Google

Exam Questions Professional-Cloud-DevOps-Engineer

Google Cloud Certified - Professional Cloud DevOps Engineer Exam

About Exambible

Your Partner of IT Exam

Found in 1998

Exambible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, Exambible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

    All examinations will be up to date.

* 24/7 Quality Support

    We will provide service round the clock.

* 100% Pass Rate

    Our guarantee that you will pass the exam.

* Unique Gurantee

    If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

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 Spinnaker to deploy your application and have created a canary deployment stage in the pipeline. Your application has an in-memory cache that loads objects at start time. You want to automate the comparison of the canary version against the production version. How should you configure the canary analysis?

A. Compare the canary with a new deployment of the current production version.
B. Compare the canary with a new deployment of the previous production version.
C. Compare the canary with the existing deployment of the current production version.
D. Compare the canary with the average performance of a sliding window of previous production versions.

Answer: A

Explanation:
https://cloud.google.com/architecture/automated-canary-analysis-kubernetes-engine-spinnaker https://spinnaker.io/guides/user/canary/best-practices/#compare-canary-against-baseline-not-against-production

NEW QUESTION 3
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. Reference the image digest in the source control tag.
B. Supply the source control tag as a parameter within the image name.
C. Use Cloud Build to include the release version tag in the application image.
D. Use GCR digest versioning to match the image to the tag in source control.

Answer: B

Explanation:
https://cloud.google.com/container-registry/docs/pushing-and-pulling

NEW QUESTION 4
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. Buckelize Ihe request latencies into ranges, and then compute the percentile at 100 ms.
B. Bucketize the request latencies into ranges, and then compute the median and 90th percentiles.
C. Count the number of home page requests that load in under 100 ms, and then divide by the total number of home page requests.
D. Count the number of home page requests that load in under 100 m
E. and then divide by the total number of all web application requests.

Answer: C

Explanation:
https://sre.google/workbook/implementing-slos/
In the SRE principles book, it's recommended treating the SLI as the ratio of two numbers: the number of good events divided by the total number of events. For example: Number of successful HTTP requests / total HTTP requests (success rate)

NEW QUESTION 5
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. All PII entries begin with the text userinfo. You want to capture these log entries in a secure location for later review and prevent them from leaking to Stackdriver Logging. What should you do?

A. Create a basic log filter matching userinfo, and then configure a log export in the Stackdriver console with Cloud Storage as a sink.
B. Use a Fluentd filter plugin with the Stackdriver Agent to remove log entries containing userinfo, and then copy the entries to a Cloud Storage bucket.
C. Create an advanced log filter matching userinfo, configure a log export in the Stackdriver console with Cloud Storage as a sink, and then configure a tog exclusion with userinfo as a filter.
D. Use a Fluentd filter plugin with the Stackdriver Agent to remove log entries containing userinfo, create an advanced log filter matching userinfo, and then configure a log export in the Stackdriver console with Cloud Storage as a sink.

Answer: B

Explanation:
https://medium.com/google-cloud/fluentd-filter-plugin-for-google-cloud-data-loss-prevention-api-42bbb1308e7

NEW QUESTION 6

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

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

NEW QUESTION 7

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: C

Explanation:
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 8

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

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

NEW QUESTION 9

You need to define Service Level Objectives (SLOs) for a high-traffic multi-region web application. Customers expect the application to always be available and have fast response times. Customers are currently happy with the application performance and availability. Based on current measurement, you observe that the 90th percentile of latency is 120ms and the 95th percentile of latency is 275ms over a 28-day window. What latency SLO would you recommend to the team to publish?

A. 90th percentile – 100ms 95th percentile – 250ms
B. 90th percentile – 120ms 95th percentile – 275ms
C. 90th percentile – 150ms 95th percentile – 300ms
D. 90th percentile – 250ms 95th percentile – 400ms

Answer: C

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

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 10

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

Answer: D

NEW QUESTION 11
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: CE

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

NEW QUESTION 12
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 14
You are on-call for an infrastructure service that has a large number of dependent systems. You receive an alert indicating that the service is failing to serve most of its requests and all of its dependent systems with hundreds of thousands of users are affected. As part of your Site Reliability Engineering (SRE) incident management protocol, you declare yourself Incident Commander (IC) and pull in two experienced people from your team as Operations Lead (OLJ and Communications Lead (CL). What should you do next?

A. Look for ways to mitigate user impact and deploy the mitigations to production.
B. Contact the affected service owners and update them on the status of the incident.
C. Establish a communication channel where incident responders and leads can communicate with each other.
D. Start a postmortem, add incident information, circulate the draft internally, and ask internal stakeholders for input.

Answer: A

Explanation:
https://sre.google/sre-book/managing-incidents/

NEW QUESTION 18
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 22
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 24
You are responsible for the reliability of a high-volume enterprise application. A large number of users report that an important subset of the application's functionality – a data intensive reporting feature – is consistently failing with an HTTP 500 error. When you investigate your application's dashboards, you notice a strong correlation between the failures and a metric that represents the size of an internal queue used for generating reports. You trace the failures to a reporting backend that is experiencing high I/O wait times. You quickly fix the issue by resizing the backend's persistent disk (PD). How you need to create an availability Service Level Indicator (SLI) for the report generation feature. How would you define it?

A. As the I/O wait times aggregated across all report generation backends
B. As the proportion of report generation requests that result in a successful response
C. As the application's report generation queue size compared to a known-good threshold
D. As the reporting backend PD throughout capacity compared to a known-good threshold

Answer: B

Explanation:
According to SRE Workbook, one of potential SLI is as below:
* Type of service: Request-driven
* Type of SLI: Availability
* Description: The proportion of requests that resulted in a successful response. https://sre.google/workbook/implementing-slos/

NEW QUESTION 28
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 31
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 minute
B. Commit to "next update" times.
C. Provide periodic updates to all stakeholders in a timely manne
D. Commit to a "next update" time in all communications.
E. Delegate the responding to internal stakeholder emails to another member of the Incident Response Tea
F. Focus on providing responses directly to customers.
G. Provide all internal stakeholder emails to the Incident Commander, and allow them to manage internal communication
H. Focus on providing responses directly to customers.

Answer: B

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 34
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 39
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 40
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 45
You use Cloud Build to build and deploy your application. You want to securely incorporate database credentials and other application secrets into the build pipeline. You also want to minimize the development effort. What should you do?

A. Create a Cloud Storage bucket and use the built-in encryption at res
B. Store the secrets in the bucket and grant Cloud Build access to the bucket.
C. Encrypt the secrets and store them in the application repositor
D. Store a decryption key in a separate repository and grant Cloud Build access to the repository.
E. Use client-side encryption to encrypt the secrets and store them in a Cloud Storage bucke
F. Store a decryption key in the bucket and grant Cloud Build access to the bucket.
G. Use Cloud Key Management Service (Cloud KMS) to encrypt the secrets and include them in your Cloud Build deployment configuratio
H. Grant Cloud Build access to the KeyRing.

Answer: D

Explanation:
https://cloud.google.com/build/docs/securing-builds/use-encrypted-credentials

NEW QUESTION 47
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 51
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 56
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 59
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 duster as a Private Cluster.
D. Configure a Google Cloud HTTP Load Balancer as Ingress.

Answer: D

Explanation:
Costs associated with a load balancer are charged to the project containing the load balancer components. Because of these benefits, container-native load balancing is the recommended solution for load balancing through Ingress. When NEGs are used with GKE Ingress, the Ingress controller facilitates the creation of all aspects of the L7 load balancer. This includes creating the virtual IP address, forwarding rules, health checks, firewall rules, and more.
https://cloud.google.com/architecture/best-practices-for-running-cost-effective-kubernetes-applications-on-gke

NEW QUESTION 63
You support a multi-region web service running on Google Kubernetes Engine (GKE) behind a Global HTTP'S Cloud Load Balancer (CLB). For legacy reasons, user requests first go through a third-party Content Delivery Network (CDN). which then routes traffic to the CLB. You have already implemented an availability Service Level Indicator (SLI) at the CLB level. However, you want to increase coverage in case of a potential load balancer misconfiguration. CDN failure, or other global networking catastrophe. Where should you measure this new SLI?
Choose 2 answers

A. Your application servers' logs
B. Instrumentation coded directly in the client
C. Metrics exported from the application servers
D. GKE health checks for your application servers
E. A synthetic client that periodically sends simulated user requests

Answer: BE

NEW QUESTION 66
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 Slacked 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 metricKin
E. set toMETRIc_KIND_UNSPECIFIEDanda valueType set to INT64.• In Stackdriver's Metrics Explorer, use a Stacked Area graph to visualize the metric.

Answer: C

Explanation:
 https://sre.google/workbook/implementing-slos/ https://cloud.google.com/architecture/adopting-slos/
Latency is commonly measured as a distribution. Given a distribution, you can measure various percentiles.
For example, you might measure the number of requests that are slower than the historical 99th percentile.

NEW QUESTION 68
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: AC

Explanation:
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 73
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 78
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 79
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 80
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. Verity the maximum node pool size, enable a horizontal pod autoscaler, and then perform a load test to verity your expected resource needs.
B. Because you are deployed on GKE and are using a cluster autoscale
C. your GKE cluster will scale automatically, regardless of growth rate.
D. 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.
E. 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

Explanation:
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 83
......

Relate Links

100% Pass Your Professional-Cloud-DevOps-Engineer Exam with Exambible Prep Materials

https://www.exambible.com/Professional-Cloud-DevOps-Engineer-exam/

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - https://www.exambible.com/

Relate Links