous deployment.
- D.** Use canary testing with rolling updates deployment,

Answer: ([SHOW ANSWER](#))

Explanation

The correct answer is B. Use shadow testing with continuous deployment.

Shadow testing is a deployment technique that involves routing a copy of the live traffic to a new version of the application, without affecting the production environment. This way, you can gather performance metrics and compare them with the current version, without exposing the new version to the users. Shadow testing can help you test against the full production load and identify any issues or bottlenecks before launching the new version. You can use continuous deployment to automate the process of deploying the new version after it passes the shadow testing.

NEW QUESTION: 41

You are performing a semiannual capacity planning exercise for your flagship service. You expect a service user growth rate of 10% month-over-month over the next six months. Your service is fully containerized and runs on Google Cloud Platform (GCP). using a

Google Kubernetes Engine (GKE) Standard regional cluster on three zones with cluster autoscaler enabled. You currently consume about 30% of your total deployed CPU capacity, and you require resilience against the failure of a zone. You want to ensure that your users experience minimal negative impact as a result of this growth or as a result of zone failure, while avoiding unnecessary costs. How should you prepare to handle the predicted growth?

- A.** Verify the maximum node pool size, enable a horizontal pod autoscaler, and then perform a load test to verify your expected resource needs.
- B.** Because you are deployed on GKE and are using a cluster autoscaler, your GKE cluster will scale automatically, regardless of growth rate.
- C.** Because you are at only 30% utilization, you have significant headroom and you won't need to add any additional capacity for this rate of growth.
- D.** Proactively add 60% more node capacity to account for six months of 10% growth rate, and then perform a load test to make sure you have enough capacity.

Answer: A (LEAVE A REPLY)

<https://cloud.google.com/kubernetes-engine/docs/concepts/horizontalpodautoscaler> The Horizontal Pod Autoscaler changes the shape of your Kubernetes workload by automatically increasing or decreasing the number of Pods in response to the workload's CPU or memory consumption

NEW QUESTION: 42

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

Answer: C (LEAVE A REPLY)

The best option for implementing the rollback actions is to update the Kubernetes Service to point to the previous Kubernetes Deployment. A Kubernetes Service is a resource that defines how to access a set of Pods. A Kubernetes Deployment is a resource that manages the creation and update of Pods. By using the blue/green deployment methodology, you can create two Deployments, one for the current version (blue) and one for the new version (green), and use a Service to switch traffic between them. If you need to rollback, you can update the Service to point to the previous Deployment (blue) and stop sending traffic to the new Deployment (green).

NEW QUESTION: 43

Your application images are built using Cloud Build and pushed to Google Container Registry (GCR). You want to be able to specify a particular version of your application for deployment based on the release version tagged in source control. What should you do when you push the image?

- A.** Use GCR digest versioning to match the image to the tag in source control.
- B.** Reference the image digest in the source control tag.
- C.** Use Cloud Build to include the release version tag in the application image.
- D.** Supply the source control tag as a parameter within the image name.

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 44

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 Cloud Storage encrypted with a key from Cloud KMS. Provide the CI/CD pipeline with access to Cloud KMS via IAM.
- C.** 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
- D.** Store secrets in a separate configuration file on Git. Provide select developers with access to the configuration file.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 45

Your company follows Site Reliability Engineering practices. You are the person in charge of Communications for a large, ongoing incident affecting your customer-facing applications. There is still no estimated time for a resolution of the outage. You are receiving emails from internal stakeholders who want updates on the outage, as well as emails from customers who want to know what is happening. You want to efficiently provide updates to everyone affected by the outage. What should you do?

- A.** Focus on responding to internal stakeholders at least every 30 minutes. Commit to "next update" times.
- B.** Provide periodic updates to all stakeholders in a timely manner. Commit to a "next update" time in all communications.
- C.** Delegate the responding to internal stakeholder emails to another member of the Incident Response Team. Focus on providing responses directly to customers.
- D.** Provide all internal stakeholder emails to the Incident Commander, and allow them to manage internal communications. Focus on providing responses directly to customers.

Answer: ([SHOW ANSWER](#))

When disaster strikes, the person who declares the incident typically steps into the IC role and directs the high-level state of the incident. The IC concentrates on the 3Cs and does the following: Commands and coordinates the incident response, delegating roles as needed. By default, the IC assumes all roles that have not been delegated yet. Communicates effectively. Stays in control of the incident response. Works with other responders to resolve the incident. <https://sre.google/workbook/incident-response/>

NEW QUESTION: 46

Your development team has created a new version of their service's API. You need to deploy the new versions of the API with the least disruption to third-party developers and end users of third-party installed applications. What should you do?

A. Introduce the new version of the API.

Announce deprecation of the old version of the API.

Deprecate the old version of the API.

Contact remaining users of the old API.

Provide best effort support to users of the old API.

Turn down the old version of the API.

B. Announce deprecation of the old version of the API.

Contact remaining users on the old API.

Introduce the new version of the API.

Deprecate the old version of the API.

Provide best effort support to users of the old API.

Turn down the old version of the API.

C. Introduce the new version of the API.

Contact remaining users of the old API.

Announce deprecation of the old version of the API.

Deprecate the old version of the API.

Turn down the old version of the API.

Provide best effort support to users of the old API.

D. Announce deprecation of the old version of the API.

Introduce the new version of the API.

Contact remaining users on the old API.

Deprecate the old version of the API.

Turn down the old version of the API.

Provide best effort support to users of the old API.

Answer: A (LEAVE A REPLY)

the **newest Professional-Cloud-DevOps-Engineer exam dumps**, the PrepPdf.com Professional-Cloud-DevOps-Engineer exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com Professional-Cloud-DevOps-Engineer dumps with Test Engine here:

<https://www.preppdf.com/Google/Professional-Cloud-DevOps-Engineer-prepaway-exam-dumps.html> (200 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 47

You support a large service with a well-defined Service Level Objective (SLO). The development team deploys new releases of the service multiple times a week. If a major incident causes the service to miss its SLO, you want the development team to shift its focus from working on features to improving service reliability. What should you do before a major incident occurs?

- A. Develop an appropriate error budget policy in cooperation with all service stakeholders.
- B. Negotiate with the product team to always prioritize service reliability over releasing new features.
- C. Negotiate with the development team to reduce the release frequency to no more than once a week.
- D. Add a plugin to your Jenkins pipeline that prevents new releases whenever your service is out of SLO.

Answer: A (LEAVE A REPLY)

Reason : Incident has not occurred yet, even when development team is already pushing new features multiple times a week. The option A says, to define an error budget "policy", not to define error budget(It is already present). Just simple means to bring in all stakeholders, and decide how to consume the error budget effectively that could bring balance between feature deployment and reliability.

The goals of this policy are to: -- Protect customers from repeated SLO misses -- Provide an incentive to balance reliability with other features <https://sre.google/workbook/error-budget-policy/>

NEW QUESTION: 48

You manage an application that runs in Google Kubernetes Engine (GKE) and uses the blue/green deployment methodology. Extracts of the Kubernetes manifests are shown below.

The Deployment app-green was updated to use the new version of the application. During post-deployment monitoring, you notice that the majority of user requests are failing. You did not observe this behavior in the testing environment. You need to mitigate the incident impact on users and enable the developers to troubleshoot the issue. What should you do?

- A. Update the Deployment app-blue to use the new version of the application.
- B. Update the Deployment app-green to use the previous version of the application.
- C. Change the selector on the Service app-2vc to app: my-app.

D. Change the selector on the Service app-svc to app: my-app, version: blue

Answer: D (LEAVE A REPLY)

Explanation

The best option for mitigating the incident impact on users and enabling the developers to troubleshoot the issue is to change the selector on the Service app-svc to app: my-app, version: blue. A Service is a resource that defines how to access a set of Pods. A selector is a field that specifies which Pods are selected by the Service. By changing the selector on the Service app-svc to app: my-app, version: blue, you can ensure that the Service only routes traffic to the Pods that have both labels app: my-app and version: blue. These Pods belong to the Deployment app-blue, which uses the previous version of the application. This way, you can mitigate the incident impact on users by switching back to the working version of the application. You can also enable the developers to troubleshoot the issue with the new version of the application in the Deployment app-green without affecting users.

NEW QUESTION: 49

You work for a global organization and run a service with an availability target of 99% with limited engineering resources. For the current calendar month you noticed that the service has 99.5% availability. You must ensure that your service meets the defined availability goals and can react to business changes including the upcoming launch of new features. You also need to reduce technical debt while minimizing operational costs. You want to follow Google-recommended practices. What should you do?

- A.** Identify, measure and eliminate toil by automating repetitive tasks
- B.** Add N+1 redundancy to your service by adding additional compute resources to the service
- C.** Allocate available engineers to the feature backlog while you ensure that the service remains within the availability target
- D.** Define an error budget for your service level availability and minimize the remaining error budget

Answer: D (LEAVE A REPLY)

NEW QUESTION: 50

You are running an application on Compute Engine and collecting logs through Stackdriver. You discover that some personally identifiable information (PII) is leaking into certain log entry fields. You want to prevent these fields from being written in new log entries as quickly as possible. What should you do?

- A.** Wait for the application developers to patch the application, and then verify that the log entries are no longer exposing PII.
- B.** Use the filter-record-transformer Fluentd filter plugin to remove the fields from the log entries in flight.

C. Stage log entries to Cloud Storage, and then trigger a Cloud Function to remove the fields and write the entries to Stackdriver via the Stackdriver Logging API.

D. Use the fluent-plugin-record-reformer Fluentd output plugin to remove the fields from the log entries in flight.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 51

You recently migrated an ecommerce application to Google Cloud. You now need to prepare the application for the upcoming peak traffic season. You want to follow Google-recommended practices. What should you do first to prepare for the busy season?

A. Migrate the application to Cloud Run, and use autoscaling.

B. Load test the application to profile its performance for scaling.

C. Create a Terraform configuration for the application's underlying infrastructure to quickly deploy to additional regions.

D. Pre-provision the additional compute power that was used last season, and expect growth.

Answer: **B** ([LEAVE A REPLY](#))

The first thing you should do to prepare your ecommerce application for the upcoming peak traffic season is to load test the application to profile its performance for scaling. Load testing is a process of simulating high traffic or user demand on your application and measuring how it responds. Load testing can help you identify any bottlenecks, errors, or performance issues that might affect your application during the busy season¹. Load testing can also help you determine the optimal scaling strategy for your application, such as horizontal scaling (adding more instances) or vertical scaling (adding more resources to each instance)².

There are different tools and methods for load testing your ecommerce application on Google Cloud, depending on the type and complexity of your application. For example, you can use Cloud Load Balancing to distribute traffic across multiple instances of your application, and use Cloud Monitoring to measure the latency, throughput, and error rate of your application³. You can also use Cloud Functions or Cloud Run to create serverless load generators that can simulate user requests and send them to your application⁴. Alternatively, you can use third-party tools such as Apache JMeter or Locust to create and run load tests on your application.

By load testing your ecommerce application before the peak traffic season, you can ensure that your application is ready to handle the expected load and provide a good user experience. You can also use the results of your load tests to plan and implement other steps to prepare your application for the busy season, such as migrating to a more scalable platform, creating a Terraform configuration for deploying to additional regions, or pre-provisioning additional compute power.

Reference:

1: Load Testing 101: How To Test Website Performance | BlazeMeter

2: Scaling applications | Google Cloud

3: Load testing using Google Cloud | Solutions | Google Cloud

4: Serverless load testing using Cloud Functions | Solutions | Google Cloud

NEW QUESTION: 52

You use a multiple step Cloud Build pipeline to build and deploy your application to Google Kubernetes Engine (GKE). You want to integrate with a third-party monitoring platform by performing a HTTP POST of the build information to a webhook. You want to minimize the development effort. What should you do?

- A.** Create a Cloud Pub/Sub push subscription to the Cloud Build cloud-builds PubSub topic to HTTP POST the build information to a webhook.
- B.** Add logic to each Cloud Build step to HTTP POST the build information to a webhook.
- C.** Add a new step at the end of the pipeline in Cloud Build to HTTP POST the build information to a webhook.
- D.** Use Stackdriver Logging to create a logs-based metric from the Cloud Build logs. Create an Alert with a Webhook notification type.

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 53

You support an e-commerce application that runs on a large Google Kubernetes Engine (GKE) cluster deployed on-premises and on Google Cloud Platform. The application consists of microservices that run in containers. You want to identify containers that are using the most CPU and memory. What should you do?

- A.** Use Stackdriver Kubernetes Engine Monitoring.
- B.** Use Stackdriver Logging to export application logs to BigQuery, aggregate logs per container, and then analyze CPU and memory consumption.
- C.** Use Prometheus to collect and aggregate logs per container, and then analyze the results in Grafana.
- D.** Use the Stackdriver Monitoring API to create custom metrics, and then organize your containers using groups.

Answer: ([SHOW ANSWER](#)**)**

NEW QUESTION: 54

Your company follows Site Reliability Engineering principles. You are writing a postmortem for an incident, triggered by a software change, that severely affected users. You want to prevent severe incidents from happening in the future. What should you do?

- A.** Identify engineers responsible for the incident and escalate to their senior management.
- B.** Ensure that test cases that catch errors of this type are run successfully before new software releases.
- C.** Follow up with the employees who reviewed the changes and prescribe practices they should follow in the future.

D. Design a policy that will require on-call teams to immediately call engineers and management to discuss a plan of action if an incident occurs.

Answer: B ([LEAVE A REPLY](#))

The best way to prevent severe incidents from happening in the future is to ensure that test cases that catch errors of this type are run successfully before new software releases. This is aligned with the Site Reliability Engineering principle of testing for reliability.

NEW QUESTION: 55

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.

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?

- 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: ([SHOW ANSWER](#)**)**

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.

NEW QUESTION: 56

You are creating Cloud Logging sinks to export log entries from Cloud Logging to BigQuery for future analysis. Your organization has a Google Cloud folder named Dev that contains development projects and a folder named Prod that contains production projects. Log entries for development projects must be exported to dev_dataset, and log entries for production projects must be exported to prod_dataset. You need to minimize the number of log sinks created and you want to ensure that the log sinks apply to future projects. What should you do?

- A. Create a single aggregated log sink at the organization level.
- B. Create a log sink in each project
- C. Create two aggregated log sinks at the organization level, and filter by project ID
- D. Create an aggregated log sink in the Dev and Prod folders

Answer: D (LEAVE A REPLY)

The best option for minimizing the number of log sinks created and ensuring that the log sinks apply to future projects is to create an aggregated log sink in the Dev and Prod folders. An aggregated log sink is a log sink that collects logs from multiple sources, such as projects, folders, or organizations. By creating an aggregated log sink in each folder, you can export log entries for development projects to dev_dataset and log entries for production projects to prod_dataset. You can also use filters to specify which logs you want to export. Additionally, by creating an aggregated log sink at the folder level, you can ensure that the log sink applies to future projects that are created under that folder.

NEW QUESTION: 57

You are configuring connectivity across Google Kubernetes Engine (GKE) clusters in different VPCs. You notice that the nodes in Cluster A are unable to access the nodes in Cluster B. You suspect that the workload access issue is due to the network configuration. You need to troubleshoot the issue but do not have execute access to workloads and nodes. You want to identify the layer at which the network connectivity is broken. What should you do?

- A. Use Network Connectivity Center to perform a Connectivity Test from Cluster A to Cluster B
- B. Use a debug container to run the traceroute command from Cluster A to Cluster B and from Cluster B to Cluster A. Identify the common failure point.
- C. Install a toolbox container on the node in Cluster A. Confirm that the routes to Cluster B are configured appropriately.
- D. Enable VPC Flow Logs in both VPCs and monitor packet drops.

Answer: A (LEAVE A REPLY)

Explanation

The best option for troubleshooting the issue without having execute access to workloads and nodes is to use Network Connectivity Center to perform a Connectivity Test from Cluster A to Cluster B. Network Connectivity Center is a service that allows you to create, manage, and monitor network connectivity across Google Cloud, hybrid, and multi-cloud environments. You can use Network Connectivity Center to perform a Connectivity Test, which is a feature that allows you to test the reachability and latency between two endpoints, such as GKE clusters, VM instances, or IP addresses. By using Network Connectivity Center to perform a Connectivity Test from Cluster A to Cluster B, you can identify the layer at which the network connectivity is broken, such as the firewall, routing, or load balancing.

NEW QUESTION: 58

Your company is developing applications that are deployed on Google Kubernetes Engine (GKE). Each team manages a different application. You need to create the development and production environments for each team while you minimize costs. Different teams should not be able to access other teams' environments. You want to follow Google-recommended practices. What should you do?

- A.** Create one Google Cloud project per team. In each project, create a cluster for development and one for production. Grant the teams Identity and Access Management (IAM) access to their respective clusters.
- B.** Create one Google Cloud project per team. In each project, create a cluster with a Kubernetes namespace for development and one for production. Grant the teams Identity and Access Management (IAM) access to their respective clusters.
- C.** Create a development and a production GKE cluster in separate projects. In each cluster, create a Kubernetes namespace per team and then configure Identity-Aware Proxy so that each team can only access its own namespace.
- D.** Create a development and a production GKE cluster in separate projects. In each cluster, create a Kubernetes namespace per team and then configure Kubernetes role-based access control (RBAC) so that each team can only access its own namespace.

Answer: D (LEAVE A REPLY)

Explanation

The best option for creating the development and production environments for each team while minimizing costs and ensuring isolation is to create a development and a production GKE cluster in separate projects, in each cluster create a Kubernetes namespace per team, and then configure Kubernetes role-based access control (RBAC) so that each team can only access its own namespace. This option allows you to use fewer clusters and projects than creating one project or cluster per team, which reduces costs and complexity. It also allows you to isolate each team's environment by using namespaces and RBAC, which prevents teams from accessing other teams' environments.

NEW QUESTION: 59

You support a high-traffic web application that runs on Google Cloud Platform (GCP). You need to measure application reliability from a user perspective without making any engineering changes to it. What should you do?

Choose 2 answers

- A.** Review current application metrics and add new ones as needed.
- B.** Modify the code to capture additional information for user interaction.
- C.** Analyze the web proxy logs only and capture response time of each request.
- D.** Create new synthetic clients to simulate a user journey using the application.
- E.** Use current and historic Request Logs to trace customer interaction with the application.

Answer: (SHOW ANSWER)

<https://cloud.google.com/architecture/adopting-slos?hl=en>

NEW QUESTION: 60

You need to run a business-critical workload on a fixed set of Compute Engine instances for several months. The workload is stable with the exact amount of resources allocated to it. You want to lower the costs for this workload without any performance implications. What should you do?

- A. Purchase Committed Use Discounts.
- B. Create an Unmanaged Instance Group for the instances used to run the workload.
- C. Convert the instances to preemptible virtual machines.
- D. Migrate the instances to a Managed Instance Group.

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 61

You are monitoring a service that uses n2-standard-2 Compute Engine instances that serve large files. Users have reported that downloads are slow. Your Cloud Monitoring dashboard shows that your VMS are running at peak network throughput. You want to improve the network throughput performance. What should you do?

- A. Deploy a Cloud NAT gateway and attach the gateway to the subnet of the VMS.
- B. Add additional network interface controllers (NICs) to your VMS.
- C. Change the machine type for your VMS to n2-standard-8.
- D. Deploy the Ops Agent to export additional monitoring metrics.

Answer: (SHOW ANSWER)

Explanation

The correct answer is C. Change the machine type for your VMs to n2-standard-8. According to the Google Cloud documentation, the network throughput performance of a Compute Engine VM depends on its machine type¹. The n2-standard-2 machine type has a maximum egress bandwidth of 4 Gbps, which can be a bottleneck for serving large files. By changing the machine type to n2-standard-8, you can increase the maximum egress bandwidth to 16 Gbps, which can improve the network throughput performance and reduce the download time for users. You also need to enable per VM Tier_1 networking performance, which is a feature that allows VMs to achieve higher network performance than the default settings².

The other options are incorrect because they do not improve the network throughput performance of your VMs. Option A is incorrect because Cloud NAT is a service that allows private IP addresses to access the internet, but it does not increase the network bandwidth or speed³. Option B is incorrect because adding additional network interfaces (NICs) or IP addresses per NIC does not increase ingress or egress bandwidth for a VM¹. Option D is incorrect because deploying the Ops Agent can help you monitor and troubleshoot your VMs, but it does not affect the network throughput performance⁴.

Valid Professional-Cloud-DevOps-Engineer Dumps shared by PrepPdf.com for Helping Passing Professional-Cloud-DevOps-Engineer Exam! PrepPdf.com now offer the **newest Professional-Cloud-DevOps-Engineer exam dumps**, the PrepPdf.com Professional-Cloud-DevOps-Engineer exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com Professional-Cloud-DevOps-Engineer dumps with Test Engine here:
<https://www.preppdf.com/Google/Professional-Cloud-DevOps-Engineer-prepaway-exam-dumps.html> (200 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 62

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?

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 (LEAVE A REPLY)

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.

NEW QUESTION: 63

Your CTO has asked you to implement a postmortem policy on every incident for internal use. You want to define what a good postmortem is to ensure that the policy is successful at your company. What should you do?

Choose 2 answers

- A.** Ensure that all postmortems include what caused the incident, identify the person or team responsible for causing the incident. and how to prevent a future occurrence of the incident.
- B.** Ensure that all postmortems include what caused the incident, how the incident could have been worse, and how to prevent a future occurrence of the incident.
- C.** Ensure that all postmortems include the severity of the incident, how to prevent a future occurrence of the incident. and what caused the incident without naming internal system components.
- D.** Ensure that all postmortems include how the incident was resolved and what caused the incident without naming customer information.
- E.** Ensure that all postmortems include all incident participants in postmortem authoring and share postmortems as widely as possible,

Answer: B,E (LEAVE A REPLY)

The correct answers are B and E.

A good postmortem should include what caused the incident, how the incident could have been worse, and how to prevent a future occurrence of the incident¹. This helps to identify the root cause of the problem, the impact of the incident, and the actions to take to mitigate or eliminate the risk of recurrence.

A good postmortem should also include all incident participants in postmortem authoring and share postmortems as widely as possible². This helps to foster a culture of learning and collaboration, as well as to increase the visibility and accountability of the incident response process.

Answer A is incorrect because it assigns blame to a person or team, which goes against the principle of blameless postmortems². Blameless postmortems focus on finding solutions rather than pointing fingers, and encourage honest and constructive feedback without fear of punishment.

Answer C is incorrect because it omits how the incident could have been worse, which is an important factor to consider when evaluating the severity and impact of the incident¹. It also avoids naming internal system components, which makes it harder to understand the technical details and root cause of the problem.

Answer D is incorrect because it omits how to prevent a future occurrence of the incident, which is the main goal of a postmortem¹. It also avoids naming customer information, which may be relevant for understanding the impact and scope of the incident.

NEW QUESTION: 64

You currently store the virtual machine (VM) utilization logs in Stackdriver. You need to provide an easy-to-share interactive VM utilization dashboard that is updated in real time and contains information aggregated on a quarterly basis. You want to use Google Cloud Platform solutions. What should you do?

- A.** 1. Export VM utilization logs from Stackdriver to a Cloud Storage bucket.
2. Enable the Cloud Storage API to pull the logs programmatically.
3. Build a custom data visualization application.
4. Display the pulled logs in a custom dashboard.
- B.** 1. Export VM utilization logs from Stackdriver to Cloud Pub/Sub.
2. From Cloud Pub/Sub, send the logs to a Security Information and Event Management (SIEM) system.
3. Build the dashboards in the SIEM system and share with your stakeholders.
- C.** 1. Export VM utilization logs from Stackdriver to BigQuery.
2. From BigQuery, export the logs to a CSV file.
3. Import the CSV file into Google Sheets.
4. Build a dashboard in Google Sheets and share it with your stakeholders.
- D.** 1. Export VM utilization logs from Stackdriver to BigQuery.
2. Create a dashboard in Data Studio.
3. Share the dashboard with your stakeholders.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 65

You are running a real-time gaming application on Compute Engine that has a production and testing environment. Each environment has their own Virtual Private Cloud (VPC) network. The application frontend and backend servers are located on different subnets in the environment's VPC. You suspect there is a malicious process communicating intermittently in your production frontend servers. You want to ensure that network traffic is captured for analysis. What should you do?

- A.** Enable VPC Flow Logs on the testing and production VPC network frontend and backend subnets with a volume scale of 0.5. Apply changes in testing before production.
- B.** Enable VPC Flow Logs on the production VPC network frontend and backend subnets only with a sample volume scale of 0.5.
- C.** Enable VPC Flow Logs on the production VPC network frontend and backend subnets only with a sample volume scale of 1.0.
- D.** Enable VPC Flow Logs on the testing and production VPC network frontend and backend subnets with a volume scale of 1.0. Apply changes in testing before production.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 66

You have an application that runs in Google Kubernetes Engine (GKE). The application consists of several microservices that are deployed to GKE by using Deployments and Services. One of the microservices is experiencing an issue where a Pod returns 403 errors after the Pod has been running for more than five hours. Your development team is working on a solution but the issue will not be resolved for a month. You need to ensure continued operations until the microservice is fixed. You want to follow Google-recommended practices and use the fewest number of steps. What should you do?

- A.** Create a cron job to terminate any Pods that have been running for more than five hours.
- B.** Add a HTTP liveness probe to the microservice's deployment.
- C.** Monitor the Pods and terminate any Pods that have been running for more than five hours.
- D.** Configure an alert to notify you whenever a Pod returns 403 errors.

Answer: ([SHOW ANSWER](#))

Explanation

The best option for ensuring continued operations until the microservice is fixed is to add a HTTP liveness probe to the microservice's deployment. A HTTP liveness probe is a type of probe that checks if a Pod is alive by sending an HTTP request and expecting a success response code. If the probe fails, Kubernetes will restart the Pod. You can add a HTTP liveness probe to your microservice's deployment by using a `livenessProbe` field in your Pod spec. This way, you can ensure that any Pod that returns 403 errors after running for more than five hours will be restarted automatically and resume normal operations.

NEW QUESTION: 67

You support a stateless web-based API that is deployed on a single Compute Engine instance in the `eu-west-2-a` zone. The Service Level Indicator (SLI) for service availability is below the specified Service Level Objective (SLO). A postmortem has revealed that requests to the API regularly time out. The time outs are due to the API having a high number of requests and running out of memory. You want to improve service availability. What should you do?

- A.** Change the specified SLO to match the measured SLI.
- B.** Set up additional service instances in other zones and load balance the traffic between all instances.
- C.** Move the service to higher-specification compute instances with more memory.
- D.** Set up additional service instances in other zones and use them as a failover in case the primary instance is unavailable.

Answer: **B** ([LEAVE A REPLY](#))

NEW QUESTION: 68

You need to create a Cloud Monitoring SLO for a service that will be published soon. You want to verify that requests to the service will be addressed in fewer than 300 ms at least 90% Of the time per calendar month.

You need to identify the metric and evaluation method to use. What should you do?

- A. Select a latency metric for a request-based method of evaluation.
- B. Select a latency metric for a window-based method of evaluation.
- C. Select an availability metric for a request-based method of evaluation.
- D. Select an availability metric for a window-based method Of evaluation.

Answer: A (LEAVE A REPLY)

Explanation

The correct answer is A. Select a latency metric for a request-based method of evaluation. A latency metric measures how responsive your service is to users. For example, you can use the `cloud.googleapis.com/http/server/response_latencies` metric to measure the latency of HTTP requests to your service¹. A request-based method of evaluation counts the number of successful requests that meet a certain criterion, such as being below a latency threshold, and compares it to the number of all requests. For example, you can define an SLI as the ratio of requests with latency below 300 ms to all requests². A request-based method of evaluation is suitable for measuring performance over time, such as per calendar month. You can set an SLO for the SLI to be at least 90%, which means that you expect 90% of the requests to have latency below 300 ms in a month³.

NEW QUESTION: 69

You are using Stackdriver to monitor applications hosted on Google Cloud Platform (GCP). You recently deployed a new application, but its logs are not appearing on the Stackdriver dashboard.

You need to troubleshoot the issue. What should you do?

- A. Confirm that the Stackdriver agent has been installed in the hosting virtual machine.
- B. Confirm that your account has the proper permissions to use the Stackdriver dashboard.
- C. Confirm that port 25 has been opened in the firewall to allow messages through to Stackdriver.
- D. Confirm that the application is using the required client library and the service account key has proper permissions.

Answer: A (LEAVE A REPLY)

<https://cloud.google.com/monitoring/agent/monitoring/troubleshooting#checklist>

NEW QUESTION: 70

Your company is developing applications that are deployed on Google Kubernetes Engine (GKE) Each team manages a different application You need to create the development and production environments for each team while you minimize costs Different teams

should not be able to access other teams environments You want to follow Google-recommended practices What should you do?

- A.** Create one Google Cloud project per team In each project create a cluster for development and one for production Grant the teams Identity and Access Management (IAM) access to their respective clusters
- B.** Create one Google Cloud project per team In each project create a cluster with a Kubernetes namespace for development and one for production Grant the teams Identity and Access Management (IAM) access to their respective clusters.
- C.** Create a development and a production GKE cluster in separate projects In each cluster create a Kubernetes namespace per team and then configure Identity-Aware Proxy so that each team can only access its own namespace
- D.** Create a development and a production GKE cluster in separate projects In each cluster create a Kubernetes namespace per team and then configure Kubernetes role-based access control (RBAC) so that each team can only access its own namespace

Answer: D (LEAVE A REPLY)

The best option for creating the development and production environments for each team while minimizing costs and ensuring isolation is to create a development and a production GKE cluster in separate projects, in each cluster create a Kubernetes namespace per team, and then configure Kubernetes role-based access control (RBAC) so that each team can only access its own namespace. This option allows you to use fewer clusters and projects than creating one project or cluster per team, which reduces costs and complexity. It also allows you to isolate each team's environment by using namespaces and RBAC, which prevents teams from accessing other teams' environments.

NEW QUESTION: 71

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.** Install an Application Performance Monitoring (APM) tool in both locations, and configure an export to a central data storage location for analysis.
- B.** Import the Stackdriver Profiler package, and configure it to relay function timing data to Stackdriver for further analysis.
- C.** Instrument the code using a timing library, and publish the metrics via a health check endpoint that is scraped by Stackdriver.
- D.** Import the Stackdriver Debugger package, and configure the application to emit debug messages with timing information.

Answer: D (LEAVE A REPLY)

NEW QUESTION: 72

Your application services run in Google Kubernetes Engine (GKE). You want to make sure that only images from your centrally-managed Google Container Registry (GCR) image registry in the altostrat-images project can be deployed to the cluster while minimizing development time. What should you do?

- A. Add a tag to each image in gcr.io/altostrat-images and check that this tag is present when the image is deployed.
- B. Add logic to the deployment pipeline to check that all manifests contain only images from gcr.io/altostrat-images.
- C. Use a Binary Authorization policy that includes the whitelist name pattern gcr.io/altostrat-images/.
- D. Create a custom builder for Cloud Build that will only push images to gcr.io/altostrat-images.

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 73

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.

Answer: ([SHOW ANSWER](#)**)**

Hostname Storage location gcr.io Stores images in data centers in the United States
asia.gcr.io Stores images in data centers in Asia eu.gcr.io Stores images in data centers within member states of the European Union us.gcr.io Stores images in data centers in the United States

NEW QUESTION: 74

You need to create a Cloud Monitoring SLO for a service that will be published soon. You want to verify that requests to the service will be addressed in fewer than 300 ms at least 90% Of the time per calendar month. You need to identify the metric and evaluation method to use. What should you do?

- A. Select a latency metric for a request-based method of evaluation.
- B. Select a latency metric for a window-based method of evaluation.
- C. Select an availability metric for a request-based method of evaluation.
- D. Select an availability metric for a window-based method Of evaluation.

Answer: A ([LEAVE A REPLY](#))

The correct answer is

A, Select a latency metric for a request-based method of evaluation.

A latency metric measures how responsive your service is to users. For example, you can use the `cloud.googleapis.com/http/server/response_latencies` metric to measure the latency of HTTP requests to your service¹. A request-based method of evaluation counts the number of successful requests that meet a certain criterion, such as being below a latency threshold, and compares it to the number of all requests. For example, you can define an SLI as the ratio of requests with latency below 300 ms to all requests². A request-based method of evaluation is suitable for measuring performance over time, such as per calendar month. You can set an SLO for the SLI to be at least 90%, which means that you expect 90% of the requests to have latency below 300 ms in a month³.

Reference:

Creating an SLO | Operations Suite | Google Cloud, Choosing a metric, Latency metric. Concepts in service monitoring | Operations Suite | Google Cloud, Service-level indicators, Request-based SLIs. Learn how to set SLOs - SRE tips | Google Cloud Blog, Setting SLOs.

NEW QUESTION: 75

You are deploying an application to Cloud Run. The application requires a password to start. Your organization requires that all passwords are rotated every 24 hours, and your application must have the latest password. You need to deploy the application with no downtime. What should you do?

- A.** Store the password in Secret Manager and send the secret to the application by using environment variables.
- B.** Store the password in Secret Manager and mount the secret as a volume within the application.
- C.** Use Cloud Build to add your password into the application container at build time. Ensure that Artifact Registry is secured from public access.
- D.** Store the password directly in the code. Use Cloud Build to rebuild and deploy the application each time the password changes.

Answer: B (LEAVE A REPLY)

The correct answer is B, Store the password in Secret Manager and mount the secret as a volume within the application.

Secret Manager is a service that allows you to securely store and manage sensitive data such as passwords, API keys, certificates, and tokens. You can use Secret Manager to rotate your secrets automatically or manually, and access them from your Cloud Run applications¹.

There are two ways to use secrets from Secret Manager in Cloud Run:

As environment variables: You can set environment variables that point to secrets in Secret Manager. Cloud Run will resolve the secrets at runtime and inject them into the environment of your application. However, this method has some limitations, such as:

The environment variables are cached for up to 10 minutes, so you may not get the latest version of the secret immediately.

The environment variables are visible in plain text in the Cloud Console and the Cloud SDK, which may expose sensitive information.

The environment variables are limited to 4 KB of data, which may not be enough for some secrets.² As file system volumes: You can mount secrets from Secret Manager as files in a volume within your application. Cloud Run will create a tmpfs volume and write the secrets as files in it. This method has some advantages, such as:

The files are updated every 30 seconds, so you can get the latest version of the secret faster.

The files are not visible in the Cloud Console or the Cloud SDK, which provides better security.

The files can store up to 64 KB of data, which allows for larger secrets.³ Therefore, for your use case, it is better to use the second method and mount the secret as a file system volume within your application. This way, you can ensure that your application has the latest password, and you can deploy it with no downtime.

To mount a secret as a file system volume in Cloud Run, you can use the following command:

```
gcloud beta run deploy SERVICE --image IMAGE_URL --update-secrets=/path/to/file=secretName:version where:
```

SERVICE is the name of your Cloud Run service.

IMAGE_URL is the URL of your container image.

/path/to/file is the path where you want to mount the secret file in your application.

secretName is the name of your secret in Secret Manager.

You can also use the Cloud Console to mount secrets as file system volumes. For more details, see [Mounting secrets from Secret Manager](#).

Reference:

1: [Overview | Secret Manager Documentation | Google Cloud](#)

2: [Using secrets as environment variables | Cloud Run Documentation | Google Cloud](#)

3: [Mounting secrets from Secret Manager | Cloud Run Documentation | Google Cloud](#)

NEW QUESTION: 76

You are ready to deploy a new feature of a web-based application to production. You want to use Google Kubernetes Engine (GKE) to perform a phased rollout to half of the web server pods.

What should you do?

- A. Use a partitioned rolling update.
- B. Use Node taints with NoExecute.
- C. Use a replica set in the deployment specification.
- D. Use a stateful set with parallel pod management policy.

Answer: ([SHOW ANSWER](#))

Explanation

<https://medium.com/velotio-perspectives/exploring-upgrade-strategies-for-stateful-sets-in-kubernetes-c02b8286f>

Valid Professional-Cloud-DevOps-Engineer Dumps shared by PrepPdf.com for Helping Passing Professional-Cloud-DevOps-Engineer Exam! PrepPdf.com now offer the **newest Professional-Cloud-DevOps-Engineer exam dumps**, the PrepPdf.com Professional-Cloud-DevOps-Engineer exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com Professional-Cloud-DevOps-Engineer dumps with Test Engine here:

<https://www.preppdf.com/Google/Professional-Cloud-DevOps-Engineer-prepaway-exam-dumps.html> (200 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 77

You built a serverless application by using Cloud Run and deployed the application to your production environment. You want to identify the resource utilization of the application for cost optimization. What should you do?

- A. Use Cloud Trace with distributed tracing to monitor the resource utilization of the application
- B. Use Cloud Profiler with Ops Agent to monitor the CPU and memory utilization of the application
- C. Use Cloud Monitoring to monitor the container CPU and memory utilization of the application
- D. Use Cloud Ops to create logs-based metrics to monitor the resource utilization of the application

Answer: A ([LEAVE A REPLY](#))

Explanation

The best option for giving developers the ability to test the latest revisions of the service before the service is exposed to customers is to run the `gcloud run deploy booking-engine --no-traffic --tag dev` command and use the `https://dev----booking-engine-abcdef.a.run.app` URL for testing. The `gcloud run deploy` command is a command that deploys a new revision of your service or updates an existing service. By using the `--no-traffic` flag, you can prevent any traffic from being sent to the new revision. By using the `--tag` flag, you can assign a tag to the new revision, such as `dev`. This way, you can create a new revision of your service without affecting your customers. You can also use the tag-based URL (e.g., `https://dev----booking-engine-abcdef.a.run.app`) to access and test the new revision.

NEW QUESTION: 78

You recently noticed that one of your services has exceeded the error budget for the current rolling window period. Your company's product team is about to launch a new feature. You want to follow Site Reliability Engineering (SRE) practices.

What should you do?

- A.** Notify the team that their error budget is used up. Negotiate with the team for a launch freeze or tolerate a slightly worse user experience.
- B.** Look through other metrics related to the product and find SLOs with remaining error budget. Reallocate the error budgets and allow the feature launch.
- C.** Escalate the situation and request additional error budget.
- D.** Notify the team about the lack of error budget and ensure that all their tests are successful so the launch will not further risk the error budget.

Answer: A ([LEAVE A REPLY](#))

The correct answer is

A, Notify the team that their error budget is used up. Negotiate with the team for a launch freeze or tolerate a slightly worse user experience.

According to the Site Reliability Engineering (SRE) practices, an error budget is the amount of unreliability that a service can tolerate without harming user satisfaction¹. An error budget is derived from the service-level objectives (SLOs), which are the measurable goals for the service quality². When a service exceeds its error budget, it means that it has violated its SLOs and may have negatively impacted the users. In this case, the SRE team should notify the product team that their error budget is used up and negotiate with them for a launch freeze or a lower SLO³. A launch freeze means that no new features are deployed until the service reliability is restored. A lower SLO means that the product team accepts a slightly worse user experience in exchange for launching new features. Both options require a trade-off between reliability and innovation, and should be agreed upon by both teams.

The other options are incorrect because they do not follow the SRE practices. Option B is incorrect because it violates the principle of error budget autonomy, which means that each service should have its own error budget and SLOs, and should not borrow or reallocate them from other services⁴. Option C is incorrect because it does not address the root cause of the error budget overspend, and may create unrealistic expectations for the service reliability. Option D is incorrect because it does not prevent the possibility of introducing new errors or bugs with the feature launch, which may further degrade the service quality and user satisfaction.

Reference:

Error Budgets, Error Budgets. Service Level Objectives, Service Level Objectives. Error Budget Policies, Error Budget Policies. Error Budget Autonomy, Error Budget Autonomy.

NEW QUESTION: 79

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 80

You are performing a semi-annual capacity planning exercise for your flagship service. You expect a service user growth rate of 10% month-over-month for the next six months. Your service is fully containerized and runs on a Google Kubernetes Engine (GKE) standard cluster across three zones with cluster autoscaling enabled. You currently consume about 30% of your total deployed CPU capacity and you require resilience against the failure of a zone. You want to ensure that your users experience minimal negative impact as a result of this growth or as a result of zone failure while you avoid unnecessary costs. How should you prepare to handle the predicted growth?

- A.** Verify the maximum node pool size, enable a Horizontal Pod Autoscaler, and then perform a load test to verify your expected resource needs.
- B.** Because you deployed the service on GKE and are using a cluster autoscaler, your GKE cluster will scale automatically regardless of growth rate.
- C.** Because you are only using 30% of deployed CPU capacity, there is significant headroom and you do not need to add any additional capacity for this rate of growth.
- D.** Proactively add 80% more node capacity to account for six months of 10% growth rate and then perform a load test to ensure that you have enough capacity.

Answer: ([SHOW ANSWER](#))

Explanation

The best option for preparing to handle the predicted growth is to verify the maximum node pool size, enable a Horizontal Pod Autoscaler, and then perform a load test to verify your expected resource needs. The maximum node pool size is a parameter that specifies the maximum number of nodes that can be added to a node pool by the cluster autoscaler. You should verify that the maximum node pool size is sufficient to accommodate your expected growth rate and avoid hitting any quota limits. The Horizontal Pod Autoscaler is a feature that automatically adjusts the number of Pods in a deployment or replica set based on observed CPU utilization or custom metrics. You should enable a Horizontal Pod Autoscaler for your application to ensure that it runs enough Pods to handle the load. A

load test is a test that simulates high user traffic and measures the performance and reliability of your application. You should perform a load test to verify your expected resource needs and identify any bottlenecks or issues.

NEW QUESTION: 81

You are running an experiment to see whether your users like a new feature of a web application. Shortly after deploying the feature as a canary release, you receive a spike in the number of 500 errors sent to users, and your monitoring reports show increased latency. You want to quickly minimize the negative impact on users. What should you do first?

- A.** Record data for the postmortem document of the incident.
- B.** Roll back the experimental canary release.
- C.** Trace the origin of 500 errors and the root cause of increased latency.
- D.** Start monitoring latency, traffic, errors, and saturation.

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 82

You are the Operations Lead for an ongoing incident with one of your services. The service usually runs at around 70% capacity. You notice that one node is returning 5xx errors for all requests. There has also been a noticeable increase in support cases from customers. You need to remove the offending node from the load balancer pool so that you can isolate and investigate the node. You want to follow Google-recommended practices to manage the incident and reduce the impact on users. What should you do?

- A.** 1. Communicate your intent to the incident team.
2. Perform a load analysis to determine if the remaining nodes can handle the increase in traffic offloaded from the removed node, and scale appropriately.
3. When any new nodes report healthy, drain traffic from the unhealthy node, and remove the unhealthy node from service.
- B.** 1. Communicate your intent to the incident team.
2. Add a new node to the pool, and wait for the new node to report as healthy.
3. When traffic is being served on the new node, drain traffic from the unhealthy node, and remove the old node from service.
- C.** 1. Drain traffic from the unhealthy node and remove the node from service.
2. Monitor traffic to ensure that the error is resolved and that the other nodes in the pool are handling the traffic appropriately.
3. Scale the pool as necessary to handle the new load.
4. Communicate your actions to the incident team.
- D.** 1. Drain traffic from the unhealthy node and remove the old node from service.
2. Add a new node to the pool, wait for the new node to report as healthy, and then serve traffic to the new node.
3. Monitor traffic to ensure that the pool is healthy and is handling traffic appropriately.

4. Communicate your actions to the incident team.

Answer: (SHOW ANSWER)

The correct answer is A, Communicate your intent to the incident team. Perform a load analysis to determine if the remaining nodes can handle the increase in traffic offloaded from the removed node, and scale appropriately. When any new nodes report healthy, drain traffic from the unhealthy node, and remove the unhealthy node from service.

This answer follows the Google-recommended practices for incident management, as described in the Chapter 9 - Incident Response, Google SRE Book¹. According to this source, some of the best practices are:

Maintain a clear line of command. Designate clearly defined roles. Keep a working record of debugging and mitigation as you go. Declare incidents early and often.

Communicate your intent before taking any action that might affect the service or the incident response. This helps to avoid confusion, duplication of work, or unintended consequences.

Perform a load analysis before removing a node from the load balancer pool, as this might affect the capacity and performance of the service. Scale the pool as necessary to handle the expected load.

Drain traffic from the unhealthy node before removing it from service, as this helps to avoid dropping requests or causing errors for users.

Answer A follows these best practices by communicating the intent to the incident team, performing a load analysis and scaling the pool, and draining traffic from the unhealthy node before removing it.

Answer B does not follow the best practice of performing a load analysis before adding or removing nodes, as this might cause overloading or underutilization of resources.

Answer C does not follow the best practice of communicating the intent before taking any action, as this might cause confusion or conflict with other responders.

Answer D does not follow the best practice of draining traffic from the unhealthy node before removing it, as this might cause errors for users.

Reference:

1: Chapter 9 - Incident Response, Google SRE Book

NEW QUESTION: 83

Your company's security team needs to have read-only access to Data Access audit logs in the _Required bucket. You want to provide your security team with the necessary permissions following the principle of least privilege and Google-recommended practices. What should you do?

- A.** Assign the roles/logging.viewer role to each member of the security team
- B.** Assign the roles/logging.viewer role to a group with all the security team members
- C.** Assign the roles/logging.privateLogViewer role to each member of the security team
- D.** Assign the roles/logging.privateLogviewer role to a group with all the security team members

Answer: D (LEAVE A REPLY)

The best option for providing your security team with the necessary permissions following the principle of least privilege and Google-recommended practices is to assign the roles/logging.privateLogViewer role to a group with all the security team members. The roles/logging.privateLogViewer role is a predefined role that grants read-only access to Data Access audit logs and other private logs in Cloud Logging. A group is a collection of users that can be assigned roles and permissions as a single unit. You can assign the roles/logging.privateLogViewer role to a group with all the security team members by using IAM policies. This way, you can provide your security team with the minimum level of access they need to view Data Access audit logs in the _Required bucket.

NEW QUESTION: 84

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?

- 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 (LEAVE A REPLY)

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 entities^{3,4}. 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.

Reference:

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.

NEW QUESTION: 85

You are running an application in a virtual machine (VM) using a custom Debian image. The image has the Stackdriver Logging agent installed. The VM has the cloud-platform scope. The application is logging information via syslog. You want to use Stackdriver Logging in the Google Cloud Platform Console to visualize the logs. You notice that syslog is not showing up in the "All logs" dropdown list of the Logs Viewer. What is the first thing you should do?

- A. Look for the agent's test log entry in the Logs Viewer.
- B. Install the most recent version of the Stackdriver agent.
- C. Verify the VM service account access scope includes the monitoring.write scope.
- D. SSH to the VM and execute the following commands on your VM: `ps ax | grep fluentd`

Answer: D (LEAVE A REPLY)

Explanation

https://cloud.google.com/compute/docs/access/service-accounts#associating_a_service_account_to_an_instance

NEW QUESTION: 86

You are using Terraform to manage infrastructure as code within a CI/CD pipeline. You notice that multiple copies of the entire infrastructure stack exist in your Google Cloud project, and a new copy is created each time a change to the existing infrastructure is made. You need to optimize your cloud spend by ensuring that only a single instance of your infrastructure stack exists at a time. You want to follow Google-recommended practices. What should you do?

- A. Create a new pipeline to delete old infrastructure stacks when they are no longer needed.

- B.** Confirm that the pipeline is storing and retrieving the terraform. if state file from Cloud Storage with the Terraform gcs backend
- C.** Verify that the pipeline is storing and retrieving the terraform.tfstate* file from a source control
- D.** Update the pipeline to remove any existing infrastructure before you apply the latest configuration

Answer: B (LEAVE A REPLY)

Explanation

The best option for optimizing your cloud spend by ensuring that only a single instance of your infrastructure stack exists at a time is to confirm that the pipeline is storing and retrieving the terraform.tfstate file from Cloud Storage with the Terraform gcs backend. The terraform.tfstate file is a file that Terraform uses to store the current state of your infrastructure. The Terraform gcs backend is a backend type that allows you to store the terraform.tfstate file in a Cloud Storage bucket. By using the Terraform gcs backend, you can ensure that your pipeline has access to the latest state of your infrastructure and avoid creating multiple copies of the entire infrastructure stack.

NEW QUESTION: 87

You are running an application on Compute Engine and collecting logs through Stackdriver. You discover that some personally identifiable information (PII) is leaking into certain log entry fields. You want to prevent these fields from being written in new log entries as quickly as possible. What should you do?

- A.** Stage log entries to Cloud Storage, and then trigger a Cloud Function to remove the fields and write the entries to Stackdriver via the Stackdriver Logging API.
- B.** Use the filter-record-transformer Fluentd filter plugin to remove the fields from the log entries in flight.
- C.** Wait for the application developers to patch the application, and then verify that the log entries are no longer exposing PII.
- D.** Use the fluent-plugin-record-reformer Fluentd output plugin to remove the fields from the log entries in flight.

Answer: B (LEAVE A REPLY)

NEW QUESTION: 88

You support a production service that runs on a single Compute Engine instance. You regularly need to spend time on recreating the service by deleting the crashing instance and creating a new instance based on the relevant image. You want to reduce the time spent performing manual operations while following Site Reliability Engineering principles. What should you do?

- A.** Create a Stackdriver Monitoring dashboard with SMS alerts to be able to start recreating the crashed instance promptly after it has crashed.

- B.** Create a Managed Instance Group with a single instance and use health checks to determine the system status.
- C.** Add a Load Balancer in front of the Compute Engine instance and use health checks to determine the system status.
- D.** File a bug with the development team so they can find the root cause of the crashing instance.

Answer: B (LEAVE A REPLY)

NEW QUESTION: 89

You built a serverless application by using Cloud Run and deployed the application to your production environment. You want to identify the resource utilization of the application for cost optimization. What should you do?

- A.** Use Cloud Trace with distributed tracing to monitor the resource utilization of the application.
- B.** Use Cloud Profiler with Ops Agent to monitor the CPU and memory utilization of the application.
- C.** Use Cloud Monitoring to monitor the container CPU and memory utilization of the application.
- D.** Use Cloud Ops to create logs-based metrics to monitor the resource utilization of the application.

Answer: A (LEAVE A REPLY)

The best option for giving developers the ability to test the latest revisions of the service before the service is exposed to customers is to run the `gcloud run deploy booking-engine --no-traffic --tag dev` command and use the `https://dev----booking-engine-abcdef.a.run.app` URL for testing. The `gcloud run deploy` command is a command that deploys a new revision of your service or updates an existing service. By using the `--no-traffic` flag, you can prevent any traffic from being sent to the new revision. By using the `--tag` flag, you can assign a tag to the new revision, such as `dev`. This way, you can create a new revision of your service without affecting your customers. You can also use the tag-based URL (e.g., `https://dev----booking-engine-abcdef.a.run.app`) to access and test the new revision.

NEW QUESTION: 90

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

- 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 ([LEAVE A REPLY](#))

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.

NEW QUESTION: 91

You need to run a business-critical workload on a fixed set of Compute Engine instances for several months. The workload is stable with the exact amount of resources allocated to it. You want to lower the costs for this workload without any performance implications.

What should you do?

- A.** Convert the instances to preemptible virtual machines.
- B.** Migrate the instances to a Managed Instance Group.
- C.** Purchase Committed Use Discounts.
- D.** Create an Unmanaged Instance Group for the instances used to run the workload.

Answer: A ([LEAVE A REPLY](#))

Valid Professional-Cloud-DevOps-Engineer Dumps shared by PrepPdf.com for Helping Passing Professional-Cloud-DevOps-Engineer Exam! PrepPdf.com now offer the **newest Professional-Cloud-DevOps-Engineer exam dumps**, the PrepPdf.com Professional-Cloud-DevOps-Engineer exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com Professional-Cloud-DevOps-Engineer dumps with Test Engine here:

<https://www.preppdf.com/Google/Professional-Cloud-DevOps-Engineer-prepaway-exam-dumps.html> (200 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 92

You support the backend of a mobile phone game that runs on a Google Kubernetes Engine (GKE) cluster.

The application is serving HTTP requests from users. You need to implement a solution that will reduce the network cost. What should you do?

- A. Configure the VPC as a Shared VPC Host project.
- B. Configure your network services on the Standard Tier.
- C. Configure your Kubernetes cluster as a Private Cluster.
- D. Configure a Google Cloud HTTP Load Balancer as Ingress.

Answer: B ([LEAVE A REPLY](#))

Explanation

The Standard Tier network service offers lower network costs than the Premium Tier. This is the correct option to reduce the network cost for the application.

NEW QUESTION: 93

You support an application deployed on Compute Engine. The application connects to a Cloud SQL instance to store and retrieve data. After an update to the application, users report errors showing database timeout messages. The number of concurrent active users remained stable. You need to find the most probable cause of the database timeout. What should you do?

- A. Determine whether there is an increased number of connections to the Cloud SQL instance.
- B. Use Cloud Security Scanner to see whether your Cloud SQL is under a Distributed Denial of Service (DDoS) attack.
- C. Use Stackdriver Profiler to visualize the resources utilization throughout the application.
- D. Check the serial port logs of the Compute Engine instance.

Answer: ([SHOW ANSWER](#)**)**

NEW QUESTION: 94

You are the Operations Lead for an ongoing incident with one of your services. The service usually runs at around 70% capacity. You notice that one node is returning 5xx errors for all requests. There has also been a noticeable increase in support cases from customers. You need to remove the offending node from the load balancer pool so that you can isolate and investigate the node. You want to follow Google-recommended practices to manage the incident and reduce the impact on users. What should you do?

- A. 1. Communicate your intent to the incident team.
2. Perform a load analysis to determine if the remaining nodes can handle the increase in traffic offloaded from the removed node, and scale appropriately.
3. When any new nodes report healthy, drain traffic from the unhealthy node, and remove the unhealthy node from service.
- B. 1. Communicate your intent to the incident team.
2. Add a new node to the pool, and wait for the new node to report as healthy.
3. When traffic is being served on the new node, drain traffic from the unhealthy node, and remove the old node from service.
- C. 1. Drain traffic from the unhealthy node and remove the node from service.

2. Monitor traffic to ensure that the error is resolved and that the other nodes in the pool are handling the traffic appropriately.

3. Scale the pool as necessary to handle the new load.

4. Communicate your actions to the incident team.

D. 1 . Drain traffic from the unhealthy node and remove the old node from service.

2. Add a new node to the pool, wait for the new node to report as healthy, and then serve traffic to the new node.

3. Monitor traffic to ensure that the pool is healthy and is handling traffic appropriately.

4. Communicate your actions to the incident team.

Answer: A (LEAVE A REPLY)

Explanation

The correct answer is A. Communicate your intent to the incident team. Perform a load analysis to determine if the remaining nodes can handle the increase in traffic offloaded from the removed node, and scale appropriately. When any new nodes report healthy, drain traffic from the unhealthy node, and remove the unhealthy node from service.

This answer follows the Google-recommended practices for incident management, as described in the Chapter

9 - Incident Response, Google SRE Book¹. According to this source, some of the best practices are:

Maintain a clear line of command. Designate clearly defined roles. Keep a working record of debugging and mitigation as you go. Declare incidents early and often.

Communicate your intent before taking any action that might affect the service or the incident response.

This helps to avoid confusion, duplication of work, or unintended consequences.

Perform a load analysis before removing a node from the load balancer pool, as this might affect the capacity and performance of the service. Scale the pool as necessary to handle the expected load.

Drain traffic from the unhealthy node before removing it from service, as this helps to avoid dropping requests or causing errors for users.

Answer A follows these best practices by communicating the intent to the incident team, performing a load analysis and scaling the pool, and draining traffic from the unhealthy node before removing it.

Answer B does not follow the best practice of performing a load analysis before adding or removing nodes, as this might cause overloading or underutilization of resources.

Answer C does not follow the best practice of communicating the intent before taking any action, as this might cause confusion or conflict with other responders.

Answer D does not follow the best practice of draining traffic from the unhealthy node before removing it, as this might cause errors for users.

References:

1: Chapter 9 - Incident Response, Google SRE Book

NEW QUESTION: 95

You need to run a business-critical workload on a fixed set of Compute Engine instances for several months.

The workload is stable with the exact amount of resources allocated to it. You want to lower the costs for this workload without any performance implications. What should you do?

- A. Convert the instances to preemptible virtual machines.
- B. Migrate the instances to a Managed Instance Group.
- C. Purchase Committed Use Discounts.
- D. Create an Unmanaged Instance Group for the instances used to run the workload.

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 96

Your team is designing a new application for deployment into Google Kubernetes Engine (GKE). You need to set up monitoring to collect and aggregate various application-level metrics in a centralized location. You want to use Google Cloud Platform services while minimizing the amount of work required to set up monitoring.

What should you do?

- A. Publish various metrics from the application directly to the Stackdriver Monitoring API, and then observe these custom metrics in Stackdriver.
- B. Install the Cloud Pub/Sub client libraries, push various metrics from the application to various topics, and then observe the aggregated metrics in Stackdriver.
- C. Install the OpenTelemetry client libraries in the application, configure Stackdriver as the export destination for the metrics, and then observe the application's metrics in Stackdriver.
- D. Emit all metrics in the form of application-specific log messages, pass these messages from the containers to the Stackdriver logging collector, and then observe metrics in Stackdriver.

Answer: A ([LEAVE A REPLY](#))

Explanation

https://cloud.google.com/kubernetes-engine/docs/concepts/custom-and-external-metrics#custom_metrics

<https://github.com/GoogleCloudPlatform/k8s-stackdriver/blob/master/custom-metrics-stackdriver-adapter/README> Your application can report a custom metric to Cloud Monitoring. You can configure Kubernetes to respond to these metrics and scale your workload automatically. For example, you can scale your application based on metrics such as queries per second, writes per second, network performance, latency when communicating with a different application, or other metrics that make sense for your workload.

<https://cloud.google.com/kubernetes-engine/docs/concepts/custom-and-external-metrics>

NEW QUESTION: 97

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.)

- 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) from 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 (LEAVE A REPLY)

Explanation

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

You want the application's SLO to more closely reflect its 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.

NEW QUESTION: 98

You are writing a postmortem for an incident that severely affected users. You want to prevent similar incidents in the future. Which two of the following sections should you include in the postmortem? (Choose two.)

- A. An explanation of the root cause of the incident
- B. A list of employees responsible for causing the incident
- C. A list of action items to prevent a recurrence of the incident
- D. Your opinion of the incident's severity compared to past incidents
- E. Copies of the design documents for all the services impacted by the incident

Answer: A,C (LEAVE A REPLY)

For a postmortem to be truly blameless, it must focus on identifying the contributing causes of the incident without indicting any individual or team for bad or inappropriate behavior.

NEW QUESTION: 99

Your company experiences bugs, outages, and slowness in its production systems.

Developers use the production environment for new feature development and bug fixes.

Configuration and experiments are done in the production environment, causing outages for users. Testers use the production environment for load testing, which often slows the production systems. You need to redesign the environment to reduce the number of bugs

and outages in production and to enable testers to load test new features. What should you do?

- A.** Secure the production environment to ensure that developers can't change it and set up one controlled update per year.
- B.** Create a development environment for writing code and a test environment for configurations, experiments, and load testing.
- C.** Create an automated testing script in production to detect failures as soon as they occur.
- D.** Create a development environment with smaller server capacity and give access only to developers and testers.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 100

Your organization recently adopted a container-based workflow for application development. Your team develops numerous applications that are deployed continuously through an automated build pipeline to a Kubernetes cluster in the production environment. The security auditor is concerned that developers or operators could circumvent automated testing and push code changes to production without approval. What should you do to enforce approvals?

- A.** Leverage Kubernetes Role-Based Access Control (RBAC) to restrict access to only approved users.
- B.** Use an Admission Controller to verify that incoming requests originate from approved sources.
- C.** Configure the build system with protected branches that require pull request approval.
- D.** Enable binary authorization inside the Kubernetes cluster and configure the build pipeline as an attestor.

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 101

You use Terraform to manage an application deployed to a Google Cloud environment. The application runs on instances deployed by a managed instance group. The Terraform code is deployed by using a CI/CD pipeline. When you change the machine type on the instance template used by the managed instance group, the pipeline fails at the terraform apply stage with the following error message:

You need to update the instance template and minimize disruption to the application and the number of pipeline runs. What should you do?

- A.** Delete the managed instance group and recreate it after updating the instance template.
- B.** Add a new instance template, update the managed instance group to use the new instance template, and delete the old instance template.
- C.** Remove the managed instance group from the Terraform state file, update the instance template, and reimport the managed instance group.

D. Set the `create_before_destroy` meta-argument to true in the lifecycle block on the instance template

Answer: [\(SHOW ANSWER\)](#)

Explanation

The best option for updating the instance template and minimizing disruption to the application and the number of pipeline runs is to set the `create_before_destroy` meta-argument to true in the lifecycle block on the instance template. The `create_before_destroy` meta-argument is a Terraform feature that specifies that a new resource should be created before destroying an existing one during an update. This way, you can avoid downtime and errors when updating a resource that is in use by another resource, such as an instance template that is used by a managed instance group. By setting the `create_before_destroy` meta-argument to true in the lifecycle block on the instance template, you can ensure that Terraform creates a new instance template with the updated machine type, updates the managed instance group to use the new instance template, and then deletes the old instance template.

NEW QUESTION: 102

You have a CI/CD pipeline that uses Cloud Build to build new Docker images and push them to Docker Hub. You use Git for code versioning. After making a change in the Cloud Build YAML configuration, you notice that no new artifacts are being built by the pipeline. You need to resolve the issue following Site Reliability Engineering practices. What should you do?

- A.** Run a Git compare between the previous and current Cloud Build Configuration files to find and fix the bug.
- B.** Change the CI pipeline to push the artifacts to Container Registry instead of Docker Hub.
- C.** Disable the CI pipeline and revert to manually building and pushing the artifacts.
- D.** Upload the configuration YAML file to Cloud Storage and use Error Reporting to identify and fix the issue.

Answer: **B** [\(LEAVE A REPLY\)](#)

NEW QUESTION: 103

You recently deployed your application in Google Kubernetes Engine (GKE) and now need to release a new version of the application. You need the ability to instantly roll back to the previous version of the application in case there are issues with the new version. Which deployment model should you use?

- A.** Perform a rolling deployment and test your new application after the deployment is complete
- B.** Perform a testing, and test your application periodically after the deployment is complete

- C. Perform a canary deployment, and test your new application periodically after the new version is deployed
- D. Perform a blue/green deployment and test your new application after the deployment is complete

Answer: D (LEAVE A REPLY)

The best deployment model for releasing a new version of your application in GKE with the ability to instantly roll back to the previous version is to perform a blue/green deployment and test your new application after the deployment is complete. A blue/green deployment is a deployment strategy that involves creating two identical environments, one running the current version of the application (blue) and one running the new version of the application (green). The traffic is switched from blue to green after testing the new version, and if any issues are discovered, the traffic can be switched back to blue instantly. This way, you can minimize downtime and risk during deployment.

NEW QUESTION: 104

You are reviewing your deployment pipeline in Google Cloud Deploy. You must reduce toil in the pipeline and you want to minimize the amount of time it takes to complete an end-to-end deployment. What should you do?

Choose 2 answers

- A. Create a trigger to notify the required team to complete the next step when manual intervention is required
- B. Divide the automation steps into smaller tasks
- C. Use a script to automate the creation of the deployment pipeline in Google Cloud Deploy
- D. Add more engineers to finish the manual steps.
- E. Automate promotion approvals from the development environment to the test environment

Answer: A,E (LEAVE A REPLY)

The best options for reducing toil in the pipeline and minimizing the amount of time it takes to complete an end-to-end deployment are to create a trigger to notify the required team to complete the next step when manual intervention is required and to automate promotion approvals from the development environment to the test environment. A trigger is a resource that initiates a deployment when an event occurs, such as a code change, a schedule, or a manual request. You can create a trigger to notify the required team to complete the next step when manual intervention is required by using Cloud Build or Cloud Functions. This way, you can reduce the waiting time and human errors in the pipeline. A promotion approval is a process that allows you to approve or reject a deployment from one environment to another, such as from development to test. You can automate promotion approvals from the development environment to the test environment by using Google Cloud Deploy or Cloud Build. This way, you can speed up the deployment process and avoid manual steps.

NEW QUESTION: 105

You manage an application that is writing logs to Stackdriver Logging. You need to give some team members the ability to export logs. What should you do?

- A.** Grant the team members the IAM role of logging.configWriter on Cloud IAM.
- B.** Configure Access Context Manager to allow only these members to export logs.
- C.** Create and grant a custom IAM role with the permissions logging.sinks.list and logging.sink.get.
- D.** Create an Organizational Policy in Cloud IAM to allow only these members to create log exports.

Answer: A ([LEAVE A REPLY](#))

Explanation/Reference: <https://cloud.google.com/logging/docs/access-control>

NEW QUESTION: 106

You are configuring your CI/CD pipeline natively on Google Cloud. You want builds in a pre-production Google Kubernetes Engine (GKE) environment to be automatically load-tested before being promoted to the production GKE environment. You need to ensure that only builds that have passed this test are deployed to production. You want to follow Google-recommended practices. How should you configure this pipeline with Binary Authorization?

- A.** Create an attestation for the builds that pass the load test by requiring the lead quality assurance engineer to sign the attestation by using a key stored in Cloud Key Management Service (Cloud KMS).
- B.** Create an attestation for the builds that pass the load test by using a private key stored in Cloud Key Management Service (Cloud KMS) authenticated through Workload Identity.
- C.** Create an attestation for the builds that pass the load test by using a private key stored in Cloud Key Management Service (Cloud KMS) with a service account JSON key stored as a Kubernetes Secret.
- D.** Create an attestation for the builds that pass the load test by requiring the lead quality assurance engineer to sign the attestation by using their personal private key.

Answer: B ([LEAVE A REPLY](#))

The correct answer is B, Create an attestation for the builds that pass the load test by using a private key stored in Cloud Key Management Service (Cloud KMS) authenticated through Workload Identity.

According to the Google Cloud documentation, Binary Authorization is a deploy-time security control that ensures only trusted container images are deployed on Google Kubernetes Engine (GKE) or Cloud Run1. Binary Authorization uses attestations to certify that a specific image has completed a previous stage in the CI/CD pipeline, such as passing a load test2. Attestations are signed by private keys that are associated with attestors, which are entities that verify the attestations3. To follow Google-recommended practices, you should store your private keys in Cloud Key Management Service (Cloud

KMS), which is a secure and scalable service for managing cryptographic keys⁴. You should also use Workload Identity, which is a feature that allows Kubernetes service accounts to act as Google service accounts, to authenticate to Cloud KMS and sign attestations without having to manage or expose service account keys⁵.

The other options are incorrect because they do not follow Google-recommended practices. Option A and option D require human intervention to sign the attestations, which is not scalable or automated. Option C exposes the service account JSON key as a Kubernetes Secret, which is less secure than using Workload Identity.

Reference:

Creating an attestor, Creating an attestor. Cloud Key Management Service

Documentation, Overview. Attestations overview, Attestations overview. Using Workload Identity with Binary Authorization, Using Workload Identity with Binary Authorization.

Binary Authorization, Binary Authorization.

Valid Professional-Cloud-DevOps-Engineer Dumps shared by PrepPdf.com for Helping Passing Professional-Cloud-DevOps-Engineer Exam! PrepPdf.com now offer the **newest Professional-Cloud-DevOps-Engineer exam dumps**, the PrepPdf.com Professional-Cloud-DevOps-Engineer exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com Professional-Cloud-DevOps-Engineer dumps with Test Engine here:

<https://www.preppdf.com/Google/Professional-Cloud-DevOps-Engineer-prepaway-exam-dumps.html> (200 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 107

You have an application running in Google Kubernetes Engine. The application invokes multiple services per request but responds too slowly. You need to identify which downstream service or services are causing the delay. What should you do?

- A. Analyze VPC flow logs along the path of the request.
- B. Create a Dataflow pipeline to analyze service metrics in real time.
- C. Investigate the Liveness and Readiness probes for each service.
- D. Use a distributed tracing framework such as OpenTelemetry or Stackdriver Trace.

Answer: ([SHOW ANSWER](#))

Valid Professional-Cloud-DevOps-Engineer Dumps shared by PrepPdf.com for Helping Passing Professional-Cloud-DevOps-Engineer Exam! PrepPdf.com now offer the **newest Professional-Cloud-DevOps-Engineer exam dumps**, the PrepPdf.com Professional-Cloud-DevOps-Engineer exam **questions have been updated** and

answers have been corrected get the **newest** PrepPdf.com Professional-Cloud-DevOps-Engineer dumps with Test Engine here:

<https://www.preppdf.com/Google/Professional-Cloud-DevOps-Engineer-prepaway-exam-dumps.html> (200 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)