Professional-Cloud-DevOps-Engineer Dumps

Google Cloud Certified - Professional Cloud DevOps Engineer Exam

https://www.certleader.com/Professional-Cloud-DevOps-Engineer-dumps.html

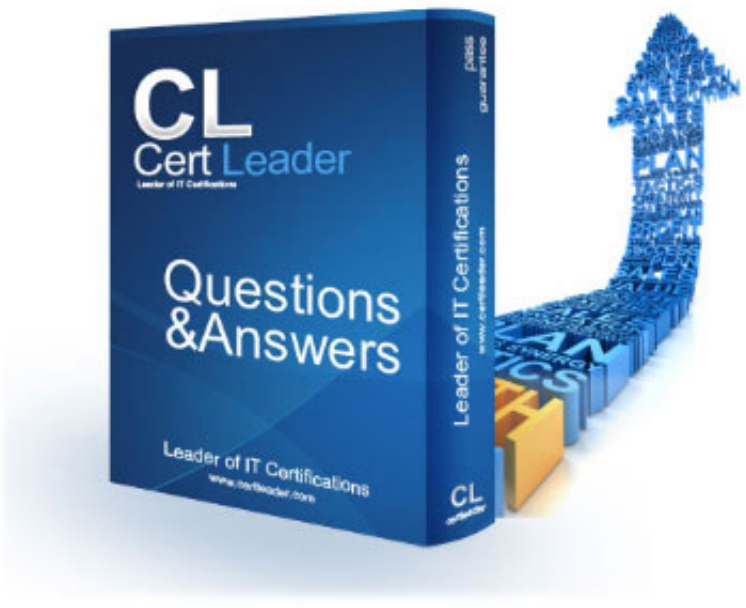

NEW QUESTION 1
You have an application running in Google Kubernetes Engine. The application invokes multiple services per request but responds too slowly. You need to identify which downstream service or services are causing the delay. What should you do?

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

Answer: C

NEW QUESTION 2
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 Cloud Storage to cache intermediate artifacts.
B. Run multiple Jenkins agents to parallelize the build.
C. Use multiple smaller build steps to minimize execution time.
D. Use larger Cloud Build virtual machines (VMs) by using the machine-type option.

Answer: C

Explanation:
https://cloud.google.com/storage/docs/best-practices
https://cloud.google.com/build/docs/speeding-up-builds#caching_directories_with_google_cloud_storage Caching directories with Google Cloud Storage To increase the speed of a build, reuse the results from a
previous build. You can copy the results of a previous build to a Google Cloud Storage bucket, use the results for faster calculation, and then copy the new results back to the bucket. Use this method when your build takes a long time and produces a small number of files that does not take time to copy to and from Google Cloud Storage.
upvoted 2 times

NEW QUESTION 3
You encounter a large number of outages in the production systems you support. You receive alerts for all the outages that wake you up at night. The alerts are due to unhealthy systems that are automatically restarted within a minute. You want to set up a process that would prevent staff burnout while following Site Reliability Engineering practices. What should you do?

A. Eliminate unactionable alerts.
B. Create an incident report for each of the alerts.
C. Distribute the alerts to engineers in different time zones.
D. Redefine the related Service Level Objective so that the error budget is not exhausted.

Answer: A

Explanation:
Eliminate bad monitoring : Unactionable alerts (i.e., spam) https://cloud.google.com/blog/products/management-tools/meeting-reliability-challenges-with-sre-principles
agree with kyubiblaze about having to remove unactionable items aka spam: "good monitoring alerts on actionable problems" @ https://cloud.google.com/blog/products/management-tools/meeting-reliability-challenges-with-sre-principles

NEW QUESTION 4
You support an application running on App Engine. The application is used globally and accessed from various device types. You want to know the number of connections. You are using Stackdriver Monitoring for App Engine. What metric should you use?

A. flex/connections/current
B. tcp_ssl_proxy/new_connections
C. tcp_ssl_proxy/open_connections
D. flex/instance/connections/current

Answer: A

Explanation:
https://cloud.google.com/monitoring/api/metrics_gcp#gcp-appengine

NEW QUESTION 5
You support a popular mobile game application deployed on Google Kubernetes Engine (GKE) across several Google Cloud regions. Each region has multiple Kubernetes clusters. You receive a report that none of the users in a specific region can connect to the application. You want to resolve the incident while following Site Reliability Engineering practices. What should you do first?

A. Reroute the user traffic from the affected region to other regions that don’t report issues.
B. Use Stackdriver Monitoring to check for a spike in CPU or memory usage for the affected region.
C. Add an extra node pool that consists of high memory and high CPU machine type instances to the cluster.
D. Use Stackdriver Logging to filter on the clusters in the affected region, and inspect error messages in the logs.

Answer: A

Explanation:
Google always aims to first stop the impact of an incident, and then find the root cause (unless the root cause just happens to be identified early on).

NEW QUESTION 6
Your application runs on Google Cloud Platform (GCP). You need to implement Jenkins for deploying application releases to GCP. You want to streamline the release process, lower operational toil, and keep user data secure. What should you do?

A. Implement Jenkins on local workstations.
B. Implement Jenkins on Kubernetes on-premises
C. Implement Jenkins on Google Cloud Functions.
D. Implement Jenkins on Compute Engine virtual machines.

Answer: D

Explanation:
Your application runs on Google Cloud Platform (GCP). You need to implement Jenkins for deploying application releases to GCP. You want to streamline the release process, lower operational toil, and keep user data secure. What should you do?
https://plugins.jenkins.io/google-compute-engine/

NEW QUESTION 7
You support a trading application written in Python and hosted on App Engine flexible environment. You want to customize the error information being sent to Stackdriver Error Reporting. What should you do?

A. Install the Stackdriver Error Reporting library for Python, and then run your code on a Compute Engine VM.
B. Install the Stackdriver Error Reporting library for Python, and then run your code on Google Kubernetes Engine.
C. Install the Stackdriver Error Reporting library for Python, and then run your code on App Engine flexible environment.
D. Use the Stackdriver Error Reporting API to write errors from your application to ReportedErrorEvent, and then generate log entries with properly formatted error messages in Stackdriver Logging.

Answer: D

Explanation:
https://cloud.google.com/error-reporting/docs/formatting-error-messages https://cloud.google.com/error-reporting/docs/reference/libraries#client-libraries-install-python no need to install error reporting library on App Engine Flex.

NEW QUESTION 8
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. Check the serial port logs of the Compute Engine instance.
B. Use Stackdriver Profiler to visualize the resources utilization throughout the application.
C. Determine whether there is an increased number of connections to the Cloud SQL instance.
D. Use Cloud Security Scanner to see whether your Cloud SQL is under a Distributed Denial of Service (DDoS) attack.

Answer: B

NEW QUESTION 9
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. File a bug with the development team so they can find the root cause of the crashing instance.
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. Create a Stackdriver Monitoring dashboard with SMS alerts to be able to start recreating the crashed instance promptly after it has crashed.

Answer: B

NEW QUESTION 10
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 logging agent to the application servers.• Give the developers the IAM Logs Private Logs Viewer role to access Stackdriver and view logs.
C. • Deploy the Stackdriver monitoring agent to the application servers.• Give the developers the IAM Monitoring Viewer role to access Stackdriver and view metrics.
D. • 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 IAM Object Viewer access to view the logs in the specified bucket.

Answer: A

Explanation:
https://cloud.google.com/logging/docs/audit#access-control

NEW QUESTION 10
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. Migrate the instances to a Managed Instance Group.
C. Convert the instances to preemptible virtual machines.
D. Create an Unmanaged Instance Group for the instances used to run the workload.

Answer: A

NEW QUESTION 12
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. Configure the build system with protected branches that require pull request approval.
B. Use an Admission Controller to verify that incoming requests originate from approved sources.
C. Leverage Kubernetes Role-Based Access Control (RBAC) to restrict access to only approved users.
D. Enable binary authorization inside the Kubernetes cluster and configure the build pipeline as an attestor.

Answer: D

Explanation:
The keywords here is "developers or operators". Option A the operators could push images to production without approval (operators could touch the cluster directly and the cluster cannot do any action against them). Rest same as francisco_guerra.

NEW QUESTION 14
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 minimizing costs. Different teams should not be able to access other teams' environments. What should you do?

A. Create one GCP Project per tea
B. In each project, create a cluster for Development and one for Productio
C. Grant the teams IAM access to their respective clusters.
D. Create one GCP Project per tea
E. In each project, create a cluster with a Kubernetes namespace for Development and one for Productio
F. Grant the teams IAM access to their respective clusters.
G. Create a Development and a Production GKE cluster in separate project
H. 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.
I. Create a Development and a Production GKE cluster in separate project
J. 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

Explanation:
https://cloud.google.com/architecture/prep-kubernetes-engine-for-prod#roles_and_groups

NEW QUESTION 15
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. Disable the CI pipeline and revert to manually building and pushing the artifacts.
B. Change the CI pipeline to push the artifacts to Container Registry instead of Docker Hub.
C. Upload the configuration YAML file to Cloud Storage and use Error Reporting to identify and fix the issue.
D. Run a Git compare between the previous and current Cloud Build Configuration files to find and fix the bug.

Answer: D

Explanation:
"After making a change in the Cloud Build YAML configuration, you notice that no new artifacts are being built by the pipeline"- means something wrong on the recent change not with the image registry.

NEW QUESTION 20
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 BigOuery.* 2. Create a dashboard in Data Studio.* 3. Share the dashboard with your stakeholders.
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 (rom Stackdriver to BigQuery.* 2. From BigQuer
D. 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.
E. * 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.

Answer: A

NEW QUESTION 21
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 I grep fluentd

Answer: D

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

NEW QUESTION 23
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 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 lo a versioned Cloud Storage bucket as the latest master version.
B. • 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.
C. • Store your code as text files in Google Drive in a defined folder structure that organizes the files.• At the end of each da
D. 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.
E. • 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.

Answer: B

NEW QUESTION 25
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 Studi
B. Filler (he logs on the cache misses.
C. Configure Stackdriver Profiler to identify and visualize when the cache misses occur based on the logs.
D. Create a logs-based metric in Stackdriver Logging and a dashboard for that metric in Stackdriver Monitoring.
E. Configure BigOuery as a sink for Stackdriver Loggin
F. Create a scheduled query to filter the cache miss logs and write them to a separate table

Answer: C

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

NEW QUESTION 26
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. Create a custom builder for Cloud Build that will only push images to gcr.io/altostrat-images.
B. Use a Binary Authorization policy that includes the whitelist name pattern gcr.io/attostrat-images/.
C. Add logic to the deployment pipeline to check that all manifests contain only images from 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: B

NEW QUESTION 27
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

Explanation:
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 28
You support a Node.js application running on Google Kubernetes Engine (GKE) in production. The application makes several HTTP requests to dependent applications. You want to anticipate which dependent applications might cause performance issues. What should you do?

A. Instrument all applications with Stackdriver Profiler.
B. Instrument all applications with Stackdriver Trace and review inter-service HTTP requests.
C. Use Stackdriver Debugger to review the execution of logic within each application to instrument all applications.

D. Modify the Node.js application to log HTTP request and response times to dependent application
E. Use Stackdriver Logging to find dependent applications that are performing poorly.

Answer: B

NEW QUESTION 30
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 an automated testing script in production to detect failures as soon as they occur.
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 a development environment for writing code and a test environment for configurations, experiments, and load testing.

Answer: D

NEW QUESTION 31
You support an application running on GCP and want to configure SMS notifications to your team for the most critical alerts in Stackdriver Monitoring. You have already identified the alerting policies you want to configure this for. What should you do?

A. Download and configure a third-party integration between Stackdriver Monitoring and an SMS gateway.Ensure that your team members add their SMS/phone numbers to the external tool.
B. Select the Webhook notifications option for each alerting policy, and configure it to use a third-party integration too
C. Ensure that your team members add their SMS/phone numbers to the external tool.
D. Ensure that your team members set their SMS/phone numbers in their Stackdriver Profil
E. Select the SMS notification option for each alerting policy and then select the appropriate SMS/phone numbers from the list.
F. Configure a Slack notification for each alerting polic
G. Set up a Slack-to-SMS integration to send SMS messages when Slack messages are receive
H. Ensure that your team members add their SMS/phone numbers to the external integration.

Answer: C

Explanation:
https://cloud.google.com/monitoring/support/notification-options#creating_channels To configure SMS notifications, do the following:
In the SMS section, click Add new and follow the instructions. Click Save. When you set up your alerting policy, select the SMS notification type and choose a verified phone number from the list.

NEW QUESTION 32
You are deploying an application that needs to access sensitive information. You need to ensure that this information is encrypted and the risk of exposure is minimal if a breach occurs. What should you do?

A. Store the encryption keys in Cloud Key Management Service (KMS) and rotate the keys frequently
B. Inject the secret at the time of instance creation via an encrypted configuration management system.
C. Integrate the application with a Single sign-on (SSO) system and do not expose secrets to the application
D. Leverage a continuous build pipeline that produces multiple versions of the secret for each instance of the application.

Answer: A

Explanation:
https://cloud.google.com/security-key-management

NEW QUESTION 34
Your product is currently deployed in three Google Cloud Platform (GCP) zones with your users divided between the zones. You can fail over from one zone to another, but it causes a 10-minute service disruption for the affected users. You typically experience a database failure once per quarter and can detect it within five minutes. You are cataloging the reliability risks of a new real-time chat feature for your product. You catalog the following information for each risk:
• Mean Time to Detect (MUD} in minutes
• Mean Time to Repair (MTTR) in minutes
• Mean Time Between Failure (MTBF) in days
• User Impact Percentage
The chat feature requires a new database system that takes twice as long to successfully fail over between zones. You want to account for the risk of the new database failing in one zone. What would be the values for the risk of database failover with the new system?

A. MTTD: 5MTTR: 10MTBF: 90Impact: 33%
B. MTTD:5 MTTR: 20MTBF: 90Impact: 33%
C. MTTD:5 MTTR: 10MTBF: 90Impact 50%
D. MTTD:5 MTTR: 20MTBF: 90Impact: 50%

Answer: B

Explanation:
https://www.atlassian.com/incident-management/kpis/common-metrics https://linkedin.github.io/school-of-sre/

NEW QUESTION 39
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 item
B. Share it with the manager only.
C. Develop a postmortem that includes the root causes, resolution, lessons learned, and a prioritized list of action item
D. Share it on the engineering organization's document portal.
E. Develop a postmortem that includes the root causes, resolution, lessons learned, the list of people responsible, and a list of action items for each perso
F. Share it with the manager only.
G. Develop a postmortem that includes the root causes, resolution, lessons learned, the list of people responsible, and a list of action items for each perso
H. Share it on the engineering organization's document portal.

Answer: B

NEW QUESTION 43
You are managing the production deployment to a set of Google Kubernetes Engine (GKE) clusters. You want to make sure only images which are successfully built by your trusted CI/CD pipeline are deployed to production. What should you do?

A. Enable Cloud Security Scanner on the clusters.
B. Enable Vulnerability Analysis on the Container Registry.
C. Set up the Kubernetes Engine clusters as private clusters.
D. Set up the Kubernetes Engine clusters with Binary Authorization.

Answer: D

Explanation:
https://cloud.google.com/binary-authorization/docs/overview

NEW QUESTION 45
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. Configure the horizontal pod autoscaler to use the average response time from the Liveness and Readiness probes.
B. Configure the vertical pod autoscaler in GKE and enable the cluster autoscaler to scale the cluster as pods expand.
C. Install the Stackdriver custom metrics adapter and configure a horizontal pod autoscaler to use the number of requests provided by the GCLB.
D. Expose the NGINX stats endpoint and configure the horizontal pod autoscaler to use the request metrics exposed by the NGINX deployment.

Answer: C

Explanation:
https://cloud.google.com/kubernetes-engine/docs/tutorials/autoscaling-metrics

NEW QUESTION 47
Your team of Infrastructure DevOps Engineers is growing, and you are starting to use Terraform to manage infrastructure. You need a way to implement code versioning and to share code with other team members. What should you do?

A. Store the Terraform code in a version-control syste
B. Establish procedures for pushing new versions and merging with the master.
C. Store the Terraform code in a network shared folder with child folders for each version releas
D. Ensure that everyone works on different files.
E. Store the Terraform code in a Cloud Storage bucket using object versionin
F. Give access to the bucket to every team member so they can download the files.
G. Store the Terraform code in a shared Google Drive folder so it syncs automatically to every team member’s compute
H. Organize files with a naming convention that identifies each new version.

Answer: A

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

NEW QUESTION 52
You manage several production systems that run on Compute Engine in the same Google Cloud Platform (GCP) project. Each system has its own set of dedicated Compute Engine instances. You want to know how must it costs to run each of the systems. What should you do?

A. In the Google Cloud Platform Console, use the Cost Breakdown section to visualize the costs per system.
B. Assign all instances a label specific to the system they ru
C. Configure BigQuery billing export and query costs per label.
D. Enrich all instances with metadata specific to the system they ru
E. Configure Stackdriver Logging to export to BigQuery, and query costs based on the metadata.
F. Name each virtual machine (VM) after the system it run
G. Set up a usage report export to a Cloud Storage bucke
H. Configure the bucket as a source in BigQuery to query costs based on VM name.

Answer: B

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

NEW QUESTION 55

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 AP
B. Deprecate the old version of the API.Contact remaining users of the old API.Provide best effort support to users of the old AP
C. Turn down the old version of the API.
D. Announce deprecation of the old version of the AP
E. Introduce the new version of the API.Contact remaining users on the old AP
F. Deprecate the old version of the AP
G. Turn down the old version of the API.Provide best effort support to users of the old API.
H. Announce deprecation of the old version of the AP
I. Contact remaining users on the old API.Introduce the new version of the AP
J. Deprecate the old version of the API.Provide best effort support to users of the old AP
K. Turn down the old version of the API.
L. Introduce the new version of the AP
M. Contact remaining users of the old API.Announce deprecation of the old version of the AP
N. 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

NEW QUESTION 57
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. Before merging new code, require 2 different peers to review the code changes.
B. Adopt the blue/green deployment strategy when releasing new code via a CD server.
C. Integrate a code linting tool to validate coding standards before any code is accepted into the repository.
D. Require developers to run automated integration tests on their local development environments before release.
E. Configure a CI serve
F. Add a suite of unit tests to your code and have your CI server run them on commit and verify any changes.

Answer: BE

NEW QUESTION 61
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

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

NEW QUESTION 63
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 production VPC network frontend and backend subnets only with a sample volume scale of 0.5.
B. Enable VPC Flow Logs on the production VPC network frontend and backend subnets only with a sample volume scale of 1.0.
C. Enable VPC Flow Logs on the testing and production VPC network frontend and backend subnets with a volume scale of 0.5. Apply changes intesting before production.
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: D

NEW QUESTION 68
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 tim
B. Instruct developers to not store secrets at rest.
C. Store secrets in a separate configuration file on Gi
D. Provide select developers with access to the configuration file.
E. Store secrets in Cloud Storage encrypted with a key from Cloud KM
F. Provide the CI/CD pipeline with access to Cloud KMS via IAM.
G. Encrypt the secrets and store them in the source code repositor
H. Store a decryption key in a separate repository and grant your pipeline access to it

Answer: C

NEW QUESTION 70

You support a service that recently had an outage. The outage was caused by a new release that exhausted the service memory resources. You rolled back the release successfully to mitigate the impact on users. You are now in charge of the post-mortem for the outage. You want to follow Site Reliability Engineering practices when developing the post-mortem. What should you do?

A. Focus on developing new features rather than avoiding the outages from recurring.
B. Focus on identifying the contributing causes of the incident rather than the individual responsible for the cause.
C. Plan individual meetings with all the engineers involve
D. Determine who approved and pushed the new release to production.
E. Use the Git history to find the related code commi
F. Prevent the engineer who made that commit from working on production services.

Answer: B

NEW QUESTION 73
You are developing a strategy for monitoring your Google Cloud Platform (GCP) projects in production using Stackdriver Workspaces. One of the requirements is to be able to quickly identify and react to production environment issues without false alerts from development and staging projects. You want to ensure that you adhere to the principle of least privilege when providing relevant team members with access to Stackdriver Workspaces. What should you do?

A. Grant relevant team members read access to all GCP production project
B. Create Stackdriver workspaces inside each project.
C. Grant relevant team members the Project Viewer IAM role on all GCP production project
D. Create Slackdriver workspaces inside each project.
E. Choose an existing GCP production project to host the monitoring workspac
F. Attach the production projects to this workspac
G. Grant relevant team members read access to the Stackdriver Workspace.
H. Create a new GCP monitoring project, and create a Stackdriver Workspace inside i
I. Attach the production projects to this workspac
J. Grant relevant team members read access to the Stackdriver Workspace.

Answer: D

Explanation:
"A Project can host many Projects and appear in many Projects, but it can only be used as the scoping project once. We recommend that you create a new Project for the purpose of having multiple Projects in the same scope."

NEW QUESTION 75
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) fro the application.
E. Announce planned downtime to consume more error budget, and ensure that users are not depending on a tighter SLO.

Answer: DE

Explanation:
https://sre.google/sre-book/service-level-objectives/
You want the application's SLO to more closely reflect it's observed reliability. The key here is error budget never goes over 5%. This means they can have additional downtime and still stay within their budget.

NEW QUESTION 80
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

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 83

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your Professional-Cloud-DevOps-Engineer Exam with Our Prep Materials Via below:

https://www.certleader.com/Professional-Cloud-DevOps-Engineer-dumps.html