Google.Professional-Cloud-DevOps-Engineer.v2022-08-09.q73

Exam Code:	Professional-Cloud-DevOps-Engineer
Exam Name:	Google Cloud Certified - Professional Cloud DevOps Engineer Exam
Certification Provider:	Google
Free Question Number:	73
Version:	v2022-08-09
# of views:	2156
# of Questions views:	730
https://www.freepdfdumps.com/Google.Professional-Cloud-DevOps-Engineer.v2022-08-09.q73.html	

NEW QUESTION: 1
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: C (LEAVE A REPLY)
Explanation/Reference: https://cloud.google.com/compute/docs/faq

NEW QUESTION: 2
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. 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. Focus on responding to internal stakeholders at least every 30 minutes. Commit to "next update" times.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 3
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 (LEAVE A REPLY)
https://cloud.google.com/compute/docs/access/service-accounts#associating_a_service_account_to_an_instance

NEW QUESTION: 4
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 metricKind. set toMETRlc_KIND_UNSPECIFIEDanda valueType set to INT64.
* In Stackdriver's Metrics Explorer, use a Stacked Area graph to visualize the metric.

Answer: C (LEAVE A REPLY)
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: 5
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. Set up the Kubernetes Engine clusters as private clusters.
B. Enable Cloud Security Scanner on the clusters.
C. Set up the Kubernetes Engine clusters with Binary Authorization.
D. Enable Vulnerability Analysis on the Container Registry.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 6
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 bigger a Spinnaker pipeline.
B. Use a custom builder in Cloud Build to trigger Jenkins pipeline.
C. Use Cloud Pub/Sub to trigger a custom deployment service running in Google Kubernetes Engine (GKE).
D. Use Cloud Build to trigger a Spinnaker pipeline.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 7
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 a Google Cloud HTTP Load Balancer as Ingress.
D. Configure your Kubernetes cluster as a Private Cluster.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 8

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: (SHOW ANSWER)
https://cloud.google.com/monitoring/api/metrics_gcp#gcp-appengine

NEW QUESTION: 9
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. Notify the development team that they will have to provide production support for the service.
B. Bring the service into production with no SLOs and build them when you have collected operational data.
C. Identify recommended reliability improvements to the service to be completed before handover.
D. djust the SLO targets to be achievable by the service so you can bring it into production.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 10
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 (LEAVE A REPLY)

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

NEW QUESTION: 11
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 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: 12
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. Redefine the related Service Level Objective so that the error budget is not exhausted.
B. Eliminate unactionable alerts.
C. Distribute the alerts to engineers in different time zones.
D. Create an incident report for each of the alerts.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 13
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. Choose an existing GCP production project to host the monitoring workspace. Attach the production projects to this workspace. Grant relevant team members read access to the Stackdriver Workspace.
B. Grant relevant team members read access to all GCP production projects. Create Stackdriver workspaces inside each project.
C. Grant relevant team members the Project Viewer IAM role on all GCP production projects. Create Slackdriver workspaces inside each project.
D. Create a new GCP monitoring project, and create a Stackdriver Workspace inside it. Attach the production projects to this workspace. Grant relevant team members read access to the Stackdriver Workspace.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 14
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. 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.
B. Publish various metrics from the application directly to the Stackdriver Monitoring API, and then observe these custom metrics in Stackdriver.
C. Install the Cloud Pub/Sub client libraries, push various metrics from the application to various topics, and then observe the aggregated 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: (SHOW ANSWER)

NEW QUESTION: 15
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:5
MTTR: 10
MTBF: 90
Impact 50%

B. MTTD:5
MTTR: 20
MTBF: 90
Impact: 33%

C. MTTD:5
MTTR: 20
MTBF: 90
Impact: 50%

D. MTTD: 5
MTTR: 10
MTBF: 90
Impact: 33%

Answer: (SHOW ANSWER)

NEW QUESTION: 16

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 ms. and then divide by the total number of all web application requests.

Answer: C (LEAVE A REPLY)

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: 17
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 (LEAVE A REPLY)
https://sre.google/sre-book/managing-incidents/

NEW QUESTION: 18
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. 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.
B. Install the Stackdriver Error Reporting library for Python, and then run your code on a Compute Engine VM.

C. Install the Stackdriver Error Reporting library for Python, and then run your code on Google Kubernetes Engine.
D. Install the Stackdriver Error Reporting library for Python, and then run your code on App Engine flexible environment.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 19
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 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.
B. 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.
C. 1. Export VM utilization logs from Stackdriver to BigQuery.
2. Create a dashboard in Data Studio.
3. Share the dashboard with your stakeholders.
D. 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.
Answer: (SHOW ANSWER)

NEW QUESTION: 20
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. Delegate the responding to internal stakeholder emails to another member of the Incident Response Team. Focus on providing responses directly to customers.
B. Focus on responding to internal stakeholders at least every 30 minutes. Commit to "next update" times.
C. Provide all internal stakeholder emails to the Incident Commander, and allow them to manage internal communications. Focus on providing responses directly to customers.

D. Provide periodic updates to all stakeholders in a timely manner. Commit to a "next update" time in all communications.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 21
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 (LEAVE A REPLY)
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: 22
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. Instrumentation coded directly in the client
B. A synthetic client that periodically sends simulated user requests
C. Your application servers' logs
D. Metrics exported from the application servers
E. GKE health checks for your application servers
Answer: A,B (LEAVE A REPLY)

NEW QUESTION: 23
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 (LEAVE A REPLY)
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: 24
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: B (LEAVE A REPLY)

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 Studio. Filler (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 BigOuery as a sink for Stackdriver Logging. Create a scheduled query to filter the cache miss logs and write them to a separate table
Answer: (SHOW ANSWER)
https://cloud.google.com/logging/docs/logs-based-metrics#counter-metric

NEW QUESTION: 26
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 Node taints with NoExecute.
B. Use a partitioned rolling update.
C. Use a stateful set with parallel pod management policy.
D. Use a replica set in the deployment specification.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 27
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 Managed Instance Group with a single instance and use health checks to determine the system status.
B. Create a Stackdriver Monitoring dashboard with SMS alerts to be able to start recreating the crashed instance promptly after it has crashed.
C. File a bug with the development team so they can find the root cause of the crashing instance.

D. Add a Load Balancer in front of the Compute Engine instance and use health checks to determine the system status.
Answer: (SHOW ANSWER)

NEW QUESTION: 28
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: B (LEAVE A REPLY)

NEW QUESTION: 29
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. SSH to the VM and execute the following commands on your VM: ps ax I grep fluentd
C. Install the most recent version of the Stackdriver agent.
D. Verify the VM service account access scope includes the monitoring.write scope.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 30
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 Slackdriver 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)
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.md 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: 31
You need to deploy a new service to production. The service needs to automatically scale using a Managed Instance Group (MIG) and should be deployed over multiple regions. The service needs a large number of resources for each instance and you need to plan for capacity. What should you do?
A. Use the n1-highcpu-96 machine type in the configuration of the MIG.
B. Monitor results of Stackdriver Trace to determine the required amount of resources.
C. Validate that the resource requirements are within the available quota limits of each region.
D. Deploy the service in one region and use a global load balancer to route traffic to this region.
Answer: C (LEAVE A REPLY)
https://cloud.google.com/compute/quotas#understanding_quotas
https://cloud.google.com/compute/quotas

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 (LEAVE A REPLY)
https://cloud.google.com/security-key-management

NEW QUESTION: 33
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. Enable binary authorization inside the Kubernetes cluster and configure the build pipeline as an attestor.
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. Configure the build system with protected branches that require pull request approval.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 34
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 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. Because you are deployed on GKE and are using a cluster autoscaler, your GKE cluster will scale automatically regardless of growth rate.
B. Verify the maximum node pool size, enable a horizontal pod autoscaler, and then perform a load test to verify your expected resource needs.
C. 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.
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.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 35
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 is Container Registry instead of Docker Hub.
C. Run a Git compare between the previous and current Cloud Build Configuration files to find and fix the bug.
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: 36
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. Create a new GCP monitoring project, and create a Stackdriver Workspace inside it. Attach the production projects to this workspace. Grant relevant team members read access to the Stackdriver Workspace.
B. Choose an existing GCP production project to host the monitoring workspace. Attach the production projects to this workspace. Grant relevant team members read access to the Stackdriver Workspace.

C. Grant relevant team members read access to all GCP production projects. Create Stackdriver workspaces inside each project.
D. Grant relevant team members the Project Viewer IAM role on all GCP production projects. Create Slackdriver workspaces inside each project.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 37
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 your account has the proper permissions to use the Stackdriver dashboard.
B. Confirm that port 25 has been opened in the firewall to allow messages through to Stackdriver.
C. Confirm that the application is using the required client library and the service account key has proper permissions.
D. Confirm that the Stackdriver agent has been installed in the hosting virtual machine.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 38
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: (SHOW ANSWER)
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: 39
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. Install the Stackdriver custom metrics adapter and configure a horizontal pod autoscaler to use the number of requests provided by the GCLB.
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. Expose the NGINX stats endpoint and configure the horizontal pod autoscaler to use the request metrics exposed by the NGINX deployment.
Answer: (SHOW ANSWER)

NEW QUESTION: 40
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 Prometheus to collect and aggregate logs per container, and then analyze the results in Grafana.
B. Use Stackdriver Logging to export application logs to BigOuery. aggregate logs per container, and then analyze CPU and memory consumption.
C. Use Stackdriver Kubernetes Engine Monitoring.
D. Use the Stackdriver Monitoring API to create custom metrics, and then organize your containers using groups.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 41
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. Use Stackdriver Logging to create a logs-based metric from the Cloud Buitd logs. Create an Alert with a Webhook notification type.
B. Add a new step at the end of the pipeline in Cloud Build to HTTP POST the build information to a webhook.
C. Create a Cloud Pub/Sub push subscription to the Cloud Build cloud-builds PubSub topic to HTTP POST the build information to a webhook.
D. Add logic to each Cloud Build step to HTTP POST the build information to a webhook.

Answer: B (LEAVE A REPLY)

NEW QUESTION: 42
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 rest. Store the secrets in the bucket and grant Cloud Build access to the bucket.
B. Encrypt the secrets and store them in the application repository. Store a decryption key in a separate repository and grant Cloud Build access to the repository.
C. Use Cloud Key Management Service (Cloud KMS) to encrypt the secrets and include them in your Cloud Build deployment configuration. Grant Cloud Build access to the KeyRing.
D. Use client-side encryption to encrypt the secrets and store them in a Cloud Storage bucket. Store a decryption key in the bucket and grant Cloud Build access to the bucket.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 43
You have migrated an e-commerce application to Google Cloud Platform (GCP). You want to prepare the application for the upcoming busy season. What should you do first to prepare for the busy season?
A. Create a runbook on inflating the disaster recovery (DR) environment if there is growth.
B. Pre-provision double the compute power used last season, expecting growth.
C. Enable AutoScaling on the production clusters, in case there is growth.
D. Load teat the application to profile its performance for scaling.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 44
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. Count the number of home page requests that load in under 100 ms, and then divide by the total number of home page requests.
B. 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.
C. Buckelize Ihe request latencies into ranges, and then compute the percentile at 100 ms.
D. Bucketize the request latencies into ranges, and then compute the median and 90th percentiles.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 45
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 tool. Ensure that your team members add their SMS/phone numbers to the external tool.
C. Ensure that your team members set their SMS/phone numbers in their Stackdriver Profile. Select the SMS notification option for each alerting policy and then select the appropriate SMS/phone numbers from the list.
D. Configure a Slack notification for each alerting policy. Set up a Slack-to-SMS integration to send SMS messages when Slack messages are received. Ensure that your team members add their SMS/phone numbers to the external integration.
Answer: (SHOW ANSWER)
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: 46
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 Private 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 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: D (LEAVE A REPLY)

NEW QUESTION: 47
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 1.0.
B. 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.
C. 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.
D. Enable VPC Flow Logs on the production VPC network frontend and backend subnets only with a sample volume scale of 0.5.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 48
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. Trace the origin of 500 errors and the root cause of increased latency.
C. Roll back the experimental canary release.
D. Start monitoring latency, traffic, errors, and saturation.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 49
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. (lex/jnstance/connections/current
C. tcp_ssl_proxy/new_connections
D. tcp_ssl_proxy/open_connections
Answer: A (LEAVE A REPLY)

NEW QUESTION: 50
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/open_connections
C. flex/instance/connections/current
D. tcp_ssl_proxy/new_connections
Answer: C (LEAVE A REPLY)

NEW QUESTION: 51
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. Enable binary authorization inside the Kubernetes cluster and configure the build pipeline as an attestor.
B. Configure the build system with protected branches that require pull request approval.
C. Use an Admission Controller to verify that incoming requests originate from approved sources.
D. Leverage Kubernetes Role-Based Access Control (RBAC) to restrict access to only approved users.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 52
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: 5
MTTR: 10
MTBF: 90
Impact: 33%
B. MTTD:5
MTTR: 10
MTBF: 90
Impact 50%
C. MTTD:5
MTTR: 20
MTBF: 90
Impact: 50%
D. MTTD:5
MTTR: 20
MTBF: 90
Impact: 33%
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 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 BigOuery. 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: 54
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 Stackdriver Monitoring dashboard with SMS alerts to be able to start recreating the crashed instance promptly after it has crashed.
C. Add a Load Balancer in front of the Compute Engine instance and use health checks to determine the system status.
D. Create a Managed Instance Group with a single instance and use health checks to determine the system status.
Answer: (SHOW ANSWER)

NEW QUESTION: 55
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. 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.
C. Use the fluent-plugin-record-reformer Fluentd output plugin to remove the fields from the log entries in flight.
D. Use the filter-record-transformer Fluentd filter plugin to remove the fields from the log entries in flight.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 56
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: 57
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. Plan individual meetings with all the engineers involved. Determine who approved and pushed the new release to production.
B. Focus on developing new features rather than avoiding the outages from recurring.
C. Focus on identifying the contributing causes of the incident rather than the individual responsible for the cause.
D. Use the Git history to find the related code commit. Prevent the engineer who made that commit from working on production services.
Answer: (SHOW ANSWER)

NEW QUESTION: 58
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. A list of employees responsible for causing the incident
B. An explanation of the root cause of the incident
C. Your opinion of the incident's severity compared to past incidents
D. Copies of the design documents for all the services impacted by the incident
E. A list of action items to prevent a recurrence of the incident
Answer: A,B (LEAVE A REPLY)

NEW QUESTION: 59
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 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 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 (LEAVE A REPLY)

NEW QUESTION: 60
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 - 150ms
95th percentile - 300ms
B. 90th percentile - 250ms
95th percentile - 400ms
C. 90th percentile - 100ms
95th percentile - 250ms
D. 90th percentile - 120ms
95th percentile - 275ms
Answer: (SHOW ANSWER)

NEW QUESTION: 61

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. Store secrets in a separate configuration file on Git. Provide select developers with access to the configuration file.
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. Prompt developers for secrets at build time. Instruct developers to not store secrets at rest.
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: B (LEAVE A REPLY)

NEW QUESTION: 62
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 team. In each project, create a cluster for Development and one for Production. Grant the teams IAM access to their respective clusters.
B. Create one GCP Project per team. In each project, create a cluster with a Kubernetes namespace for Development and one for Production. Grant the teams 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)

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

NEW QUESTION: 63
You created a Stackdriver chart for CPU utilization in a dashboard within your workspace project. You want to share the chart with your Site Reliability Engineering (SRE) team only. You want to ensure you follow the principle of least privilege. What should you do?
A. Share the workspace Project ID with the SRE team. Assign the SRE team the Dashboard Viewer IAM role in the workspace project.
B. Share the workspace Project ID with the SRE team. Assign the SRE team the Monitoring Viewer IAM role in the workspace project.
C. Click "Share chart by URL" and provide the URL to the SRE team. Assign the SRE team the Dashboard Viewer IAM role in the workspace project.
D. Click "Share chart by URL" and provide the URL to the SRE team. Assign the SRE team the Monitoring Viewer IAM role in the workspace project.
Answer: (SHOW ANSWER)

NEW QUESTION: 64
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: 5
MTTR: 10
MTBF: 90
Impact: 33%
B. MTTD:5
MTTR: 20
MTBF: 90
Impact: 33%
C. MTTD:5
MTTR: 10
MTBF: 90

Impact 50%
D. MTTD:5
MTTR: 20
MTBF: 90
Impact: 50%
Answer: B (LEAVE A REPLY)
https://www.atlassian.com/incident-management/kpis/common-metrics
https://linkedin.github.io/school-of-sre/

NEW QUESTION: 65
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. Announce planned downtime to consume more error budget, and ensure that users are not depending on a tighter SLO.
B. Have more frequent or potentially risky application releases.
C. Add more serving capacity to all of your application's zones.
D. Implement and measure additional Service Level Indicators (SLIs) fro the application.
E. Tighten the SLO match the application's observed reliability.
Answer: C,D (LEAVE A REPLY)

NEW QUESTION: 66
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. Roll back the experimental canary release.
B. Record data for the postmortem document of the incident.
C. Start monitoring latency, traffic, errors, and saturation.
D. Trace the origin of 500 errors and the root cause of increased latency.
Answer: (SHOW ANSWER)

NEW QUESTION: 67
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 (LEAVE A REPLY)
https://cloud.google.com/binary-authorization/docs/overview

NEW QUESTION: 68
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. Secure the production environment to ensure that developers can't change it and set up one controlled update per year.
C. Create a development environment for writing code and a test environment for configurations, experiments, and load testing.
D. Create a development environment with smaller server capacity and give access only to developers and testers.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 69
You support a service with a well-defined Service Level Objective (SLO). Over the previous 6 months, your service has consistently met its SLO and customer satisfaction has been consistently high. Most of your service's operations tasks are automated and few repetitive tasks occur frequently. You want to optimize the balance between reliability and deployment velocity while following site reliability engineering best practices. What should you do? (Choose two.)
A. Shift engineering time to other services that need more reliability.
B. Make the service's SLO more strict.
C. Get the product team to prioritize reliability work over new features.
D. Change the implementation of your Service Level Indicators (SLIs) to increase coverage.
E. Increase the service's deployment velocity and/or risk.
Answer: (SHOW ANSWER)

NEW QUESTION: 70

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. Create an Unmanaged Instance Group for the instances used to run the workload.
D. Purchase Committed Use Discounts.
Answer: (SHOW ANSWER)

NEW QUESTION: 71
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 vertical pod autoscaler in GKE and enable the cluster autoscaler to scale the cluster as pods expand.
B. Expose the NGINX stats endpoint and configure the horizontal pod autoscaler to use the request metrics exposed by the NGINX deployment.
C. Install the Stackdriver custom metrics adapter and configure a horizontal pod autoscaler to use the number of requests provided by the GCLB.
D. Configure the horizontal pod autoscaler to use the average response time from the Liveness and Readiness probes.
Answer: (SHOW ANSWER)

NEW QUESTION: 72
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 custom deployment service running in Google Kubernetes Engine (GKE).
B. Use Cloud Pub/Sub to trigger a Spinnaker pipeline.
C. Use a custom builder in Cloud Build to trigger a Jenkins pipeline.
D. Use Cloud Build to trigger a Spinnaker pipeline.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 73
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. 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.
B. Follow up with the employees who reviewed the changes and prescribe practices they should follow in the future.
C. Ensure that test cases that catch errors of this type are run successfully before new software releases.
D. Identify engineers responsible for the incident and escalate to their senior management.
Answer: B (LEAVE A REPLY)