EXAMTOPICS

-Expert Verified, Online, **Free**.



CERTIFICATION TEST

-CertificationTest.net - Cheap & Quality Resources With Best Support

Question #1 *Topic 1*

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 applications. Use Stackdriver Logging to find dependent applications that are performing poorly.

Suggested Answer: *B*
Official Answer: B

Community vote distribution
B (100%)

Manh Highly Voted 4 years, 9 months ago
Answer is B. The keyword is "make several requests to dependent app". So you need trace for it. Cloud Trace Find performance bottlenecks in production. Cloud Profiler Continuous CPU and heap profiling to improve performance and reduce costs.
upvoted 12 times

Charun Highly Voted 5 years, 1 month ago
I have submitted B answer
upvoted 6 times

nhatne Most Recent 7 months, 3 weeks ago
it's B
upvoted 1 times

mississippi__ 1 year, 4 months ago
b IS CORRECT
upvoted 1 times

kirthi0771 1 year, 4 months ago
answer is B
upvoted 1 times

quismorioej 2 years ago
correct
upvoted 1 times

dija123 2 years, 3 months ago
Agree with B
upvoted 1 times

Jason_Cloud_at 2 years, 9 months ago
Right Answer, Cleared the exam yesterday , All the questions were from this study material , Always go with community answers and prepare well.
upvoted 3 times

LowJi 2 years, 11 months ago
B is right answer.
upvoted 1 times

tuanuv1 3 years ago
Answer is B.
upvoted 1 times

Hiren_Meghnathi 3 years ago
B is correct answer.
upvoted 1 times

Pulk 3 years ago
B is the answer

upvoted 1 times

floppino 3 years, 7 months ago

Ans: B Keyword: HTTP requests -> Trace

upvoted 1 times

chelbsik 3 years, 8 months ago

Submitted B on the exam

upvoted 1 times

AzureDP900 3 years, 10 months ago

B is right answer.

upvoted 1 times

GCP72 4 years ago

The answer is B.

upvoted 1 times

Sav94 4 years, 3 months ago

The answer is B.

upvoted 1 times

Question #2 *Topic 1*

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 Monitoring Viewer IAM role in the workspace project.

B. Share the workspace Project ID with the SRE team. Assign the SRE team the Dashboard Viewer IAM role in the workspace project.

C. Click ג€Share chart by URLג€ and provide the URL to the SRE team. Assign the SRE team the Monitoring Viewer IAM role in the workspace project.

D. Click ג€Share chart by URLג€ and provide the URL to the SRE team. Assign the SRE team the Dashboard Viewer IAM role in the workspace project.

Suggested Answer: *A*
Official Answer: A

Community vote distribution



syslog Highly Voted 5 years, 2 months ago
I think it's C, because dashboard viewer "Read-only access to dashboard configurations." SRE team wants to view data, not configurations.
upvoted 18 times

syslog Highly Voted 5 years, 2 months ago
I think it's C, because dashboard viewer "Read-only access to dashboard configurations." SRE team wants to view data, not configurations.correct there is no such role - "dashboard viewer" the correct name is monitoring dashboard configuration viewer (and the permission is - Read-only access to dashboard configurations). so correct answer should be - Cbut it gives access to all list in monitoring of the project.C is correctThere is, but only the single role won't give sufficient access. https://cloud.google.com/monitoring/access-control?hl=de#monitoring.dashboardViewer
upvoted 18 times

akg001 5 years, 2 months ago
correct there is no such role - "dashboard viewer" the correct name is monitoring dashboard configuration viewer (and the permission is - Read-only access to dashboard configurations).

so correct answer should be - C
upvoted 9 times

akg001 5 years, 2 months ago
correct there is no such role - "dashboard viewer" the correct name is monitoring dashboard configuration viewer (and the permission is - Read-only access to dashboard configurations). so correct answer should be - Cbut it gives access to all list in monitoring of the project.C is correctThere is, but only the single role won't give sufficient access. https://cloud.google.com/monitoring/access-control?hl=de#monitoring.dashboardViewer
upvoted 9 times

gcpengineer 5 years ago
but it gives access to all list in monitoring of the project.
upvoted 1 times

AzureDP900 3 years, 9 months ago
C is correct
upvoted 3 times

Raz0r 2 years, 8 months ago
There is, but only the single role won't give sufficient access. https://cloud.google.com/monitoring/access-control?hl=de#monitoring.dashboardViewer
upvoted 2 times

danchoif2 4 years, 11 months ago
I did a test and see that: The valid roles are: Monitoring Viewer and Monitoring Dashboard Configuration Viewer. You can only share chart by URL using Metrics explorer. With only Monitoring Dashboard Configuration Viewer role, user cannot see anything in Monitoring page. I create a custom role from Monitoring Dashboard Configuration Viewer role and add resourcemanager.projects.get permission. Now user can see list of custom

dashboards and the charts in these custom dashboards. User cannot see standard GCP dashboards. User cannot see the the chart in Metrics explorer (using the shared URL). Opening the URL, user will see errors: "Invalid resource type" and "Invalid metric type". So even if I ignore the typo of Dashboard Viewer role in B and D, they are still incorrect answers. So only A and C are valid. But I think C is better because the question is: "You want to share the chart" not the whole dashboard.

upvoted 13 times

Srikantha Most Recent 3 months, 2 weeks ago

Based on Principle of Least Privilege, option C which gives public url is not the best practice

upvoted 1 times

HunkyBunky 11 months ago

Thats defenetly C, becuase there is no Dashboard Viewer iam role at GCP IAM. And sharing chart by URL is more secure than a whole project. GCP ducumentation suggest that: https://cloud.google.com/monitoring/charts/share-dashboards

upvoted 1 times

harsh5kalsait 1 year, 8 months ago

The most restrictive and correct option is B. Sharing the workspace Project ID and assigning the roles/monitoring.dashboardViewer role grants the SRE team only the ability to view dashboards. They won't have access to other Monitoring features or resources within the project. This adheres to the principle of least privilege. Option A is too broad. The roles/monitoring.viewer role grants read access to all monitoring data, not just the specific chart. Options C and D are unsuitable because "Share chart by URL" makes the chart publicly accessible to anyone with the link, violating the requirement of sharing only with the SRE team and least privilege principle. Even with a Viewer role assigned, the URL method bypasses specific IAM controls.

upvoted 1 times

Khumo1 1 year, 9 months ago

100% C

upvoted 1 times

francisco_guerra 1 year, 10 months ago

Ans: C this question is kind of weird because the role is: roles/monitoring.dashboardViewer "Monitoring Dashboard Configuration Viewer" but monitoring viewer its too much access so I tested in console https://cloud.google.com/iam/docs/understanding-roles with Dashboard Viewer I would ´nt be allowed to see the chart URL with monitoring works fine Dashboard Viewer Provides read-only access to get and list information about all monitoring data and configurations.

upvoted 5 times

AzureDP900 1 year, 10 months ago

C There are a number of IAM security roles related to monitoring. The big three are viewer, editor, and admin. To create the monitoring Workspace initially, a user will need the Monitoring Editor or Admin role in the Workspace's host project. The Monitoring Viewer can get read-only access to the Monitoring console and API. The Monitoring Editor has read-write access to the Monitoring console and APIs and can write monitoring data and configurations into the Workspace. And the Monitoring Admin has full access to, and control over, all monitoring resources. Past these big three roles, monitoring roles exist to provide and limit access to alert policies, dashboards, notification channels, service monitoring, and uptime checks.

upvoted 1 times

PhilipKoku 2 years, 1 month ago

D) Dashboard viewer follows the principle of least privileges.

upvoted 1 times

andersonbispos42 2 years, 2 months ago

My vote is C

upvoted 1 times

thewalker 2 years, 2 months ago

Selected Answer: D

Option A grants the SRE team more access than necessary. The Monitoring Viewer role allows them to view all monitoring data in the project, not just the chart you want to share. Option B is a good option, but it is not as secure as option D. The Dashboard Viewer role allows the SRE team to view all dashboards in the project, not just the chart you want to share. Option C is not secure. Anyone who has the URL can view the chart, even if they are not a member of the SRE team. Option D is the most secure option. It allows the SRE team to view the chart without giving them access to any other data in the project. Here are the steps on how to share the chart with the SRE team:

upvoted 2 times

thewalker 2 years, 2 months ago

Option A grants the SRE team more access than necessary. The Monitoring Viewer role allows them to view all monitoring data in the project, not just the chart you want to share. Option B is a good option, but it is not as secure as option D. The Dashboard Viewer role allows the SRE team to view all dashboards in the project, not just the chart you want to share. Option C is not secure. Anyone who has the URL can view the chart, even if they are not a member of the SRE team. Option D is the most secure option. It allows the SRE team to view the chart without giving them access to any other

data in the project. Here are the steps on how to share the chart with the SRE team:Open the Stackdriver dashboard that contains the chart you want to share. Click the Share button in the top right corner of the dashboard. Click Share chart by URL . Copy the URL and share it with the SRE team. In the IAM & Admin section of the Google Cloud Console, navigate to the workspace project. Click on IAM & Admin > IAM . Click Add . In the New members field, enter the email address of the SRE team. In the Select a role dropdown, select Dashboard Viewer . Click Save . The SRE team will now be able to view the chart by clicking on the URL you shared with them. They will not have access to any other data in the project.Additional Considerations It is important to note that the principle of least privilege is not just about security. It is also about efficiency. By giving users only the access they need, you can reduce the risk of errors and make it easier to manage your resources. Here are some additional tips for following the principle of least privilege: Review your IAM roles regularly. Make sure that users only have the access they need. Use groups to manage access. This can make it easier to grant and revoke access to multiple users at once. Use temporary access when possible. This can help to reduce the risk of unauthorized access. By following these tips, you can help to ensure that your

upvoted 2 times

thewalker 2 years, 2 months ago

Open the Stackdriver dashboard that contains the chart you want to share. Click the Share button in the top right corner of the dashboard. Click Share chart by URL . Copy the URL and share it with the SRE team. In the IAM & Admin section of the Google Cloud Console, navigate to the workspace project. Click on IAM & Admin > IAM . Click Add . In the New members field, enter the email address of the SRE team. In the Select a role dropdown, select Dashboard Viewer . Click Save . The SRE team will now be able to view the chart by clicking on the URL you shared with them. They will not have access to any other data in the project.

upvoted 1 times

thewalker 2 years, 2 months ago

Open the Stackdriver dashboard that contains the chart you want to share. Click the Share button in the top right corner of the dashboard. Click Share chart by URL . Copy the URL and share it with the SRE team. In the IAM & Admin section of the Google Cloud Console, navigate to the workspace project. Click on IAM & Admin > IAM . Click Add . In the New members field, enter the email address of the SRE team. In the Select a role dropdown, select Dashboard Viewer . Click Save . The SRE team will now be able to view the chart by clicking on the URL you shared with them. They will not have access to any other data in the project.Additional Considerations It is important to note that the principle of least privilege is not just about security. It is also about efficiency. By giving users only the access they need, you can reduce the risk of errors and make it easier to manage your resources. Here are some additional tips for following the principle of least privilege: Review your IAM roles regularly. Make sure that users only have the access they need. Use groups to manage access. This can make it easier to grant and revoke access to multiple users at once. Use temporary access when possible. This can help to reduce the risk of unauthorized access. By following these tips, you can help to ensure that your Google Cloud resources are secure and efficient.

upvoted 1 times

thewalker 2 years, 2 months ago

Additional Considerations It is important to note that the principle of least privilege is not just about security. It is also about efficiency. By giving users only the access they need, you can reduce the risk of errors and make it easier to manage your resources. Here are some additional tips for following the principle of least privilege: Review your IAM roles regularly. Make sure that users only have the access they need. Use groups to manage access. This can make it easier to grant and revoke access to multiple users at once. Use temporary access when possible. This can help to reduce the risk of unauthorized access. By following these tips, you can help to ensure that your Google Cloud resources are secure and efficient.

upvoted 1 times

jinaldesailive 2 years, 5 months ago

There is no such role as "Dashboard Viewer"

upvoted 1 times

maxdanny 2 years, 9 months ago

The answer is C because : Dashboard viewer role not exists , monitor-viewer only https://cloud.google.com/iam/docs/understanding-roles#monitoring.viewer and Cloud Monitoring allows you to share the URL of the individual Dashboard and not of the entire project ID

upvoted 1 times

Jason_Cloud_at 2 years, 9 months ago

It is the right answer

upvoted 1 times

floppino 3 years, 7 months ago

Ans: C Exam take on 19/12/2022, 50/50 from this dump without buying the full access.

upvoted 2 times

MTS88 3 years, 8 months ago

I think is A, because there isn't any button "Share chart by URL", and least role is Monitoring Viewer (Dashboard viewer can see only dashboard configurations)

upvoted 2 times

chelbsik 3 years, 8 months ago

Selected Answer: A

Decided to go for A on the exam because there is no "Share chart by URL" button, it's just "Share by URL".

upvoted 1 times

chelbsik 3 years, 8 months ago

Decided to go for A on the exam because there is no "Share chart by URL" button, it's just "Share by URL".But "Share by URL" is in chart? I remember it's only logging explorer

upvoted 1 times

MTS88 3 years, 8 months ago

But "Share by URL" is in chart? I remember it's only logging explorer

upvoted 1 times

Question #3 *Topic 1*

Your organization wants to implement Site Reliability Engineering (SRE) culture and principles. Recently, a service that you support had a limited outage. A manager on another team asks you to provide a formal explanation of what happened so they can action remediations. What should you do?

A. Develop a postmortem that includes the root causes, resolution, lessons learned, and a prioritized list of action items. Share it with the manager only.

B. Develop a postmortem that includes the root causes, resolution, lessons learned, and a prioritized list of action items. Share it on the engineering organization's document portal.

C. Develop a postmortem that includes the root causes, resolution, lessons learned, the list of people responsible, and a list of action items for each person. Share it with the manager only.

D. Develop a postmortem that includes the root causes, resolution, lessons learned, the list of people responsible, and a list of action items for each person. Share it on the engineering organization's document portal.

Suggested Answer: *B*
Official Answer: B

Community vote distribution
B (100%)

devopsbatch Highly Voted 4 years, 8 months ago
B it could be based on this In order to maintain a healthy postmortem culture within an organization, it's important to share postmortems as widely as possible.agree with you for B based on blameless postmortem culture. C and D promote blame by adding people responsible.
upvoted 25 times

devopsbatch Highly Voted 4 years, 8 months ago
B it could be based on this In order to maintain a healthy postmortem culture within an organization, it's important to share postmortems as widely as possible.
upvoted 25 times

akg001 4 years, 8 months ago
agree with you for B based on blameless postmortem culture. C and D promote blame by adding people responsible.
upvoted 6 times

Charun 4 years, 7 months ago
I have submitted B answer
upvoted 6 times

andersonbispos42 Most Recent 1 year, 8 months ago
Vote on B
upvoted 1 times

Jason_Cloud_at 2 years, 3 months ago
Right answer
upvoted 1 times

floppino 3 years, 1 month ago
Ans: B But, it doesn't show answers after 40. Can you please send me the questions list somehow?
upvoted 1 times

floppino 3 years, 1 month ago
Selected Answer: B
Ans: B
upvoted 1 times

ChiggaBoy 3 years, 1 month ago
But, it doesn't show answers after 40. Can you please send me the questions list somehow?
upvoted 1 times

chelbsik 3 years, 2 months ago
Submitted B

upvoted 1 times

DoodleDo 3 years, 2 months ago

Ans B. as the application is not writing to STDOUT or STDERR. Source for the answer - https://cloud.google.com/stackdriver/docs/solutions/gke/managing-logs Extract from this source relevant for the answer. GKE's default logging agent provides a managed solution to deploy and manage the agents that send the logs for your clusters to Cloud Logging. Depending on your GKE cluster master version, either fluentd or fluentbit are used to collect logs. Starting from GKE 1.17, logs are collected using a fluentbit-based agent. GKE clusters using versions prior to GKE 1.17 use a fluentd-based agent. If you want to alter the default behavior of the fluentdagents, then you can run a customized fluentd agent or a customized fluentbit agent.

Common use cases include: removing sensitive data from your logs collecting additional logs not written to STDOUT or STDERR using specific performance-related settings customized log formatting

upvoted 1 times

DoodleDo 3 years, 2 months ago

Ans B. as the application is not writing to STDOUT or STDERR. Source for the answer - https://cloud.google.com/stackdriver/docs/solutions/gke/managing-logs Extract from this source relevant for the answer. GKE's default logging agent provides a managed solution to deploy and manage the agents that send the logs for your clusters to Cloud Logging. Depending on your GKE cluster master version, either fluentd or fluentbit are used to collect logs. Starting from GKE 1.17, logs are collected using a fluentbit-based agent. GKE clusters using versions prior to GKE 1.17 use a fluentd-based agent. If you want to alter the default behavior of the fluentdagents, then you can run a customized fluentd agent or a customized fluentbit agent. Common use cases include: removing sensitive data from your logs collecting additional logs not written to STDOUT or STDERR using specific performance-related settings customized log formattingThis response was for Question 4 and not for Question 3. I posted it here by mistake.

upvoted 1 times

DoodleDo 3 years, 2 months ago

This response was for Question 4 and not for Question 3. I posted it here by mistake.

upvoted 1 times

AzureDP900 3 years, 3 months ago

B is right answer

upvoted 1 times

AzureDP900 3 years, 3 months ago

B is looks perfect

upvoted 1 times

amanp 3 years, 6 months ago

B is correct. C and D is wrong as it is against the SRE policy of Blameless postmortem. A is wrong as these actions cannot be shared with just manager

upvoted 3 times

GCP72 3 years, 6 months ago

The answer is B

upvoted 2 times

Ananda 3 years, 9 months ago

Submitted B in the exam. I cleared the exam today

upvoted 3 times

shivanisarthi 3 years, 9 months ago

B should be the right one

upvoted 1 times

lakiluk 3 years, 11 months ago

Selected B

upvoted 2 times

alaahakim 4 years, 2 months ago

Ans: B

upvoted 2 times

Question #4 *Topic 1*

You have a set of applications running on a Google Kubernetes Engine (GKE) cluster, and you are using Stackdriver Kubernetes Engine Monitoring. You are bringing a new containerized application required by your company into production. This application is written by a third party and cannot be modified or reconfigured. The application writes its log information to /var/log/app_messages.log, and you want to send these log entries to Stackdriver Logging. What should you do?

A. Use the default Stackdriver Kubernetes Engine Monitoring agent configuration.

B. Deploy a Fluentd daemonset to GKE. Then create a customized input and output configuration to tail the log file in the application's pods and write to Stackdriver Logging.

C. Install Kubernetes on Google Compute Engine (GCE) and redeploy your applications. Then customize the built-in Stackdriver Logging configuration to tail the log file in the application's pods and write to Stackdriver Logging.

D. Write a script to tail the log file within the pod and write entries to standard output. Run the script as a sidecar container with the application's pod. Configure a shared volume between the containers to allow the script to have read access to /var/log in the application container.

Suggested Answer: *B*

Reference: https://cloud.google.com/solutions/customizing-stackdriver-logs-fluentd

Community vote distribution

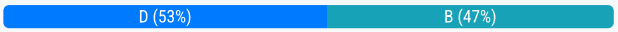

muhasinem Highly Voted 5 years ago

B is correct.

upvoted 6 times

francisco_guerra Highly Voted 1 year, 11 months ago

Ans is B Besides the list of default logs that the Logging agent streams by default, you can customize the Logging agent to send additional logs to Logging or to adjust agent settings by adding input configurations. The configuration definitions in these sections apply to the fluent-plugin-google-cloud output plugin only and specify how logs are transformed and ingested into Cloud Logging.

upvoted 6 times

BlackCoffee Most Recent 1 month, 1 week ago

Why B is the correct approach: By default, Google Kubernetes Engine (GKE) and Stackdriver (Google Cloud Logging) automatically collect everything written to the standard output (stdout) and standard error (stderr) streams of your containers. They do not scan random file paths inside a container like /var/log/app_messages.log. The Sidecar Pattern: Because you cannot modify the third-party application to write to stdout, the standard Kubernetes solution is to use a logging agent sidecar container.

upvoted 1 times

Srikantha 3 months, 2 weeks ago

B is over engineering and D is easy and maintainable

upvoted 1 times

HunkyBunky 11 months ago

I guess that this should be B, because it looks much robust that a D option

upvoted 1 times

knightus 1 year ago

a sidecar container is the cleanest and most GKE-native solution to forward application logs from a file to Stackdriver.

upvoted 1 times

intre 1 year ago

That is the use case for sidecar container

upvoted 1 times

bartrek 1 year, 6 months ago

So in general we have a process running in a pod that opens a local file (inside the pod). Answer B does not mention volume (regardless of a type). As far as i know fluentd is not capable of accessing other pod's files unless hostpath or rwx storage is configured (shared storage). The answer does not mention that however even if this is mentioned this would not be a good choice as this way would easily lead into a log corruption when application is started in more than one replica. You may also ask yourself a question: is it right to deploy a daemonset in order to cover an edge-case?

Is it good to taint the hosts with a hostpath volume? In my opinion D is the least effort solution, utilizes existing mechanisms, does not break scaling nor taints the gke hosts. Application would have to support external log truncation (open the logfile in append mode) in order to prevent infinite file growth. Regardless if you decide to use tail -f or fluentd sidecar this method will prevent log from being corrupted and will play nice with horizontal pod autoscaling. It would be interesting if someone decides to use init container - symlinking logfile to /dev/stdout (or similar) before the main container pod is started.

upvoted 1 times

mohan999 1 year, 9 months ago

I think it should be D, because its says you have a set of applications already running on a Google Kubernetes Engine (GKE) cluster, and you are using Stackdriver Kubernetes Engine Monitoring. I believe we should not enable the workload logging if we were to use the custom agent - So, now just for one new application, it does not make any sense to disable the default stackdriver setting on GKE and use the custom agent for all application logging. So, D would be correct option in this context since we are only customizing this for one application.

upvoted 1 times

WakandaF 1 year, 10 months ago

I believe that is B, This tutorial describes how to customize Fluentd logging for a Google Kubernetes Engine cluster. You'll learn how to host your own configurable Fluentd daemonset to send logs to Cloud Logging, instead of selecting the cloud logging option when creating the Google Kubernetes Engine (GKE) cluster, which does not allow configuration of the Fluentd daemon

upvoted 4 times

DoodleDo 1 year, 10 months ago

Ans B. as the application is not writing to STDOUT or STDERR. Source for the answer - https://cloud.google.com/stackdriver/docs/solutions/gke/managing-logs Extract from this source relevant for the answer. GKE's default logging agent provides a managed solution to deploy and manage the agents that send the logs for your clusters to Cloud Logging. Depending on your GKE cluster master version, either fluentd or fluentbit are used to collect logs. Starting from GKE 1.17, logs are collected using a fluentbit-based agent. GKE clusters using versions prior to GKE 1.17 use a fluentd-based agent. If you want to alter the default behavior of the fluentdagents, then you can run a customized fluentd agent or a customized fluentbit agent. Common use cases include: removing sensitive data from your logs collecting additional logs not written to STDOUT or STDERR using specific performance-related settings customized log formatting

upvoted 1 times

Watcharin_start 1 year, 10 months ago

Ans. is B. This solution is want to send log msg to the Stackdriver logging. However, C is possible IF we have a third party log collector like FluentD that monitoring STDOUT log from the cluster, but this question does not. Therefore, the answer must is B, because any content of container that write in files, we can scrape it in hosted worker node. It store in path /var/containers/* . If we create a customize tail path in the FluentD to monitor any path that store log content, we could use a STAR(*) operator to tell a plugins to find all folders that contain file named app_message.log .

upvoted 1 times

galkin 1 year, 11 months ago

Sidecar is expected here

upvoted 2 times

dija123 2 years, 3 months ago

Deploy a Fluentd daemonset to GKE

upvoted 1 times

Jason_Cloud_at 2 years, 9 months ago

To collect log entries from a specific file within each node in your GKE cluster, you can use a DaemonSet, Fluentd is often used as the logging agent for log forwarding in GKE.

upvoted 2 times

singularis 3 years, 7 months ago

The answer is D. Fluent D cannot read logs from files. https://kubernetes.io/docs/concepts/cluster-administration/logging/#sidecar-container-with-logging-agent By having your sidecar containers write to their own stdout and stderr streams, you can take advantage of the kubelet and the logging agent that already run on each node. The sidecar containers read logs from a file, a socket, or journald. Each sidecar container prints a log to its own stdout or stderr stream.

upvoted 3 times

hanweiCN 3 years, 8 months ago

i will go with B. flentd daemonset cannot collect logs from other pods , it collects logs from host filesystem (stdout, stderr) https://cloud.google.com/architecture/best-practices-for-operating-containers#use_the_native_logging_mechanisms_of_containers

upvoted 1 times

chelbsik 3 years, 8 months ago

Submitted B

upvoted 1 times

Question #5 *Topic 1*

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

A. Look for the agent's test log entry in the Logs Viewer.

B. Install the most recent version of the Stackdriver agent.

C. Verify the VM service account access scope includes the monitoring.write scope.

D. SSH to the VM and execute the following commands on your VM: ps ax | grep fluentd.

Suggested Answer: *D*

Reference: https://groups.google.com/g/google-stackdriver-discussion/c/FXehB9a-5Vk?pli=1

Community vote distribution

D (100%)

raf2121 Highly Voted 4 years, 6 months ago

Between C and D, I think D Reason : When an instance is created, we can specify which service account the instance uses when calling Google Cloud APIs. The instance is automatically configured with access scope and one such access scope is monitoring.write (Link : https://cloud.google.com/compute/docs/access/service- read is to publish metric data and logging.write is to write compute engine logs. Considering above, I believe D as the answer (check whether the agent is running)Agree. If you check here: https://cloud.google.com/logging/docs/agent/logging/troubleshooting#checklist You will see that first recommended troubleshooting step is to check if the agent is running or not... So it should be D. Also if you refer to Google Groups link provided as the answer, you will see that they first checked if the agent is running/installed.

upvoted 12 times

raf2121 Highly Voted 4 years, 6 months ago

Between C and D, I think D

Reason : When an instance is created, we can specify which service account the instance uses when calling Google Cloud APIs. The instance is automatically configured with access scope and one such access scope is monitoring.write (Link : https://cloud.google.com/compute/docs/access/service- read is to publish metric data and logging.write is to write compute engine logs.

Considering above, I believe D as the answer (check whether the agent is running)

upvoted 12 times

cetanx 4 years, 5 months ago

Agree. If you check here: https://cloud.google.com/logging/docs/agent/logging/troubleshooting#checklist You will see that first recommended troubleshooting step is to check if the agent is running or not... So it should be D. Also if you refer to Google Groups link provided as the answer, you will see that they first checked if the agent is running/installed.

upvoted 3 times

desertlotus1211 Most Recent 1 year, 10 months ago

Why nnot Answer A as a starting point?

upvoted 1 times

floppino 3 years, 1 month ago

Ans: D

upvoted 2 times

DoodleDo 3 years, 2 months ago

D. Check if fluentd is running is the right answer. C is incorrect as monitoring.write scope is for monitoring agent and not logging agent.

upvoted 2 times

GCP72 3 years, 6 months ago

Ans: D

upvoted 1 times

Devtestnew 3 years, 8 months ago

Ans: D

upvoted 1 times

Ananda 3 years, 9 months ago

Submitted D

upvoted 1 times

gomezzang 3 years, 10 months ago

answer is D

upvoted 1 times

alexweberlopes 4 years ago

The VM has cloud-platform scope on the VM means Full access if you using the default compute account. I would say answer D is right check is the fluentd daemon running.

upvoted 2 times

pddddd 4 years ago

https://cloud.google.com/logging/docs/agent/logging/troubleshooting#test-agent

upvoted 1 times

alaahakim 4 years, 2 months ago

Ans: D

upvoted 2 times

Wwhite44 4 years, 2 months ago

C is not good, “monitoring.write" is use for metric not log

upvoted 2 times

danchoif2 4 years, 5 months ago

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

When you create an instance using the gcloud command-line tool or the Google Cloud Console, you can specify which service account the instance uses when calling Google Cloud APIs. The instance is automatically configured with the following access scopes:

* ...

So D is better than C.

upvoted 1 times

danchoif2 4 years, 5 months ago

https://cloud.google.com/compute/docs/access/service-accounts#associating_a_service_account_to_an_instance When you create an instance using the gcloud command-line tool or the Google Cloud Console, you can specify which service account the instance uses when calling Google Cloud APIs. The instance is automatically configured with the following access scopes: * ... So D is better than C.i don't think so , reason is option-D is looking for fluent d which is specifically for customized logs , also by default the syslog will be captured with default monitoring agent and no need for fluent-dhttps://cloud.google.com/logging/docs/agent/logging/installation Logging agent is fluentd based (now fluentbit). The Logging agent streams logs from your VM instances and from selected third-party software packages to Cloud Logging. It is a best practice to run the Logging agent on all your VM instances. The VM images for Compute Engine and Amazon Elastic Compute Cloud (EC2) don't include the Logging agent, so you must complete these steps to install it on those instances. The agent runs under both Linux and Windows. If your VMs are running in Google Kubernetes Engine or App Engine, the agent is already included in the VM image, so you can skip this page.

upvoted 1 times

Manumj 4 years, 4 months ago

i don't think so , reason is option-D is looking for fluent d which is specifically for customized logs , also by default the syslog will be captured with default monitoring agent and no need for fluent-d

upvoted 1 times

Manumj 4 years, 4 months ago

i don't think so , reason is option-D is looking for fluent d which is specifically for customized logs , also by default the syslog will be captured with default monitoring agent and no need for fluent-dhttps://cloud.google.com/logging/docs/agent/logging/installation Logging agent is fluentd based (now fluentbit). The Logging agent streams logs from your VM instances and from selected third-party software packages to Cloud Logging. It is a best practice to run the Logging agent on all your VM instances. The VM images for Compute Engine and Amazon Elastic Compute Cloud (EC2) don't include the Logging agent, so you must complete these steps to install it on those instances. The agent runs under both Linux and Windows. If your VMs are running in Google Kubernetes Engine or App Engine, the agent is already included in the VM image, so you can skip this page.

upvoted 1 times

danchoif2 4 years, 3 months ago

https://cloud.google.com/logging/docs/agent/logging/installation Logging agent is fluentd based (now fluentbit). The Logging agent streams logs from your VM instances and from selected third-party software packages to Cloud Logging. It is a best practice to run the Logging agent on all your VM instances. The VM images for Compute Engine and Amazon Elastic Compute Cloud (EC2) don't include the Logging agent, so you must complete these steps to install it on those instances. The agent runs under both Linux and Windows. If your VMs are running in Google Kubernetes Engine or App Engine, the agent is already included in the VM image, so you can skip this page.

upvoted 1 times

guruguru 4 years, 7 months ago

C, first thing to check should be the access scope. There are 3 types of scopes, default, full access and access for each API. Even thought the default scope have the Stackdriver write access, it doesn't mean the instance has the default scope. It could be access for each API. Hence, first thing to check is the scope. After that, you can check the fluentd service in the system. C over D.Monitoring.write is for metric, so C is not correct. In that case logging.write was the right command. So D is the correct ANS ;)

upvoted 2 times

guruguru 4 years, 7 months ago

C, first thing to check should be the access scope. There are 3 types of scopes, default, full access and access for each API. Even thought the default scope have the Stackdriver write access, it doesn't mean the instance has the default scope. It could be access for each API. Hence, first thing to check is the scope. After that, you can check the fluentd service in the system. C over D.

upvoted 2 times

Wwhite44 4 years, 2 months ago

Monitoring.write is for metric, so C is not correct. In that case logging.write was the right command. So D is the correct ANS ;)

upvoted 1 times

TNT87 4 years, 8 months ago

Ans C https://cloud.google.com/logging/docs/agent/logging/troubleshooting

upvoted 2 times

TNT87 4 years, 8 months ago

Ans C https://cloud.google.com/logging/docs/agent/logging/troubleshootingSorry Ans is D

upvoted 2 times

TNT87 4 years, 4 months ago

Sorry Ans is D

upvoted 2 times

francisco_guerra 4 years, 8 months ago

Ans is C Why?: A: Does not make sense at all B: the version it's not possible? C: Its OK because both (scope and service account) need access to write to logging D: Really? the question clearly says: "The image has the Stackdriver Logging agent installed."

Explanation: Generally, you can just set the cloud-platform access scope to allow access to most of the Cloud APIs, then grant the service account only relevant IAM roles. The combination of access scopes granted to the virtual machine instance and the IAM roles granted to the service account determines the amount of access the service account has for that instance. The service account can execute API methods only if they are allowed by both the access scope and its IAM roles.

https://cloud.google.com/compute/docs/access/create-enable-service-accounts-for-instances

Please do not try to confuse people if you put your explanation add the link

upvoted 3 times

francisco_guerra 4 years, 8 months ago

Ans is C Why?: A: Does not make sense at all B: the version it's not possible? C: Its OK because both (scope and service account) need access to write to logging D: Really? the question clearly says: "The image has the Stackdriver Logging agent installed." Explanation: Generally, you can just set the cloud-platform access scope to allow access to most of the Cloud APIs, then grant the service account only relevant IAM roles. The combination of access scopes granted to the virtual machine instance and the IAM roles granted to the service account determines the amount of access the service account has for that instance. The service account can execute API methods only if they are allowed by both the access scope and its IAM roles. https://cloud.google.com/compute/docs/access/create-enable-service-accounts-for-instances Please do not try to confuse people if you put your explanation add the linkI think it's D: https://cloud.google.com/logging/docs/agent/logging/troubleshooting#test-agent. Maybe the agent is installed but not running.I think the answer is D. C couldn't be right because it suggest to verify the monitoring.write scope. But this kind of scope is for monitoring agent not for logging agent. If you see here: https://cloud.google.com/logging/docs/agent/logging/troubleshooting?hl=en they say: If you do not have suitable access scopes in your Compute Engine instance, add the needed access scopes to your instance. The following table shows

the scopes relevant to the Logging and Monitoring agents: Access scope Agent permissions https://www.googleapis.com/auth/logging.write Adequate for the Logging agent https://www.googleapis.com/auth/monitoring.write Adequate for the Monitoring agent what do you think about it?
upvoted 3 times

syslog 4 years, 7 months ago

I think it's D: https://cloud.google.com/logging/docs/agent/logging/troubleshooting#test-agent. Maybe the agent is installed but not running.
upvoted 5 times

burndayl 4 years, 2 months ago

I think the answer is D. C couldn't be right because it suggest to verify the monitoring.write scope. But this kind of scope is for monitoring agent not for logging agent. If you see here: https://cloud.google.com/logging/docs/agent/logging/troubleshooting?hl=en they say: If you do not have suitable access scopes in your Compute Engine instance, add the needed access scopes to your instance. The following table shows the scopes relevant to the Logging and Monitoring agents: Access scope Agent permissions https://www.googleapis.com/auth/logging.write Adequate for the Logging agent https://www.googleapis.com/auth/monitoring.write Adequate for the Monitoring agent what do you think about it?
upvoted 1 times

devopsbatch 4 years, 8 months ago

C because vm need writer access and i feel D as well because fluent need to be available to export ,aggregate and process logs into stack driver monitoring tool... suggestions needed here
upvoted 3 times

devopsbatch 4 years, 8 months ago

C because vm need writer access and i feel D as well because fluent need to be available to export ,aggregate and process logs into stack driver monitoring tool... suggestions needed hereDefault: read-only access to Storage and Service Management, write access to Stackdriver Logging and Monitoring, read/write access to Service Control. So answer should be DCloud-platform scope covers all the GCP service scope so no need to add any other scope. hence answer should be Di think it's D, cause ans C is less precise. on question mention "logging agent" not "monitoring agent" and from here :https://cloud.google.com/logging/docs/agent/logging/troubleshooting#verifying_access_scopes it should be "logging.write" not "monitoring.write"
upvoted 3 times

akg001 4 years, 8 months ago

Default: read-only access to Storage and Service Management, write access to Stackdriver Logging and Monitoring, read/write access to Service Control. So answer should be D
upvoted 1 times

rinkeshgala1 4 years, 8 months ago

Cloud-platform scope covers all the GCP service scope so no need to add any other scope. hence answer should be D
upvoted 1 times

Zuy01 4 years, 7 months ago

i think it's D, cause ans C is less precise. on question mention "logging agent" not "monitoring agent" and from here :https://cloud.google.com/logging/docs/agent/logging/troubleshooting#verifying_access_scopes it should be "logging.write" not "monitoring.write"
upvoted 2 times

Question #6 *Topic 1*

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 Build logs. Create an Alert with a Webhook notification type.

D. Create a Cloud Pub/Sub push subscription to the Cloud Build cloud-builds PubSub topic to HTTP POST the build information to a webhook.

Suggested Answer: *D*

Official Answer: D

Community vote distribution

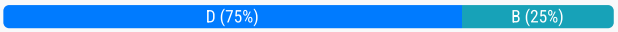

Charun Highly Voted 3 years, 7 months ago

I have submitted D answer

upvoted 17 times

francisco_guerra Highly Voted 3 years, 8 months ago

Ans: D Pub/Sub A: No becauseThere is not Structure attribute to create a http request in the steps and remember you want minimize the development effort. B: The same A C: minimize the development effort D: Its OK To receive messages from push subscriptions, use a webhook and process the POST requests that Pub/Sub sends to the push endpoint. For more information about processing these POST requests in App Engine, see Writing and responding to Pub/Sub messages. https://cloud.google.com/pubsub/docs/push https://cloud.google.com/build/docs/subscribe-build-notifications

upvoted 17 times

Ankitkb2011 Most Recent 1 year, 3 months ago

A is correct because Easy to implement and D is not correct because Scalable but complex

upvoted 1 times

Ankitkb2011 1 year, 3 months ago

A is correct because Easy to implement and not correct because Scalable but complex

upvoted 1 times

floppino 2 years, 1 month ago

Ans: D

upvoted 2 times

GCP72 2 years, 6 months ago

Ans: D

upvoted 1 times

Ananda 2 years, 9 months ago

Submitted D in exam.

upvoted 1 times

akshay_jadhav 2 years, 9 months ago

Cloud Build -> Pubsub -> HTTP Builder.....SO ans is D

upvoted 2 times

pddddd 3 years ago

Not sure on D. Pub/Sub posts into specific format and the actual payload in is message.data. Third party system API (where you post status) will have its own POST format. I feel B is better - you add a step in CLoud Build to do the POST.

upvoted 3 times

alaahakim 3 years, 2 months ago

Ans : D

upvoted 1 times

muk5658 3 years, 2 months ago

Its D. Cloud Build -> Pubsub -> HTTP Builder

upvoted 1 times

rinkeshgala1 3 years, 8 months ago

D is correct answer https://cloud.google.com/build/docs/subscribe-build-notificationsAgree D looks correct to me.

upvoted 2 times

rinkeshgala1 3 years, 8 months ago

D is correct answer https://cloud.google.com/build/docs/subscribe-build-notifications

upvoted 2 times

akg001 3 years, 8 months ago

Agree D looks correct to me.

upvoted 2 times

devopsbatch 3 years, 8 months ago

no idea on this

upvoted 1 times

Question #7 *Topic 1*

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.

Suggested Answer: *D*

Reference: https://cloud.google.com/solutions/automated-canary-analysis-kubernetes-engine-spinnaker

Community vote distribution

A (100%)

rinkeshgala1 Highly Voted 5 years, 2 months ago

i guess A

upvoted 16 times

rinkeshgala1 Highly Voted 5 years, 2 months ago

i guess A A is correct. 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-productionAgree with - A is the correct one. https://cloud.google.com/architecture/automated-canary-analysis-kubernetes-engine-spinnakerCorrect in the link: "You might be tempted to compare the canary deployment against your current production deployment. Instead always compare the canary against an equivalent baseline, deployed at the same time."

upvoted 16 times

DucLee3110 5 years, 2 months ago

A is correct. 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

upvoted 3 times

DucLee3110 5 years, 2 months ago

A is correct. 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-productionAgree with - A is the correct one. https://cloud.google.com/architecture/automated-canary-analysis-kubernetes-engine-spinnaker

upvoted 3 times

akg001 5 years, 2 months ago

Agree with - A is the correct one. https://cloud.google.com/architecture/automated-canary-analysis-kubernetes-engine-spinnaker

upvoted 2 times

giammydell 4 years, 9 months ago

Correct in the link: "You might be tempted to compare the canary deployment against your current production deployment. Instead always compare the canary against an equivalent baseline, deployed at the same time."

upvoted 3 times

francisco_guerra 5 years, 2 months ago

Ans A https://spinnaker.io/guides/user/canary/best-practices/#compare-canary-against-baseline-not-against-production You might be tempted to compare the canary deployment against your current production deployment. Instead always compare the canary against an equivalent baseline, deployed at the same time. The baseline uses the same version and configuration that is currently running in production, but is otherwise identical to the canary: Same time of deployment Same size of deployment Same type and amount of traffic In this way, you control for version and configuration only, and you reduce factors that could affect the analysis, like the cache warmup time, the heap size, and so on.yes correctAgree with A

upvoted 10 times

francisco_guerra 5 years, 2 months ago

Ans A

https://spinnaker.io/guides/user/canary/best-practices/#compare-canary-against-baseline-not-against-production You might be tempted to compare the canary deployment against your current production deployment. Instead always compare the canary against an equivalent baseline, deployed at the same time.

The baseline uses the same version and configuration that is currently running in production, but is otherwise identical to the canary:

Same time of deployment Same size of deployment Same type and amount of traffic In this way, you control for version and configuration only, and you reduce factors that could affect the analysis, like the cache warmup time, the heap size, and so on.

upvoted 10 times

MF2C 4 years, 9 months ago

yes correct

upvoted 1 times

AzureDP900 3 years, 9 months ago

Agree with A

upvoted 1 times

RuchiMishra Most Recent 1 year, 10 months ago

Explanation taking into account the in-memory cache: In-memory Cache Impact: The in-memory cache loading at start time means that both the canary and the current production version could have different data in their caches. This difference could skew the comparison if you compare the canary directly to the existing production version (Option C). Leveling the Playing Field: By deploying a fresh instance of the current production version (Option A), you ensure that both the canary and the comparison baseline have the same starting point in terms of cache data. This makes the comparison more accurate and reliable.

upvoted 1 times

floppino 3 years, 7 months ago

Ans: A

upvoted 1 times

hanweiCN 3 years, 9 months ago

and apply " a new deployment production " for analysis is against canary test purpose.

upvoted 1 times

hanweiCN 3 years, 9 months ago

i think i will go with D the baseline version must have : Ensure that the baseline and production versions of your application are identical. Deploy the baseline version at the same time that you deploy the canary. Ensure that the baseline deployment (such as the number of application instances and autoscaling policies) matches the canary deployment. Use the baseline version to serve the same traffic as the canary. in this case , the " new deployment of production" as a baseline version is violating the third rules ? the new deployment of production is not matching canary deployment

upvoted 1 times

abhishrek 3 years, 9 months ago

A is correct

upvoted 1 times

AzureDP900 3 years, 9 months ago

A is right You might be tempted to compare the canary deployment against your current production deployment. Instead always compare the canary against an equivalent baseline, deployed at the same time. The baseline uses the same version and configuration that is currently running in production, but is otherwise identical to the canary: Same time of deployment Same size of deployment Same type and amount of traffic In this way, you control for version and configuration only, and you reduce factors that could affect the analysis, like the cache warmup time, the heap size, and so on.

upvoted 2 times

Vlad_dryha 3 years, 10 months ago

answer is A

upvoted 1 times

GCP72 4 years ago

Ans : A

upvoted 1 times

Halimb 4 years, 1 month ago

Correct with most comments here. It must be A.

upvoted 1 times

americoleonardo 4 years, 2 months ago

A is ok

upvoted 1 times

gomezzang 4 years, 4 months ago

A is correct!!!

upvoted 1 times

gcpz 4 years, 7 months ago

A is the correct answer

upvoted 1 times

alaahakim 4 years, 8 months ago

Ans : A

upvoted 1 times

lucabinc 4 years, 8 months ago

100% A

upvoted 1 times

Charun 5 years, 1 month ago

I have submitted A

upvoted 6 times

Question #8 *Topic 1*

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

A. Bucketize the request latencies into ranges, and then compute the 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 request that load in under 100 ms, and then divide by the total number of all web application requests.

Suggested Answer: *C*
Reference: https://sre.google/workbook/implementing-slos/

Community vote distribution
C (100%)

Charun Highly Voted 4 years, 7 months ago
I have submitted C answerhttps://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)C is right
upvoted 17 times

Charun Highly Voted 4 years, 7 months ago
I have submitted C answer
upvoted 17 times

raf2121 4 years, 6 months ago
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)
upvoted 8 times

raf2121 4 years, 6 months ago
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)C is right
upvoted 8 times

AzureDP900 3 years, 3 months ago
C is right
upvoted 1 times

danchoif2 4 years, 5 months ago
A, B and D: aren't specific to home page, so they are incorrect. C is correct.
upvoted 9 times

Raz0r Most Recent 2 years, 1 month ago
C is the way to go.
upvoted 1 times

floppino 3 years, 1 month ago
Ans: C Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/
upvoted 1 times

GCP72 3 years, 6 months ago
The correct answer is 'C'
upvoted 3 times

Alpha121 3 years, 8 months ago

C is the answer
upvoted 1 times

alaahakim 4 years, 2 months ago
Ans: C
upvoted 1 times

WakandaF 4 years, 8 months ago
B or C?
upvoted 1 times

TNT87 4 years, 8 months ago
Ans C
upvoted 6 times

francisco_guerra 4 years, 8 months ago
Ans is B After a lot of research i find out this:

The SLI for availability indicates whether the service is working. The SLI for availability is defined as follows: The proportion of valid requests served successfully. (good request / total request)

The SLI for latency (sometimes called speed) indicates whether the service is fast enough. The proportion of valid requests served faster than a threshold. 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. In this case, we consider good events to be events that are faster than this threshold, which was set by examining the historical distribution. You can also set this threshold based on product requirements. You can even set multiple latency SLOs, for example typical latency versus tail latency.
upvoted 4 times

francisco_guerra 4 years, 8 months ago
Ans is B After a lot of research i find out this: The SLI for availability indicates whether the service is working. The SLI for availability is defined as follows: The proportion of valid requests served successfully. (good request / total request) The SLI for latency (sometimes called speed) indicates whether the service is fast enough. The proportion of valid requests served faster than a threshold. 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. In this case, we consider good events to be events that are faster than this threshold, which was set by examining the historical distribution. You can also set this threshold based on product requirements. You can even set multiple latency SLOs, for example typical latency versus tail latency.Still not sure about Band if this is not enough for you check how google calculate Availability and latency Availability sum(rate(http_requests_total{host="api", status!~"5.."}[7d])) / sum(rate(http_requests_total{host="api"}[7d]) Latency histogram_quantile(0.9, rate(http_request_duration_seconds_bucket[7d])) histogram_quantile(0.99, rate(http_request_duration_seconds_bucket[7d])) https://sre.google/workbook/implementing-slos/ https://cloud.google.com/architecture/adopting-slos/The commands shared create a histogram that demonstrates the latency distribution for 90 and 99 % of the requests. Does not define how to measure the SLI.The correct answer is "C", you can confirm it here, https://sre.google/workbook/slo-document/ Latency The proportion of sufficiently fast requests, as measured from the load balancer metrics. “Sufficiently fast” is defined as < 400 ms, or < 850 ms. count of "api" http_requests with a duration less than or equal to "0.4" seconds divided by count of all "api" http_requests
upvoted 4 times

Dark_Lord 4 years, 7 months ago
Still not sure about B
upvoted 1 times

francisco_guerra 4 years, 8 months ago
and if this is not enough for you check how google calculate Availability and latency

Availability

sum(rate(http_requests_total{host="api", status!~"5.."}[7d])) / sum(rate(http_requests_total{host="api"}[7d])

Latency histogram_quantile(0.9, rate(http_request_duration_seconds_bucket[7d])) histogram_quantile(0.99, rate(http_request_duration_seconds_bucket[7d]))

https://sre.google/workbook/implementing-slos/ https://cloud.google.com/architecture/adopting-slos/
upvoted 1 times

francisco_guerra 4 years, 8 months ago

and if this is not enough for you check how google calculate Availability and latency Availability sum(rate(http_requests_total{host="api", status!~"5.."}[7d])) / sum(rate(http_requests_total{host="api"}[7d]) Latency histogram_quantile(0.9, rate(http_request_duration_seconds_bucket[7d])) histogram_quantile(0.99, rate(http_request_duration_seconds_bucket[7d])) https://sre.google/workbook/implementing-slos/ https://cloud.google.com/architecture/adopting-slos/The commands shared create a histogram that demonstrates the latency distribution for 90 and 99 % of the requests. Does not define how to measure the SLI.

upvoted 1 times

pddddd 4 years ago

The commands shared create a histogram that demonstrates the latency distribution for 90 and 99 % of the requests. Does not define how to measure the SLI.

upvoted 2 times

Dark_Lord 4 years, 7 months ago

The correct answer is "C", you can confirm it here, https://sre.google/workbook/slo-document/ Latency The proportion of sufficiently fast requests, as measured from the load balancer metrics. “Sufficiently fast” is defined as < 400 ms, or < 850 ms. count of "api" http_requests with a duration less than or equal to "0.4" seconds divided by count of all "api" http_requests

upvoted 3 times

akg001 4 years, 8 months ago

Answer -C SLI = good events/valid events X 100

upvoted 5 times

akg001 4 years, 8 months ago

Answer -C SLI = good events/good events X 100

upvoted 4 times

devopsbatch 4 years, 8 months ago

answer C

upvoted 5 times

Question #9 *Topic 1*

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

A. Before merging new code, require 2 different peers to review the code changes.

B. Adopt the blue/green deployment strategy when releasing new code via a CD server.

C. Integrate a code linting tool to validate coding standards before any code is accepted into the repository.

D. Require developers to run automated integration tests on their local development environments before release.

E. Configure a CI server. Add a suite of unit tests to your code and have your CI server run them on commit and verify any changes.

Suggested Answer: *A, C*
Official Answer: AC

Community vote distribution
BE (94%)

Charun Highly Voted 4 years, 7 months ago
ans BE
upvoted 31 times

francisco_guerra Highly Voted 4 years, 8 months ago
Ans: B & E A: No, More peers to review dont automate anything B: Ok CD C: No, Linting is for code format D: No, Integration test are needed but its better automatically E: Ok CI CI/CD its OK
upvoted 28 times

edd004 Most Recent 2 years, 2 months ago
E) CI D) CD
upvoted 1 times

theseawillclaim 2 years, 2 months ago
Pretty sure some answers are wrong on purpose. B&E are the only suitable options, as they talk about CI/CD.
upvoted 1 times

jomonkp 2 years, 2 months ago
B and E
upvoted 1 times

maxdanny 2 years, 3 months ago
C because Linting is important to reduce errors and improve the overall quality of your code E for testing the single unit and so reduce inattentive behavior
upvoted 1 times

rseibel 2 years, 5 months ago
B: reduces roll back time, E: reduces failure rate
upvoted 1 times

carloscorreia 2 years, 5 months ago
Ans: B and E A: I understand the importance, but is not it that will avoid errors because we are humans B: Blue/Green will ensure that part of the users were not affected by the bug C: Lint code is just one to validate a code format and indentation D: To avoid errors is important that tests be run in the pipeline E: CI it's a good moment to run all checks
upvoted 2 times

Parteekbhardwaj 3 years, 2 months ago
ANS IS
upvoted 2 times

AzureDP900 3 years, 3 months ago
BE is right

upvoted 1 times

kapara 3 years, 6 months ago

testing comment dont approve please

upvoted 1 times

GCP72 3 years, 6 months ago

Ans BE

upvoted 2 times

Rupo7 3 years, 10 months ago

Definitely B, as this will reduce recovery time as the question asks for. I would think D over E, as the problem states 'not working as expected in production'. Unit tests might not spot an environment specific or application behaviour related issue, whereas automated integration tests are more likely too e.g. might test with DB integration, mocked services.

upvoted 1 times

simbu1299 4 years, 1 month ago

Answer is B & E

upvoted 4 times

dturan 4 years, 2 months ago

ans BE

upvoted 1 times

PJ_2021 4 years, 3 months ago

Answer is B E

upvoted 3 times

Manh 4 years, 3 months ago

BE is good for CI/CD practices

upvoted 1 times

Question #10 *Topic 1*

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. ג€¢ Deploy the Stackdriver logging agent to the application servers. ג€¢ Give the developers the IAM Logs Viewer role to access Stackdriver and view logs.

B. ג€¢ Deploy the Stackdriver logging agent to the application servers. ג€¢ Give the developers the IAM Logs Private Logs Viewer role to access Stackdriver and view logs.

C. ג€¢ Deploy the Stackdriver monitoring agent to the application servers. ג€¢ Give the developers the IAM Monitoring Viewer role to access Stackdriver and view metrics.

D. ג€¢ Install the gsutil command line tool on your application servers. ג€¢ 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. ג€¢ Give the developers the IAM Object Viewer access to view the logs in the specified bucket.

Suggested Answer: *B*
Official Answer: B

Community vote distribution
A (92%) B (8%)

devopsbatch Highly Voted 4 years, 8 months ago
A roles/logging.viewer (Logs Viewer) gives you read-only access to all features of Logging, except Access Transparency logs and Data Access audit logs.
upvoted 27 times

devopsbatch Highly Voted 4 years, 8 months ago
A roles/logging.viewer (Logs Viewer) gives you read-only access to all features of Logging, except Access Transparency logs and Data Access audit logs.correct - A . least privilege principleA is right
upvoted 27 times

akg001 4 years, 8 months ago
correct - A . least privilege principle
upvoted 9 times

AzureDP900 3 years, 3 months ago
A is right
upvoted 2 times

Charun 4 years, 7 months ago
correct A
upvoted 11 times

jinaldesailive 1 year, 11 months ago
A is the correct answer. As B is talking about private logs viewer, there is nothing like that role in GCP.For access to all logs in the _Required and _Default buckets, including data access logs, grant the Private Logs Viewer (roles/logging.privateLogViewer) role.
https://cloud.google.com/logging/docs/access-control#logging.privateLogViewer
upvoted 1 times

jinaldesailive 1 year, 11 months ago
Selected Answer: A
A is the correct answer. As B is talking about private logs viewer, there is nothing like that role in GCP.
upvoted 1 times

habla2019pasta Most Recent 1 year, 9 months ago
For access to all logs in the _Required and _Default buckets, including data access logs, grant the Private Logs Viewer (roles/logging.privateLogViewer) role. https://cloud.google.com/logging/docs/access-control#logging.privateLogViewer
upvoted 1 times

fdsfsdgsdfxcvxcv 2 years, 1 month ago
B, https://cloud.google.com/logging/docs/routing/overview

upvoted 1 times

jomonkp 2 years, 2 months ago

A - correct option

upvoted 1 times

maxdanny 2 years, 3 months ago

the correct answer is A, the privateLogViewer gives extra access to Data Access Logs that's is not required
https://cloud.google.com/logging/docs/view/logs-explorer-interface

upvoted 1 times

mohan999 3 years, 3 months ago

A is correct, Private Logs Viewer gives you extra access to Data access logs and the question was about viewing application logs.

upvoted 2 times

GCP72 3 years, 6 months ago

The correct answer is "A"

upvoted 2 times

mgm7 3 years, 7 months ago

Default tier is premium. There is NO mention of the load balancer being used and there is no default for this.

upvoted 1 times

prasathdv 3 years, 8 months ago

Ans: Option A. :Logs Viewer role. Least config setup (as per question). Option B is incorrect due to additional audit log viewing access which is inappropriate to this question. ref: https://cloud.google.com/logging/docs/access-control

upvoted 1 times

prasathdv 3 years, 8 months ago

Option A (Least config settings). Option B - Private viewer log is for viewing data audit logs. "The Logs Viewer role doesn't let principals read the Data Access audit logs that are in the _Default bucket. To read these Data Access audit logs, principals need the Private Logs Viewer role (roles/logging.privateLogViewer) for the appropriate log view." ref: https://cloud.google.com/logging/docs/access-control

upvoted 2 times

buldas 4 years ago

A. OK B. Logs Private Logs is for Data Logs C. Nope D. what?

upvoted 3 times

PhilipKoku 4 years ago

A) You only need logging.viewer https://cloud.google.com/logging/docs/access-control

upvoted 2 times

roastc 4 years, 1 month ago

Looks like answer A is correct. A logging agent is required to enable the custom logs pushed to Stackdriver https://cloud.google.com/logging/docs/agent/logging . Developers need only Log Viewer permission, which is enough in this case and Private Log viewer is a superset of log viewer permission with elevated permission to view the private data in logs. Which is not needed in this case.

upvoted 4 times

simbu1299 4 years, 1 month ago

Correct Answer is A

upvoted 1 times

scjs 4 years, 1 month ago

B is correct as it talks about application logs https://cloud.google.com/logging/docs/access-control The Logs Viewer role doesn't let you read the Data Access audit logs that are in the _Default bucket roles/logging.privateLogViewer (Private Logs Viewer) includes all the permissions contained by roles/logging.viewer, plus the ability to read Data Access audit logs in the _Defaultand why do you need the Data Access Logs?

upvoted 1 times

scjs 4 years, 1 month ago

B is correct as it talks about application logs https://cloud.google.com/logging/docs/access-control

The Logs Viewer role doesn't let you read the Data Access audit logs that are in the _Default bucket

roles/logging.privateLogViewer (Private Logs Viewer) includes all the permissions contained by roles/logging.viewer, plus the ability to read Data Access audit logs in the _Default

upvoted 1 times

pddddd 4 years ago

and why do you need the Data Access Logs?

upvoted 1 times

alaahakim 4 years, 2 months ago

Ans : A

upvoted 1 times

Question #11 *Topic 1*

You support the backend of a mobile phone game that runs on a Google Kubernetes Engine (GKE) cluster. The application is serving HTTP requests from users.
You need to implement a solution that will reduce the network cost. What should you do?

A. Configure the VPC as a Shared VPC Host project.

B. Configure your network services on the Standard Tier.

C. Configure your Kubernetes cluster as a Private Cluster.

D. Configure a Google Cloud HTTP Load Balancer as Ingress.

Suggested Answer: *C*
Reference: https://cloud.google.com/solutions/prep-kubernetes-engine-for-prod

Community vote distribution
B (64%) D (33%) (3

francisco_guerra Highly Voted 5 years, 2 months ago
Ans is D A: No, Doest make sense B: Who says that we are using a premium tier? C: This does not help with the network cost? D: Ok :) 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-gkeD is rightIt is B because premium tier is by default everywhere. Standard tier is cheaper. Why not D? This is single HTTP application. You don't know how many HTTP endpoints this app expose to Internet. If it is a single endpoint it does not matter if you create Service type Loadbalancer or Cluster IP-> Ingress. You still have one CLB in GCP, so cost is the sameWho says that we aren't use the load balancer? I think it's B anyway, because you have to choose standard tier when "Cost is your main consideration, and you're willing to trade-off some network performance" https://cloud.google.com/network-tiersThe doc didn't says that Container-native load balancing will decrees the spend in network costs. It's only refers to performance. Switching to Standard tier will be the most cost-specific answer
upvoted 28 times

francisco_guerra Highly Voted 5 years, 2 months ago
Ans is D

A: No, Doest make sense B: Who says that we are using a premium tier? C: This does not help with the network cost? D: Ok :) 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
upvoted 28 times

AzureDP900 3 years, 9 months ago
D is right
upvoted 1 times

lxs 4 years, 2 months ago
It is B because premium tier is by default everywhere. Standard tier is cheaper. Why not D? This is single HTTP application. You don't know how many HTTP endpoints this app expose to Internet. If it is a single endpoint it does not matter if you create Service type Loadbalancer or Cluster IP-> Ingress. You still have one CLB in GCP, so cost is the same
upvoted 9 times

syslog 5 years, 1 month ago
Who says that we aren't use the load balancer? I think it's B anyway, because you have to choose standard tier when "Cost is your main consideration, and you're willing to trade-off some network performance" https://cloud.google.com/network-tiers
upvoted 15 times

Meyucho 3 years, 2 months ago

The doc didn't says that Container-native load balancing will decrees the spend in network costs. It's only refers to performance. Switching to Standard tier will be the most cost-specific answer

upvoted 4 times

Charun 5 years, 1 month ago

D 100%

upvoted 9 times

gaurav517 Most Recent 8 months, 3 weeks ago

Directly reduces network costs in the easiest way.

upvoted 1 times

HunkyBunky 11 months ago

I thinks that right answer is B, becuase D makes no sense at from a price reduction costs

upvoted 1 times

luangomesbr 1 year, 2 months ago

Ans shoulb be B

upvoted 1 times

BIGQUERY_ALT_ALT 1 year, 8 months ago

Only B can reduce cost not other options.

upvoted 1 times

mouthwash 1 year, 10 months ago

Should be D. Mobile game apps may not be the right fit for standard tier

upvoted 1 times

6a8c7ad 2 years ago

B reduces cost

upvoted 1 times

quismorioej 2 years ago

answer should be D

upvoted 2 times

dija123 2 years, 3 months ago

Agree with D over B

upvoted 1 times

jinaldesailive 2 years, 5 months ago

I'll go with D. AS HTTP LB will help to reduce the costs. In question it is not mentioned that premium tier is used, so we can't assume that premium tier is used (eventhough it is default network tier) and go with B. So D is the correct answer.

upvoted 2 times

vladik820 2 years, 8 months ago

Answer B

upvoted 3 times

theseawillclaim 2 years, 8 months ago

I go for B, as "Premium" is the default tier. D could be nice, but totally relies on the architecture of the application. If the application needs to be global or re-route lots of requests, a LB makes no miracles.

upvoted 2 times

jomonkp 2 years, 8 months ago

D -configuring the load balancer as ingress

upvoted 2 times

bhunias 2 years, 8 months ago

I am confused between B and D. Default network is Premium which is costly, so changing to Standard also make sense. D is also correct when you'll have multiple endpoints and Google also suggested to LB. But, I'll go with (B) as it'll reduce some cost.

upvoted 1 times

Cassim 3 years, 2 months ago

To reduce network costs for the mobile phone game running on a Google Kubernetes Engine (GKE) cluster, the recommended approach is to use a Google Cloud HTTP Load Balancer as Ingress. This option is the best as it provides several advantages such as scalability, cost-effectiveness, and security.

upvoted 1 times

Question #12 *Topic 1*

You encountered a major service outage that affected all users of the service for multiple hours. After several hours of incident management, the service returned to normal, and user access was restored. You need to provide an incident summary to relevant stakeholders following the Site Reliability Engineering recommended practices. What should you do first?

A. Call individual stakeholders to explain what happened.

B. Develop a post-mortem to be distributed to stakeholders.

C. Send the Incident State Document to all the stakeholders.

D. Require the engineer responsible to write an apology email to all stakeholders.

Suggested Answer: *A*

Official Answer: A

Community vote distribution

B (100%)

devopsbatch Highly Voted 4 years, 8 months ago

B postmortem analysis report to stakeholders

upvoted 25 times

devopsbatch Highly Voted 4 years, 8 months ago

B postmortem analysis report to stakeholderscorrect.

upvoted 25 times

akg001 4 years, 8 months ago

correct.

upvoted 3 times

francisco_guerra 4 years, 8 months ago

A: No B: Correct C: No, Incident State document is used for consultation with incident participants. D: No, Blameless

upvoted 7 times

jinaldesailive Most Recent 1 year, 11 months ago

Obviously B is the answer.

upvoted 1 times

jomonkp 2 years, 2 months ago

option B -Develop a post-mortem to be distributed to stakeholders.

upvoted 1 times

carloscorreia 2 years, 5 months ago

Ans : A I don't know why A is correct Is important to write a postmortem to ensure that the incident is documented

upvoted 2 times

Rog_4444 2 years, 9 months ago

I choose option D. Option B, which suggests configuring network services on the Standard Tier, can help reduce network costs to some extent, but it may not be the most effective solution. The Standard Tier offers lower-cost network services compared to the Premium Tier, which can result in lower network costs. However, the Standard Tier may not offer the same level of functionality and features as the Premium Tier, which can limit your ability to optimize your network for your specific needs. In contrast, using a Google Cloud HTTP Load Balancer as Ingress, as suggested in option D, can provide more fine-grained control over network traffic, which can help you optimize your network usage and reduce network costs. Additionally, a load balancer can distribute incoming traffic across multiple backend instances, which can help reduce the amount of network bandwidth used by each instance.Your suggestion belongs to previous question Sir. In this question it is option B

upvoted 1 times

Rog_4444 2 years, 9 months ago

I choose option D.

Option B, which suggests configuring network services on the Standard Tier, can help reduce network costs to some extent, but it may not be the most effective solution. The Standard Tier offers lower-cost network services compared to the Premium Tier, which can result in lower network costs. However, the Standard Tier may not offer the same level of functionality and features as the Premium Tier, which can limit your ability to optimize

your network for your specific needs.

In contrast, using a Google Cloud HTTP Load Balancer as Ingress, as suggested in option D, can provide more fine-grained control over network traffic, which can help you optimize your network usage and reduce network costs. Additionally, a load balancer can distribute incoming traffic across multiple backend instances, which can help reduce the amount of network bandwidth used by each instance.

upvoted 1 times

Jason_Cloud_at 2 years, 3 months ago

Your suggestion belongs to previous question Sir. In this question it is option B

upvoted 2 times

abhishrek 3 years, 3 months ago

why is A showing as correct ans.

upvoted 4 times

GCP72 3 years, 6 months ago

The correct answer is "C"Plz ignore my previous ans ,the correct answer is "B"

upvoted 2 times

GCP72 3 years, 6 months ago

Selected Answer: B

The correct answer is "C"

upvoted 2 times

GCP72 3 years, 6 months ago

Plz ignore my previous ans ,the correct answer is "B"

upvoted 2 times

skm4690 3 years, 10 months ago

B postmortem report to be share with stakeholders as per SRE.

upvoted 3 times

chidambaram 4 years, 2 months ago

As per SRE culture, postmortem is the correct one.

upvoted 5 times

alaahakim 4 years, 2 months ago

Ans : B

upvoted 2 times

Wwhite44 4 years, 2 months ago

B for sure

upvoted 2 times

edner 4 years, 3 months ago

B, blameless postmortem https://sre.google/sre-book/postmortem-culture/

upvoted 3 times

Charun 4 years, 7 months ago

B correct

upvoted 5 times

Question #13 *Topic 1*

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. Verify the maximum node pool size, enable a horizontal pod autoscaler, and then perform a load test to verify your expected resource needs.

B. Because you are deployed on GKE and are using a cluster autoscaler, your GKE cluster will scale automatically, regardless of growth rate.

C. Because you are at only 30% utilization, you have significant headroom and you won't need to add any additional capacity for this rate of growth.

D. Proactively add 60% more node capacity to account for six months of 10% growth rate, and then perform a load test to make sure you have enough capacity.

Suggested Answer: *B*
Official Answer: B

Community vote distribution

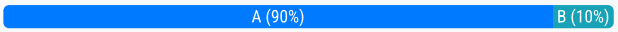

Charun Highly Voted 5 years, 1 month ago
answer A
upvoted 25 times

francisco_guerra Highly Voted 5 years, 2 months ago
A: Correct. 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 B: Incorrect. It is not based on the CPU its based on the workload C: No, Hope is not an strategy D: No, have more resource than needed
upvoted 25 times

francisco_guerra 5 years, 2 months ago
A: Correct. 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 B: Incorrect. It is not based on the CPU its based on the workload C: No, Hope is not an strategy D: No, have more resource than neededAgreed
upvoted 25 times

AzureDP900 3 years, 9 months ago
Agreed
upvoted 1 times

Kedar1234 Most Recent 2 years, 1 month ago
Examtopics suggest Ans B
upvoted 1 times

jomonkp 2 years, 8 months ago
option A
upvoted 1 times

Hiren_Meghnathi 2 years, 9 months ago
Answer A
upvoted 1 times

umesh0u 3 years, 5 months ago
i think A
upvoted 2 times

akarayol 3 years, 7 months ago
A and B confusing
upvoted 1 times

mohan999 3 years, 9 months ago

B is incorrect, because the cluster autoscaler doesn't work based on CPU/Memory usage on the node. The cluster will be scaled on the basis of resources requested by the workloads. So A is more relevant since HPA will automatically scales up/down based on the workload usage.

upvoted 1 times

GCP72 4 years ago

The correct answer is "A"

upvoted 1 times

Ananda 4 years, 3 months ago

Submit the option provide in A

upvoted 1 times

gomezzang 4 years, 4 months ago

A is correct

upvoted 1 times

skm4690 4 years, 4 months ago

A seems to be correct.

upvoted 1 times

Goatee_McGee 4 years, 5 months ago

Answer is A

upvoted 1 times

asdqweasd 4 years, 5 months ago

Even with autoscaling maximul limit effect scaling

upvoted 1 times

burndayl 4 years, 8 months ago

C for me: 30% of total deployed CPU consumption mean 10% of CUP on each zone; 10% of increment month-over-month mean at the beginning of the second mont you will have 11% of CPU, at the tird 12,1% and so on; you can easly accomodate a fault of a zone; you don't have additional cost

upvoted 2 times

burndayl 4 years, 8 months ago

C for me: 30% of total deployed CPU consumption mean 10% of CUP on each zone; 10% of increment month-over-month mean at the beginning of the second mont you will have 11% of CPU, at the tird 12,1% and so on; you can easly accomodate a fault of a zone; you don't have additional cost10% growth is only predicted value, but you should also be prepared for scenarios if the growth at real time is higher. In any case, using HPA and scaling based on it would be an optimal choice in my opinion.

upvoted 2 times

mohan999 3 years, 9 months ago

10% growth is only predicted value, but you should also be prepared for scenarios if the growth at real time is higher. In any case, using HPA and scaling based on it would be an optimal choice in my opinion.

upvoted 2 times

alaahakim 4 years, 8 months ago

Ans: A

upvoted 1 times

Bala999 4 years, 9 months ago

Correct answer A.

upvoted 1 times

Question #14 *Topic 1*

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 Build to trigger a Spinnaker pipeline.

B. Use Cloud Pub/Sub to trigger a Spinnaker pipeline.

C. Use a custom builder in Cloud Build to trigger Jenkins pipeline.

D. Use Cloud Pub/Sub to trigger a custom deployment service running in Google Kubernetes Engine (GKE).

Suggested Answer: *D*
Official Answer: D

Community vote distribution

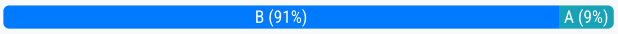

MF2C Highly Voted 4 years, 5 months ago
B is correct : https://cloud.google.com/architecture/continuous-delivery-toolchain-spinnaker-cloud#triggering_a_spinnaker_pipeline_when_a_docker_image_is_pushed_to_container_registryAgree B is right
upvoted 21 times

MF2C Highly Voted 4 years, 5 months ago
B is correct : https://cloud.google.com/architecture/continuous-delivery-toolchain-spinnaker-cloud#triggering_a_spinnaker_pipeline_when_a_docker_image_is_pushed_to_container_registry
upvoted 21 times

AzureDP900 3 years, 3 months ago
Agree B is right
upvoted 1 times

Charun 4 years, 7 months ago
B is correct
upvoted 8 times

manishk39 Most Recent 1 year, 5 months ago
Must be a A. Pub/Sub is for notification, cloud build for trigger the change
upvoted 1 times

jomonkp 2 years, 2 months ago
option B
upvoted 1 times

floppino 3 years, 1 month ago
Ans: B Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/
upvoted 3 times

Trinabh 3 years, 4 months ago
B is correct
upvoted 2 times

GCP72 3 years, 6 months ago
B is correct
upvoted 3 times

rfd 4 years ago
most sensible choice
upvoted 4 times

alaahakim 4 years, 2 months ago
Ans : B
upvoted 1 times

Luvero 4 years, 3 months ago

Its B Google Cloud Build vs Spinnaker: What are the differences? What is Google Cloud Build? Continuously build, test, and deploy. Cloud Build lets you build software quickly across all languages. Get complete control over defining custom workflows for building, testing, and deploying across multiple environments such as VMs, serverless, Kubernetes, or Firebase. What is Spinnaker? Multi-cloud continuous delivery platform for releasing software changes with high velocity and confidence. Created at Netflix, it has been battle-tested in production by hundreds of teams over millions of deployments. It combines a powerful and flexible pipeline management system with integrations to the major cloud providers. Google Cloud Build and Spinnaker belong to "Continuous Deployment" category of the tech stack.

upvoted 7 times

gcpengineer 4 years, 6 months ago

A is the answer

upvoted 2 times

gcpengineer 4 years, 6 months ago

A is the answerCloud Build can't deploy inform Spinnaker to initiate a pipeline as and when the GCR is updated with new image. Whereas CloudPub/Sub can do. Answer should B.

upvoted 2 times

jagadeeshs 4 years ago

Cloud Build can't deploy inform Spinnaker to initiate a pipeline as and when the GCR is updated with new image. Whereas CloudPub/Sub can do. Answer should B.

upvoted 2 times

Dark_Lord 4 years, 7 months ago

B, the pub/sub can trigger any deployment tool(spinnaker, cloud Run etc), but in platform native approach is to use spinnnaker.

You can also validate this, because they have only used platform native tools in the question like google container registry(instead of docker hub)

upvoted 6 times

Dark_Lord 4 years, 7 months ago

B, the pub/sub can trigger any deployment tool(spinnaker, cloud Run etc), but in platform native approach is to use spinnnaker. You can also validate this, because they have only used platform native tools in the question like google container registry(instead of docker hub)no no noeyyy,sorry meant BYes, Spinnaker for Google Cloud can be configured to trigger pipelines from Cloud Build when an image is updated in Google Cloud Registry

upvoted 6 times

TNT87 4 years, 5 months ago

no no no

upvoted 1 times

TNT87 4 years, 5 months ago

no no noeyyy,sorry meant B

upvoted 1 times

TNT87 4 years, 5 months ago

eyyy,sorry meant B

upvoted 2 times

desertlotus1211 1 year, 10 months ago

Yes, Spinnaker for Google Cloud can be configured to trigger pipelines from Cloud Build when an image is updated in Google Cloud Registry

upvoted 1 times

MikeFR 4 years, 7 months ago

OK for B but there is mistake by saying to bigger instead of to trigger

upvoted 2 times

holahola 4 years, 8 months ago

B is correct, because it says that image is built already, then there must be a trigger of some sort for spinnaker. Cloud Build does not trigger Spinnaker itself. Cloud Pub/Sub trigger can be configured in Spinnaker as documented: https://spinnaker.io/guides/user/pipeline/triggers/pubsub/Yes, Spinnaker for Google Cloud can be configured to trigger pipelines from Cloud Build when an image is updated in Google Cloud Registry

upvoted 2 times

holahola 4 years, 8 months ago

B is correct, because it says that image is built already, then there must be a trigger of some sort for spinnaker. Cloud Build does not trigger Spinnaker itself. Cloud Pub/Sub trigger can be configured in Spinnaker as documented: https://spinnaker.io/guides/user/pipeline/triggers/pubsub/

upvoted 2 times

devopsbatch 4 years, 8 months ago

A cloud build spinakker pipeline

upvoted 4 times

devopsbatch 4 years, 8 months ago

A cloud build spinakker pipelineCloud build directly can't trigger spinnaker pipeline. you have to use cloud pub/sub to trigger spinnaker pipeline. (In the backend it will use cloud build notifiers

upvoted 4 times

rinkeshgala1 4 years, 8 months ago

Cloud build directly can't trigger spinnaker pipeline. you have to use cloud pub/sub to trigger spinnaker pipeline. (In the backend it will use cloud build notifiers

upvoted 4 times

Question #15 *Topic 1*

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 (MTTD) 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%

Suggested Answer: *C*

Official Answer: C

Community vote distribution

B (100%)

Charun Highly Voted 4 years, 7 months ago

B 100%

upvoted 18 times

rinkeshgala1 Highly Voted 4 years, 8 months ago

B MTR will increase to 20

upvoted 11 times

alpha_canary Most Recent 2 years ago

Pretty simple maths: MTTD stays 5 as no change in it MTTR doubles to 20 minutes because question says "a new database system that takes twice as long to successfully fail over between zones." MTBF stays the same i.e., once per quarter (90 days) Impact: 33% because if one zone fails, only users in that zone is affeted and we have users spread over 3 zones.

upvoted 6 times

jomonkp 2 years, 2 months ago

option B

upvoted 2 times

Cassim 2 years, 8 months ago

Impact would be 33% since failure is one zone out of three zones (1/3)

upvoted 2 times

jaykumarjkd99 3 years ago

B 100%

upvoted 1 times

taqihurr 3 years, 1 month ago

B 100%

upvoted 1 times

Greg123123 3 years, 1 month ago

I was struggling between B and D as I couldn't find any information that indicate the impact of users. After checking out the discussion here most of people say it's B because there are 3 zones. If this is the case, there is a typo in the question: "Your product is currently deployed in three Google Cloud Platform (GCP) zones with your users divided between "the"(three) zones". So it should be B.

upvoted 4 times

Greg123123 3 years, 1 month ago

I was struggling between B and D as I couldn't find any information that indicate the impact of users. After checking out the discussion here most of people say it's B because there are 3 zones. If this is the case, there is a typo in the question: "Your product is currently deployed in three Google Cloud Platform (GCP) zones with your users divided between "the"(three) zones". So it should be B.You are right. There is a typo in the question. Nowhere does it say three zones.

upvoted 4 times

run4fun 2 years, 8 months ago

You are right. There is a typo in the question. Nowhere does it say three zones.

upvoted 1 times

floppino 3 years, 1 month ago

Ans: B Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/

upvoted 1 times

DoodleDo 3 years, 2 months ago

Can someone give a explanation on why MTTR will be 20 mins?Because "The chat feature requires a new database system that takes twice as long to successfully fail over between zones." which is 2x10min

upvoted 1 times

DoodleDo 3 years, 2 months ago

Can someone give a explanation on why MTTR will be 20 mins?

upvoted 1 times

chelbsik 3 years, 2 months ago

Because "The chat feature requires a new database system that takes twice as long to successfully fail over between zones." which is 2x10min

upvoted 6 times

GCP72 3 years, 6 months ago

B is correct

upvoted 2 times

Bill831231 3 years, 7 months ago

not sure should be "B" or "D", since description like "between Zones" and "fail from one to another" indicate that there are only 2 zones

upvoted 1 times

Ananda 3 years, 9 months ago

Submit B in exam

upvoted 1 times

PhilipKoku 4 years ago

B is the answer: MTTD: 5 minutes MTTR: 20 minutes (twice the time) Impact: 33% (1 of 3 zones impacted)B is right

upvoted 8 times

PhilipKoku 4 years ago

Selected Answer: B

B is the answer: MTTD: 5 minutes MTTR: 20 minutes (twice the time) Impact: 33% (1 of 3 zones impacted)

upvoted 8 times

AzureDP900 3 years, 3 months ago

B is right

upvoted 1 times

alaahakim 4 years, 2 months ago

Ans : B

upvoted 3 times

MF2C 4 years, 5 months ago

How to cal User Impact Percentage?

upvoted 2 times

MF2C 4 years, 5 months ago

How to cal User Impact Percentage?Because of 3 zones, one zone fail is 33%

upvoted 2 times

MF2C 4 years, 3 months ago

Because of 3 zones, one zone fail is 33%

upvoted 3 times

TNT87 4 years, 5 months ago

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

upvoted 3 times

Question #16 *Topic 1*

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.

Suggested Answer: *B*

Reference: https://codelabs.developers.google.com/codelabs/cloud-builder-gke-continuous-deploy/index.html#1

Community vote distribution

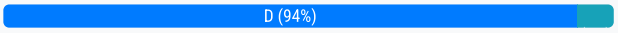

devopsbatch Highly Voted 4 years, 8 months ago

D because binary authorization is deploy time security tool and it will allow only trusted and attested containers into GKE

upvoted 29 times

devopsbatch Highly Voted 4 years, 8 months ago

D because binary authorization is deploy time security tool and it will allow only trusted and attested containers into GKEAgreed with D.

upvoted 29 times

AzureDP900 3 years, 3 months ago

Agreed with D.

upvoted 2 times

Charun 4 years, 7 months ago

D 100%

upvoted 9 times

trashbox Most Recent 1 year, 9 months ago

Exam on 2024-04-26

upvoted 2 times

desertlotus1211 1 year, 10 months ago

Answer is D: The question states: 'only images which are successfully built..' which means Vulnerability Scanns have been completed...

upvoted 2 times

jomonkp 2 years, 2 months ago

option D

upvoted 2 times

carloscorreia 2 years, 5 months ago

To ensure that only images successfully built by your trusted CI/CD pipeline are deployed to production on Google Kubernetes Engine (GKE) clusters, you should set up the Kubernetes Engine clusters with Binary Authorization. Therefore, the correct answer is: D. Set up the Kubernetes Engine clusters with Binary Authorization.

upvoted 2 times

carloscorreia 2 years, 5 months ago

Selected Answer: B

B The question approach is about a trusted image generated, Is possible to create an image binary authorized with vulnerabilities?

upvoted 1 times

carloscorreia 2 years, 5 months ago

B The question approach is about a trusted image generated, Is possible to create an image binary authorized with vulnerabilities?yes it is possible, as long as it is attested by an attestor.

upvoted 1 times

kagami2 2 years, 2 months ago

yes it is possible, as long as it is attested by an attestor.

upvoted 1 times

SarumanMX 2 years, 6 months ago

In another mock test I took this was a select-2 answers, and those were B & D

upvoted 1 times

samuelmorher 2 years, 7 months ago

Who is the one that selects the correct answers?, because it matches 1/9999, looks more like a random.

upvoted 1 times

floppino 3 years, 1 month ago

Ans: D Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/

upvoted 1 times

shivanisarthi 3 years, 2 months ago

D is the right answer. The dump is valid, got all questions from here and cleared the exam

upvoted 2 times

hanweiCN 3 years, 2 months ago

i will go with D, as there is no vulnerability analysis , it is vulnerability scan in container analysis service. and the binary authorization use metadata store to secure trusted repository.

upvoted 2 times

GCP72 3 years, 6 months ago

answer is D

upvoted 1 times

Halimb 3 years, 7 months ago

Must be D.

upvoted 1 times

xtxrtx 3 years, 8 months ago

answer is D

upvoted 1 times

PhilipKoku 4 years ago

D as Binary Authorisation enforces only trusted builds

upvoted 3 times

pondai 4 years, 1 month ago

D https://cloud.google.com/binary-authorization/?hl=zh-tw&utm_source=google&utm_medium=cpc&utm_campaign=japac-TW-all-en-dr-bkws-all-all-trial-e-dr-1009882&utm_content=text-ad-none-none-DEV_c-CRE_284325707741-ADGP_Hybrid%20%7C%20BKWS%20-%20EXA%20%7C%20Txt%20~%20Security%20and%20Identity%20~%20Binary%20Authorization_binary%20authorization-KWID_43700035200283171-kwd-473575914620&userloc_9040379-network_g&utm_term=KW_gcp%20binary%20authorization&gclid=CjwKCAiArOqOBhBmEiwAsgeLmXG4Cu1unuvnKxmO7tAykQtTp-q-jDNRi-mBpNES-S5xNMPb6qSLJBoCTxQQAvD_BwE&gclsrc=aw.ds

upvoted 2 times

Question #17 *Topic 1*

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

A. Use Stackdriver Kubernetes Engine Monitoring.

B. Use Prometheus to collect and aggregate logs per container, and then analyze the results in Grafana.

C. Use the Stackdriver Monitoring API to create custom metrics, and then organize your containers using groups.

D. Use Stackdriver Logging to export application logs to BigQuery, aggregate logs per container, and then analyze CPU and memory consumption.

Suggested Answer: *B*
Official Answer: B

Community vote distribution

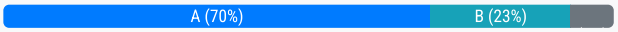

danchoif2 Highly Voted 4 years, 5 months ago

* https://cloud.google.com/anthos/clusters/docs/on-prem GKE on-prem is also called Anthos clusters on VMware

* https://cloud.google.com/anthos/clusters/docs/on-prem/concepts/logging-and-monitoring You have several logging and monitoring options for your Anthos clusters on VMware: + Cloud Logging and Cloud Monitoring, enabled by in-cluster agents deployed with Anthos clusters on VMware. + Prometheus and Grafana, disabled by default. + Validated configurations with third-party solutions.

=> it means, if not a special situation, the correct should be using the first option: Logging and Monitoring. In this case, we want metrics, so Monitoring (aka. Cloud Monitoring, Stackdriver Monitoring) should be used. We are talking about GKE, so we will use Kubernetest Engine Monitoring (https://cloud.google.com/kubernetes-engine-monitoring).

Obviously, A is correct.
upvoted 22 times

danchoif2 Highly Voted 4 years, 5 months ago

* https://cloud.google.com/anthos/clusters/docs/on-prem GKE on-prem is also called Anthos clusters on VMware * https://cloud.google.com/anthos/clusters/docs/on-prem/concepts/logging-and-monitoring You have several logging and monitoring options for your Anthos clusters on VMware: + Cloud Logging and Cloud Monitoring, enabled by in-cluster agents deployed with Anthos clusters on VMware. + Prometheus and Grafana, disabled by default. + Validated configurations with third-party solutions. => it means, if not a special situation, the correct should be using the first option: Logging and Monitoring. In this case, we want metrics, so Monitoring (aka. Cloud Monitoring, Stackdriver Monitoring) should be used. We are talking about GKE, so we will use Kubernetest Engine Monitoring (https://cloud.google.com/kubernetes-engine-monitoring). Obviously, A is correct.we are talking about GKE and GKE on Premise so you need a multi cloud monitoring option not a GKE logging option like AI agree with @helg, if you check the link: https://cloud.google.com/kubernetes-engine/docs/concepts/observability. It says: "Note: The provided GKE dashboards only display information for GKE clusters running on Google Cloud. They don't display information for GKE clusters running anywhere else, for example using on-premises or bare-metal servers."I found another link indicating B is the answer, https://cloud.google.com/anthos/clusters/docs/on-prem/1.3/concepts/logging-and-monitoring I copy a slice: You have several logging and monitoring options for your GKE on-prem clusters: - Cloud Logging and Cloud Monitoring, enabled by in-cluster agents deployed with GKE on-prem. - Prometheus and Grafana, disabled by default. - Validated configurations with third-party solutions.
upvoted 22 times

helg 4 years, 1 month ago

we are talking about GKE and GKE on Premise so you need a multi cloud monitoring option not a GKE logging option like A
upvoted 6 times

helg 4 years, 1 month ago

we are talking about GKE and GKE on Premise so you need a multi cloud monitoring option not a GKE logging option like AI agree with @helg, if you check the link: https://cloud.google.com/kubernetes-engine/docs/concepts/observability. It says: "Note: The provided GKE dashboards only display information for GKE clusters running on Google Cloud. They don't display information for GKE clusters running anywhere else, for example using on-

premises or bare-metal servers."I found another link indicating B is the answer, https://cloud.google.com/anthos/clusters/docs/on-prem/1.3/concepts/logging-and-monitoring I copy a slice: You have several logging and monitoring options for your GKE on-prem clusters: - Cloud Logging and Cloud Monitoring, enabled by in-cluster agents deployed with GKE on-prem. - Prometheus and Grafana, disabled by default. - Validated configurations with third-party solutions.

upvoted 6 times

Feliphus 2 years, 1 month ago

I agree with @helg, if you check the link: https://cloud.google.com/kubernetes-engine/docs/concepts/observability. It says: "Note: The provided GKE dashboards only display information for GKE clusters running on Google Cloud. They don't display information for GKE clusters running anywhere else, for example using on-premises or bare-metal servers."I found another link indicating B is the answer, https://cloud.google.com/anthos/clusters/docs/on-prem/1.3/concepts/logging-and-monitoring I copy a slice: You have several logging and monitoring options for your GKE on-prem clusters: - Cloud Logging and Cloud Monitoring, enabled by in-cluster agents deployed with GKE on-prem. - Prometheus and Grafana, disabled by default. - Validated configurations with third-party solutions.

upvoted 2 times

Feliphus 2 years, 1 month ago

I agree with @helg, if you check the link: https://cloud.google.com/kubernetes-engine/docs/concepts/observability. It says: "Note: The provided GKE dashboards only display information for GKE clusters running on Google Cloud. They don't display information for GKE clusters running anywhere else, for example using on-premises or bare-metal servers."

upvoted 2 times

Feliphus 2 years, 1 month ago

I found another link indicating B is the answer, https://cloud.google.com/anthos/clusters/docs/on-prem/1.3/concepts/logging-and-monitoring I copy a slice: You have several logging and monitoring options for your GKE on-prem clusters: - Cloud Logging and Cloud Monitoring, enabled by in-cluster agents deployed with GKE on-prem. - Prometheus and Grafana, disabled by default. - Validated configurations with third-party solutions.

upvoted 2 times

Charun 4 years, 7 months ago

B correct

upvoted 13 times

Charun 4 years, 7 months ago

B correctPoint for discussion : Google highly recommends Google Logging and monitoring when running workloads only on GKE on-prem and GKE. For applications with component running on GKE on-prem and traditional on-premises infrastructure, other monitoring and logging solutions for an end-to-end view of application can be considered considering what is stated in the question "GKE cluster deployed on-premises and Google Cloud Platform", should "A" be the answernor for A! STackdriver Kub Engine only supports GKE not GKE on premise!

upvoted 13 times

raf2121 4 years, 6 months ago

Point for discussion : Google highly recommends Google Logging and monitoring when running workloads only on GKE on-prem and GKE. For applications with component running on GKE on-prem and traditional on-premises infrastructure, other monitoring and logging solutions for an end-to-end view of application can be considered considering what is stated in the question "GKE cluster deployed on-premises and Google Cloud Platform", should "A" be the answernor for A! STackdriver Kub Engine only supports GKE not GKE on premise!

upvoted 13 times

raf2121 4 years, 6 months ago

Point for discussion : Google highly recommends Google Logging and monitoring when running workloads only on GKE on-prem and GKE. For applications with component running on GKE on-prem and traditional on-premises infrastructure, other monitoring and logging solutions for an end-to-end view of application can be considered

considering what is stated in the question "GKE cluster deployed on-premises and Google Cloud Platform", should "A" be the answer

upvoted 13 times

helg 4 years, 1 month ago

nor for A! STackdriver Kub Engine only supports GKE not GKE on premise!

upvoted 3 times

gaurav517 Most Recent 8 months, 3 weeks ago

GKE on prem also has all the agents required to send metrics/logs to Cloud Monitoring (previously Stackdriver). So as of now, A is the right answer.

upvoted 1 times

jinaldesailive 1 year, 11 months ago

The answer will change pre and post GKE 1.24 release. If the question is asked before the release of GKE 1.24 the answer should be A. Post GKE 1.24 release metrics deprecated, B is the answer.

upvoted 2 times

alpha_canary 2 years ago

If the on-prem was a non-GKE kubernetes solution, the the answer would be B. But it's GKE on-prem, A is the best option here

upvoted 2 times

jomonkp 2 years, 2 months ago

option A

upvoted 1 times

SarumanMX 2 years, 6 months ago

GKE workload monitoring has been deprecated, so should be Prometheus https://cloud.google.com/stackdriver/docs/solutions/gke/managing-metrics

upvoted 2 times

Ajitshetty 2 years, 7 months ago

It should be B check this link It is clearly mentioned if you are application running on GKE or anthos cluster on Vmware you use stack driver counter part but if it is on On-prem then better find other solution With respect to options we have here it is prometheus https://cloud.google.com/anthos/clusters/docs/on-prem/latest/concepts/logging-and-monitoring#:~:text=For%20applications%20with%20components%20running%20on%20Anthos%20clusters%20on%20VMware%20and%20traditional%20on%2Dpremises%20infrastructure%2C%20you%20might%20consider%20other%20solutions%20for%20an%20end%2Dto%2Dend%20view%20of%20those%20applications

upvoted 2 times

felipeschossler 2 years, 10 months ago

Prometheus is made for metrics not logs, in this way the correct answer is A, not be.Its telling to pick which uses most cpu and memory this has nothing to do with logsare these questions still relevant?

upvoted 1 times

felipeschossler 2 years, 10 months ago

Selected Answer: A

Prometheus is made for metrics not logs, in this way the correct answer is A, not be.

upvoted 1 times

sidharthwader 2 years, 7 months ago

Its telling to pick which uses most cpu and memory this has nothing to do with logsare these questions still relevant?

upvoted 1 times

sidharthwader 2 years, 7 months ago

Its telling to pick which uses most cpu and memory this has nothing to do with logs

upvoted 1 times

aswani 2 years, 7 months ago

are these questions still relevant?

upvoted 2 times

taqihurr 3 years ago

B. Use Prometheus to collect and aggregate logs per container, and then analyze the results in Grafana. Prometheus is a popular open-source monitoring and alerting system that is well-suited for monitoring containers running in a Kubernetes cluster. It can scrape metrics from the Kubernetes API server and export them to a time series database, where they can be queried and visualized in Grafana. This approach allows you to monitor CPU and memory usage of individual containers, and set up alerts if usage exceeds certain thresholds.

upvoted 2 times

taqihurr 3 years ago

You support an e-commerce application that runs on a large Google Kubernetes Engine (GKE) cluster deployed on-premises and on Google Cloud Platform. The application consists of microservices that run in containers. You want to identify containers that are using the most CPU and memory. What should you do? A. Use Stackdriver Kubernetes Engine Monitoring. B. Use Prometheus to collect and aggregate logs per container, and then analyze the results in Grafana. C. Use the Stackdriver Monitoring API to create custom metrics, and then organize your containers using groups. D. Use Stackdriver Logging to export application logs to BigQuery, aggregate logs per container, and then analyze CPU and memory consumption.

upvoted 1 times

floppino 3 years, 1 month ago

Ans: A Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/any way to take dump of the discussions?

upvoted 2 times

floppino 3 years, 1 month ago

Selected Answer: A

Ans: A

Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/

upvoted 2 times

jaykumarjkd99 3 years ago

any way to take dump of the discussions?

upvoted 1 times

DoodleDo 3 years, 2 months ago

Ans A - Anthos GKE On-Prem agent collects systems metric however it doesn't collect applications metric. Given that question asks for CPU and Memory GKE On-Prem agent should suffice of containers in on-prem. Source - https://cloud.google.com/anthos/clusters/docs/on-prem/1.3/concepts/logging-and-monitoring

upvoted 1 times

hanweiCN 3 years, 2 months ago

for anthos clusters, need cloud monitoring & cloud logging agent, not GKE monitoring.

upvoted 1 times

hanweiCN 3 years, 2 months ago

The Cloud Operations Suite for GKE only displays information for GKE clusters running on Google Cloud. It does not display information for GKE clusters running elsewhere, such as using on-premises or bare metal servers https://cloud-google-com.translate.goog/stackdriver/docs/solutions/gke?hl=fr&_x_tr_sl=fr&_x_tr_tl=en&_x_tr_hl=en&_x_tr_pto=sc" Cloud Operations Suite for GKE " is a tailored cloud monitoring & logging service for GKE. not cloud monitoring & logging service. A is correct.

upvoted 1 times

hanweiCN 3 years, 2 months ago

The Cloud Operations Suite for GKE only displays information for GKE clusters running on Google Cloud. It does not display information for GKE clusters running elsewhere, such as using on-premises or bare metal servers https://cloud-google-com.translate.goog/stackdriver/docs/solutions/gke?hl=fr&_x_tr_sl=fr&_x_tr_tl=en&_x_tr_hl=en&_x_tr_pto=sc

upvoted 1 times

hanweiCN 3 years, 2 months ago

" Cloud Operations Suite for GKE " is a tailored cloud monitoring & logging service for GKE. not cloud monitoring & logging service. A is correct.

upvoted 1 times

enter_co 3 years, 3 months ago

The eCommerce app is a workload. Since v. 1.24 GKE workload metrics were deprecated and replaced with Prometheus https://cloud.google.com/stackdriver/docs/solutions/gke/workload-metrics#gcloud https://cloud.google.com/stackdriver/docs/solutions/gke/gmp-migration

upvoted 1 times

mudot 3 years, 3 months ago

I work in the hybrid environment [anthos] and can confirm the required metric availability

upvoted 4 times

Question #18 *Topic 1*

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 toad 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.

Suggested Answer: *A*

Official Answer: A

Community vote distribution

D (100%)

devopsbatch Highly Voted 4 years, 2 months ago

D i think creating seperate env for dev and testers

upvoted 23 times

Charun Highly Voted 4 years, 1 month ago

answer is D

upvoted 11 times

jomonkp Most Recent 1 year, 8 months ago

option B

upvoted 3 times

maxdanny 1 year, 9 months ago

It is necessary to check the environment, one of development and one of test

upvoted 2 times

Jason_Cloud_at 1 year, 9 months ago

Right answer

upvoted 1 times

SarumanMX 2 years ago

A makes no sense

upvoted 1 times

Ale_valdivia_vivar 2 years, 4 months ago

i thinks is D because the questions is about how redesign the system...

upvoted 1 times

LEONIDA85 2 years, 7 months ago

why not B?Even if we let alone the env separation in D, B doesn't address all the needs in the question. For example, they want to perform load test, if B has a smaller capacity the load test isn't accurate.Please do know why not B

upvoted 3 times

LEONIDA85 2 years, 7 months ago

why not B?

upvoted 3 times

Greg123123 2 years, 7 months ago

Even if we let alone the env separation in D, B doesn't address all the needs in the question. For example, they want to perform load test, if B has a smaller capacity the load test isn't accurate.

upvoted 1 times

Swepu 2 years, 7 months ago

Please do know why not B

upvoted 1 times

floppino 2 years, 7 months ago

Ans: D Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/

upvoted 2 times

sri1016 2 years, 8 months ago

answer is D

upvoted 1 times

AzureDP900 2 years, 9 months ago

D is right

upvoted 1 times

EricaZhao 2 years, 11 months ago

D is correct

upvoted 1 times

GCP72 3 years ago

answer is D

upvoted 1 times

Ananda 3 years, 3 months ago

Submitted D in the exam

upvoted 1 times

PhilipKoku 3 years, 6 months ago

D is the answer, keep production separated, dev environment to write the code, test environment to test functionality and load before deployments to prod

upvoted 3 times

cyrus86 3 years, 6 months ago

D looks right to me

upvoted 2 times

vijaigcp 3 years, 6 months ago

Creating Separate ENVs for DEV and QA is an age old best practice. Same applies here

upvoted 3 times

Question #19 *Topic 1*

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

Suggested Answer: *D*

Reference: https://cloud.google.com/monitoring/api/metrics_gcp

Community vote distribution

A (100%)

francisco_guerra Highly Voted 4 years, 1 month ago

Ans A A: Metric for App engine & for version B: NO, Metrics for Cloud Load Balancing. C: NO, Metrics for Cloud Load Balancing. D Metric for App engine & for instance An App Engine app is made up of a single application resource that consists of one or more services. Each service can be configured to use different runtimes and to operate with different performance settings. Within each service, you deploy versions of that service. Each version then runs within one or more instances, depending on how much traffic you configured it to handle. if the version runs within one or more instances we need for the version.agree, the question is about number of connections of the app (could be multiple instances), not an instance.

upvoted 23 times

francisco_guerra Highly Voted 4 years, 1 month ago

Ans A A: Metric for App engine & for version B: NO, Metrics for Cloud Load Balancing. C: NO, Metrics for Cloud Load Balancing. D Metric for App engine & for instance

An App Engine app is made up of a single application resource that consists of one or more services. Each service can be configured to use different runtimes and to operate with different performance settings. Within each service, you deploy versions of that service. Each version then runs within one or more instances, depending on how much traffic you configured it to handle.

if the version runs within one or more instances we need for the version.

upvoted 23 times

danchoif2 3 years, 9 months ago

agree, the question is about number of connections of the app (could be multiple instances), not an instance.

upvoted 1 times

Charun 4 years, 1 month ago

A correct

upvoted 6 times

jomonkp Most Recent 1 year, 8 months ago

option A

upvoted 1 times

maxdanny 1 year, 9 months ago

A is correct answer https://cloud.google.com/monitoring/api/metrics_gcp#gcp-appengine flex/connections/current show the current active connections for all instance in appengine

upvoted 1 times

floppino 2 years, 7 months ago

Ans: A Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/

upvoted 1 times

AzureDP900 2 years, 9 months ago

A flex/connections/current GA Connections GAUGE, DOUBLE, 1 gae_app Number of current active connections per App Engine flexible environment version. Sampled every 60 seconds. After sampling, data is not visible for up to 240 seconds.

upvoted 1 times

GCP72 3 years ago

The correct answer is "A"

upvoted 2 times

Ananda 3 years, 3 months ago

Submitted A

upvoted 1 times

PhilipKoku 3 years, 6 months ago

A) Is the correct answer https://cloud.google.com/monitoring/api/metrics_gcp#gcp-appengine

upvoted 3 times

alaahakim 3 years, 8 months ago

Ans : A

upvoted 1 times

TNT87 3 years, 11 months ago

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

upvoted 5 times

ralf_cc 4 years, 2 months ago

A it is

upvoted 5 times

devopsbatch 4 years, 2 months ago

flex/instance/connections/current GA Connections GAUGE, DOUBLE, 1 gae_instance Number of current active connections per App Engine flexible environment instance. Sampled every 60 seconds. After sampling, data is not visible for up to 240 seconds.

upvoted 1 times

devopsbatch 4 years, 2 months ago

flex/instance/connections/current GA Connections GAUGE, DOUBLE, 1 gae_instance Number of current active connections per App Engine flexible environment instance. Sampled every 60 seconds. After sampling, data is not visible for up to 240 seconds.what about option A we need to count application connection count and the count at instance level. no. of instance will keep on changing in app engine in the backend.Agree for the answer -A

upvoted 1 times

rinkeshgala1 4 years, 2 months ago

what about option A we need to count application connection count and the count at instance level. no. of instance will keep on changing in app engine in the backend.Agree for the answer -A

upvoted 2 times

rinkeshgala1 4 years, 2 months ago

what about option A we need to count application connection count and the count at instance level. no. of instance will keep on changing in app engine in the backend.

upvoted 2 times

akg001 4 years, 2 months ago

Agree for the answer -A

upvoted 2 times

Question #20 *Topic 1*

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

A. Check the serial port logs of the Compute Engine instance.

B. Use Stackdriver Profiler to visualize the resources utilization throughout the application.

C. Determine whether there is an increased number of connections to the Cloud SQL instance.

D. Use Cloud Security Scanner to see whether your Cloud SQL is under a Distributed Denial of Service (DDoS) attack.

Suggested Answer: *C*
Official Answer: C

Community vote distribution
B (70%) C (30%)

francisco_guerra Highly Voted 5 years, 1 month ago
Ans: B I don't find anything about ddos attacks and Cloud security Scanner & databases; its most used with App engine and compute so.

I go with Stackdriver profiler
upvoted 24 times

francisco_guerra Highly Voted 5 years, 1 month ago
Ans: B I don't find anything about ddos attacks and Cloud security Scanner & databases; its most used with App engine and compute so. I go with Stackdriver profilerAgree. Ans B
upvoted 24 times

Manh 4 years, 9 months ago
Agree. Ans B
upvoted 3 times

ralf_cc 5 years, 2 months ago
B - getting to know the app perf
upvoted 10 times

Srikantha Most Recent 3 months, 3 weeks ago
In this scenario, where an application update has triggered database timeouts despite a stable user count, the most probable cause is a change in the application's code that affects how it interacts with the database. The correct action is: C. Determine whether there is an increased number of connections to the Cloud SQL instance.
upvoted 1 times

gaurav517 8 months, 3 weeks ago
Most probable cause: Application update introduced connection leak. A performance issue on application is not likely to cause database timeout issues.
upvoted 1 times

HunkyBunky 10 months, 4 weeks ago
In my opinion - C looks better that application utilization
upvoted 1 times

mouthwash 1 year, 10 months ago
Can't be B as I wonder what use resource utilization metrics are going to be. Going with C
upvoted 1 times

winston9 2 years, 2 months ago
Cannot be B, Profiler, it's used to check how an app consumes RAM and CPU, not but database connections
upvoted 2 times

jomonkp 2 years, 8 months ago
option C
upvoted 2 times

alve 3 years, 5 months ago

Ans: B I don't find anything about ddos attacks and Cloud security Scanner & databases; its most used with App engine and compute so.

upvoted 1 times

izekc 3 years, 5 months ago

C would be correct

upvoted 2 times

07397 3 years, 5 months ago

I would go for B: As it says, "The **number** of concurrent active users remained stable." Since users report database timeout errors, the newly deployed version can only service a particular load and that load is **lower than** expected number. It is likely that the new version has changes that lead to high CPU utilization or any other bottleneck, so we should use the profiler to locate the issue.

upvoted 4 times

Greg123123 3 years, 7 months ago

I would go for C. C. Determine whether there is an increased number of connections to the Cloud SQL instance. 1. Because users report the error message saying "database timeout". If the errors is from the application then it wouldn't be "database timeout", so rather than running Profiler in your application, go check whether there is an increased number of connections is more intuitive. 2. Other than that, the question mentions "After an update to the application" which imply that there could be a bug regarding to SQL so that's why if "concurrent active users remained stable", the database could still timeout because of bugs. Please be aware that "concurrent active users remained stable" doesn't mean qps (query-per-second) is still the same.

upvoted 8 times

floppino 3 years, 7 months ago

Ans: B Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/

upvoted 2 times

GCP72 4 years ago

Ans: B

upvoted 2 times

Ananda 4 years, 3 months ago

Submit B

upvoted 2 times

paul223344 4 years, 4 months ago

Answer B: Use Stackdriver Profiler to visualize the resources utilization throughout the application. High CPU usage can most definitely cause dropped or ignored connections. The database engine and underlying OS are fighting for resources and aren't able to respond to the connection in time. Finding out why the query is eating so much CPU usage and optimizing it. https://stackoverflow.com/questions/69919454/high-cpu-usage-on-cloud-sql-causing-timeouts # Cloud Profiler is a statistical, low-overhead profiler that continuously gathers CPU usage and memory-allocation information (supported profile types: CPU time, Heap, Allocated heap, Contention, Threads, Wall time) from your production applications. It attributes that information to the source code that generated it, helping you identify the parts of your application that are consuming the most resources, and otherwise illuminating your applications performance characteristics. https://cloud.google.com/profiler/docs/about-profiler

upvoted 5 times

sankofa 4 years, 6 months ago

C as app update could have introduced a bug where although app user concurrency remain stable, could have introduced more db connections

upvoted 3 times

Question #21 *Topic 1*

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.

Suggested Answer: *C*

Official Answer: C

Community vote distribution

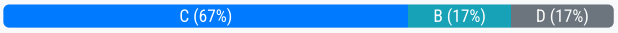

francisco_guerra Highly Voted 4 years, 7 months ago

Ans C Cloud Build provides the following default substitutions: $TAG_NAME: build.Source.RepoSource.Revision.TagNameOK but doesn't it make more sense to use this $TAG to include the release version tag in the application image? I would go with B.Sorry, I meant "to use this $TAG as a parameter within the image name?"Agree. Cloud Build provides a variable for the tag, when we push the image, we must supply the variable to tag the image.

upvoted 23 times

francisco_guerra Highly Voted 4 years, 7 months ago

Ans C Cloud Build provides the following default substitutions: $TAG_NAME: build.Source.RepoSource.Revision.TagName

upvoted 23 times

cetanx 4 years, 5 months ago

OK but doesn't it make more sense to use this $TAG to include the release version tag in the application image?

I would go with B.

upvoted 5 times

cetanx 4 years, 5 months ago

OK but doesn't it make more sense to use this $TAG to include the release version tag in the application image? I would go with B.Sorry, I meant "to use this $TAG as a parameter within the image name?"Agree. Cloud Build provides a variable for the tag, when we push the image, we must supply the variable to tag the image.

upvoted 5 times

cetanx 4 years, 5 months ago

Sorry, I meant "to use this $TAG as a parameter within the image name?"

upvoted 2 times

danchoif2 4 years, 3 months ago

Agree. Cloud Build provides a variable for the tag, when we push the image, we must supply the variable to tag the image.

upvoted 2 times

devopsbatch 4 years, 8 months ago

C inside cloudbuil.yml file can include release version

upvoted 5 times

Srikantha Most Recent 3 months, 3 weeks ago

he correct approach to ensure your deployments are tied to specific releases is: B. Supply the source control tag as a parameter within the image name. Why this is the standard practice In Docker and Google Container Registry (GCR/Artifact Registry), the "image name" includes the tag (e.g., gcr.io/project-id/app:v1.2.3). By tagging the image with the same version string used in your source control (like Git), you create a direct, human-readable link between your code and your artifact.

upvoted 1 times

leeban99 1 year, 7 months ago

C seems ok

upvoted 1 times

BleHi 1 year, 11 months ago

imho D cause Using GCR digest versioning involves using the SHA256 digest of the image as the unique identifier for the image. This ensures that the image can be reliably referenced and identified regardless of any tags that may be applied to it.

upvoted 1 times

jomonkp 2 years, 2 months ago

option C

upvoted 1 times

floppino 3 years, 1 month ago

Ans: C Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/

upvoted 1 times

GCP72 3 years, 6 months ago

The correct answer is C

upvoted 2 times

nuthanj800 3 years, 8 months ago

C is a right option

upvoted 2 times

meetplanet 4 years, 2 months ago

I go for B https://cloud.google.com/container-registry/docs/pushing-and-pulling

upvoted 2 times

alaahakim 4 years, 2 months ago

Ans: C

upvoted 3 times

guruguru 4 years, 7 months ago

C. tag your images with the sourced code release version.

upvoted 3 times

JuanitoNN 4 years, 7 months ago

I think is B, is the normal process in this caseshttps://cloud.google.com/container-registry/docs/pushing-and-pulling

upvoted 1 times

JuanitoNN 4 years, 7 months ago

I think is B, is the normal process in this cases

upvoted 1 times

JuanitoNN 4 years, 7 months ago

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

upvoted 1 times

Question #22 *Topic 1*

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 (OL) 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.

Suggested Answer: *C*
Official Answer: C

Community vote distribution



Charun Highly Voted 4 years, 7 months ago
C is the answer
upvoted 15 times

francisco_guerra Highly Voted 4 years, 7 months ago
Ans: C Prepare Beforehand In addition to incident response training, it helps to prepare for an incident beforehand. Use the following tips and strategies to be better prepared. Decide on a communication channel Decide and agree on a communication channel (Slack, a phone bridge, IRC, HipChat, etc.) beforehand. Keep your audience informed Unless you acknowledge that an incident is happening and actively being addressed, people will automatically assume nothing is being done to resolve the issue. Similarly, if you forget to call off the response once the issue has been mitigated or resolved, people will assume the incident is ongoing. You can preempt this dynamic by keeping your audience informed throughout the incident with regular status updates. Having a prepared list of contacts (see the next tip) saves valuable time and ensures you don't miss anyone. https://sre.google/workbook/incident-response/Agreed, C is right.
upvoted 14 times

francisco_guerra 4 years, 7 months ago
Ans: C Prepare Beforehand In addition to incident response training, it helps to prepare for an incident beforehand. Use the following tips and strategies to be better prepared.

Decide on a communication channel Decide and agree on a communication channel (Slack, a phone bridge, IRC, HipChat, etc.) beforehand.

Keep your audience informed Unless you acknowledge that an incident is happening and actively being addressed, people will automatically assume nothing is being done to resolve the issue. Similarly, if you forget to call off the response once the issue has been mitigated or resolved, people will assume the incident is ongoing. You can preempt this dynamic by keeping your audience informed throughout the incident with regular status updates. Having a prepared list of contacts (see the next tip) saves valuable time and ensures you don't miss anyone.

https://sre.google/workbook/incident-response/
upvoted 14 times

AzureDP900 3 years, 3 months ago
Agreed, C is right.
upvoted 2 times

jomonkp 2 years, 2 months ago
Option C
upvoted 1 times

floppino 3 years, 1 month ago
Ans: C Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/

upvoted 1 times

moitsu 3 years, 3 months ago

"The Ops lead works with the incident commander to respond to the incident by applying operational tools to the task at hand. The operations team should be the 'only' group modifying the system during an incident." Has to be C

upvoted 1 times

AzureDP900 3 years, 3 months ago

based on my experience with daily incidents it should be C followed by A

upvoted 1 times

atkhan 3 years, 4 months ago

C is correct answer

upvoted 1 times

EricaZhao 3 years, 5 months ago

Correct answer is A

upvoted 2 times

EricaZhao 3 years, 5 months ago

the correct Answer is A, i took a google training of this professional DevOps Engineer, the teacher said C communication channel is important, but the most of import thing is A. Look for ways to mitigate user impact and deploy the mitigations to production. solve the issue and reduce the impact of the users.

upvoted 2 times

GCP72 3 years, 6 months ago

The correct answer is C

upvoted 1 times

gomezzang 3 years, 10 months ago

You are IC. https://sre.google/sre-book/managing-incidents/

upvoted 3 times

ric79 3 years, 11 months ago

In this case there are 3 different peoples. The mitigation is for the OL. The IC must coordinate and setup a living Incident State Document which is the the communication channel for the operational group

upvoted 1 times

zygomar 3 years, 12 months ago

agree with rinkeshgala1: "mitigating the issue is not the job of Incident commander but of Operation Lead. Incident commander should setup common communication channel , so option C"

upvoted 4 times

buldas 4 years ago

Selected Answer: C

Ans A is a task for Operation Lead

upvoted 3 times

buldas 4 years ago

Ans A is a task for Operation LeadI Agree. This statement is asking the roles of an incident team, not which it is role more important, which it would be recovering the service (A) asap

upvoted 3 times

Feliphus Most Recent 2 years, 2 months ago

I Agree. This statement is asking the roles of an incident team, not which it is role more important, which it would be recovering the service (A) asap

upvoted 1 times

PhilipKoku 4 years ago

A - Stop the impact to Production and end users ASAP.

upvoted 2 times

pondai 4 years, 1 month ago

I will pick C https://sre.google/sre-book/managing-incidents/ In Unmanaged Incidents exmple: on-call engineer Mary try to fix problem, but the technical task at hand was overwhelming. Poor Communication:other engineers who could have lent a hand in debugging or fixing the issue weren't used effectively

upvoted 2 times

Wwhite44 4 years, 2 months ago

Ans A is not, cause it means:"man solve the problem", it is not your task solve it, you are IC not technician, C for sure

upvoted 4 times

Question #23 *Topic 1*

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 projects. Create Stackdriver workspaces inside each project.

B. Grant relevant team members the Project Viewer IAM role on all GCP production projects. Create Stackdriver workspaces inside each project.

C. 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.

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.

Suggested Answer: *C*
Official Answer: C

Community vote distribution
D (100%)

Charun Highly Voted 4 years, 7 months ago
D 100%
upvoted 25 times

vijaigcp Highly Voted 4 years ago
Answer - D When you want to manage metrics for multiple projects, we recommend that you create a project to be the scoping project for that metrics scope. https://cloud.google.com/monitoring/settings/multiple-projectsDefinitely D is right without any second thought.
upvoted 13 times

vijaigcp 4 years ago
Selected Answer: D
Answer - D When you want to manage metrics for multiple projects, we recommend that you create a project to be the scoping project for that metrics scope. https://cloud.google.com/monitoring/settings/multiple-projects
upvoted 13 times

AzureDP900 3 years, 3 months ago
Definitely D is right without any second thought.
upvoted 3 times

habla2019pasta Most Recent 1 year, 9 months ago
D because the recommendation is that you create a separate project for monitoring your GCP projects.
upvoted 1 times

jomonkp 2 years, 2 months ago
option D
upvoted 2 times

feholen210 2 years, 10 months ago
D for sure
upvoted 1 times

singularis 3 years, 1 month ago
D 100%
upvoted 1 times

floppino 3 years, 1 month ago
Ans: D Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/
upvoted 2 times

atkhan 3 years, 4 months ago

D is correct answer

upvoted 1 times

GCP72 3 years, 6 months ago

D for sure

upvoted 1 times

Bill831231 3 years, 7 months ago

sounds like D is the best

upvoted 1 times

Ananda 3 years, 9 months ago

Submitted D

upvoted 1 times

helg 4 years, 1 month ago

D- new project with a stackdriver workspace (clean) + least privilege when giving read access to relevant members ..

upvoted 1 times

ralf_cc 4 years, 8 months ago

D it is

upvoted 7 times

devopsbatch 4 years, 8 months ago

C answerI think D because when I choose to add other project to mine, without create a new project, Google tells me: "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."

upvoted 2 times

devopsbatch 4 years, 8 months ago

C answer

upvoted 2 times

syslog 4 years, 8 months ago

I think D because when I choose to add other project to mine, without create a new project, Google tells me: "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."

upvoted 9 times

Question #24 *Topic 1*

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. 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 from Stackdriver to BigQuery. 2. From BigQuery, export the logs to a CSV file. 3. Import the CSV file into Google Sheets. 4. Build a dashboard in Google Sheets and share it with your stakeholders.

D. 1. Export VM utilization logs from Stackdriver to 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.

Suggested Answer: *A*
Official Answer: A

Community vote distribution
A (100%)

Charun Highly Voted 4 years, 7 months ago
A is correct
upvoted 22 times

kai324 Highly Voted 4 years, 6 months ago
A for sure
upvoted 12 times

habla2019pasta Most Recent 1 year, 9 months ago
A for sure
upvoted 1 times

jomonkp 2 years, 2 months ago
option A
upvoted 3 times

singularis 3 years, 1 month ago
A for sure
upvoted 2 times

floppino 3 years, 1 month ago
Ans: A Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/
upvoted 4 times

atkhan 3 years, 4 months ago
A is correct because google cloud studio is build to provide real time metric data and it can be directly integrated with BigQuery
upvoted 3 times

atkhan 3 years, 4 months ago
A is correct because google cloud studio is build to provide real time metric data and it can be directly integrated with BigQueryAgreed, A is right
upvoted 3 times

AzureDP900 3 years, 3 months ago
Agreed, A is right
upvoted 2 times

GCP72 3 years, 6 months ago
A is correct
upvoted 1 times

Ananda 3 years, 9 months ago

A is correct

upvoted 1 times

vijaigcp 4 years ago

Answer - A B & C are ruled out straight away. Between A & D, as the ask is real time, D can be ruled out.
https://cloud.google.com/logging/docs/export/configure_export_v2

upvoted 3 times

helg 4 years, 1 month ago

A Stackdriver -- BQ -- DataStudio

upvoted 3 times

alaahakim 4 years, 2 months ago

Ans: A

upvoted 2 times

ralf_cc 4 years, 8 months ago

A - no doubt

upvoted 4 times

devopsbatch 4 years, 8 months ago

answer B. because SEIM helps realtime analysis of security and also export logs into pubsub is realtime transactions.In the question it is asking for VM utilisation not the security events. so answer is AA is correct. It mentioned clearly in the questions: "You want to use Google Cloud Platform solutions"

upvoted 3 times

devopsbatch 4 years, 8 months ago

answer B. because SEIM helps realtime analysis of security and also export logs into pubsub is realtime transactions.

upvoted 3 times

akg001 4 years, 8 months ago

In the question it is asking for VM utilisation not the security events. so answer is A

upvoted 2 times

DucLee3110 4 years, 8 months ago

A is correct. It mentioned clearly in the questions: "You want to use Google Cloud Platform solutions"

upvoted 8 times

Question #25 *Topic 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.

Suggested Answer: *C*

Reference: https://cloud.google.com/compute/docs/faq

Community vote distribution

A (100%)

Charun Highly Voted 4 years, 7 months ago

A is correct

upvoted 26 times

habla2019pasta Most Recent 1 year, 9 months ago

A is correct.

upvoted 1 times

jomonkp 2 years, 2 months ago

Option A

upvoted 1 times

foreigncar99 2 years, 5 months ago

Answer is A

upvoted 1 times

floppino 3 years, 1 month ago

Ans: A Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/

upvoted 1 times

GCP72 3 years, 6 months ago

Answer is A

upvoted 2 times

ghadxx 3 years, 9 months ago

C. Since the requirement is to run “business-critical workloads”, preemptible instances not ideal since they can be stopped randomly.

upvoted 3 times

PhilipKoku 4 years ago

A - Committed usage discounts is the right answer as we have a consistent workload for a known duration.

upvoted 2 times

vijaigcp 4 years ago

A is correct, committed usage discounts

upvoted 2 times

cyrus86 4 years ago

A is correct

upvoted 2 times

Wwhite44 4 years, 2 months ago

In study guide you can find “user discount” and “preemptable”. In this case performance are qequired, so C is not possible. A for sure

upvoted 2 times

not_thanos 4 years, 2 months ago

Answer is A => several months can be up to or more than a year C is wrong because 1. Business critical workload, you don't want that stopped randomly 2. Preemptible requires the workload to be resumable -> not stated in the question

upvoted 3 times

meetplanet 4 years, 2 months ago

please look at the question: "You need to run a business-critical workload". this is not a good idea to go for preemtible vms as it can effect performance and bussiness. I go for A.

upvoted 1 times

danchoif2 4 years, 3 months ago

B and D is not related to cost because of stable workload. Now, really unsure between A and C. A: several months vs. at least one year C: a fixed set of Compute Engine instances vs. preemptable I think I will choose C because using preemptable, we can still try to find a way to create a new preemptable VM (same or different zones) if existing one is terminated.

upvoted 2 times

tycho 4 years, 4 months ago

A) not correct committed discounts 1-3 years, we have several months B) MIG - not correct - stable workloads C) Preemptible vm's - not correct they last only 24 hours D) not sure--Unmanaged instance groups-- do they offer cost savings?

upvoted 2 times

daas01 4 years, 4 months ago

C is correct, committed use discounts must be purchased for a minimum of one year.

upvoted 2 times

gcpengineer 4 years, 6 months ago

C is the answer

upvoted 2 times

Question #26 *Topic 1*

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. Adjust the SLO targets to be achievable by the service so you can bring it into production.

B. Notify the development team that they will have to provide production support for the service.

C. Identify recommended reliability improvements to the service to be completed before handover.

D. Bring the service into production with no SLOs and build them when you have collected operational data.

Suggested Answer: *B*
Official Answer: B

Community vote distribution
C (100%)

Charun Highly Voted 4 years, 7 months ago
C is correct
upvoted 16 times

devopsbatch Highly Voted 4 years, 8 months ago
C identify improvements before handover with no meet SLO
upvoted 10 times

devopsbatch 4 years, 8 months ago
C identify improvements before handover with no meet SLOI also agree for the answer - C
upvoted 10 times

akg001 4 years, 8 months ago
I also agree for the answer - C
upvoted 3 times

habla2019pasta Most Recent 1 year, 9 months ago
C should be correct
upvoted 1 times

jomonkp 2 years, 2 months ago
Option C
upvoted 1 times

raghu09 2 years, 10 months ago
Code has to be optimised well before submitting to the production or pre-production so that it can meet the SLO/
upvoted 1 times

floppino 3 years, 1 month ago
Ans: C Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/
upvoted 1 times

AzureDP900 3 years, 3 months ago
C is the best
upvoted 1 times

ejlp 3 years, 5 months ago
I choose C
upvoted 1 times

GCP72 3 years, 6 months ago
Answer is C
upvoted 1 times

PhilipKoku 4 years ago

C - If the app doesn't meet the set SLOs, improvements need to be met or a revised SLO need to be agreed with relevant stakeholders.

upvoted 3 times

vijaigcp 4 years ago

Answer C makes sense

upvoted 1 times

cyrus86 4 years ago

C is correct,

upvoted 1 times

alaahakim 4 years, 2 months ago

Ans : C

upvoted 1 times

giammydell 4 years, 3 months ago

Why B is Not correct? SLO is used to understand if give priority to Develop or to OperationAccording to SRE book, next phase of conducting PRR in Simple PRR model is to select items in PRR to improve before hand over the service to SRE team. https://sre.google/sre-book/evolving-sre-engagement-model/#improvements-and-refactoring-xqsrUdcyO So C is correct.

upvoted 1 times

giammydell 4 years, 3 months ago

Why B is Not correct? SLO is used to understand if give priority to Develop or to Operation

upvoted 1 times

danchoif2 4 years, 3 months ago

According to SRE book, next phase of conducting PRR in Simple PRR model is to select items in PRR to improve before hand over the service to SRE team. https://sre.google/sre-book/evolving-sre-engagement-model/#improvements-and-refactoring-xqsrUdcyO So C is correct.

upvoted 4 times

francisco_guerra 4 years, 7 months ago

I don't find anything about this but C its the only one that ensure that the service can meet its SLOs in production so I go with C

upvoted 5 times

Question #27 *Topic 1*

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. Start monitoring latency, traffic, errors, and saturation.

C. Record data for the postmortem document of the incident.

D. Trace the origin of 500 errors and the root cause of increased latency.

Suggested Answer: *D*

Reference: https://cloud.google.com/solutions/automated-canary-analysis-kubernetes-engine-spinnaker

Community vote distribution

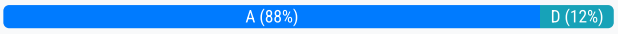

devopsbatch Highly Voted 4 years, 8 months ago

A rollback experiment

upvoted 30 times

devopsbatch Highly Voted 4 years, 8 months ago

A rollback experimentAgree with AAgree with A, that is even why Spinnaker has "Manual judgment" stage that if Canary deployment seems dangerous, it can be immediately cancelled.

upvoted 30 times

akg001 4 years, 8 months ago

Agree with AAgree with A, that is even why Spinnaker has "Manual judgment" stage that if Canary deployment seems dangerous, it can be immediately cancelled.

upvoted 5 times

akg001 4 years, 8 months ago

Agree with A

upvoted 5 times

holahola 4 years, 8 months ago

Agree with A, that is even why Spinnaker has "Manual judgment" stage that if Canary deployment seems dangerous, it can be immediately cancelled.

upvoted 4 times

habla2019pasta Most Recent 1 year, 9 months ago

A. First and foremost, reduce user impact.

upvoted 1 times

jomonkp 2 years, 2 months ago

Option A

upvoted 2 times

Joey1 2 years, 8 months ago

Why are there so many wrong answers listed as "Correct Answers" ? For example for this question A is obviously correct and the most voted. Why does it say D is correct?

upvoted 3 times

Joey1 2 years, 8 months ago

Why are there so many wrong answers listed as "Correct Answers" ? For example for this question A is obviously correct and the most voted. Why does it say D is correct?It's because the community answers are in 95% of all cases the right ones.

upvoted 3 times

Raz0r 2 years, 1 month ago

It's because the community answers are in 95% of all cases the right ones.

upvoted 1 times

raghu09 2 years, 10 months ago

We need to rollback the current deployment so that user are not negatively impacted, then we can go with post-mortem, root cause alaysis and fixes.

upvoted 2 times

floppino 3 years, 1 month ago

Ans: A Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/

upvoted 1 times

sri1016 3 years, 2 months ago

The first step that you do to reduce issue is to rollback, then comes the identifying and fixing the errors part

upvoted 1 times

AzureDP900 3 years, 3 months ago

Rollback is the best option to avoid customer impact. Latency issue can be triaged in staging env with performance test hence my answer is A.

upvoted 1 times

GCP72 3 years, 6 months ago

Agree with A

upvoted 1 times

PhilipKoku 4 years ago

A - Rollback the canary to bring back stability to production; then review logs to find out what caused the issues.

upvoted 3 times

vijaigcp 4 years ago

Answer - A for quickly minimizing the negative impact on users

upvoted 1 times

cyrus86 4 years ago

Selected Answer: D

The correct answer is D, the deployment is Canary, hence a very small amount of user traffic, the whole purpose of Canary here will be to find the origin or errors and fix it before releasing to all the users.

upvoted 2 times

cyrus86 4 years ago

The correct answer is D, the deployment is Canary, hence a very small amount of user traffic, the whole purpose of Canary here will be to find the origin or errors and fix it before releasing to all the users.Nope! You have logs to analyze the issue."You want to quickly minimize the negative impact on users" states that the you need to make the service available ASAP. Doesn't matter about the percent of users using the service. You can't keep waiting until you do the trace and fix. So answer should be A, per SRE principle.

upvoted 2 times

buldas 4 years ago

Nope! You have logs to analyze the issue.

upvoted 1 times

jagadeeshs 4 years ago

"You want to quickly minimize the negative impact on users" states that the you need to make the service available ASAP. Doesn't matter about the percent of users using the service. You can't keep waiting until you do the trace and fix. So answer should be A, per SRE principle.

upvoted 4 times

Sekierer 4 years, 1 month ago

A is correct

upvoted 1 times

cloudbee 4 years, 1 month ago

A should be the answer. error happens after canary release so we are not sure for actual cause. so first we should rollback canary release.

upvoted 1 times

alaahakim 4 years, 2 months ago

Ans : A

upvoted 1 times

driftwood 4 years, 6 months ago

Has to be A

upvoted 3 times

Question #28 *Topic 1*

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. ↄ€¢ Store your code in a Git-based version control system. ↄ€¢ Establish a process that allows developers to merge their own changes at the end of each day. ↄ€¢ Package and upload code to a versioned Cloud Storage basket as the latest master version.

B. ↄ€¢ Store your code in a Git-based version control system. ↄ€¢ Establish a process that includes code reviews by peers and unit testing to ensure integrity and functionality before integration of code. ↄ€¢ Establish a process where the fully integrated code in the repository becomes the latest master version.

C. ↄ€¢ Store your code as text files in Google Drive in a defined folder structure that organizes the files. ↄ€¢ At the end of each day, confirm that all changes have been captured in the files within the folder structure. ↄ€¢ Rename the folder structure with a predefined naming convention that increments the version.

D. ↄ€¢ Store your code as text files in Google Drive in a defined folder structure that organizes the files. ↄ€¢ 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. ↄ€¢ Upload the .zip archive to a versioned Cloud Storage bucket and accept it as the latest version.

Suggested Answer: *A*
Official Answer: A

Community vote distribution
B (100%)

devopsbatch Highly Voted 4 years, 8 months ago
B peer review and source code management tool required
upvoted 31 times

devopsbatch Highly Voted 4 years, 8 months ago
B peer review and source code management tool requiredAgree with you. Answer is B
upvoted 31 times

akg001 4 years, 8 months ago
Agree with you. Answer is B
upvoted 5 times

Charun 4 years, 7 months ago
B for sure
upvoted 6 times

habla2019pasta Most Recent 1 year, 9 months ago
B. Terraform state should never be on a dev machine.. Instead it should be hosted on a sever that is access controlled and updated from latest git commit
upvoted 1 times

jomonkp 2 years, 2 months ago
Option B
upvoted 2 times

maxdanny 2 years, 3 months ago
B is correct because the terraform source file they must be versioned with Git control, Furthermore, changes must be merged and approved before going into production
upvoted 1 times

raghu09 2 years, 10 months ago
A can't be the option why the code be pushed to any other storage when it is Git based?
upvoted 1 times

floppino 3 years, 1 month ago
Ans: B Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/

upvoted 1 times

GCP72 3 years, 6 months ago

B is correct ans

upvoted 1 times

Learning_GCP 3 years, 7 months ago

The important thing here is to avoid overwriting the code, not unit testing and other things. I am confused.

upvoted 1 times

Ananda 3 years, 9 months ago

submitted B

upvoted 1 times

Aiffone 3 years, 9 months ago

Does code review guarantee codes wont be overwritten or its merging of codes that guarantees this?

upvoted 1 times

Aiffone 3 years, 9 months ago

Does code review guarantee codes wont be overwritten or its merging of codes that guarantees this?This is one of the main role in a GIT based system: to allow some peers/coders working together and avoid overwritten the code between them. The first thing you do in the morning is to fetch your local copy of the code, to be sure you are updated previous you start coding. Although the better option, if the project allows it, that you were the only person working on a specific branch. In this case, there is not going to be overwritten problems

upvoted 1 times

Feliphus 2 years, 1 month ago

This is one of the main role in a GIT based system: to allow some peers/coders working together and avoid overwritten the code between them. The first thing you do in the morning is to fetch your local copy of the code, to be sure you are updated previous you start coding. Although the better option, if the project allows it, that you were the only person working on a specific branch. In this case, there is not going to be overwritten problems

upvoted 1 times

Aiffone 3 years, 10 months ago

I wanto go with A because B (code review) does not guarantee the developers won't over write each other's codes.

upvoted 1 times

PhilipKoku 4 years ago

B - Git based repository + Peer review and Unit testing

upvoted 2 times

cyrus86 4 years ago

gitflow

upvoted 1 times

Murty549 4 years, 1 month ago

Selected Answer: B

B is correct. A is false because the code is already stored in a got repo, so there is no need to upload the code to the cloud storage bucket.

upvoted 3 times

Murty549 4 years, 1 month ago

B is correct. A is false because the code is already stored in a got repo, so there is no need to upload the code to the cloud storage bucket.yes, exactly. B is perfect

upvoted 3 times

AzureDP900 3 years, 3 months ago

yes, exactly. B is perfect

upvoted 1 times

Sekierer 4 years, 1 month ago

B is correct

upvoted 1 times

alaahakim 4 years, 2 months ago

Agree with B

upvoted 1 times

Question #29 *Topic 1*

You support a high-traffic web application with a microservice architecture. The home page of the application displays multiple widgets containing content such as the current weather, stock prices, and news headlines. The main serving thread makes a call to a dedicated microservice for each widget and then lays out the homepage for the user. The microservices occasionally fail; when that happens, the serving thread serves the homepage with some missing content. Users of the application are unhappy if this degraded mode occurs too frequently, but they would rather have some content served instead of no content at all. You want to set a Service Level Objective (SLO) to ensure that the user experience does not degrade too much. What Service Level Indicator (SLI) should you use to measure this?

A. A quality SLI: the ratio of non-degraded responses to total responses.

B. An availability SLI: the ratio of healthy microservices to the total number of microservices.

C. A freshness SLI: the proportion of widgets that have been updated within the last 10 minutes.

D. A latency SLI: the ratio of microservice calls that complete in under 100 ms to the total number of microservice calls.

Suggested Answer: *D*

Reference: https://cloud.google.com/stackdriver/docs/solutions/slo-monitoring

Community vote distribution



francisco_guerra Highly Voted 4 years, 1 month ago

Ans: A Quality as an SLI Quality is a helpful SLI for complex services that are designed to fail gracefully by degrading when dependencies are slow or unavailable. The SLI for quality is defined as follows: The proportion of valid requests served without degradation of service.

https://cloud.google.com/architecture/adopting-slos

upvoted 34 times

Charun Highly Voted 4 years, 1 month ago

Ans: B

upvoted 10 times

Charun 4 years, 1 month ago

Ans: BNope A

upvoted 10 times

TNT87 3 years, 7 months ago

Nope A

upvoted 2 times

jomonkp Most Recent 1 year, 8 months ago

Option A

upvoted 1 times

floppino 2 years, 7 months ago

Ans: A Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/

upvoted 1 times

sadheeshs 2 years, 8 months ago

Confusion between ET and user answers. How should I know which answer is correct? Anyone can address this?

upvoted 1 times

GCP72 3 years ago

Answer A is looks correct for me

upvoted 1 times

rajatjain12 3 years, 1 month ago

I think D is the correct

upvoted 2 times

Ananda 3 years, 3 months ago

Submitted A

upvoted 3 times

gomezzang 3 years, 3 months ago

I think B.

upvoted 1 times

PhilipKoku 3 years, 6 months ago

A - Quality SLI: Number of correct pages / Total number of pages served.

upvoted 5 times

cyrus86 3 years, 6 months ago

https://cloud.google.com/architecture/adopting-slos#sli-quality

upvoted 2 times

guid1984 3 years, 8 months ago

Answer A - Quality as an SLI suits well for degrading dependencies

upvoted 1 times

not_thanos 3 years, 8 months ago

https://cloud.google.com/architecture/adopting-slos?hl=en#sli-quality Quality is a helpful SLI for complex services that are designed to fail gracefully by degrading when dependencies are slow or unavailable. The SLI for quality is defined as follows: The proportion of valid requests served without degradation of service. For example, a web page might load its main content from one datastore and load ancillary, optional assets from 100 other services and datastores. If one optional service is out of service or too slow, the page can still be rendered without the ancillary elements. By measuring the number of requests that are served a degraded response (that is, a response missing at least one backend service's response), you can report the ratio of requests that were bad. You might even track how many responses to the user were missing a response from a single backend, or were missing responses from multiple backends.

upvoted 6 times

alaahakim 3 years, 8 months ago

Ans : B

upvoted 2 times

Manh 3 years, 9 months ago

Ans is A https://sre.google/workbook/implementing-slos/ Request-driven: Quality If the service degrades gracefully when overloaded or when backends are unavailable, you need to measure the proportion of responses that were served in an undegraded state. For example, if the User Data store is unavailable, the game is still playable but uses generic imagery.Ans B https://cloud.google.com/blog/products/gcp/available-or-not-that-is-the-question-cre-life-lessons, your answer is wrong

upvoted 2 times

Manh 3 years, 9 months ago

Ans is A https://sre.google/workbook/implementing-slos/ Request-driven: Quality

If the service degrades gracefully when overloaded or when backends are unavailable, you need to measure the proportion of responses that were served in an undegraded state. For example, if the User Data store is unavailable, the game is still playable but uses generic imagery.

upvoted 2 times

TNT87 3 years, 9 months ago

Ans B https://cloud.google.com/blog/products/gcp/available-or-not-that-is-the-question-cre-life-lessons, your answer is wrong

upvoted 3 times

giammydell 3 years, 9 months ago

I think D is the correct, B is not time-bound. B have to be writen as an availability probability in a time window

upvoted 1 times

MF2C 3 years, 11 months ago

A Users may not be consciously aware of the degradation in quality until it becomes severe, but their subconscious perceptions may still have an impact on your business if e.g. degrading quality means serving less relevant ads to users, reducing click-through rates.

upvoted 2 times

Question #30 *Topic 1*

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 two.)

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.

Suggested Answer: *C, D*
Official Answer: CD

Community vote distribution
BE (93%)

celia20200410 Highly Voted 4 years, 7 months ago
https://cloud.google.com/architecture/adopting-slos#choosing_a_measurement_method B > Using client instrumentation. E > Implementing synthetic testing.How you code instrumentation at each client side?Client instrumentation can be easily done using Browser instrumentation using various UI frameworks you can do that. One of the most popular framework and SDK is OpenTelemetry. So the answer is B and E.
upvoted 23 times

celia20200410 Highly Voted 4 years, 7 months ago
https://cloud.google.com/architecture/adopting-slos#choosing_a_measurement_method B > Using client instrumentation. E > Implementing synthetic testing.
upvoted 23 times

cloudbee 4 years, 1 month ago
How you code instrumentation at each client side?
upvoted 2 times

cloudbee 4 years, 1 month ago
How you code instrumentation at each client side?Client instrumentation can be easily done using Browser instrumentation using various UI frameworks you can do that. One of the most popular framework and SDK is OpenTelemetry. So the answer is B and E.
upvoted 2 times

jinaldesailive 1 year, 11 months ago
Client instrumentation can be easily done using Browser instrumentation using various UI frameworks you can do that. One of the most popular framework and SDK is OpenTelemetry. So the answer is B and E.
upvoted 1 times

Charun 4 years, 7 months ago
BE is correct
upvoted 11 times

zanhsieh Most Recent 1 year ago
BE. A: No. Application server logs are too localized and do not detect global issues like CDN failures. B: Yes. Instrumenting the client directly provides visibility into real end-user experience and can capture issues before traffic reaches your infrastructure. C: No. Metrics from servers only show internal health and do not reflect upstream failures like those from the CDN or load balancer. D: No. GKE health checks are internal to the cluster and can't detect problems in the CDN or CLB layers. E: Yes. Synthetic clients simulate end-user traffic and can detect issues in the CDN, load balancer, or global network. https://sre.google/sre-book/monitoring-distributed-systems/#implementing-slis-1 https://cloud.google.com/monitoring/synthetic-monitoring
upvoted 1 times

habla2019pasta 1 year, 9 months ago
BE A - Is inside the GKE .. which is useless at this point C - Same as A D - Same as it too late in the pipeline to detect these issues

upvoted 1 times

alpha_canary 2 years ago

I dont still get why the answers are B & E. Option A & C also appear in below documentation along with B & E so everyone is only chosing B & E
https://cloud.google.com/architecture/framework/reliability/adopting-slos#choose_a_measurement_method

upvoted 1 times

jomonkp 2 years, 2 months ago

Option B and E

upvoted 1 times

maxdanny 2 years, 3 months ago

https://cloud.google.com/architecture/framework/reliability/adopting-slos#choose_a_measurement_method

upvoted 1 times

floppino 3 years, 1 month ago

Ans: BE Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here:
https://www.examtopics.com/discussions/google/1/

upvoted 1 times

hanweiCN 3 years, 2 months ago

from my opinion, network issue will not be detected in app server side. so any SLI in application wont work.

upvoted 1 times

Bill831231 3 years, 6 months ago

If need something beyong CDN and CLB, seems only option is on client side directly

upvoted 3 times

GCP72 3 years, 6 months ago

For me, B and E is correct

upvoted 2 times

Pankul 3 years, 9 months ago

For me, it is B and E

upvoted 2 times

zygomar 3 years, 12 months ago

nobody submitted a voting comment but B&E is the consensus on most feedback

upvoted 1 times

guid1984 4 years, 1 month ago

IMHO -> (E and C)A synthetic client that periodically sends simulated user requests could be used for checking CDN connection is up and running at frequent intervals and other one C for metrics exposed from Application servers. Since as per the question CLB SLI is already defined.

upvoted 1 times

alaahakim 4 years, 2 months ago

Ans : BE

upvoted 1 times

muk5658 4 years, 2 months ago

Its B and E, when there is a networking catastrophe we cannot rely on Application as traffic may not even come there.

upvoted 1 times

Question #31 *Topic 1*

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

A. Publish various metrics from the application directly to the Stackdriver Monitoring API, and then observe these custom metrics in Stackdriver.

B. Install the Cloud Pub/Sub client libraries, push various metrics from the application to various topics, and then observe the aggregated metrics in Stackdriver.

C. Install the OpenTelemetry client libraries in the application, configure Stackdriver as the export destination for the metrics, and then observe the application's metrics in Stackdriver.

D. Emit all metrics in the form of application-specific log messages, pass these messages from the containers to the Stackdriver logging collector, and then observe metrics in Stackdriver.

Suggested Answer: *C*
Official Answer: C

Community vote distribution
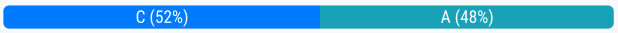

giammydell Highly Voted 4 years, 3 months ago
C is correct, for A the installing of stackdriver library in the application is missing, without the lbrary custom metric cannot be created by the application
upvoted 19 times

Luvero Highly Voted 4 years, 3 months ago
You want to use Google Cloud Platform services while minimizing the amount of work required to set up monitoring. Hence A
upvoted 14 times

Srikantha Most Recent 2 months, 3 weeks ago
The correct answer is 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.
upvoted 1 times

gaurav517 8 months, 3 weeks ago
A requires writing custom code to call API directly. C has the least effort.
upvoted 1 times

Luis_Ehlen 1 year, 5 months ago
Open telemetry is the less effort option
upvoted 1 times

habla2019pasta 1 year, 9 months ago
C should be correct
upvoted 1 times

chidambaram 1 year, 10 months ago
C. OpenTelemetry with Stackdriver: OpenTelemetry is a vendor-neutral instrumentation framework. By configuring Stackdriver as the export destination, you leverage a standardized approach with minimal coding effort. OpenTelemetry offers an easy way to collect and export metrics from your application to Stackdriver, minimizing setup work. Answer : C
upvoted 1 times

jinaldesailive 1 year, 11 months ago
OpenTelemetry provides standardized metrics collection, with it's SDK and various client libraries. So using OpenTelemetry is much faster and easier than using Cloud Monitoring API directly. So it's C.
upvoted 1 times

jomonkp 2 years, 2 months ago
Option C
upvoted 1 times

macqly 2 years, 3 months ago

You can create user-defined metrics by using the Cloud Monitoring API directly. However, we recommend that you use OpenTelemetry. https://cloud.google.com/monitoring/custom-metricsI think coding raw API calls is more effort than configure Open Telemetry. But maybe I am wrong.

upvoted 2 times

macqly 2 years, 3 months ago

You can create user-defined metrics by using the Cloud Monitoring API directly. However, we recommend that you use OpenTelemetry. https://cloud.google.com/monitoring/custom-metrics

upvoted 2 times

macqly 2 years, 3 months ago

I think coding raw API calls is more effort than configure Open Telemetry. But maybe I am wrong.

upvoted 1 times

ManishKS 2 years, 4 months ago

C should be correct Answer

upvoted 1 times

Cassim 2 years, 6 months ago

C is correct. Installing Open Telemetry client libraries in the application. setting stackdriver as export destination and then observing the app metrics.

upvoted 1 times

Cassim 2 years, 6 months ago

C is correct. Installing Open Telemetry client libraries in the application. setting stackdriver as export destination and then observing the app metrics.Hey Cassim ? Are these questions still relevant ? Thanks in advance for your answer!yes they are, just replace the name stack driver with Cloud Operations Suite and you will be fine https://cloud.google.com/products/operations?hl=en

upvoted 1 times

B_Ahmed 2 years, 6 months ago

Hey Cassim ? Are these questions still relevant ? Thanks in advance for your answer!

upvoted 2 times

B_Ahmed 2 years, 6 months ago

Hey Cassim ? Are these questions still relevant ? Thanks in advance for your answer!yes they are, just replace the name stack driver with Cloud Operations Suite and you will be fine https://cloud.google.com/products/operations?hl=en

upvoted 2 times

Orzechowski 2 years, 5 months ago

yes they are, just replace the name stack driver with Cloud Operations Suite and you will be fine https://cloud.google.com/products/operations?hl=en

upvoted 1 times

boras 3 years, 1 month ago

You should use a library/SDK instead of coding raw API requests all over the place. Hence I go with C - OpenTelemetry.

upvoted 3 times

Greg123123 3 years, 1 month ago

The ans should be A. 3 main points here: application-level metrics use Google Cloud Platform services minimizing the amount of work required B: b is not correct as google suggest only external metrics are pushed using pub/sub. For custom metrics use stackdriver api. C: OpenTelemetry is not a gcp service D: too much work

upvoted 4 times

floppino 3 years, 1 month ago

Ans: A Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/

upvoted 2 times

hanweiCN 3 years, 2 months ago

openTelemetry library only appear in Cloud Trace , and openTelemetry is not a GCP service ? https://cloud.google.com/trace/docs/setup

upvoted 2 times

hanweiCN 3 years, 2 months ago

stackdriver API can not monitor application-level metrics https://cloud.google.com/stackdriver/docs/solutions/gke/installing

upvoted 2 times

Question #32 *Topic 1*

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

A. File a bug with the development team so they can find the root cause of the crashing instance.

B. Create a Managed instance Group with a single instance and use health checks to determine the system status.

C. Add a Load Balancer in front of the Compute Engine instance and use health checks to determine the system status.

D. Create a Stackdriver Monitoring dashboard with SMS alerts to be able to start recreating the crashed instance promptly after it was crashed.

Suggested Answer: *A*
Official Answer: A

Community vote distribution
B (85%) A (15%)

devopsbatch Highly Voted 4 years, 8 months ago
B managed instance groups can be handled when a vm crashed and immediately created new oneAgree. correct Answer : B
upvoted 26 times

devopsbatch Highly Voted 4 years, 8 months ago
B managed instance groups can be handled when a vm crashed and immediately created new one
upvoted 26 times

akg001 4 years, 8 months ago
Agree. correct Answer : B
upvoted 6 times

Charun 4 years, 7 months ago
B is correct
upvoted 8 times

habla2019pasta Most Recent 1 year, 9 months ago
MIG is the right way.
upvoted 1 times

jomonkp 2 years, 2 months ago
option B
upvoted 1 times

maxdanny 2 years, 3 months ago
B https://cloud.google.com/compute/docs/instance-groups?hl=en#autohealing
upvoted 1 times

calex1 2 years, 5 months ago
Correct B
upvoted 1 times

sipekmichal 3 years, 1 month ago
definitely B
upvoted 1 times

floppino 3 years, 1 month ago
Ans: B Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/
upvoted 1 times

kapara 3 years, 6 months ago
I think that the answer is B, Although SRE principles guide you to find the root cause and post-mortem, the question clearly asks you to: Reduce time spent on manual operations. Therefore the answer is B (Although deep in my heart I would personally combine A and B)

upvoted 4 times

GCP72 3 years, 6 months ago

B is correct

upvoted 1 times

Ananda 3 years, 9 months ago

Submit B in the exam

upvoted 1 times

PhilipKoku 4 years ago

B - Managed Instance Group with one Compute Engine instance will take care of this use case.

upvoted 1 times

hantanbl 4 years ago

answer is A. If the root cause of crashing is not found and fixed, the crashing will still happen frequently even if we go MIG. Waste of money and time and does not follows SRE principles

upvoted 4 times

hantanbl 4 years ago

answer is A. If the root cause of crashing is not found and fixed, the crashing will still happen frequently even if we go MIG. Waste of money and time and does not follows SRE principlesbut if the question is saying that "You want to reduce the time spent performing manual operations while following Site Reliability Engineering principles",is not improvement and fix the issue, I think if we consider reducing the time spent performing manual operations, cloud be Option B

upvoted 4 times

jeffersonkozak 2 years, 8 months ago

but if the question is saying that "You want to reduce the time spent performing manual operations while following Site Reliability Engineering principles",is not improvement and fix the issue, I think if we consider reducing the time spent performing manual operations, cloud be Option B

upvoted 1 times

Nalex9ja 4 years ago

B because it automates toil. I can't think of a better answer.

upvoted 1 times

vijaigcp 4 years ago

Answer B, SRE principles says remove the toil where ever possible. MIG will fit here by automating the VM creation.

upvoted 2 times

cyrus86 4 years ago

IMO - SRE practice should be to get the root cause of crashing of VM and fix it. MIG will not help here.

upvoted 3 times

garmstrong 4 years, 2 months ago

I'd say B, as a managed instance group can recreate crashed VMs

upvoted 1 times

Question #33 *Topic 1*

Your application artifacts are being built and deployed via a CI/CD pipeline. You want the CI/CD pipeline to securely access application secrets. You also want to more easily rotate secrets in case of a security breach. What should you do?

A. Prompt developers for secrets at build time. Instruct developers to not store secrets at rest.

B. Store secrets in a separate configuration file on Git. Provide select developers with access to the configuration file.

C. Store secrets in Cloud Storage encrypted with a key from Cloud KMS. Provide the CI/CD pipeline with access to Cloud KMS via IAM.

D. Encrypt the secrets and store them in the source code repository. Store a decryption key in a separate repository and grant your pipeline access to it.

Suggested Answer: *C*

Official Answer: C

Community vote distribution



driftwood Highly Voted 4 years, 6 months ago

C is the "best" from the choices given

upvoted 13 times

Charun Highly Voted 4 years, 7 months ago

Answer C

upvoted 6 times

habla2019pasta Most Recent 1 year, 9 months ago

Option C

upvoted 1 times

jomonkp 2 years, 2 months ago

Option C

upvoted 1 times

calex1 2 years, 5 months ago

C is the only viable option...

upvoted 1 times

floppino 3 years, 1 month ago

Ans: C Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/

upvoted 1 times

GCP72 3 years, 6 months ago

Answer is C

upvoted 1 times

vijaigcp 4 years ago

Answer is C , https://cloud.google.com/security-key-management

upvoted 2 times

cyrus86 4 years ago

C is the answer

upvoted 1 times

alaahakim 4 years, 2 months ago

Ans : C

upvoted 1 times

akg001 4 years, 8 months ago

C is the best option.

upvoted 4 times

devopsbatch 4 years, 8 months ago

answer C storing secrets in cloud is better option

upvoted 4 times

Question #34 *Topic 1*

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

A. Focus on responding to internal stakeholders at least every 30 minutes. Commit to ג€next updateג€ times.

B. Provide periodic updates to all stakeholders in a timely manner. Commit to a ג€next updateג€ time in all communications.

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

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

Suggested Answer: *C*
Official Answer: C

Community vote distribution
B (89%) C (11%)

francisco_guerra Highly Voted 4 years, 1 month ago
Ans: B (the communication lead CAN'T delegate) 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/
upvoted 24 times

Charun Highly Voted 4 years, 1 month ago
B is correct
upvoted 10 times

jomonkp Most Recent 1 year, 8 months ago
Option B
upvoted 1 times

floppino 2 years, 7 months ago
Ans: B Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/
upvoted 1 times

mudot 2 years, 9 months ago
Appeard in 7/11/2022 exam Keys : Large incident, you need to respond to stakeholder + users, manage it effectively
upvoted 1 times

GCP72 3 years ago
Answer is B
upvoted 1 times

vijaigcp 3 years, 6 months ago
Ans B , The CL's main duties include providing periodic updates to the incident response team and stakeholders, and managing inquiries about the incident.
upvoted 1 times

dalecolor 3 years, 7 months ago
B was my first answer, but turns out that CL can delegate, according to SRE workbook (https://sre.google/workbook/incident-response/): "Both the CL and OL may lead a team of people to help manage their specific areas of incident response. These teams can expand or contract as needed. If the incident becomes small enough, the CL role can be subsumed back into the IC role." In that case, taking into account that this is a major incidente, C seems to be the most accurate answer.
upvoted 1 times

garmstrong 3 years, 8 months ago

B - providing updates to all stakeholders is your job as being in charge of communications. Incident Commander delegates.

upvoted 2 times

meetplanet 3 years, 8 months ago

B is correct. everyone is saying B is correct. not sure why final answer is C!!

upvoted 2 times

danchoif2 3 years, 9 months ago

B is correct C is wrong. If CL want to delegate, s/he must talk with IC that we need another internal CL. IC will delegate if needed.

upvoted 3 times

ralf_cc 4 years, 2 months ago

B is most correct one, with quite a lot of details missing, but the direction is right

upvoted 2 times

rinkeshgala1 4 years, 2 months ago

we can have external communication lead. as there is a large outage we can choose to have one more role of external communication lead. hence option C

upvoted 1 times

rinkeshgala1 4 years, 2 months ago

we can have external communication lead. as there is a large outage we can choose to have one more role of external communication lead. hence option Cfor answer C. one doubt does Communication Leads (CL) can delegate the task or this is Incident commander (IC) can do... B looks more closer.Looks correct to me as well.

upvoted 1 times

akg001 4 years, 2 months ago

for answer C. one doubt does Communication Leads (CL) can delegate the task or this is Incident commander (IC) can do... B looks more closer.

upvoted 2 times

akg001 4 years, 2 months ago

Looks correct to me as well.

upvoted 1 times

devopsbatch 4 years, 2 months ago

could be B

upvoted 1 times

Question #35 *Topic 1*

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

A. Assign the Container Developer role to the Cloud Build service account.

B. Specify the Container Developer role for Cloud Build in the cloudbuild.yaml file.

C. Create a new service account with the Container Developer role and use it to run Cloud Build.

D. Create a separate step in Cloud Build to retrieve service account credentials and pass these to kubectl.

Suggested Answer: *C*
Official Answer: C

Community vote distribution
A (67%) C (33%)

TNT87 Highly Voted 4 years, 5 months ago
A https://cloud.google.com/build/docs/securing-builds/configure-user-specified-service-accounts
upvoted 11 times

Charun Highly Voted 4 years, 7 months ago
A is correct
upvoted 6 times

gaurav517 Most Recent 8 months, 3 weeks ago
C: Follows principle of least privilege Dedicated SA only for GKE deployment pipelines Isolates permissions - only specific builds use this SA Other Cloud Build pipelines don't get GKE access Better security posture and access control Industry best practice for production environments ⚠ Slightly more initial setup, but worth it for security
upvoted 1 times

jomonkp 2 years, 2 months ago
Option C
upvoted 4 times

Orzechowski 2 years, 5 months ago
A https://cloud.google.com/build/docs/deploying-builds/deploy-gke#required_iam_permissions
upvoted 1 times

sidharthwader 2 years, 7 months ago
Doesn't C makes more sense ? Why is it A?are these questions still relevant?I hope that are still relevant. I'll take the exam this week, can someone confirm about aswani question?can someone confirm now?
upvoted 2 times

sidharthwader 2 years, 7 months ago
Doesn't C makes more sense ? Why is it A?
upvoted 2 times

aswani 2 years, 7 months ago
are these questions still relevant?I hope that are still relevant. I'll take the exam this week, can someone confirm about aswani question?can someone confirm now?
upvoted 2 times

Magist3r 2 years, 7 months ago
I hope that are still relevant. I'll take the exam this week, can someone confirm about aswani question?
upvoted 1 times

Magist3r 2 years, 7 months ago
I hope that are still relevant. I'll take the exam this week, can someone confirm about aswani question?can someone confirm now?
upvoted 1 times

joshtechgroup 2 years, 1 month ago
can someone confirm now?

upvoted 1 times

samuelmorher 2 years, 7 months ago

I think that the A is incorrect... The good practices says that the CB like the other resources should avoid to use the default SA, so the correct one is the C which creates a SA and then give the required roles.are these questions still relevant?

upvoted 1 times

samuelmorher 2 years, 7 months ago

I think that the A is incorrect... The good practices says that the CB like the other resources should avoid to use the default SA, so the correct one is the C which creates a SA and then give the required roles.

upvoted 1 times

aswani 2 years, 7 months ago

are these questions still relevant?

upvoted 1 times

felipeschossler 2 years, 10 months ago

100% A and this is the doc the proves this: https://cloud.google.com/build/docs/deploying-builds/deploy-gke#required_iam_permissions

upvoted 3 times

juliefighting 3 years, 1 month ago

Answer is A https://cloud.google.com/build/docs/deploying-builds/deploy-gke#required_iam_permissions

upvoted 1 times

floppino 3 years, 1 month ago

Ans: A Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/

upvoted 1 times

hanweiCN 3 years, 2 months ago

i think A, new service account needs " Cloud Build Service Account " role and " kubernete engine developer" role to execute the build steps for cloud build.

upvoted 2 times

AzureDP900 3 years, 3 months ago

A is more suitable for this scenario https://cloud.google.com/build/docs/securing-builds/configure-access-for-cloud-build-service-account

upvoted 1 times

okercho 3 years, 4 months ago

I think A is correct, but please note that question specify that kubectl builder (https://github.com/GoogleCloudPlatform/cloud-builders/tree/master/kubectl) and NOT gke-deploy (https://github.com/GoogleCloudPlatform/cloud-builders/tree/master/gke-deploy) is being used! https://cloud.google.com/build/docs/deploying-builds/deploy-gke In any case, as specified in kubectl builder documentation: When executed in the Cloud Build environment, commands are executed with credentials of the builder service account for the build project.

upvoted 1 times

GCP72 3 years, 6 months ago

Answer is A

upvoted 1 times

Pankul 3 years, 9 months ago

A is the correct answer

upvoted 1 times

cloudbee 4 years, 1 month ago

A should be the correct one. because assigning permission to cloud build service account will give permission to deploy while minimizing additional overhead.

upvoted 4 times

Biden 4 years, 2 months ago

Agree with A. Reference to container.developer role: https://cloud.google.com/kubernetes-engine/docs/how-to/iam

upvoted 3 times

Question #36 *Topic 1*

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

A. Link Stackdriver Logging as a source in Google Data Studio. Filter the logs on the cache misses.

B. Configure Stackdriver Profiler to identify and visualize when the cache misses occur based on the logs.

C. Create a logs-based metric in Stackdriver Logging and a dashboard for that metric in Stackdriver Monitoring.

D. Configure BigQuery as a sink for Stackdriver Logging. Create a scheduled query to filter the cache miss logs and write them to a separate table.

Suggested Answer: *C*
Official Answer: C

Community vote distribution
C (100%)

Charun Highly Voted 4 years, 7 months ago
C for sure
upvoted 18 times

syslog Highly Voted 4 years, 8 months ago
I think C: https://cloud.google.com/logging/docs/logs-based-metrics#counter-metric
upvoted 11 times

habla2019pasta Most Recent 1 year, 9 months ago
Option C
upvoted 1 times

jomonkp 2 years, 2 months ago
Option C
upvoted 1 times

ngz123 2 years, 5 months ago
are these questions still valid?
upvoted 1 times

floppino 3 years, 1 month ago
Ans: C Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/
upvoted 2 times

GCP72 3 years, 6 months ago
Answer is C
upvoted 1 times

Ananda 3 years, 9 months ago
Submitted C
upvoted 1 times

vijaigcp 4 years ago
Answer C
upvoted 2 times

TNT87 4 years, 5 months ago
C https://cloud.google.com/logging/docs/logs-based-metrics#counter-metric
upvoted 4 times

devopsbatch 4 years, 8 months ago
D export to biguery as sink and analyse logs
upvoted 1 times

devopsbatch 4 years, 8 months ago
D export to biguery as sink and analyse logsand how would you visualize the data in bq?

upvoted 1 times

tycho 4 years, 4 months ago

and how would you visualize the data in bq?

upvoted 2 times

Question #37 *Topic 1*

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.

Suggested Answer: *D*

Official Answer: D

Community vote distribution



Charun Highly Voted 4 years, 7 months ago

C 100%

upvoted 15 times

TNT87 Highly Voted 4 years, 5 months ago

C https://cloud.google.com/compute/quotas

upvoted 12 times

trashbox Most Recent 1 year, 9 months ago

Exam on 2024-04-26

upvoted 1 times

jomonkp 2 years, 2 months ago

option c

upvoted 2 times

calex1 2 years, 5 months ago

C 100%

upvoted 1 times

floppino 3 years, 1 month ago

Ans: C Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/

upvoted 1 times

GCP72 3 years, 6 months ago

Answer is C

upvoted 1 times

ghadxx 3 years, 9 months ago

Knowing available quota limits allows you to plan for capacity

upvoted 2 times

asdqweasd 3 years, 12 months ago

You need to make sure of quota limit other things will be handled by MIG

upvoted 1 times

cyrus86 4 years ago

C shud be the answer

upvoted 1 times

meetplanet 4 years, 2 months ago

Correct answer is C. this service is deployed into MIG and we only have to take care of quota

upvoted 2 times

Shasha1 4 years, 2 months ago

Answer is D Networking and load balancing quotas are not regional quotes, those are global quotas. Any region can use a global quota. For example, in-use and static external IP addresses assigned to load balancers and HTTP(S) proxies consume global quotas. reference https://cloud.google.com/compute/quotas

upvoted 1 times

Shasha1 4 years, 2 months ago

Answer is D Networking and load balancing quotas are not regional quotes, those are global quotas. Any region can use a global quota. For example, in-use and static external IP addresses assigned to load balancers and HTTP(S) proxies consume global quotas. reference https://cloud.google.com/compute/quotasD says "one" region, the solution is meant to be multi regional

upvoted 1 times

Aiffone 3 years, 9 months ago

D says "one" region, the solution is meant to be multi regional

upvoted 1 times

kubosuke 4 years, 6 months ago

absolutely C 'cause this service is deployed into MIG and you only have to take care of quota.

upvoted 3 times

ralf_cc 4 years, 8 months ago

C - it is asking for quota

upvoted 4 times

DucLee3110 4 years, 8 months ago

I think C is the answer. https://cloud.google.com/compute/quotas#understanding_quotas https://cloud.google.com/compute/quotasAgree for the answer : C

upvoted 5 times

DucLee3110 4 years, 8 months ago

I think C is the answer. https://cloud.google.com/compute/quotas#understanding_quotas https://cloud.google.com/compute/quotas

upvoted 5 times

akg001 4 years, 8 months ago

Agree for the answer : C

upvoted 2 times

devopsbatch 4 years, 8 months ago

A can be the answer

upvoted 1 times

devopsbatch 4 years, 8 months ago

A can be the answeryou have to check the required resources available on those particular regions or not. As you know not all type of VM or a resource available in all regions. There is not surety that n1-highcpu-96 machine type available in all regions.And nobody is saying you need 96 CPU.

upvoted 1 times

akg001 4 years, 8 months ago

you have to check the required resources available on those particular regions or not. As you know not all type of VM or a resource available in all regions. There is not surety that n1-highcpu-96 machine type available in all regions.

upvoted 1 times

akg001 4 years, 8 months ago

you have to check the required resources available on those particular regions or not. As you know not all type of VM or a resource available in all regions. There is not surety that n1-highcpu-96 machine type available in all regions.And nobody is saying you need 96 CPU.

upvoted 1 times

syslog 4 years, 8 months ago

And nobody is saying you need 96 CPU.

upvoted 4 times

Question #38 *Topic 1*

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 log 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.

Suggested Answer: *A*
Official Answer: A

Community vote distribution
B (67%) C (33%)

DucLee3110 Highly Voted 5 years, 2 months ago
looks like it is B. https://medium.com/google-cloud/fluentd-filter-plugin-for-google-cloud-data-loss-prevention-api-42bbb1308e76
upvoted 19 times

DucLee3110 Highly Voted 5 years, 2 months ago
looks like it is B. https://medium.com/google-cloud/fluentd-filter-plugin-for-google-cloud-data-loss-prevention-api-42bbb1308e76to me , looks B is the correct answer .Why not D?prevent them from leaking to Stackdriver Logging. If you need to create a log export & log filter so the information is leaking to logging.Agree with BNot available any more.B to me as well. Because fluentd can filter the logs quite nicely before passing information to Stackdriver. It can cober sensitive information such as credit card details, social security numbers, etc. Once the filtering is done, then the log can be passed to Cloud Storage, but the unfiltered information should not even reach stackdriver, so most of the answers are wrong.
upvoted 19 times

akg001 5 years, 2 months ago
to me , looks B is the correct answer .Why not D?prevent them from leaking to Stackdriver Logging. If you need to create a log export & log filter so the information is leaking to logging.
upvoted 2 times

akg001 5 years, 2 months ago
to me , looks B is the correct answer .
upvoted 2 times

syslog 5 years, 2 months ago
Why not D?
upvoted 2 times

syslog 5 years, 2 months ago
Why not D?prevent them from leaking to Stackdriver Logging. If you need to create a log export & log filter so the information is leaking to logging.
upvoted 2 times

francisco_guerra 5 years, 1 month ago
prevent them from leaking to Stackdriver Logging. If you need to create a log export & log filter so the information is leaking to logging.
upvoted 6 times

Manh 4 years, 9 months ago
Agree with B
upvoted 1 times

irmingard_examtopics 2 years, 8 months ago
Not available any more.
upvoted 1 times

holahola 5 years, 1 month ago

B to me as well. Because fluentd can filter the logs quite nicely before passing information to Stackdriver. It can cober sensitive information such as credit card details, social security numbers, etc. Once the filtering is done, then the log can be passed to Cloud Storage, but the unfiltered information should not even reach stackdriver, so most of the answers are wrong.

upvoted 5 times

francisco_guerra 5 years, 1 month ago

Im not pretty sure but Ans B Prevent them form leaking to Stackdriver logging A: Incorrect, Leaking to Stackdriver B: Correct, not leaking to Stackdriver & fluentD C: Incorrect, Leaking D: If we removed why we need to create a filter matching there will not be logs with userinfo?

upvoted 9 times

Srikantha Most Recent 2 months, 3 weeks ago

Correct Answer: C Why C is the best solution: Create an advanced log filter matching userinfo* → This identifies the PII logs accurately. Configure a log export (sink) to Cloud Storage → This securely captures the sensitive logs in a dedicated, controlled location for later review (e.g., security/compliance team). Configure a log exclusion with the same filter → This prevents the PII logs from being stored in the default _Default log bucket in Cloud Logging, stopping the leakage. This is the standard, recommended Google Cloud Logging pattern for handling sensitive data (PII) in logs.

upvoted 1 times

gaurav517 8 months, 3 weeks ago

 Log export (sink): Captures userinfo logs to Cloud Storage (secure review) Log exclusion: Prevents userinfo logs from being stored in Stackdriver Two-step process works together: Export filter sends PII logs to Cloud Storage Exclusion filter prevents same logs from staying in Stackdriver Uses native GCP logging features Meets both requirements perfectly No custom development needed

upvoted 1 times

gaurav517 8 months, 3 weeks ago

It should be C. 1. Sink (Include): Create a high-priority sink to export logs matching userinfo to a private Cloud Storage bucket. 2. Exclusion Filter (Exclude): Create an exclusion filter to drop logs matching userinfo from all Cloud Logging storage buckets.

upvoted 1 times

account123 2 years ago

Why not D? Use fluentd filter to remove logs with userinfo and export existing logs for review?

upvoted 1 times

jomonkp 2 years, 8 months ago

option b

upvoted 1 times

jeffersonkozak 3 years, 2 months ago

A suggests creating a basic log filter and configuring a log export to Cloud Storage, but it does not address preventing the PII entries from leaking to Stackdriver Logging. By creating only a basic log filter, the PII data would still be accessible within Stackdriver Logging.are these questions still relevant? @jeffersonkozak

upvoted 1 times

jeffersonkozak 3 years, 2 months ago

Selected Answer: C

A suggests creating a basic log filter and configuring a log export to Cloud Storage, but it does not address preventing the PII entries from leaking to Stackdriver Logging. By creating only a basic log filter, the PII data would still be accessible within Stackdriver Logging.

upvoted 1 times

aswani 3 years, 1 month ago

are these questions still relevant? @jeffersonkozak

upvoted 1 times

Greg123123 3 years, 7 months ago

All A,C and D are leaking to stackdriver. So the ans has to be B

upvoted 1 times

floppino 3 years, 7 months ago

Ans: B Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/

upvoted 1 times

raghupothula 3 years, 8 months ago

will go with B

upvoted 1 times

saiprasathdv 3 years, 11 months ago

Option B: https://cloud.google.com/logging/docs/agent/logging/configuration. Custom defined log entries has this structure "[TAG_NAME]+Payload+timestamp+Severity+labels". Here "Userinfo" is the TAG_NAME. Fluentd filter plugins used to filter out logs based on TAG_NAME. finally this could be stored in Cloud storage.

upvoted 3 times

GCP72 4 years ago

Answer is B

upvoted 1 times

mgm7 4 years, 1 month ago

Selected Answer: C

B is a poor answer IMHO because it seems to include a manual task of copying the PII data to the bucket at some later date (cron job or personal intervention, ugly either way). https://cloud.google.com/logging/docs/routing/overview makes it quite clear that this should be taken care by routing and the exclusion filter which would imply that C is the correct answer. "Cloud logging" is only one of the possible sink choices.

upvoted 1 times

mgm7 4 years, 1 month ago

B is a poor answer IMHO because it seems to include a manual task of copying the PII data to the bucket at some later date (cron job or personal intervention, ugly either way). https://cloud.google.com/logging/docs/routing/overview makes it quite clear that this should be taken care by routing and the exclusion filter which would imply that C is the correct answer. "Cloud logging" is only one of the possible sink choices.Nevermind... I though B is ugly, if the PII data is removed with the fluentd filter, it will never arrive at the advanced filter.

upvoted 1 times

mgm7 4 years, 1 month ago

Nevermind... I though B is ugly, if the PII data is removed with the fluentd filter, it will never arrive at the advanced filter.

upvoted 1 times

TNT87 4 years, 7 months ago

Ans B https://medium.com/google-cloud/fluentd-filter-plugin-for-google-cloud-data-loss-prevention-api-42bbb1308e76

upvoted 2 times

garmstrong 4 years, 8 months ago

Agree with B

upvoted 1 times

Shasha1 4 years, 8 months ago

A question is about capture logs entries in a secure location for later review, not removing the log sensitive data before store then in a secure location. so answer is A

upvoted 4 times

Question #39 *Topic 1*

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. Upload the configuration YAML file to Cloud Storage and use Error Reporting to identify and fix the issue.

D. Run a Git compare between the previous and current Cloud Build Configuration files to find and fix the bug.

Suggested Answer: *B*
Official Answer: B

Community vote distribution
D (100%)

Charun Highly Voted 4 years, 7 months ago
D is correct
upvoted 21 times

ralf_cc Highly Voted 4 years, 8 months ago
D - find out what's changed in the build spec
upvoted 12 times

habla2019pasta Most Recent 1 year, 9 months ago
Option D - Figure out what was the recent change
upvoted 1 times

jomonkp 2 years, 2 months ago
Option D
upvoted 2 times

aswani 2 years, 7 months ago
are these questions relevant yet? with the new exam version launched in may
upvoted 3 times

Greg123123 3 years, 1 month ago
D is correct. Honestly I couldn't relate this to any SRE practices. Anyone can give me a hints? I could deduce the answer though. A: It didn't mention anything urgent, so no need to mitigate in this ugly way. B: miss the point. It worked perfectly without CR, why change it? C: upload the yaml to CS won't identify the issue for you. D: this has to be correct. The question emphasize "code versioning" and "change in the Cloud Build YAML". We can easily identify the problem through comparison.
upvoted 2 times

floppino 3 years, 1 month ago
Ans: D Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/
upvoted 2 times

EvelynSU 3 years, 2 months ago
D is correct
upvoted 1 times

GCP72 3 years, 6 months ago
Answer is D
upvoted 1 times

Ananda 3 years, 9 months ago
Submitted D in the examHow many questions are from this site for exam?
upvoted 2 times

Ananda 3 years, 9 months ago

Selected Answer: D

Submitted D in the exam

upvoted 2 times

akshay_jadhav 3 years, 9 months ago

How many questions are from this site for exam?

upvoted 2 times

vijaigcp 4 years ago

Would go with D

upvoted 3 times

Shasha1 4 years, 2 months ago

B is correct artifact can not push to Docker hub only docker images possible. therefore, need to push the artifacts to google cloud container registry not to the public docker hub. https://cloud.google.com/build/docs/interacting-with-dockerhub-images

upvoted 5 times

devopsbatch 4 years, 8 months ago

B push images into container registry instead docker hub

upvoted 1 times

devopsbatch 4 years, 8 months ago

B push images into container registry instead docker hub"After making a change in the Cloud Build YAML configuration, you notice that no new artifacts are being built by the pipeline"- means something wrong on the recent change not with the image registry. correct answer should be - D

upvoted 1 times

akg001 4 years, 8 months ago

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

upvoted 9 times

Question #40 *Topic 1*

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

A. Identify engineers responsible for the incident and escalate to their senior management.

B. Ensure that test cases that catch errors of this type are run successfully before new software releases.

C. Follow up with the employees who reviewed the changes and prescribe practices they should follow in the future.

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

Suggested Answer: *C*

Official Answer: C

Community vote distribution

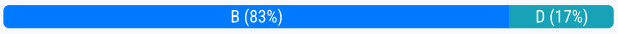

Charun Highly Voted 4 years, 7 months ago

B is correctAgree with B. I find this answer in "Site Reliability Engineering: How Google Runs Production Systems".

upvoted 26 times

Charun Highly Voted 4 years, 7 months ago

B is correct

upvoted 26 times

looseboy 4 years, 6 months ago

Agree with B. I find this answer in "Site Reliability Engineering: How Google Runs Production Systems".

upvoted 7 times

devopsbatch 4 years, 8 months ago

B make automation better

upvoted 10 times

devopsbatch 4 years, 8 months ago

B make automation betterAgree with you. IMO - B is the correct answer.

upvoted 10 times

akg001 4 years, 8 months ago

Agree with you. IMO - B is the correct answer.

upvoted 7 times

thewalker Most Recent 1 year, 8 months ago

B. Ensure that test cases that catch errors of this type are run successfully before new software releases. Here's why: Option A focuses on blaming individuals, which is not helpful in a postmortem. The goal is to learn from the incident and prevent future occurrences, not to assign fault. Option C is a good step, but it's not enough. While following up with reviewers is important, the primary focus should be on improving the testing process to catch errors before they reach production. Option D might be helpful in some cases, but it's not a general solution. Not all incidents require immediate escalation, and relying solely on on-call teams to make that decision can lead to unnecessary escalations.Option B addresses the root cause of the incident by ensuring that similar errors are caught during testing. This will help to prevent future incidents and improve the overall reliability of the system. Here are some additional steps you can take to prevent future incidents: Implement a strong continuous integration and continuous delivery (CI/CD) pipeline. This will help to automate the testing process and ensure that new code is tested thoroughly before it is deployed to production. Use a blameless postmortem process. This will encourage engineers to be open and honest about their mistakes, which will help to identify and fix problems more quickly. Invest in training and education for your engineers. This will help them to understand the importance of reliability and how to write code that is less prone to errors. By taking these steps, you can help to prevent future incidents and improve the overall reliability of your system.Additional Considerations It's important to note that there is no single solution that will prevent all incidents. However, by taking a proactive approach and focusing on improving your testing and development processes, you can significantly reduce the risk of future incidents. Here are some additional resources that you may find helpful: Site Re

upvoted 1 times

thewalker 1 year, 8 months ago

Selected Answer: B

B. Ensure that test cases that catch errors of this type are run successfully before new software releases.

Here's why:

Option A focuses on blaming individuals, which is not helpful in a postmortem. The goal is to learn from the incident and prevent future occurrences, not to assign fault. Option C is a good step, but it's not enough. While following up with reviewers is important, the primary focus should be on improving the testing process to catch errors before they reach production. Option D might be helpful in some cases, but it's not a general solution. Not all incidents require immediate escalation, and relying solely on on-call teams to make that decision can lead to unnecessary escalations.
upvoted 1 times

thewalker 1 year, 8 months ago

Option B addresses the root cause of the incident by ensuring that similar errors are caught during testing. This will help to prevent future incidents and improve the overall reliability of the system.

Here are some additional steps you can take to prevent future incidents:

Implement a strong continuous integration and continuous delivery (CI/CD) pipeline. This will help to automate the testing process and ensure that new code is tested thoroughly before it is deployed to production. Use a blameless postmortem process. This will encourage engineers to be open and honest about their mistakes, which will help to identify and fix problems more quickly. Invest in training and education for your engineers. This will help them to understand the importance of reliability and how to write code that is less prone to errors. By taking these steps, you can help to prevent future incidents and improve the overall reliability of your system.
upvoted 1 times

thewalker 1 year, 8 months ago

Option B addresses the root cause of the incident by ensuring that similar errors are caught during testing. This will help to prevent future incidents and improve the overall reliability of the system. Here are some additional steps you can take to prevent future incidents: Implement a strong continuous integration and continuous delivery (CI/CD) pipeline. This will help to automate the testing process and ensure that new code is tested thoroughly before it is deployed to production. Use a blameless postmortem process. This will encourage engineers to be open and honest about their mistakes, which will help to identify and fix problems more quickly. Invest in training and education for your engineers. This will help them to understand the importance of reliability and how to write code that is less prone to errors. By taking these steps, you can help to prevent future incidents and improve the overall reliability of your system.Additional Considerations It's important to note that there is no single solution that will prevent all incidents. However, by taking a proactive approach and focusing on improving your testing and development processes, you can significantly reduce the risk of future incidents. Here are some additional resources that you may find helpful: Site Reliability Engineering (SRE) Book: https://landing.google.com/sre/book.html Google Cloud Reliability Engineering: https://cloud.google.com/solutions/reliability-engineering Postmortem Best Practices: https://sre.google/sre-book/postmortems/
upvoted 1 times

thewalker 1 year, 8 months ago

Additional Considerations It's important to note that there is no single solution that will prevent all incidents. However, by taking a proactive approach and focusing on improving your testing and development processes, you can significantly reduce the risk of future incidents. Here are some additional resources that you may find helpful: Site Reliability Engineering (SRE) Book: https://landing.google.com/sre/book.html Google Cloud Reliability Engineering: https://cloud.google.com/solutions/reliability-engineering Postmortem Best Practices: https://sre.google/sre-book/postmortems/
upvoted 1 times

habla2019pasta 1 year, 9 months ago

B - Blameless postmortem
upvoted 1 times

alpha_canary 2 years ago

The focus should be on improving processes and systems, not blaming individuals.
upvoted 1 times

jomonkp 2 years, 2 months ago

Option B
upvoted 1 times

Greg123123 3 years, 1 month ago

It is B. The main point here is "trigged by software change". This make B make sense and right to the point. D is not correct because it doesn't prevent this from happen again...
upvoted 1 times

floppino 3 years, 1 month ago

Ans: B Exam passed and taken on 19/12/2022, 50/50 from this dump without buying the full access and looking for 'devops' word here: https://www.examtopics.com/discussions/google/1/

upvoted 1 times

GCP72 3 years, 6 months ago

Answer is B

upvoted 2 times

devOpsForFun 3 years, 10 months ago

D is correct. B is too narrow

upvoted 2 times

PhilipKoku 4 years ago

B - Blameless port-mortens. Focus on the process and not in the people.

upvoted 3 times

cyrus86 4 years ago

I will go with C here

upvoted 1 times

meetplanet 4 years, 2 months ago

I go for D based on the on-call doc

upvoted 1 times

Vallal 4 years, 2 months ago

c is the correct answer

upvoted 4 times

muk5658 4 years, 2 months ago

I would go with B as they are saying incident is triggered by a software change, if this has been tested thoroughly it would have been avoided.

upvoted 2 times

MF2C 4 years, 5 months ago

B is better over D because D is incident management, B is RCA.

upvoted 2 times

kubosuke 4 years, 6 months ago

I think D is correct. > you want to prevent severe incidents from happening in the future. we need to chose the answer that 'prevent severe incidents'. B: create test cases can reduce the possibility of incidents, but it doesn't ease the 'severe incidents'. to minimize the effect of incidents, we need to make a consensus to manage incidents. for these backgrounds, the answer is D.

upvoted 4 times

Question #41 *Topic 1*

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 two.)

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.

Suggested Answer: *B, D*

Official Answer: BD

Community vote distribution

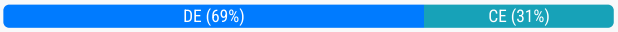

ralf_cc Highly Voted 4 years, 2 months ago

DE - synthetic transactions

upvoted 29 times

ralf_cc Highly Voted 4 years, 2 months ago

DE - synthetic transactionsA&D. Why E isnt correct - Synthetic transaction already provides the capability mentioned in E - "Use current and historic Request Logs to trace customer interaction with the application". Instead, option A - to review and add additional metrics makes sense !!It says "measure application reliability from a user perspective" - you need synthetic measurements from outside the infrastructure (D). To simulate the synthetic journey mentioned, you first need to understand the current user behaviours using the 'request' logs (E).

upvoted 29 times

Biden 3 years, 8 months ago

A&D. Why E isnt correct - Synthetic transaction already provides the capability mentioned in E - "Use current and historic Request Logs to trace customer interaction with the application". Instead, option A - to review and add additional metrics makes sense !!It says "measure application reliability from a user perspective" - you need synthetic measurements from outside the infrastructure (D). To simulate the synthetic journey mentioned, you first need to understand the current user behaviours using the 'request' logs (E).

upvoted 3 times

Biden 3 years, 8 months ago

A&D. Why E isnt correct - Synthetic transaction already provides the capability mentioned in E - "Use current and historic Request Logs to trace customer interaction with the application". Instead, option A - to review and add additional metrics makes sense !!

upvoted 3 times

Rupo7 3 years, 1 month ago

It says "measure application reliability from a user perspective" - you need synthetic measurements from outside the infrastructure (D). To simulate the synthetic journey mentioned, you first need to understand the current user behaviours using the 'request' logs (E).

upvoted 5 times

Charun 4 years, 1 month ago

CE ans

upvoted 12 times

Charun 4 years, 1 month ago

CE ansWeb proxy is not a reverse proxy, it is a forward proxy - a type of server that runs at the client side. If you have clients from all over the world, how will you collect their proxy logs? So it cannot be C - D&E should be the answer which don't require any engineering changes "in the" application.

upvoted 12 times

cetanx 3 years, 11 months ago

Web proxy is not a reverse proxy, it is a forward proxy - a type of server that runs at the client side. If you have clients from all over the world, how will you collect their proxy logs? So it cannot be C - D&E should be the answer which don't require any engineering changes "in the" application.

upvoted 10 times

gaurav517 Most Recent 8 months, 3 weeks ago

Option D: Create synthetic clients to simulate user journey No application changes required Synthetic monitoring/active probing from user perspective Simulates real user interactions and journeys Measures end-to-end user experience Can test from multiple geographic locations Proactive monitoring (detects issues before users report) Strong option for user perspective Option E: Use current and historic Request Logs to trace customer interaction Uses existing logs (no code changes) Request logs capture actual user behavior Can trace real user journeys through the application Historic data shows trends and patterns Strong option for real user monitoring

upvoted 1 times

jomonkp 1 year, 8 months ago

Option C and E

upvoted 1 times

enter_co 2 years, 9 months ago

A - Adding metrics doesn't necessarily reflect reliability from a user perspective' B - Modifying the code is against the requirement C - Analyzing proxy logs doesn't connect the findings to a user perspective From the request logs, user journeys can be recreated (E) and the replays can be fed to synthetic clients to simulate/replay (D).

upvoted 4 times

paul223344 3 years, 4 months ago

Answer: CE A. Review current application metrics and add new ones as needed. ==> (x) "add new ones" needs engineering changes B. Modify the code to capture additional information for user interaction. ==> (x) "Modify the code" needs engineering changes C. Analyze the web proxy logs only and capture response time of each request. ==> (O) no engineering changes D. Create new synthetic clients to simulate a user journey using the application. => (x) "Create new synthetic client" needs engineering changes. E. Use current and historic Request Logs to trace customer interaction with the application. ==> (O) no engineering changes

upvoted 5 times

PhilipKoku 3 years, 6 months ago

D & E - Reliability review using synthetic transactions and customer journeys from logs.

upvoted 3 times

cloudbee 3 years, 7 months ago

This two option doesn't require engineering changes into the application. Web Proxy logs is a forward proxy thing so it present in client side. others need changes

upvoted 3 times

guid1984 3 years, 8 months ago

C& E. C-> a web proxy relays URL requests from clients to a server. Analyzing web proxy logs can give unobtrusive insights into the browsing behavior of users

upvoted 1 times

rspstudent 3 years, 8 months ago

Selected Answer: DE

upvoted 2 times

Freemiums 3 years, 8 months ago

Selected Answer: CE

C and E should be the correct answer, since the question state that "without making ANY engineering changes"

upvoted 4 times

Freemiums 3 years, 8 months ago

C and E should be the correct answer, since the question state that "without making ANY engineering changes"with making any engineering changes to IT (i.e the application itself)

upvoted 4 times

not_thanos 3 years, 8 months ago

with making any engineering changes to IT (i.e the application itself)

upvoted 4 times

Biden 3 years, 9 months ago

A&D should be the correct answer: D - Easy/default choice since this doesnt need any changes on the app A - Since the ask is to "measure reliability" - review and add the appropriate SLIs. Option E -tracing is needed only for collecting/troubleshooting latency issues - hence not the correct choice

upvoted 4 times

raf2121 4 years ago

To me , it looks C & E B & D, need engineering changes and investment, hence I ruled these two out This leaves us with A, C and E - The question is asking for "Need to measure Application Reliability from User perspective without Engineering Changes" - This rules out A, as talks about adding new metric but not stating which one Considering high-traffic web application as "Request Driven Services" - Two of the suggested SLI's are Availability and Latency - C & E to me covers Latency https://cloud.google.com/architecture/adopting-slos?hl=en

upvoted 4 times

gcpengineer 4 years ago

A & D is the answer

upvoted 2 times

kubosuke 4 years ago

maybe C,D,E, and it looks like create synthetic client looks like "engineering change", so I choose C and E .

upvoted 1 times

kubosuke 4 years ago

maybe C,D,E, and it looks like create synthetic client looks like "engineering change", so I choose C and E .umm but > C. Analyze the web proxy logs 'only' and the phrase 'only' bothers me... we can know the reliability by getting the status of web applications by web proxy logs(I mean HTTP status is good/bad), but we don't need to view these logs 'only'...C cant be the ans , as it talks about response time. we need answer for reliability

upvoted 1 times

kubosuke 4 years ago

umm but > C. Analyze the web proxy logs 'only' and the phrase 'only' bothers me... we can know the reliability by getting the status of web applications by web proxy logs(I mean HTTP status is good/bad), but we don't need to view these logs 'only'...C cant be the ans , as it talks about response time. we need answer for reliability

upvoted 2 times

kubosuke 4 years ago

umm but > C. Analyze the web proxy logs 'only' and the phrase 'only' bothers me... we can know the reliability by getting the status of web applications by web proxy logs(I mean HTTP status is good/bad), but we don't need to view these logs 'only'...

upvoted 2 times

gcpengineer 4 years ago

C cant be the ans , as it talks about response time. we need answer for reliability

upvoted 4 times

akg001 4 years, 2 months ago

for me it looks like C,E the correct answer. E for sure.Same here, because no engineering effort should be required. All others definitely need engineering effort.

upvoted 5 times

akg001 4 years, 2 months ago

for me it looks like C,E the correct answer. E for sure.

upvoted 5 times

holahola 4 years, 1 month ago

Same here, because no engineering effort should be required. All others definitely need engineering effort.

upvoted 2 times

rinkeshgala1 4 years, 2 months ago

may be A&E as these options may not need any additional engineering efforts

upvoted 2 times

Question #42 *Topic 1*

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

A. Grant the team members the IAM role of logging.configWriter on Cloud IAM.

B. Configure Access Context Manager to allow only these members to export logs.

C. Create and grant a custom IAM role with the permissions logging.sinks.list and logging.sink.get.

D. Create an Organizational Policy in Cloud IAM to allow only these members to create log exports.

Suggested Answer: *A*

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

Community vote distribution

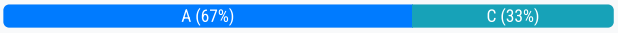

rinkeshgala1 Highly Voted 4 years, 2 months ago

option A

upvoted 27 times

rinkeshgala1 Highly Voted 4 years, 2 months ago

option Aagree for the A.I understand that option A gives the ability to export logs, but isn't C the best option following the least privilege principle since the question only says that the team members needs to export logs and not to write them?

upvoted 27 times

akg001 4 years, 2 months ago

agree for the A.I understand that option A gives the ability to export logs, but isn't C the best option following the least privilege principle since the question only says that the team members needs to export logs and not to write them?

upvoted 5 times

akg001 4 years, 2 months ago

agree for the A.

upvoted 5 times

irocketsoldier 3 years, 9 months ago

I understand that option A gives the ability to export logs, but isn't C the best option following the least privilege principle since the question only says that the team members needs to export logs and not to write them?

upvoted 6 times

Manh 3 years, 9 months ago

It's should be C. least privilege

The question is ask about export log and does not mention about read and write log Option A give too many permission Logs Configuration Writer (roles/logging.configWriter) Provides permissions to read and write the configurations of logs-based metrics and sinks for exporting logs.

logging.buckets.create logging.buckets.delete logging.buckets.get logging.buckets.list logging.buckets.undelete logging.buckets.update logging.cmekSettings.* logging.exclusions.* logging.locations.* logging.logMetrics.* logging.logServiceIndexes.* logging.logServices.* logging.logs.list logging.notificationRules.* logging.operations.* logging.sinks.* logging.views.create logging.views.delete logging.views.get logging.views.list logging.views.update resourcemanager.projects.get resourcemanager.projects.list

upvoted 8 times

Manh 3 years, 9 months ago

It's should be C. least privilege The question is ask about export log and does not mention about read and write log Option A give too many permission Logs Configuration Writer (roles/logging.configWriter) Provides permissions to read and write the configurations of logs-based metrics and sinks for exporting logs. logging.buckets.create logging.buckets.delete logging.buckets.get logging.buckets.list logging.buckets.undelete logging.buckets.update logging.cmekSettings.* logging.exclusions.* logging.locations.* logging.logMetrics.* logging.logServiceIndexes.* logging.logServices.* logging.logs.list logging.notificationRules.* logging.operations.* logging.sinks.* logging.views.create logging.views.delete logging.views.get logging.views.list logging.views.update resourcemanager.projects.get resourcemanager.projects.listability to use sinks Add logging.sinks.{list, create, get, update, delete} , the list, get function can only have view permission. can not create sinks to export logs. u need create

sink to export logs.After review again, Ans A had enough permission to export log
https://cloud.google.com/logging/docs/routing/overviewlogging.sinks.create is needed to export logs - this is why C is wrong
upvoted 8 times

hanweiCN 2 years, 8 months ago
ability to use sinks Add logging.sinks.{list, create, get, update, delete} , the list, get function can only have view permission. can not create sinks to export logs. u need create sink to export logs.
upvoted 1 times

Manh 3 years, 9 months ago
After review again, Ans A had enough permission to export log https://cloud.google.com/logging/docs/routing/overviewlogging.sinks.create is needed to export logs - this is why C is wrong
upvoted 2 times

Manh 3 years, 9 months ago
After review again, Ans A had enough permission to export log https://cloud.google.com/logging/docs/routing/overview
upvoted 2 times

Goram113 3 years, 8 months ago
logging.sinks.create is needed to export logs - this is why C is wrong
upvoted 11 times

BlackCoffee Most Recent 1 month, 1 week ago
Not C Custom roles can be used, but the permissions logging.sinks.list and logging.sink.get only allow team members to view existing sinks, not create or modify them to set up a new export. To create exports, permissions like logging.sinks.create are required.
upvoted 1 times

kleranti 1 year, 1 month ago
option c
upvoted 1 times

jomonkp 1 year, 8 months ago
option C
upvoted 2 times

GaneshSurwase 2 years, 7 months ago
correct ans is A Logs Configuration Writer (roles/logging.configWriter) Provides permissions to read and write the configurations of logs-based metrics and sinks for exporting logs.
upvoted 1 times

Ananda 3 years, 3 months ago
Option A
upvoted 1 times

Sreedharveluru 3 years, 3 months ago
There is no such thing called logging.sink.get
upvoted 3 times

Epic_rose 3 years, 4 months ago
roles/logging.configWriter (Logs Configuration Writer) gives you the permissions to create log-based metrics, exclusions, buckets, and views, and to use sinks. To use the Logs Explorer (console) for these actions, add roles/logging.viewer.
upvoted 1 times

Shasha1 3 years, 5 months ago
A Logs configuration writer can access to configure log exporting and metrics
upvoted 1 times

ric79 3 years, 5 months ago
What is the minimum set of privs in order to export logs?
upvoted 1 times

zygomar 3 years, 5 months ago
in addition to other comments here, C would be too restrictive. User new to have logs list permissions at least to know which logs to export. Goram113 also indicate that logging.sinks.create is needed to export logs hence why C is wrong
upvoted 1 times

PhilipKoku 3 years, 6 months ago
C - Use principle of minimum access required to fulfill the requirements
upvoted 1 times

vijaigcp 3 years, 6 months ago

Agree with A

upvoted 1 times

cloudbee 3 years, 7 months ago

https://cloud.google.com/logging/docs/export/configure_export_v2#before-you-begin A is the answer

upvoted 3 times

giammydell 3 years, 9 months ago

but C could follow the least privilege priciple

upvoted 3 times

not_thanos 3 years, 8 months ago

Write answer is A as stated in the documentation here https://cloud.google.com/logging/docs/export/configure_export_v2#before-you-begin "Note that this guide describes creating and managing sinks at the Cloud project level, but you can create sinks (non-aggregated) for billing accounts, folders, and organizations. As you get started, ensure the following: You have a Google Cloud project with logs that you can see in the Logs Explorer. You have one of the following IAM roles for the source Cloud project from which you're routing logs. Owner (roles/owner) Logging Admin (roles/logging.admin) Logs Configuration Writer (roles/logging.configWriter) The permissions contained in these roles allow you to create, delete, or modify sinks. For information on setting IAM roles, see the Logging Access control guide."

upvoted 4 times

giammydell 3 years, 9 months ago

but C could follow the least privilege priciplel agree. logging.configWriter (answer A) gives too much power to the team members. We only need to give them the rights to export, not change the whole logging configuration. C is ok.

upvoted 3 times

sticky 3 years, 10 months ago

A is correct

Logs Configuration Writer (roles/logging.configWriter) - Provides permissions to read and write the configurations of logs-based metrics and sinks for exporting logs. https://cloud.google.com/logging/docs/access-control

upvoted 4 times

Trony 3 years, 9 months ago

I agree. logging.configWriter (answer A) gives too much power to the team members. We only need to give them the rights to export, not change the whole logging configuration. C is ok.

upvoted 1 times

Question #43 *Topic 1*

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

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

B. Use a Binary Authorization policy that includes the whitelist name pattern gcr.io/altostrat-images/.

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

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

Suggested Answer: *D*
Official Answer: D

Community vote distribution
B (100%)

rinkeshgala1 Highly Voted 4 years, 8 months ago
option B
upvoted 25 times

rinkeshgala1 Highly Voted 4 years, 8 months ago
option Bagree for answer-B
upvoted 25 times

akg001 4 years, 8 months ago
agree for answer-B
upvoted 5 times

Charun 4 years, 7 months ago
B is correct
upvoted 12 times

alpha_canary Most Recent 2 years ago
https://cloud.google.com/binary-authorization/docs/key-concepts#allowlist_patterns
upvoted 1 times

jomonkp 2 years, 2 months ago
Option B
upvoted 2 times

zanhsieh 2 years, 7 months ago
D. See admissionWhitelistPatterns in: https://cloud.google.com/binary-authorization/docs/example-policies#add_exempt_imagesShould be B not D.are these questions still valid?
upvoted 1 times

zanhsieh 2 years, 7 months ago
Selected Answer: B
D. See admissionWhitelistPatterns in: https://cloud.google.com/binary-authorization/docs/example-policies#add_exempt_images
upvoted 1 times

zanhsieh 2 years, 7 months ago
Should be B not D.are these questions still valid?
upvoted 1 times

zanhsieh 2 years, 7 months ago
Should be B not D.
upvoted 1 times

aswani 2 years, 7 months ago
are these questions still valid?
upvoted 1 times

Halimb 3 years, 7 months ago

See https://cloud.google.com/binary-authorization/docs/example-policies

upvoted 2 times

Ananda 3 years, 9 months ago

B is the answer

upvoted 2 times

PhilipKoku 4 years ago

B - Binary authorisation is the answer…

upvoted 2 times

vijaigcp 4 years ago

Answer B - Binary Authorization

upvoted 1 times

Shasha1 4 years, 2 months ago

D https://cloud.google.com/container-registry/docs/using-with-google-cloud-platform

upvoted 3 times

tycho 4 years, 4 months ago

B) https://cloud.google.com/binary-authorization/docs/cloud-build

upvoted 4 times

Question #44 *Topic 1*

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.

Suggested Answer: *B*
Official Answer: B

Community vote distribution
C (100%)

Charun Highly Voted 4 years, 1 month ago
Option C
upvoted 13 times

kubosuke Highly Voted 4 years ago
C is correct A. Configure the horizontal pod autoscaler to use the average response time from the Liveness and Readiness Probes. --> using health check as a trigger of scaling is weird. if the response time of the health check is delayed, it may be caused by resources issues such as CPU, memories, and so on. so you should use such values as SLIs. B. Configure the vertical pod autoscaler in GKE and enable the cluster autoscaler to scale the cluster as pods expand. --> it doesn't referred to pod autoscaling. D. Expose the NGINX stats endpoint and configure the horizontal pod autoscaler to use the request metrics exposed by the NGINX deployment. --> if you use request metrics as SLIs, you should use custom metrics as SLIs. it is a little bit redundant.
upvoted 13 times

jomonkp Most Recent 1 year, 8 months ago
option C
upvoted 1 times

Wael216 2 years, 5 months ago
C. Install the Stackdriver custom metrics adapter and configure a horizontal pod autoscaler to use the number of requests provided by the GCLB. To scale the deployment of the application's frontend using an appropriate Service Level Indicator (SLI), we need to monitor the traffic coming to the application. One way to do this is to install the Stackdriver custom metrics adapter, which provides visibility into GCLB metrics such as request counts, bytes sent and received, and active connections. We can then configure a horizontal pod autoscaler (HPA) to scale the number of pods based on the request count coming through the GCLB, which will help to ensure that our application is always available to handle the incoming traffic.
upvoted 3 times

pradoUA 2 years, 9 months ago
I will go with C
upvoted 1 times

ramzez4815 2 years, 11 months ago
C is the correct answer as per Google documentation
upvoted 2 times

khushboo93s 3 years, 2 months ago
C is correct
upvoted 1 times

Murty549 3 years, 7 months ago
Option B is incorrect because there will be no benefit in vertically scaling the front end if the requests are very high. Because in that case the network is the bottleneck and not the instance resources. In my opinion the Answer is C.

upvoted 2 times

cloudbee 3 years, 7 months ago

c looks more feasible to me. but also B is in small favour. Horizontal pod autoscaler will scale based on the custom metrics such as requests per second(i.e. no. of requests). but vertical autoscaling is also useful feature for the frontend application. as it requires more resource needs if the traffic is high. Since vertical autoscaling will first delete the pod and adjust the cpu and memory to recreate the pod, It can cause downtime for that duration and not recommended. Hence, More inclined towards answer C. Still, reply my answer with your explanation please https://cloud.google.com/kubernetes-engine/docs/concepts/horizontalpodautoscaler#overview

upvoted 1 times

Shasha1 3 years, 8 months ago

B is correct front-end web applications Scale based on the number of incoming request. so need vertical scaling Back-end Batch Processing (Scale Horizontally) reference https://docs.rightscale.com/faq/What_is_auto-scaling.html

upvoted 3 times

TNT87 3 years, 11 months ago

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

upvoted 2 times

IamSuren 4 years ago

According to Google https://cloud.google.com/kubernetes-engine/docs/tutorials/autoscaling-metrics#cpu_1 Horizontal Pod Autoscalers can scale based on CPU utilization natively, so the Custom Metrics Adapter is not needed, therefor C doesn't fit .

upvoted 2 times

ralf_cc 4 years, 2 months ago

C - https://cloud.google.com/kubernetes-engine/docs/tutorials/autoscaling-metrics You want to scale horizontally

upvoted 4 times

rinkeshgala1 4 years, 2 months ago

Option B

upvoted 1 times

Question #45 *Topic 1*

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

A. Operations Lead

B. Engineering Lead

C. Communications Lead

D. Customer Impact Assessor

E. External Customer Communications Lead

Suggested Answer: *A, E*
Official Answer: AE

Community vote distribution
AC (100%)

syslog Highly Voted 4 years, 2 months ago
AC

https://sre.google/workbook/incident-response/ "The main roles in incident response are the Incident Commander (IC), Communications Lead (CL), and Operations or Ops Lead (OL)."
upvoted 34 times

syslog Highly Voted 4 years, 2 months ago
AC https://sre.google/workbook/incident-response/ "The main roles in incident response are the Incident Commander (IC), Communications Lead (CL), and Operations or Ops Lead (OL)."could be. A and C. why not E? this is customer impacting incident so should be an external customer comm lead as well.
upvoted 34 times

akg001 4 years, 2 months ago
could be. A and C. why not E? this is customer impacting incident so should be an external customer comm lead as well.
upvoted 2 times

Charun 4 years, 1 month ago
AC for sure
upvoted 13 times

zanhsieh Most Recent 1 year ago
AC. External communication may be needed for customer-impacting incidents, but it's a secondary priority compared to resolving the incident and coordinating internally. This role is typically activated later, if necessary, and is not a core SRE incident management role during the initial response. That's why prefer C not E.
upvoted 1 times

jomonkp 1 year, 8 months ago
option A and C
upvoted 1 times

maxdanny 1 year, 9 months ago
AC Main Roles in Incident Response
upvoted 1 times

AzureDP900 2 years, 9 months ago
A,C is right The main roles in incident response are the Incident Commander (IC), Communications Lead (CL), and Operations or Ops Lead (OL). IMAG organizes these roles into a hierarchy: the IC leads the incident response, and the CL and OL report to the IC.
upvoted 1 times

ramzez4815 2 years, 11 months ago
AC, been there many times in real situations
upvoted 1 times

PhilipKoku 3 years, 6 months ago

A - Operations lead: fixes the issue C - Communications lead: keeps stakeholders informed

upvoted 2 times

vijaigcp 3 years, 6 months ago

Answer A & C as per Google SRE

upvoted 1 times

cloudbee 3 years, 7 months ago

It's mentioned in sre book by Google.

upvoted 1 times

Biden 3 years, 9 months ago

Who gave AE as the correct answer ? why not this tool display/reveal the correct answer from the "Highly Voted" from the discussion thread for each Q?

upvoted 2 times

Luvero 3 years, 9 months ago

A, C - other roles are appropriate ..

upvoted 3 times

Question #46 *Topic 1*

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.

Suggested Answer: *D*
Official Answer: D

Community vote distribution



kubosuke Highly Voted 4 years, 6 months ago
C 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.
upvoted 18 times

Charun Highly Voted 4 years, 7 months ago
C option
upvoted 7 times

alpha_canary 2 years ago
Stackdriver (now Operations Suite) natively supports SMS notification channel, so going through external tool route or slack route is not cloud native approach. https://cloud.google.com/monitoring/support/notification-options#sms
upvoted 1 times

jomonkp 2 years, 2 months ago
option C
upvoted 1 times

jeffersonkozak 2 years, 8 months ago
Option C suggests ensuring that team members set their SMS/phone numbers in their Stackdriver Profile. This approach is the simplest and most straightforward of the options presented. It requires no additional integration or configuration, and team members can easily manage their contact information in their Stackdriver profile. However, it does require team members to have access to and familiarity with Stackdriver, which may not be the case for all members of the team.Same argument goes for Slack integration as well. Team members needs to setup Slack access, channel, etc. And it seems redundant also. Final goal is to send alert message using SMS, which Google Cloud has in-build support.
upvoted 2 times

jeffersonkozak 2 years, 8 months ago
Selected Answer: D
Option C suggests ensuring that team members set their SMS/phone numbers in their Stackdriver Profile. This approach is the simplest and most straightforward of the options presented. It requires no additional integration or configuration, and team members can easily manage their contact information in their Stackdriver profile. However, it does require team members to have access to and familiarity with Stackdriver, which may not be the case for all members of the team.
upvoted 2 times

jinaldesailive Most Recent 1 year, 11 months ago
Same argument goes for Slack integration as well. Team members needs to setup Slack access, channel, etc. And it seems redundant also. Final goal is to send alert message using SMS, which Google Cloud has in-build support.
upvoted 1 times

pradoUA 3 years, 3 months ago

Had this question on exam 25.10.2022 and originally it says: C. Ensure that your team members set their SMS/phone numbers in their Cloud Monitoring. Select the SMS notification option for each alerting policy and then select the appropriate SMS/phone numbers from the list. Hence definitely C is the answer

upvoted 3 times

Ananda 3 years, 9 months ago

Submit C

upvoted 2 times

PhilipKoku 4 years ago

C - Google Cloud Monitoring (previously known as Stackdriver) supports SMS notifications

upvoted 3 times

cyrus86 4 years ago

I dont know if there is something called - Stackdriver Profile . The SMS can be configured in notification channel. But I dont see that in the options, hence the correct answer should be DYes, I also don't think there is anything such as stackdriver profile.

upvoted 2 times

cyrus86 4 years ago

Selected Answer: D

I dont know if there is something called - Stackdriver Profile . The SMS can be configured in notification channel. But I dont see that in the options, hence the correct answer should be D

upvoted 2 times

Aiffone 3 years, 9 months ago

Yes, I also don't think there is anything such as stackdriver profile.

upvoted 2 times

cloudbee 4 years, 1 month ago

Although stackdriver profile is not an named option in gcp. It looks like that to add phone no in stackdriver. so pretty straigtforward right answer is C

upvoted 1 times

Shasha1 4 years, 2 months ago

i think answer is D There is no Stackdriver Profile exists to add their phone number. only possible to add the phone number in the notification channel under sms category

upvoted 3 times

utwo 4 years, 2 months ago

is option C

upvoted 2 times

TNT87 4 years, 5 months ago

C https://cloud.google.com/monitoring/support/notification-options

upvoted 6 times

Question #47 *Topic 1*

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. ג€¢ In your application, create a metric with a metricKind set to DELTA and a valueType set to DOUBLE. ג€¢ In Stackdriver's Metrics Explorer, use a Stacked Bar graph to visualize the metric.

B. ג€¢ In your application, create a metric with a metricKind set to CUMULATIVE and a valueType set to DOUBLE. ג€¢ In Stackdriver's Metrics Explorer, use a Line graph to visualize the metric.

C. ג€¢ In your application, create a metric with a metricKind set to GAUGE and a valueType set to DISTRIBUTION. ג€¢ In Stackdriver's Metrics Explorer, use a Heatmap graph to visualize the metric.

D. ג€¢ In your application, create a metric with a metricKind set to METRIC_KIND_UNSPECIFIED and a valueType set to INT64. ג€¢ In Stackdriver's Metrics Explorer, use a Stacked Area graph to visualize the metric.

Suggested Answer: *A*
Official Answer: A

Community vote distribution
C (89%) A (11%)

raf2121 Highly Voted 4 years, 6 months ago
Answer C Reference : https://cloud.google.com/monitoring/api/v3/kinds-and-types?hl=en GAUGE Metric : In which value measures a specific instant in time DELTA Metric : In which the value measures the change since it was last recorded CUMULATIVE metric : In which the value constantly increases over time Question asks, "Latency of HTTP responses" - This needs to be specific instant in time , which is GAUGE, hence C
upvoted 24 times

Charun Highly Voted 4 years, 7 months ago
C is correct
upvoted 11 times

alpha_canary Most Recent 2 years ago
GAUGE metrics record a value at a particular point in time and DISTRIBUTION captures distribution statistics. A Heatmap is a good way to visualize latencies across all users. https://cloud.google.com/monitoring/api/v3/kinds-and-types?hl=en#metric-kinds
upvoted 1 times

jomonkp 2 years, 2 months ago
option c
upvoted 1 times

jeffersonkozak 2 years, 8 months ago
C is partially correct, GAUGE with DISTRIBUTUION, but "use a Heatmap graph to visualize the metric" will not help to visualiaze the Latencies of all users, will help to identify regions, but it;s not a question
upvoted 1 times

AzureDP900 3 years, 3 months ago
C is right as per GCP documentation. A gauge metric, in which the value measures a specific instant in time. For example, metrics measuring CPU utilization are gauge metrics; each point records the CPU utilization at the time of measurement. Another example of a gauge metric is the current temperature.
upvoted 1 times

Ananda 3 years, 9 months ago
Submitted C
upvoted 2 times

PhilipKoku 4 years ago
C - Gauge to see latency distribution
upvoted 2 times

cloudbee 4 years, 1 month ago
It's guage metric as it is asking latency of http responses.

upvoted 2 times

francisco_guerra 4 years, 7 months ago

Ans C: A: Incorrect: Stacked bar could be but latency its used with distributions B: Incorrect: If we used a cumulative the latency would be increasing in the graph. C: Correct: Latency is commonly measured as a distribution. D: Incorrect: METRIC_KIND_UNSPECIFIED is not the answer that gcp wanted. 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. Latency histogram_quantile(0.9, rate(http_request_duration_seconds_bucket[7d]) histogram_quantile(0.99, rate(http_request_duration_seconds_bucket[7d])) https://sre.google/workbook/implementing-slos/ https://cloud.google.com/architecture/adopting-slos/

upvoted 9 times

ralf_cc 4 years, 8 months ago

Option BAbsolutely not. Cumulative means a growing graph, whose latency values add up to the previous ones, and it doesn't make sense. It's C for sure.

upvoted 1 times

ralf_cc 4 years, 8 months ago

Option B

upvoted 1 times

syslog 4 years, 7 months ago

Absolutely not. Cumulative means a growing graph, whose latency values add up to the previous ones, and it doesn't make sense. It's C for sure.

upvoted 5 times

rinkeshgala1 4 years, 8 months ago

Option C

upvoted 5 times

Question #48 *Topic 1*

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. Import the Stackdriver Profiler package, and configure it to relay function timing data to Stackdriver for further analysis.

B. Import the Stackdriver Debugger package, and configure the application to emit debug messages with timing information.

C. Instrument the code using a timing library, and publish the metrics via a health check endpoint that is scraped by Stackdriver.

D. Install an Application Performance Monitoring (APM) tool in both locations, and configure an export to a central data storage location for analysis.

Suggested Answer: *B*
Official Answer: B

Community vote distribution
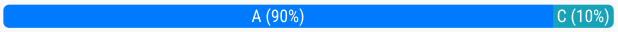

Charun Highly Voted 4 years, 7 months ago
Answer A
upvoted 23 times

rinkeshgala1 Highly Voted 4 years, 8 months ago
may be A, as profiler works both inside and outside of GCP.
upvoted 8 times

rinkeshgala1 4 years, 8 months ago
may be A, as profiler works both inside and outside of GCP.I agree. A is the right answer.
upvoted 8 times

akg001 4 years, 8 months ago
I agree. A is the right answer.
upvoted 3 times

gaurav517 Most Recent 8 months, 2 weeks ago
Given all non-ideal options, C is closer to ideal solution where "instrumentation" is meant broadly to include running a standard Prometheus system exporter (like node_exporter), which is then scraped by Cloud Monitoring, as this is the industry standard for hybrid metrics collection.
upvoted 1 times

alpha_canary 2 years ago
https://cloud.google.com/profiler/docs/profiling-external
upvoted 1 times

jomonkp 2 years, 2 months ago
option A
upvoted 1 times

PhilipKoku 4 years ago
A - Profiler for resource utilisation.
upvoted 2 times

cyrus86 4 years ago
Answer should be A
upvoted 2 times

cloudbee 4 years, 1 month ago
Profiler can be used outside GCP also it is used to measure cpu utilization by the application.
upvoted 3 times

ralf_cc 4 years, 8 months ago
It is A
upvoted 4 times

Question #49 *Topic 1*

You need to reduce the cost of virtual machines (VM) for your organization. After reviewing different options, you decide to leverage preemptible VM instances.
Which application is suitable for preemptible VMs?

A. A scalable in-memory caching system.

B. The organization's public-facing website.

C. A distributed, eventually consistent NoSQL database cluster with sufficient quorum.

D. A GPU-accelerated video rendering platform that retrieves and stores videos in a storage bucket.

Suggested Answer: *D*
Reference: https://cloud.google.com/preemptible-vms

Community vote distribution



Charun Highly Voted 5 years, 1 month ago
D is correctA A GPU-accelerated video rendering platform that retrieves and stores videos in a storage bucket: Video rendering requires a stable and powerful infrastructure with persistent storage, which is not provided by preemptible VMs. Additionally, GPUs are not available on all preemptible VM instances.https://cloud.google.com/compute/docs/instances/preemptible#preemptible-with-gpu
upvoted 17 times

Charun Highly Voted 5 years, 1 month ago
D is correct
upvoted 17 times

NiveusSol 3 years, 4 months ago
A A GPU-accelerated video rendering platform that retrieves and stores videos in a storage bucket: Video rendering requires a stable and powerful infrastructure with persistent storage, which is not provided by preemptible VMs. Additionally, GPUs are not available on all preemptible VM instances.https://cloud.google.com/compute/docs/instances/preemptible#preemptible-with-gpu
upvoted 3 times

NiveusSol 3 years, 4 months ago
A A GPU-accelerated video rendering platform that retrieves and stores videos in a storage bucket: Video rendering requires a stable and powerful infrastructure with persistent storage, which is not provided by preemptible VMs. Additionally, GPUs are not available on all preemptible VM instances.
upvoted 3 times

xarom 2 years, 12 months ago
https://cloud.google.com/compute/docs/instances/preemptible#preemptible-with-gpu
upvoted 3 times

devopsbatch 5 years, 2 months ago
ans: D
upvoted 6 times

devopsbatch 5 years, 2 months ago
ans: DAgree with you on Answer DD GPU-accelerated video rendering platform that retrieves and stores videos in a storage bucket: Video rendering requires a stable and powerful infrastructure with persistent storage, which is not provided by preemptible VMs. Additionally, GPUs are not available on all preemptible VM instances.
upvoted 6 times

akg001 5 years, 2 months ago
Agree with you on Answer D
upvoted 4 times

NiveusSol 3 years, 4 months ago
D GPU-accelerated video rendering platform that retrieves and stores videos in a storage bucket: Video rendering requires a stable and powerful infrastructure with persistent storage, which is not provided by preemptible VMs. Additionally, GPUs are not available on all preemptible VM instances.

upvoted 2 times

Srikantha Most Recent 2 months, 3 weeks ago

Correct Answer: D. A GPU-accelerated video rendering platform that retrieves and stores videos in a storage bucket. Why D is suitable for Preemptible VMs: Preemptible VMs (now called Spot VMs) are significantly cheaper (up to 80% discount) but can be terminated by Google at any time with only 30 seconds notice.

upvoted 1 times

gaurav517 8 months, 2 weeks ago

 Batch processing workload - perfect for preemptible VMs Stateless: Input from bucket, output to bucket Fault-tolerant: Can restart failed rendering jobs Job-based: Each video render is independent Interruptible: Can checkpoint progress or restart Cost-sensitive: GPU instances are expensive, 80% savings significant Not time-critical: Rendering can be queued/retried Perfect fit for preemptible VMs

upvoted 1 times

jfolleas_izivia 9 months, 2 weeks ago

Caches are stateless and disposable. If a cache instance is preempted, the data is simply loaded again from the primary data store (e.g., database) when requested, causing a slight latency spike but no Video rendering is a long-running, stateful batch job (the rendering progress). While the final output is safe in a storage bucket, preemption would waste significant GPU time and money by forcing the entire rendering job (or a large portion of it) to restart from the last checkpoint, minimizing the intended cost savings.

upvoted 1 times

dneves 1 year, 6 months ago

Preemptible VMs are short-lived and can be terminated by Google Cloud at any time if resources are needed elsewhere, so they are ideal for fault-tolerant applications that can handle interruptions without significant impact on functionality. Option C describes a distributed, eventually consistent NoSQL database. Eventually consistent systems can tolerate temporary unavailability of some nodes and reconcile the data once nodes are back online. Quorum-based systems (like many NoSQL databases) can continue to operate even if some VMs are preempted, as long as the required quorum of nodes is still available. Therefore, this setup is well-suited for preemptible VMs, as it can handle intermittent node loss.

upvoted 2 times

6a8c7ad 2 years ago

Has to be D. Definitely not A

upvoted 2 times

winston9 2 years, 2 months ago

cannot be A, if the VM is removed by the preemption process, we lose data in memory

upvoted 1 times

jinaldesailive 2 years, 5 months ago

It seems A & D both as suitable answer. But I'll go with A, as attaching GPU with preemptible VM will increase cost, and in this question the purpose of opting for preemptible VM is reducing cost. https://cloud.google.com/compute/docs/instances/preemptible#preemptible-with-gpu

upvoted 1 times

FI22 2 years, 5 months ago

Will go with D. Looking for cost effective. https://cloud.google.com/compute/docs/instances/preemptible#preemptible-with-gpu

upvoted 1 times

jomonkp 2 years, 8 months ago

Option D

upvoted 1 times

Jason_Cloud_at 2 years, 9 months ago

I would go with A compared to other options

upvoted 2 times

Watcharin_start 3 years, 3 months ago

Ans is D Video rendering service is like application type called Batch job. Therefore, we can use instance type preemptible for them. If they complete task, they could be destroy and generate new instance to work continuously next task.

upvoted 2 times

izekc 3 years, 5 months ago

C is more accurate

upvoted 3 times

eks4x 3 years, 8 months ago

Why not A? "A scalable in-memory caching system." In general a caching system is not critical to the function of an application. If the cache is down it will cause requests to have cache miss and query the DB instead. User requests will still get served albeit slower. In addition the answer specifies that the caching system is "scalable" reducing further the impact of 1 VM getting preempted, ie traffic can be automatically redirected to other cache replicas. To me all other answers seem to have a more severe impact on the user in case the VM is preempted.

upvoted 5 times

kisshs 3 years, 8 months ago

C is the right answer

upvoted 1 times

hanweiCN 3 years, 8 months ago

i think is C, database cluster is storage, and distributed, eventually consistence is resistant for the preempted. and sufficient quorum can ensure the DB transitions.

upvoted 2 times

Question #50 *Topic 1*

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

A. Configure the build system with protected branches that require pull request approval.

B. Use an Admission Controller to verify that incoming requests originate from approved sources.

C. Leverage Kubernetes Role-Based Access Control (RBAC) to restrict access to only approved users.

D. Enable binary authorization inside the Kubernetes cluster and configure the build pipeline as an attestor.

Suggested Answer: *C*
Official Answer: C

Community vote distribution
D (88%) A (12%)

francisco_guerra Highly Voted 4 years, 7 months ago
this question is a little bit strange, but first we need to remove the invalid answers B: Incorrect An admission controller is a piece of code that intercepts requests to the Kubernetes API server prior to persistence of the object, but after the request is authenticated and authorized. (its for security but not "enforce approvals") C: Incorrect, we need to "enforce approvals" roles apply in the cluster and Ops always could push to production without approval. A: Incorrect, for me this answer sound well but this does not sound that an answer for a gcp exam and this do not enforce the use of the pipeline. D: Correct, they cannot push code to production without approval because their images are not signed.
upvoted 27 times

francisco_guerra Highly Voted 4 years, 7 months ago
I win the exam today so this questions help me a lot
upvoted 15 times

jomonkp 2 years, 2 months ago
Option D
upvoted 1 times

WhyIronMan 3 years, 2 months ago
If you are not familiarised, Pull Requests are one way to bring changes from one branch (e.g. develop) into protected branches (e.g. master, main). 1. First you need to protect the production branch (e.g, master, main) 2. If a developer attempts to push new code to a production (now protected) it will trow a Permission denied error like this: remote: Permission denied to update branch master. To git.com:org/repository.git ! [remote rejected] master -> master (pre-receive hook declined) error: failed to push some refs to 'git.com:org/repository.git' 3. in order to push their code to production branch, the developers will need to open a Pull Request (If you are using GitLab, it is called Merge Request) and ask someone to Review and Approve your changes. 4. The pipeline points to the protected branch; any new code pushed trigger the pipeline, and runs the tests and then deploy it. This is how we do devops.but you missed the point here "The security auditor is concerned that developers or operators could circumvent". I see why you think pull request is the only way to circumvent the deployment process. But how about the operators? They have access to cluster and can simply redeploy it by some kubectl / cloud build command. So A is not correct.That said, I send the question back to you, how about the developers?: "is concerned that developers or operators could circumvent" Developer will push code direct to master, with bug that is not caught by the tests and here we go: you do have a signed imaged, attested, with a BUG or something like that will be deployed so the question depends on the point of view, the combination of both are right and one more thing, D prevent to push to the registry, but the operator can get the kubernetes deployment yaml and point to other GCR...
upvoted 1 times

WhyIronMan 3 years, 2 months ago
Selected Answer: A
If you are not familiarised, Pull Requests are one way to bring changes from one branch (e.g. develop) into protected branches (e.g. master, main).

1. First you need to protect the production branch (e.g, master, main) 2. If a developer attempts to push new code to a production (now protected) it will trow a Permission denied error like this: remote: Permission denied to update branch master. To git.com:org/repository.git ! [remote rejected] master -> master (pre-receive hook declined) error: failed to push some refs to 'git.com:org/repository.git'

3. in order to push their code to production branch, the developers will need to open a Pull Request (If you are using GitLab, it is called Merge Request) and ask someone to Review and Approve your changes. 4. The pipeline points to the protected branch; any new code pushed trigger the pipeline, and runs the tests and then deploy it.

This is how we do devops.

upvoted 1 times

Greg123123 3 years, 1 month ago

but you missed the point here "The security auditor is concerned that developers or operators could circumvent". I see why you think pull request is the only way to circumvent the deployment process. But how about the operators? They have access to cluster and can simply redeploy it by some kubectl / cloud build command. So A is not correct.

upvoted 2 times

Greg123123 3 years, 1 month ago

but you missed the point here "The security auditor is concerned that developers or operators could circumvent". I see why you think pull request is the only way to circumvent the deployment process. But how about the operators? They have access to cluster and can simply redeploy it by some kubectl / cloud build command. So A is not correct.That said, I send the question back to you, how about the developers?: "is concerned that developers or operators could circumvent" Developer will push code direct to master, with bug that is not caught by the tests and here we go: you do have a signed imaged, attested, with a BUG or something like that will be deployed so the question depends on the point of view, the combination of both are right and one more thing, D prevent to push to the registry, but the operator can get the kubernetes deployment yaml and point to other GCR...

upvoted 2 times

WhyIronMan 3 years, 1 month ago

That said, I send the question back to you, how about the developers?: "is concerned that developers or operators could circumvent" Developer will push code direct to master, with bug that is not caught by the tests and here we go: you do have a signed imaged, attested, with a BUG or something like that will be deployed so the question depends on the point of view, the combination of both are right and one more thing, D prevent to push to the registry, but the operator can get the kubernetes deployment yaml and point to other GCR...

upvoted 1 times

ramzez4815 3 years, 5 months ago

D is correct answer, binary auth is best practice

upvoted 3 times

lxs 3 years, 8 months ago

Correct me if I am wrong, but this question is ambiguous. You can push the code at 3 stages: 1. You can merge a branch to master without Merge Request if the master is not protected 2. You can push the image to container registry to a repository if you have role assigned (only pipeline should be privileged to do). 3. An operator can change the code altering image/yaml using kubectl cli. The ultimate question is which problem are we trying to solve?

upvoted 2 times

PhilipKoku 4 years ago

D - Binary authorisation

upvoted 3 times

DarkMatterOne 4 years, 1 month ago

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

upvoted 2 times

maddy94 4 years, 4 months ago

Questions 51-55 is not available.. can someone please help me to get 51-55 questions?

upvoted 3 times

zanhsieh 4 years, 7 months ago

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

upvoted 4 times

ralf_cc 4 years, 8 months ago

D - PR approval will ensure the automated testing etc., the question is asking how to ensure all code changes go through the pipeline, where automated tests are integrated

upvoted 1 times

akg001 4 years, 8 months ago

Answer C is the most close answer. Leverage best practice . answer A is for pulling the code but in the question , the security auditor is concern about pushing the code .

upvoted 1 times

akg001 4 years, 8 months ago

Answer C is the most close answer. Leverage best practice . answer A is for pulling the code but in the question , the security auditor is concern about pushing the code .If you are not familiarised, Pull Requests are one way to bring changes from one branch (e.g. develop) into protected branches (e.g. master, main). In order to push their code to production branch, the developers will need to open a Pull Request (If you are using GitLab, it is called Merge Request) and ask someone to Review and Approve their changes.

upvoted 1 times

WhyIronMan 3 years, 2 months ago

If you are not familiarised, Pull Requests are one way to bring changes from one branch (e.g. develop) into protected branches (e.g. master, main). In order to push their code to production branch, the developers will need to open a Pull Request (If you are using GitLab, it is called Merge Request) and ask someone to Review and Approve their changes.

upvoted 1 times

devopsbatch 4 years, 8 months ago

A could be the ans

upvoted 2 times

devopsbatch 4 years, 8 months ago

A could be the ansI think you are right : A https://github.community/t/best-practices-for-protected-branches/10204IMHO the statement is concert about runnigs unsafe workload on a cluster Kubernet, and not how to protect the code repository, in this other case, pull request approval and an Admission Controller would be fine. Then D is the ans

upvoted 2 times

akg001 4 years, 8 months ago

I think you are right : A https://github.community/t/best-practices-for-protected-branches/10204

upvoted 2 times

Feliphus Most Recent 2 years, 1 month ago

IMHO the statement is concert about runnigs unsafe workload on a cluster Kubernet, and not how to protect the code repository, in this other case, pull request approval and an Admission Controller would be fine. Then D is the ans

upvoted 1 times

Question #51 *Topic 1*

You support a stateless web-based API that is deployed on a single Compute Engine instance in the europe-west2-a zone. The Service Level Indicator (SLI) for service availability is below the specified Service Level Objective (SLO). A postmortem has revealed that requests to the API regularly time out. The time outs are due to the API having a high number of requests and running out memory. You want to improve service availability. What should you do?

A. Change the specified SLO to match the measured SLI

B. Move the service to higher-specification compute instances with more memory

C. Set up additional service instances in other zones and load balance the traffic between all instances

D. Set up additional service instances in other zones and use them as a failover in case the primary instance is unavailable

Suggested Answer: *C*
Official Answer: C

Community vote distribution
C (100%)

Alaaelanwr Highly Voted 3 years, 9 months ago
Ans: C
upvoted 12 times

jomonkp Most Recent 1 year, 8 months ago
Option C
upvoted 1 times

ramzez4815 2 years, 10 months ago
C is the correct anwer, it is required to increase reliability
upvoted 3 times

Ananda 3 years, 3 months ago
Submitted C
upvoted 2 times

Sekierer 3 years, 7 months ago
Ans: C
upvoted 1 times

muk5658 3 years, 8 months ago
C, this option will offer resilience and distribute the load, also provides ability to configure health checks at VM level and Load Balancer can send only to Healthy Instances.
upvoted 1 times

Manh 3 years, 9 months ago
ANSWER: C
upvoted 2 times

Question #52 *Topic 1*

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 in testing 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.

Suggested Answer: *D*
Official Answer: D

Community vote distribution
B (83%) D (17%)

NXD Highly Voted 4 years, 3 months ago
B https://cloud.google.com/vpc/docs/flow-logs#log-sampling
upvoted 16 times

Suraj2611 Highly Voted 4 years, 3 months ago
The Answer for this is BMost people rated D, sure that is better, but it will require re-deployment other than just changing the instance type. If a company doesn't have the capacity or resources to setup D, B is the quickly way to improve SLI. At the end of the day customer first, then you focus on how to do D or move things to serverless like Cloud Run
upvoted 6 times

Suraj2611 4 years, 3 months ago
The Answer for this is B
upvoted 6 times

4246cd7 Most Recent 2 years ago
Most people rated D, sure that is better, but it will require re-deployment other than just changing the instance type. If a company doesn't have the capacity or resources to setup D, B is the quickly way to improve SLI. At the end of the day customer first, then you focus on how to do D or move things to serverless like Cloud Run
upvoted 1 times

jomonkp 2 years, 2 months ago
Option B
upvoted 1 times

maxdanny 2 years, 3 months ago
B: enabling VPC Flow Logs on the production VPC network frontend and backend subnets only with a sample volume scale of 1.0. This means that all traffic will be logged, ensuring that all network traffic is captured for analysis.
upvoted 1 times

anshad666 2 years, 10 months ago
App Engine grants the Error Reporting Writer role by default. The Error Reporting library for Python can be used without needing to explicitly provide credentials. Error Reporting is automatically enabled for App Engine flexible environment applications. No additional setup is required
upvoted 1 times

hello_aws 3 years ago
Selected Answer: D
in real project background, first we should keep the consistency of configuration between test env and prod env. second, we should apply changes in testing before production.
upvoted 1 times

hello_aws 3 years ago

in real project background, first we should keep the consistency of configuration between test env and prod env. second, we should apply changes in testing before production.That is true, but the statement says the problem is only in the production frontend servers, maybe you are not going to find the malicious process in the testing frontend servers

upvoted 1 times

Feliphus 2 years, 1 month ago

That is true, but the statement says the problem is only in the production frontend servers, maybe you are not going to find the malicious process in the testing frontend servers

upvoted 1 times

WhyIronMan 3 years, 2 months ago

question saying test envs can be eliminated

upvoted 3 times

ramzez4815 3 years, 5 months ago

Correct answer is B

upvoted 2 times

Ananda 3 years, 9 months ago

Submitted B

upvoted 2 times

ric79 3 years, 11 months ago

B is better that A because you are filtering 1:20 instead of 1:10 and the malicious process generates very low traffic

upvoted 1 times

Sekierer 4 years, 1 month ago

Ans: B

upvoted 1 times

Alaaelanwr 4 years, 3 months ago

i think Answer: D

upvoted 2 times

Alaaelanwr 4 years, 3 months ago

i think Answer: Dif there isnt problem in test environment why log itQuestion clearly says only in Production they suspect, so we can eliminate the C and D options.

upvoted 2 times

giammydell 4 years, 3 months ago

if there isnt problem in test environment why log it

upvoted 4 times

muk5658 4 years, 2 months ago

Question clearly says only in Production they suspect, so we can eliminate the C and D options.

upvoted 3 times

Question #53 *Topic 1*

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

A. Store the Terraform code in a version-control system. Establish procedures for pushing new versions and merging with the master.

B. Store the Terraform code in a network shared folder with child folders for each version release. Ensure that everyone works on different files.

C. Store the Terraform code in a Cloud Storage bucket using object versioning. Give access to the bucket to every team member so they can download the files.

D. Store the Terraform code in a shared Google Drive folder so it syncs automatically to every team member's computer. Organize files with a naming convention that identifies each new version.

Suggested Answer: *A*

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

Community vote distribution

A (100%)

giammydell Highly Voted 4 years, 3 months ago

A for sure

upvoted 15 times

alpha_canary Most Recent 2 years ago

https://developer.hashicorp.com/terraform/cloud-docs/recommended-practices/part3.3#4-create-workspaces version-control systems like GitLab, GitHub, BitBucket, etc is the obvious choice, B, C, D are ridiculous

upvoted 3 times

jomonkp 2 years, 2 months ago

Option A

upvoted 2 times

pk1001 2 years, 3 months ago

version-control system like Git

upvoted 1 times

raghu09 2 years, 10 months ago

Version Control is the preferred way

upvoted 1 times

WhyIronMan 3 years, 2 months ago

A - Always use git

upvoted 1 times

meeku16 3 years, 5 months ago

For option C & D, downloading not at all recommended. For option B, hectic process So option A is the correct answer

upvoted 2 times

jagadeeshs 4 years ago

A is the correct Answer!

upvoted 1 times

Sekierer 4 years, 1 month ago

A for sure

upvoted 3 times

Suraj2611 4 years, 3 months ago

A is the guaranteed ans

upvoted 2 times

TNT87 4 years, 3 months ago

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

upvoted 4 times

job_search83 4 years, 3 months ago

A for sure

upvoted 3 times

Question #54 *Topic 1*

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.

Suggested Answer: *B*
Official Answer: B

Community vote distribution
A (77%) B (14%) D (9%)

KyubiBlaze Highly Voted 4 years, 1 month ago
Why do you complicate stuff. A is the answer. Why not D, because if A is not there D is useless. Question says you are using Stackdriver monitoring, not saying you have an agent installed already. You need the agent to export logs. So first thing you'll always see in the agent is there, and running. Next service account, next client libraries. I hope this clears your doubts.For "client libraries" you mean something like the fluent-logger-python ?You are assuming the new application is deployed on compute engine though there is no reference in question that says where the application is deployed.It never says that u are using a VM. What if the service is running on CloudRun? B is the answer
upvoted 13 times

KyubiBlaze Highly Voted 4 years, 1 month ago
Why do you complicate stuff. A is the answer. Why not D, because if A is not there D is useless. Question says you are using Stackdriver monitoring, not saying you have an agent installed already. You need the agent to export logs. So first thing you'll always see in the agent is there, and running. Next service account, next client libraries. I hope this clears your doubts.
upvoted 13 times

ric79 3 years, 11 months ago
For "client libraries" you mean something like the fluent-logger-python ?
upvoted 1 times

csrazdan 3 years, 4 months ago
You are assuming the new application is deployed on compute engine though there is no reference in question that says where the application is deployed.
upvoted 3 times

Meyucho 2 years, 2 months ago
It never says that u are using a VM. What if the service is running on CloudRun? B is the answer
upvoted 1 times

hfudjdn 4 years, 3 months ago
D - The question states that you are already using Stackdriver for the GCE instance (assuming the agent is already installed). However, the "new" application has issues shipping the logsThe question mention application logs, not only monitoring metrics, so I think "D" is correct. The stackdriver agent would only provide metrics of resources. I think you need to setup logging "client libraries" on the application in order to have the logs, therefore "D". https://cloud.google.com/logging/docs/reference/libraries
upvoted 6 times

hfudjdn 4 years, 3 months ago
D - The question states that you are already using Stackdriver for the GCE instance (assuming the agent is already installed). However, the "new" application has issues shipping the logs
upvoted 6 times

JohnnieWalker 4 years, 3 months ago
The question mention application logs, not only monitoring metrics, so I think "D" is correct. The stackdriver agent would only provide metrics of resources. I think you need to setup logging "client libraries" on the application in order to have the logs, therefore "D". https://cloud.google.com/logging/docs/reference/libraries

upvoted 4 times

muthucts Most Recent 1 year, 6 months ago

No specific service mentioned, so B is the best match for the question

upvoted 1 times

Meyucho 2 years, 2 months ago

It never says in which service is running so A does not consider all posibillities. B is the best match

upvoted 1 times

jomonkp 2 years, 2 months ago

Option A

upvoted 2 times

maxdanny 2 years, 3 months ago

the question asks about logs, only answer correct is A, the user already has access to the Stackdriver dashboard

upvoted 1 times

zanhsieh 2 years, 7 months ago

D. Stackdriver agent or Cloud Monitor agent only ships metrics, not log. If opt A change to Ops agent then it will be correct since Ops Agent handles both metrics and log. https://cloud.google.com/logging/docs/agent/ops-agent/configuration#default "The Ops Agent collects both metrics and logs by default. You can change this default behavior by configuring the Ops Agent." https://cloud.google.com/monitoring/agent/monitoring/installation#:~:text=The%20Ops%20Agent%20collects%20both%20metrics%20and%20logs%20by%20default.%20You%20can%20change%20this%20default%20behavior%20by%20configuring%20the%20Ops%20Agent.

upvoted 2 times

hanweiCN 3 years, 2 months ago

service account key does not have permission. service account has permission.

upvoted 3 times

WhyIronMan 3 years, 2 months ago

Key sentences for this question: 1. You recently deployed a NEW APPLICATION, 2. If it is a new application you CAN'T assume Stackdriver agent is already installed B) Wrong this is not a dashboard issue, even do, the first step is check if agent is installed C) Wrong, port 25 is used for SMTP server D) Wrong, if you are using GCP, VM, wth..

upvoted 3 times

juancambb 3 years, 2 months ago

I think a lot the best answer is B, because the question not suggest the application is running in compute engine, and if the user has dashboard viewer permissions but no logs viewer permissions the problem mentioned in the question can happen

upvoted 1 times

ssmb 3 years, 4 months ago

A should be the most likely correct answer.

upvoted 1 times

eliC 3 years, 8 months ago

A is the answer ref.: https://cloud.google.com/monitoring/agent/monitoring/troubleshooting#checklist

upvoted 3 times

ric79 3 years, 11 months ago

I think A Troubleshoot legacy monitoring - agent installed ed run: ps -ax | grep fluentd - logging test is visible? - VM API Access scope (ie Full Access scope or Logging Write for logging agent) - Service account (if not default is used) or credential key - Application Client library installed - Exclusion rule exists on Log Router? - Firewall rule?

upvoted 1 times

Sekierer 4 years, 1 month ago

A is correct

upvoted 3 times

TNT87 4 years, 1 month ago

The correct answer is A not D..D is mentioning one wrong thing thats why.
https://cloud.google.com/logging/docs/agent/logging/troubleshooting#try-installing

upvoted 1 times

not_thanos 4 years, 2 months ago

A and D seems correct. I think D is less correct because since you are hosting on GCP, best practice is to not add a service account to the application itself but instead leverage the native logging capabilities of GCP's offerings like GKE, GCP, etc.

upvoted 5 times

TNT87 4 years, 2 months ago

https://cloud.google.com/logging/docs/agent/logging/troubleshooting#verifying_default_service_account_permission Ans D

upvoted 2 times

Question #55 *Topic 1*

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

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

B. Configure the containers in the build pipeline to always update themselves before release.

C. Reconfigure the existing operating system vulnerability software to exist inside the container.

D. Implement static code analysis tooling against the Docker files used to create the containers.

Suggested Answer: *A*
Official Answer: A

Community vote distribution
A (100%)

giammydell Highly Voted 3 years, 9 months ago
A is correct for me
upvoted 10 times

jomonkp Most Recent 1 year, 8 months ago
Option A
upvoted 2 times

WhyIronMan 2 years, 8 months ago
A) for sure
upvoted 1 times

ssmb 2 years, 10 months ago
A is correct.
upvoted 2 times

ramzez4815 2 years, 10 months ago
Correct ans is A as per Google's best practices
upvoted 1 times

kapara 3 years ago
A - correct since this system would have alerted the issue after deployment in stating/dev and before deployment to production. - preventing issues in prod. https://cloud.google.com/container-analysis/docs/container-analysis B - Updates don't prevent vulnerabilities C - Not addressing the root cause. D - Static code analysis against a Docker file is useless since the code does not reside there.
upvoted 1 times

zygomar 3 years, 5 months ago
all comments so far agree that A is indeed the answer
upvoted 3 times

Sekierer 3 years, 7 months ago
A is correct
upvoted 1 times

Alaaelanwr 3 years, 9 months ago
Ans: A
upvoted 3 times

job_search83 3 years, 9 months ago
should be A
upvoted 3 times

Question #56 *Topic 1*

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

A. Use Cloud Storage to cache intermediate artifacts.

B. Run multiple Jenkins agents to parallelize the build.

C. Use multiple smaller build steps to minimize execution time.

D. Use larger Cloud Build virtual machines (VMs) by using the machine-type option.

Suggested Answer: *C*
Official Answer: C

Community vote distribution

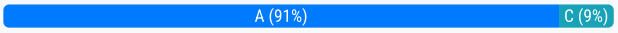

TNT87 Highly Voted 4 years, 3 months ago
Ans A https://cloud.google.com/storage/docs/best-practices
upvoted 14 times

bob_builder Highly Voted 4 years, 2 months ago
A - https://cloud.google.com/build/docs/speeding-up-builds#caching_directories_with_google_cloud_storage
upvoted 8 times

jomonkp 2 years, 2 months ago
Option A
upvoted 2 times

calex1 2 years, 5 months ago
Definitely A A. Use Cloud Storage to cache intermediate artifacts.
upvoted 1 times

jeffersonkozak 2 years, 8 months ago
https://cloud.google.com/blog/products/containers-kubernetes/kubernetes-best-practices-how-and-why-to-build-small-container-images
upvoted 1 times

Greg123123 3 years, 1 month ago
IMO this question / options could be reported incorrectly. For A, according to the doc, we can only store the resulting image in cloud storage, not intermediate layer of the image. To do so, we should use Kaniko cache instead. So A isn't correct. For D, it works but it increases the cost which violates the requirement from the question.
upvoted 1 times

Greg123123 3 years, 1 month ago
IMO this question / options could be reported incorrectly. For A, according to the doc, we can only store the resulting image in cloud storage, not intermediate layer of the image. To do so, we should use Kaniko cache instead. So A isn't correct. For D, it works but it increases the cost which violates the requirement from the question.I wonder if ans C is the correct, you can find the same answer: Use multiple smaller build steps to minimize execution time in Q-126 as well.
upvoted 1 times

Feliphus Most Recent 2 years, 1 month ago
I wonder if ans C is the correct, you can find the same answer: Use multiple smaller build steps to minimize execution time in Q-126 as well.
upvoted 1 times

ssmb 3 years, 4 months ago
Most likely A is correct, even after reading the suggested document at: https://cloud.google.com/build/docs/speeding-up-builds
upvoted 1 times

Ananda 3 years, 9 months ago
Submitted A
upvoted 2 times

TNT87 4 years, 1 month ago
Ans A and thats cheap and fast

upvoted 3 times

TNT87 4 years, 1 month ago

minimising costs...high virtual machines are expensive

upvoted 1 times

TNT87 4 years, 2 months ago

lets correct something fellows ok. A is not ok, why, we can ok only use intermediate artifacts in Kaniko cache, not in Cloud storage, in cloud storage we use results not intermediate artifacts "You can copy the results of a previous build to a Google Cloud Storage bucket, use the results for faster calculation, and then copy the new results back to the bucket." go re-read the best practices and you will see what im talking about, then this leaves D as the correct answer. https://cloud.google.com/build/docs/speeding-up-buildsif you question and point out the error, don't just give your answer so you help the discussionsorry not see ansif you question and point out the error, don't just give your answer so you help the discussionA is correct choice. Why D is not correct, because even if you break steps into small chunks, you are just increasing the number of steps, and not solving the actual problem of reducing execution time. If step 1 takes 5min it will take 5min, if you break it down into three steps, combined Step 1+2+3 = 5min again, and they don't run in parallel. What did you solve? Nothing. If you cache ie. store intermediate artefacts in cloud storage, you do not need to recompile it in next steps, saves time and cost.

upvoted 3 times

TNT87 4 years, 2 months ago

lets correct something fellows ok. A is not ok, why, we can ok only use intermediate artifacts in Kaniko cache, not in Cloud storage, in cloud storage we use results not intermediate artifacts "You can copy the results of a previous build to a Google Cloud Storage bucket, use the results for faster calculation, and then copy the new results back to the bucket." go re-read the best practices and you will see what im talking about, then this leaves D as the correct answer. https://cloud.google.com/build/docs/speeding-up-builds

upvoted 3 times

leandrojpg 3 years, 2 months ago

if you question and point out the error, don't just give your answer so you help the discussionsorry not see ansif you question and point out the error, don't just give your answer so you help the discussion

upvoted 1 times

leandrojpg 3 years, 2 months ago

sorry not see ans

upvoted 1 times

leandrojpg 3 years, 2 months ago

if you question and point out the error, don't just give your answer so you help the discussion

upvoted 1 times

KyubiBlaze 4 years, 1 month ago

A is correct choice. Why D is not correct, because even if you break steps into small chunks, you are just increasing the number of steps, and not solving the actual problem of reducing execution time. If step 1 takes 5min it will take 5min, if you break it down into three steps, combined Step 1+2+3 = 5min again, and they don't run in parallel. What did you solve? Nothing. If you cache ie. store intermediate artefacts in cloud storage, you do not need to recompile it in next steps, saves time and cost.

upvoted 4 times

ESP_SAP 4 years, 2 months ago

Correct Answer is (D): Using custom virtual machine sizes In addition to the standard machine type, Cloud Build provides four high-CPU virtual machine types to run your builds. To increase the speed of your build, select a virtual machine with a higher CPU. Requesting a high-CPU machine may increase the startup time of your build as Cloud Build only starts these machines on demand. https://cloud.google.com/build/docs/speeding-up-buildsin as much as D is correct A is also mentioned inthe Best practicesYou are correct sir....

upvoted 6 times

ESP_SAP 4 years, 2 months ago

Correct Answer is (D):

Using custom virtual machine sizes In addition to the standard machine type, Cloud Build provides four high-CPU virtual machine types to run your builds. To increase the speed of your build, select a virtual machine with a higher CPU. Requesting a high-CPU machine may increase the startup time of your build as Cloud Build only starts these machines on demand.

https://cloud.google.com/build/docs/speeding-up-builds

upvoted 6 times

TNT87 4 years, 2 months ago

in as much as D is correct A is also mentioned inthe Best practices

upvoted 1 times

TNT87 4 years, 2 months ago

You are correct sir....

upvoted 1 times

Tire 4 years, 2 months ago

Caching directories with Google Cloud Storage To increase the speed of a build, reuse the results from a previous build. You can copy the results of a previous build to a Google Cloud Storage bucket, use the results for faster calculation, and then copy the new results back to the bucket. Use this method when your build takes a long time and produces a small number of files that does not take time to copy to and from Google Cloud Storage.results not intermediate artifacts, ok

upvoted 2 times

Tire 4 years, 2 months ago

Selected Answer: A

Caching directories with Google Cloud Storage To increase the speed of a build, reuse the results from a previous build. You can copy the results of a previous build to a Google Cloud Storage bucket, use the results for faster calculation, and then copy the new results back to the bucket. Use this method when your build takes a long time and produces a small number of files that does not take time to copy to and from Google Cloud Storage.

upvoted 2 times

TNT87 4 years, 2 months ago

results not intermediate artifacts, ok

upvoted 1 times

muk5658 4 years, 2 months ago

Both A and D are valid options, but if we look in the question they are asking to Minimize Cost - Larger Compute Instances comes at price. So I will go with A (Caching from GCS bucket)

upvoted 3 times

Manh 4 years, 3 months ago

Ans is A

upvoted 1 times

Shasha1 4 years, 3 months ago

A https://cloud.google.com/build/docs/speeding-up-builds

upvoted 3 times

MBA_1 4 years, 3 months ago

A is not correct . From the Link shared -https://cloud.google.com/build/docs/speeding-up-builds - Cloud Build is to store previous builds and not intermediate. For Intermediate the ask is to use Kaniko.

If you scroll down. - you can find -- "To increase the speed of your build, select a virtual machine with a higher CPU. Requesting a high-CPU machine may increase the startup time of your build as Cloud Build only starts these machines on demand."

And hence C is the answer

upvoted 1 times

MBA_1 4 years, 3 months ago

A is not correct . From the Link shared -https://cloud.google.com/build/docs/speeding-up-builds - Cloud Build is to store previous builds and not intermediate. For Intermediate the ask is to use Kaniko. If you scroll down. - you can find -- "To increase the speed of your build, select a virtual machine with a higher CPU. Requesting a high-CPU machine may increase the startup time of your build as Cloud Build only starts these machines on demand." And hence C is the answerboth A and D seems correctbut the question is asking for "You want to reduce the build time while minimizing cost and development effort" -> only A is best optionSorry, I meant D is the answer

upvoted 1 times

Manh 4 years, 3 months ago

both A and D seems correctbut the question is asking for "You want to reduce the build time while minimizing cost and development effort" -> only A is best option

upvoted 1 times

Manh 4 years, 3 months ago

both A and D seems correct

upvoted 1 times

Manh 4 years, 3 months ago

but the question is asking for "You want to reduce the build time while minimizing cost and development effort" -> only A is best option

upvoted 2 times

MBA_1 4 years, 3 months ago

Sorry, I meant D is the answer

upvoted 1 times

Question #57 *Topic 1*

You support a web application that is hosted on Compute Engine. The application provides a booking service for thousands of users. Shortly after the release of a new feature, your monitoring dashboard shows that all users are experiencing latency at login. You want to mitigate the impact of the incident on the users of your service. What should you do first?

A. Roll back the recent release.

B. Review the Stackdriver monitoring.

C. Upsize the virtual machines running the login services.

D. Deploy a new release to see whether it fixes the problem.

Suggested Answer: *C*

Official Answer: C

Community vote distribution

A (100%)

PJ_2021 Highly Voted 4 years, 3 months ago

A Rollback is needed to mitigate the impact. Once the is done review can be done

upvoted 12 times

Meyucho Most Recent 2 years, 2 months ago

MITIGATE is the word here... increasing resources will no garantize the mitigation. Rolling back will do this

upvoted 2 times

jomonkp 2 years, 2 months ago

Option A

upvoted 1 times

ssmb 3 years, 4 months ago

A is correct.

upvoted 2 times

ramzez4815 3 years, 5 months ago

A is the correct answer, to properly mitigate the issue a rollback is required

upvoted 1 times

PhilipKoku 4 years ago

A - Rollback to previous stable version. Then you need to find what is causing the issue.

upvoted 2 times

Sekierer 4 years, 1 month ago

A is correct

upvoted 2 times

Manh 4 years, 3 months ago

A Rollback first

upvoted 1 times

MBA_1 4 years, 3 months ago

B - Find whats wrong with the system and mitigate. As the Q says only login is affected, So no point is rollback

upvoted 1 times

MBA_1 4 years, 3 months ago

B - Find whats wrong with the system and mitigate. As the Q says only login is affected, So no point is rollbackyou still use service you cannot login?

https://cloud.google.com/architecture/identifying-causes-of-app-latency-with-stackdriver-and-opencensus

upvoted 1 times

MF2C 4 years, 3 months ago

you still use service you cannot login?

upvoted 1 times

MBA_1 4 years, 3 months ago

https://cloud.google.com/architecture/identifying-causes-of-app-latency-with-stackdriver-and-opencensus

upvoted 1 times

giammydell 4 years, 3 months ago

But the question is about mitigating the problema, with the rollback the system goes back to the old version.

upvoted 3 times

giammydell 4 years, 3 months ago

But the question is about mitigating the problema, with the rollback the system goes back to the old version.roll back is the answer

upvoted 3 times

TNT87 4 years, 3 months ago

roll back is the answer

upvoted 2 times

Nik22 4 years, 3 months ago

First roll back and then use monitoring to review what went wrong.

upvoted 4 times

job_search83 4 years, 3 months ago

A rollback

upvoted 4 times

Question #58 *Topic 1*

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.

Suggested Answer: *A*
Official Answer: A

Community vote distribution
A (100%)

giammydell Highly Voted 3 years, 9 months ago
Ans: A
upvoted 11 times

jomonkp Most Recent 1 year, 8 months ago
Option A
upvoted 1 times

WhyIronMan 2 years, 8 months ago
A) should be the correct answer.
upvoted 1 times

ssmb 2 years, 10 months ago
A should be the correct answer.
upvoted 2 times

ramzez4815 2 years, 10 months ago
The clear answer is A
upvoted 1 times

emdee202 3 years, 3 months ago
Ans: A
upvoted 1 times

Sekierer 3 years, 7 months ago
A is correct
upvoted 1 times

TNT87 3 years, 9 months ago
https://cloud.google.com/security-key-management
upvoted 1 times

TNT87 3 years, 9 months ago
Ans is A.
upvoted 2 times

Alaaelanwr 3 years, 9 months ago
Ans: A
upvoted 3 times

Question #59 *Topic 1*

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.

Suggested Answer: *A*
Official Answer: A

Community vote distribution
A (100%)

AL12 Highly Voted 4 years, 3 months ago
I reckon its A, the reason is because it seems like the problem is automatically fixed with an restart of the service after a minute, therefore engineers don't really need to be woken up about these problems. If it failed multiple times or if the restart failed, then the engineer should be woken upA or C
upvoted 14 times

AL12 Highly Voted 4 years, 3 months ago
I reckon its A, the reason is because it seems like the problem is automatically fixed with an restart of the service after a minute, therefore engineers don't really need to be woken up about these problems. If it failed multiple times or if the restart failed, then the engineer should be woken up
upvoted 14 times

MF2C 4 years, 3 months ago
A or C
upvoted 1 times

09bd94b Most Recent 1 year, 6 months ago
Agree with A. It does not make sense to wake up an engineer when you know that there is no need for any remedy action
upvoted 1 times

Greg123123 3 years, 1 month ago
It should be A rather than D. To follow SRE practice, we should eliminate unactionable alert which is pointless and to increase precision. While D also looks valid, the question never say that the application is being affected (e.g. has downtime), and never says any actions are needed. As a result, there is no need to redefine SLI and since they didn't spend time to resolve it no error budget is spent.
upvoted 2 times

ssmb 3 years, 4 months ago
Between A and C, B and D answers are not good. I lean more towards A because those alerts seem unactionable a the moment alert is received, ie: machine restarted automatically already. This would be best imidiate action as per the question. Of course the source of alerts should be looked at and fixed separately from addressing the issue in question.
upvoted 2 times

zygomar 3 years, 12 months ago
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
upvoted 4 times

Sekierer 4 years, 1 month ago
A is correct
upvoted 1 times

KyubiBlaze 4 years, 1 month ago
A - You have to remove "unactionable" alerts, these alerts are useless if you can't take any action. Simple reason, C might be following SRE practice, but it is distributing the problem, not solving it. B and D, totally No.
upvoted 3 times

gcpz 4 years, 1 month ago

answer is c. it follows google SRE and prevents staff burnout. https://sre.google/workbook/team-lifecycles/

upvoted 1 times

ESP_SAP 4 years, 2 months ago

The team may continue to work on non-reliability features if: The outage was caused by a company-wide networking problem. The outage was caused by a service maintained by another team, who have themselves frozen releases to address their reliability issues. The error budget was consumed by users out of scope for the SLO (e.g., load tests or penetration testers). Miscategorized errors consume budget even though no users were impacted. https://sre.google/workbook/error-budget-policy/Correct Answer is (D):

upvoted 3 times

ESP_SAP 4 years, 2 months ago

The team may continue to work on non-reliability features if:

The outage was caused by a company-wide networking problem. The outage was caused by a service maintained by another team, who have themselves frozen releases to address their reliability issues. The error budget was consumed by users out of scope for the SLO (e.g., load tests or penetration testers). Miscategorized errors consume budget even though no users were impacted.

https://sre.google/workbook/error-budget-policy/

upvoted 3 times

ESP_SAP 4 years, 2 months ago

Correct Answer is (D):

upvoted 2 times

Manh 4 years, 3 months ago

Answer D

upvoted 1 times

NXD 4 years, 3 months ago

C follows the SRE.

upvoted 3 times

NXD 4 years, 3 months ago

C follows the SRE.The statemene says: you encounter a large number of outages in the production systems you support, then eliminating the alerts doesn't seem to be a good idea. If there is another support team in another time zone. What's happen if the server doesn't reboot or the services don't start fine?. There is not a correct answer between options, what it would be to resolve the reboot problem. I don't know which is better if A or C, I suppose we have losed some information in the statement or in the answers. But in this situation I agree @NXD and choose CSorry, but I change to ans A. I have noticed this question is repeated as Q133 but without the text: You receive alerts for all the outages that wake you up at night

upvoted 3 times

Feliphus 2 years, 1 month ago

The statemene says: you encounter a large number of outages in the production systems you support, then eliminating the alerts doesn't seem to be a good idea. If there is another support team in another time zone. What's happen if the server doesn't reboot or the services don't start fine?. There is not a correct answer between options, what it would be to resolve the reboot problem. I don't know which is better if A or C, I suppose we have losed some information in the statement or in the answers. But in this situation I agree @NXD and choose CSorry, but I change to ans A. I have noticed this question is repeated as Q133 but without the text: You receive alerts for all the outages that wake you up at night

upvoted 1 times

Feliphus 2 years, 1 month ago

The statemene says: you encounter a large number of outages in the production systems you support, then eliminating the alerts doesn't seem to be a good idea. If there is another support team in another time zone. What's happen if the server doesn't reboot or the services don't start fine?. There is not a correct answer between options, what it would be to resolve the reboot problem. I don't know which is better if A or C, I suppose we have losed some information in the statement or in the answers. But in this situation I agree @NXD and choose C

upvoted 1 times

Feliphus 2 years, 1 month ago

Sorry, but I change to ans A. I have noticed this question is repeated as Q133 but without the text: You receive alerts for all the outages that wake you up at night

upvoted 1 times

TNT87 4 years, 3 months ago

Ans D https://www.atlassian.com/incident-management/kpis/error-budget

upvoted 4 times

TNT87 4 years, 3 months ago

Ans D https://www.atlassian.com/incident-management/kpis/error-budgetAns A...point of correctionNO!D is correct

upvoted 4 times

TNT87 4 years, 1 month ago

Ans A...point of correction

upvoted 1 times

TNT87 4 years, 1 month ago

Ans A...point of correctionNO!D is correct

upvoted 1 times

TNT87 4 years, 1 month ago

NO!D is correct

upvoted 1 times

neutrino9 4 years, 3 months ago

Should be A

upvoted 3 times

job_search83 4 years, 3 months ago

D redefine SLI

upvoted 1 times

Question #60 *Topic 1*

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. Load teat the application to profile its performance for scaling.

B. Enable AutoScaling on the production clusters, in case there is growth.

C. Pre-provision double the compute power used last season, expecting growth.

D. Create a runbook on inflating the disaster recovery (DR) environment if there is growth.

Suggested Answer: *B*
Official Answer: B

Community vote distribution
A (90%) B (10%)

NXD Highly Voted 3 years, 9 months ago
https://cloud.google.com/architecture/black-friday-production-readiness#preparation_stage The objective of the preparation stage is to test the system's ability to scale for peak user traffic and to document the results. Completing the preparation stage results in architecture refinement to handle peak traffic more efficiently and increase system reliability. This stage also yields procedures for operations and support that help streamline processes for handling the peak event and any issues that might occur. Consider this stage as practice for the peak event from a system and operations perspective. A is exactly what mentioned above. B is the step after the preparation stage.Changing the architecture for scale and reliability Load testing and failure testing, along with architecture reviews, encourage limited-scope architectural changes that can enhance the scale and reliability of the system. However, introducing changes adds risk, so limit the changes to a conservative range of time.https://cloud.google.com/architecture/black-friday-production-readiness#changing_the_architecture_for_scale_and_reliability
upvoted 18 times

NXD Highly Voted 3 years, 9 months ago
https://cloud.google.com/architecture/black-friday-production-readiness#preparation_stage

The objective of the preparation stage is to test the system's ability to scale for peak user traffic and to document the results. Completing the preparation stage results in architecture refinement to handle peak traffic more efficiently and increase system reliability. This stage also yields procedures for operations and support that help streamline processes for handling the peak event and any issues that might occur. Consider this stage as practice for the peak event from a system and operations perspective.

A is exactly what mentioned above. B is the step after the preparation stage.
upvoted 18 times

TNT87 3 years, 8 months ago
Changing the architecture for scale and reliability Load testing and failure testing, along with architecture reviews, encourage limited-scope architectural changes that can enhance the scale and reliability of the system. However, introducing changes adds risk, so limit the changes to a conservative range of time.https://cloud.google.com/architecture/black-friday-production-readiness#changing_the_architecture_for_scale_and_reliability
upvoted 1 times

KyubiBlaze 3 years, 7 months ago
Come on, no brainer. A is the answer. You load test to understand how your application perform under heavy load. "Prepare for busy season" B - No, Option A will give you insight in how your applications works under load, and how do you scale, if it cannot scale, autoscaling in meaningless. So first you test your application in controlled environment. Not wait for the busy time to come and then realise autoscaling is also unable to meet demand. Or Maybe you even reach your quotas.Good analysis, before moving any new features first thing to do performance test to understand how system behaves under load, Auto scaling is later step... A is 100% correct
upvoted 5 times

KyubiBlaze 3 years, 7 months ago
Selected Answer: A
Come on, no brainer. A is the answer. You load test to understand how your application perform under heavy load. "Prepare for busy season" B - No, Option A will give you insight in how your applications works under load, and how do you scale, if it cannot scale, autoscaling in meaningless. So first

you test your application in controlled environment. Not wait for the busy time to come and then realise autoscaling is also unable to meet demand. Or Maybe you even reach your quotas.

upvoted 5 times

AzureDP900 2 years, 9 months ago

Good analysis, before moving any new features first thing to do performance test to understand how system behaves under load, Auto scaling is later step... A is 100% correct

upvoted 1 times

jomonkp Most Recent 1 year, 8 months ago

Option A

upvoted 1 times

raghu09 2 years, 4 months ago

Would go with A because you need to know the performance of your application if it scales to serve larger customers. Scaling decisions of the application comes after.

upvoted 1 times

ssmb 2 years, 10 months ago

After moving the application onto GCP, some amount of testing should be carried out before deciding what to do next.

upvoted 1 times

Sekierer 3 years, 7 months ago

A is correct

upvoted 2 times

TNT87 3 years, 7 months ago

Set up load and performance testing Load testing is the process of deploying a test version of the system and creating requests to simulate high use of the system. Load testing normally focuses on testing for sustainable user-perceived behavior at some percentile below the absolute peak. Testing for peak requires hitting that top percentile with consistent good performance.

upvoted 2 times

TNT87 3 years, 7 months ago

B is the answerSorry guys this is wrongFor Auto Scaling you need to provide parameters which will not be know if load test is not done. Otherwise you will be using arbitrary values which may not be consistent with actual load.

upvoted 1 times

TNT87 3 years, 7 months ago

Selected Answer: B

B is the answer

upvoted 1 times

TNT87 3 years, 7 months ago

Sorry guys this is wrong

upvoted 1 times

csrazdan 2 years, 10 months ago

For Auto Scaling you need to provide parameters which will not be know if load test is not done. Otherwise you will be using arbitrary values which may not be consistent with actual load.

upvoted 1 times

TNT87 3 years, 8 months ago

https://cloud.google.com/blog/topics/retail/preparing-for-peak-holiday-season-while-wfh Ans A makes sense according to this doc

upvoted 1 times

tee_dee26 3 years, 9 months ago

Passed exam a couple of days ago. Chose A

upvoted 4 times

MBA_1 3 years, 9 months ago

Should be A , To check how application performs with load

upvoted 4 times

MBA_1 3 years, 9 months ago

Should be A , To check how application performs with loadAgree with A

upvoted 4 times

Manh 3 years, 9 months ago

Agree with A

upvoted 1 times

Alaaelanwr 3 years, 9 months ago

Ans : B

upvoted 2 times

giammydell 3 years, 9 months ago

Ans: B

upvoted 2 times

TNT87 3 years, 9 months ago

Ans B https://cloud.google.com/architecture/black-friday-production-readiness

upvoted 2 times

Question #61 *Topic 1*

You support a web application that runs on App Engine and uses CloudSQL and Cloud Storage for data storage. After a short spike in website traffic, you notice a big increase in latency for all user requests, increase in CPU use, and the number of processes running the application. Initial troubleshooting reveals:

➩ After the initial spike in traffic, load levels returned to normal but users still experience high latency.

➩ Requests for content from the CloudSQL database and images from Cloud Storage show the same high latency.

➩ No changes were made to the website around the time the latency increased.

➩ There is no increase in the number of errors to the users.

You expect another spike in website traffic in the coming days and want to make sure users don't experience latency. What should you do?

A. Upgrade the GCS buckets to Multi-Regional.

B. Enable high availability on the CloudSQL instances.

C. Move the application from App Engine to Compute Engine.

D. Modify the App Engine configuration to have additional idle instances.

Suggested Answer: *B*

Official Answer: B

Community vote distribution

D (100%)

ESP_SAP Highly Voted 4 years, 2 months ago

Correct Answer is D: Scaling App Engine scales the number of instances automatically in response to processing volume. This scaling factors in the automatic_scaling settings that are provided on a per-version basis in the configuration file. A service with basic scaling is configured by setting the maximum number of instances in the max_instances parameter of the basic_scaling setting. The number of live instances scales with the processing volume. You configure the number of instances of each version in that service's configuration file. The number of instances usually corresponds to the size of a dataset being held in memory or the desired throughput for offline work. You can adjust the number of instances of a manually-scaled version very quickly, without stopping instances that are currently running, using the Modules API set_num_instances function.
https://cloud.google.com/appengine/docs/standard/python/how-instances-are-managed

upvoted 10 times

MBA_1 Highly Voted 4 years, 3 months ago

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

upvoted 6 times

MBA_1 4 years, 3 months ago

D is correct - https://cloud.google.com/appengine/docs/standard/python/config/appref max_idle_instances Optional. The maximum number of idle instances that App Engine should maintain for this version. Specify a value from 1 to 1000. If not specified, the default value is automatic, which means App Engine will manage the number of idle instances. Keep the following in mind:

A high maximum reduces the number of idle instances more gradually when load levels return to normal after a spike. This helps your application maintain steady performance through fluctuations in request load, but also raises the number of idle instances (and consequent running costs) during such periods of heavy load.

upvoted 6 times

Manh 4 years, 3 months ago

Agree with D

upvoted 1 times

alpha_canary Most Recent 2 years ago

https://cloud.google.com/appengine/docs/legacy/standard/python/config/appref#scaling_elements:~:text=traffic%20or%20not.-,max_idle_instances,-Optional.%20The%20maximum

upvoted 1 times

jomonkp 2 years, 2 months ago

Option D

upvoted 1 times

Greg123123 3 years, 1 month ago

D, no brainer. A B and C don't even make sense.A. then how about CloudSQL? B. then how about cloud storage? C. the same could also happen and even worse if it isn't MIG

upvoted 1 times

Greg123123 3 years, 1 month ago

Selected Answer: D

D, no brainer. A B and C don't even make sense.

upvoted 1 times

Greg123123 3 years, 1 month ago

A. then how about CloudSQL? B. then how about cloud storage? C. the same could also happen and even worse if it isn't MIG

upvoted 1 times

ramzez4815 3 years, 4 months ago

Vote for D

upvoted 2 times

Sekierer 4 years, 1 month ago

D is correct

upvoted 4 times

TNT87 4 years, 1 month ago

D is the answer

upvoted 2 times

Biden 4 years, 3 months ago

D looks like the nearest correct solution though this will not address with high latency with CloudSQL & storage instances

upvoted 1 times

NXD 4 years, 3 months ago

D is correct to me. A: wrong – because increase in latency for all user requests, not for specific region. So B: wrong – problem affects both database and GCS C: wrong increase in CPU use, and the number of processes running the application => problem is that CPU is not enough to run application processes.

upvoted 3 times

giammydell 4 years, 3 months ago

B could be useful but does not solve latency problem on Cloud Storage side

upvoted 3 times

TNT87 4 years, 3 months ago

Ans B https://cloud.google.com/sql/docs/mysql/high-availabilityAns D,B has latency issues

upvoted 2 times

TNT87 4 years, 3 months ago

Ans B https://cloud.google.com/sql/docs/mysql/high-availability

upvoted 2 times

TNT87 4 years, 2 months ago

Ans D,B has latency issues

upvoted 1 times

Nik22 4 years, 3 months ago

A. Upgrade the GCS buckets to Multi-Regional. B. Enable high availability on the CloudSQL instances. C. Move the application from App Engine to Compute Engine. D. Modify the App Engine configuration to have additional idle instances. There is nothing in the question that talks about region. If I assume that bucket is created in the region where the app is deployed B - Gives redundancy. I don't see how this can improve the latency C - Does not make any sense D - Additional instance, would reduce the CPU Time & memory. I would go with D. But not too sure.

upvoted 2 times

Question #62 *Topic 1*

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

A. Implement Jenkins on local workstations.

B. Implement Jenkins on Kubernetes on-premises.

C. Implement Jenkins on Google Cloud Functions.

D. Implement Jenkins on Compute Engine virtual machines.

Suggested Answer: *D*

References: https://plugins.jenkins.io/google-compute-engine/

Community vote distribution

D (100%)

enter_co Most Recent 1 year, 9 months ago

Deploying Jenkins on GKE allows reduction of toil and simplified management, because Cloud IAM will allow Jenkins access to various GCP services through GCP-managed permission model.

upvoted 1 times

jomonkp 2 years, 8 months ago

option D

upvoted 2 times

PhilipKoku 4 years, 6 months ago

D - Jenkins with Compute Engine in GCP.

upvoted 4 times

Sekierer 4 years, 7 months ago

D is correct

upvoted 4 times

muk5658 4 years, 8 months ago

I am sorry ignore my previous comment, they are saying Kubernetes on Premise, not GKE (Google Managed), so I would go with GCE.

upvoted 2 times

muk5658 4 years, 8 months ago

I believe its B - GKE since they are saying reducing the toil and secure. We can use GKE secrets and also enabling autoscaling.

Correct me if my thinking is wrong

upvoted 1 times

muk5658 4 years, 8 months ago

I believe its B - GKE since they are saying reducing the toil and secure. We can use GKE secrets and also enabling autoscaling. Correct me if my thinking is wrongOption B says deploying on Kubernetes on premises not on GKE.

upvoted 1 times

raghu09 3 years, 4 months ago

Option B says deploying on Kubernetes on premises not on GKE.

upvoted 1 times

TNT87 4 years, 9 months ago

he Google Compute Engine (GCE) Plugin allows you to use GCE virtual machines (VMs) with Jenkins to execute build tasks. GCE VMs provision quickly, are destroyed by Jenkins when idle, and offer Preemptible VMs that run at a much lower price than regular VMs.

upvoted 4 times

TNT87 4 years, 9 months ago

he Google Compute Engine (GCE) Plugin allows you to use GCE virtual machines (VMs) with Jenkins to execute build tasks. GCE VMs provision quickly, are destroyed by Jenkins when idle, and offer Preemptible VMs that run at a much lower price than regular VMs.yes, Jenkins only got 2 options for GCP as of now, GKE or compute engine

upvoted 4 times

MF2C 4 years, 9 months ago

yes, Jenkins only got 2 options for GCP as of now, GKE or compute engine

upvoted 1 times

TNT87 4 years, 9 months ago

Ans D https://plugins.jenkins.io/google-compute-engine/

upvoted 3 times

job_search83 4 years, 9 months ago

D, GCE

upvoted 3 times

Question #63 *Topic 1*

You are working with a government agency that requires you to archive application logs for seven years. You need to configure Stackdriver to export and store the logs while minimizing costs of storage. What should you do?

A. Create a Cloud Storage bucket and develop your application to send logs directly to the bucket.

B. Develop an App Engine application that pulls the logs from Stackdriver and saves them in BigQuery.

C. Create an export in Stackdriver and configure Cloud Pub/Sub to store logs in permanent storage for seven years.

D. Create a sink in Stackdriver, name it, create a bucket on Cloud Storage for storing archived logs, and then select the bucket as the log export destination.

Suggested Answer: *D*

References: https://jayendrapatil.com/google-cloud-logging/

Community vote distribution

D (100%)

Manh Highly Voted 3 years, 9 months ago

Answer D https://cloud.google.com/logging/docs/routing/overview

upvoted 11 times

job_search83 Highly Voted 3 years, 9 months ago

D, sink

upvoted 8 times

jomonkp Most Recent 1 year, 8 months ago

Option D

upvoted 1 times

ramzez4815 2 years, 10 months ago

D is correct, it's actually explaining the steps to do what the question ask.

upvoted 1 times

ramzez4815 2 years, 11 months ago

D is correct as per Googles best practices

upvoted 1 times

emdee202 3 years, 3 months ago

Ans: D

upvoted 3 times

Sekierer 3 years, 7 months ago

D is correct

upvoted 3 times

giammydell 3 years, 9 months ago

ans: D

upvoted 4 times

Alaaelanwr 3 years, 9 months ago

Ans : D

upvoted 3 times

Question #64 *Topic 1*

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

A. Install the Stackdriver Error Reporting library for Python, and then run your code on a Compute Engine VM.

B. Install the Stackdriver Error Reporting library for Python, and then run your code on Google Kubernetes Engine.

C. Install the Stackdriver Error Reporting library for Python, and then run your code on App Engine flexible environment.

D. Use the Stackdriver Error Reporting API to write errors from your application to ReportedErrorEvent, and then generate log entries with properly formatted error messages in Stackdriver Logging.

Suggested Answer: *C*
References: https://cloud.google.com/error-reporting/docs/setup/app-engine-flexible-environment

Community vote distribution
D (71%) C (29%)

muk5658 Highly Voted 4 years, 2 months ago
@TNT87 - Answer is C, the link you shared has a pip install in the beginning which means Python requires installing library first. pip install --upgrade google-cloud-error-reporting
upvoted 21 times

muk5658 Highly Voted 4 years, 2 months ago
@TNT87 - Answer is C, the link you shared has a pip install in the beginning which means Python requires installing library first. pip install --upgrade google-cloud-error-reportingAns C https://cloud.google.com/error-reporting/docs/setup/python i think you can see that 1 month ago answered it ok, the one with answer D was a mistakeApp Engine flexible environment App Engine grants the Error Reporting Writer role by default. The Error Reporting library for Python can be used without needing to explicitly provide credentials. Error Reporting is automatically enabled for App Engine flexible environment applications. No additional setup is required. Yes anser is C https://cloud.google.com/error-reporting/docs/setup/python#app-engine
upvoted 21 times

TNT87 4 years, 2 months ago
Ans C https://cloud.google.com/error-reporting/docs/setup/python i think you can see that 1 month ago answered it ok, the one with answer D was a mistakeApp Engine flexible environment App Engine grants the Error Reporting Writer role by default. The Error Reporting library for Python can be used without needing to explicitly provide credentials. Error Reporting is automatically enabled for App Engine flexible environment applications. No additional setup is required. Yes anser is C https://cloud.google.com/error-reporting/docs/setup/python#app-engine
upvoted 1 times

TNT87 4 years, 2 months ago
Ans C https://cloud.google.com/error-reporting/docs/setup/python i think you can see that 1 month ago answered it ok, the one with answer D was a mistake
upvoted 1 times

TNT87 4 years, 2 months ago
App Engine flexible environment App Engine grants the Error Reporting Writer role by default. The Error Reporting library for Python can be used without needing to explicitly provide credentials. Error Reporting is automatically enabled for App Engine flexible environment applications. No additional setup is required. Yes anser is C https://cloud.google.com/error-reporting/docs/setup/python#app-engine
upvoted 3 times

gaurav517 Most Recent 8 months, 2 weeks ago
Option C: Install Error Reporting library, run on App Engine flexible □ Keeps current deployment on App Engine flexible □ Error Reporting library allows customization of error data □ Can add custom fields, context, user information □ Minimal changes - just add library to existing app □ Works seamlessly with App Engine flexible □ Simple and direct solution □ No infrastructure changes needed
upvoted 1 times

Min3em 1 year, 7 months ago
https://cloud.google.com/error-reporting/docs/setup/python#app-engine:~:text=Error%20Reporting%20is%20automatically%20enabled%20for%20App%20Engine%20flexible%20environment%20applications.%20No%2

0additional%20setup%20is%20required. "Error Reporting is automatically enabled for App Engine flexible environment applications. No additional setup is required."

upvoted 1 times

tell3r 1 year, 9 months ago

ReportedErrorEvent is for formatting, not customization

upvoted 1 times

jinaldesailive 1 year, 11 months ago

C is the answer. As it's simplest to configure with minimal effort.

upvoted 1 times

alpha_canary 2 years ago

wy not D? While it's possible to use the Stackdriver Error Reporting API, it's easier and more straightforward to use the library designed specifically for the language of the application (Python, in this case). https://cloud.google.com/error-reporting/docs/setup/python#install_the_client_library

upvoted 1 times

jomonkp 2 years, 2 months ago

Option D

upvoted 2 times

anshad666 2 years, 10 months ago

App Engine grants the Error Reporting Writer role by default. The Error Reporting library for Python can be used without needing to explicitly provide credentials.Error Reporting is automatically enabled for App Engine flexible environment applications. No additional setup is required

upvoted 3 times

WhyIronMan 3 years, 2 months ago

C) for sure As per Error Reporting setup instructions for App Engine Flexible: https://cloud.google.com/error-reporting/docs/setup/app-engine-flexible-environment "If you'd like to customize the error information being sent to Error Reporting, you can use the instrumentation libraries available for a number of languages" Which includes a step for installing Python client libraries https://cloud.google.com/error-reporting/docs/setup/python#installing_the_client_library

upvoted 2 times

DoodleDo 3 years, 2 months ago

Ans C There are two things asked here. One is Error Reporting enabled for App Engine Flex. It is based on this source - Ans D Error Reporting Auto Enabled for App Engine Flex - Source - https://cloud.google.com/error-reporting/docs/setup/app-engine-flexible-environment Second the question asks about customize error (not formatting the error). If you'd like to customize the error information being sent to Error Reporting, you can use the instrumentation libraries available for a number of languages: Source - https://cloud.google.com/error-reporting/docs/setup/app-engine-flexible-environment For formatting the error then use option D - Source - https://cloud.google.com/error-reporting/docs/formatting-error-messages

upvoted 3 times

enter_co 3 years, 3 months ago

You want to customize the error information being sent to Stackdriver Error Reporting. As per Error Reporting setup instructions for App Engine Flexible: https://cloud.google.com/error-reporting/docs/setup/app-engine-flexible-environment "If you'd like to customize the error information being sent to Error Reporting, you can use the instrumentation libraries available for a number of languages" Which includes a step for installing Python client libraries https://cloud.google.com/error-reporting/docs/setup/python#installing_the_client_library

upvoted 2 times

mudot 3 years, 3 months ago

Appeard in 7/11/2022 exam keys: you need to customise the message + it is already on appEngine

upvoted 3 times

mudot 3 years, 3 months ago

isnt the question already says it is running on app engine flexible ?

upvoted 2 times

pradoUA 3 years, 3 months ago

Selected Answer: D

The question ask: "You want to customize the error information being sent", hence for me the answer is D

upvoted 4 times

pradoUA 3 years, 3 months ago

The question ask: "You want to customize the error information being sent", hence for me the answer is DExactly that is the keyword triggers me D as right choice

upvoted 4 times

AzureDP900 3 years, 3 months ago

Exactly that is the keyword triggers me D as right choice

upvoted 1 times

iamzeah 3 years, 6 months ago

If you're using the Error Reporting API, you can report error events from your application by writing them to ReportedErrorEvent. Doing this generates log entries with properly formatted error messages in Cloud Logging. https://cloud.google.com/error-reporting/docs/formatting-error-messages

upvoted 4 times

nhadi82 3 years, 7 months ago

If you're using the Error Reporting API, you can report error events from your application by writing them to ReportedErrorEvent. Doing this generates log entries with properly formatted error messages in Cloud Logging. The resulting logName is formatted as follows:

upvoted 3 times

eliC 3 years, 8 months ago

D is correct. There is no need to setup library Refs.: (1) https://cloud.google.com/error-reporting/docs/setup/app-engine-flexible-environment (2) https://cloud.google.com/error-reporting/docs/setup/python#app-engine

upvoted 2 times

Question #65 *Topic 1*

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
90
percentile of latency is 120ms and the 95
percentile of latency is 275ms over a 28-day window. What latency SLO would you recommend to the team to th th publish?

A. 90 percentile 100 "€ɹms th 95 percentile 250 "€ɹms th

B. 90 percentile 120 "€ɹms th 95 percentile 275 "€ɹms th

C. 90 percentile 150 "€ɹms th 95 percentile 300 "€ɹms th

D. 90 percentile 250 "€ɹms th 95 percentile 400 "€ɹms th

Suggested Answer: *B*
Official Answer: B

Community vote distribution
C (54%) B (46%)

job_search83 Highly Voted 4 years, 9 months ago
C, https://sre.google/sre-book/service-level-objectives/ "Don't pick a target based on current performance"
upvoted 25 times

Nik22 Highly Voted 4 years, 9 months ago
Picking an SLO based upon current performance can commit you to unnecessarily strict SLOs. Select slightly lower SLO. Will go wtih C
upvoted 11 times

blazt Most Recent 1 year, 11 months ago
option B, why setting lower SLO !
upvoted 1 times

jomonkp 2 years, 8 months ago
Option C
upvoted 1 times

maxdanny 2 years, 9 months ago
https://www.exam-answer.com/slos-for-high-traffic-web-application
upvoted 1 times

izekc 3 years, 5 months ago
pick the SLO a bit higher then the current value would be better
upvoted 1 times

JayDeng 3 years, 8 months ago
C is correct answer according to https://sre.google/workbook/implementing-slos/#using-the-slis-to-calculate-starter-slos "Using the SLIs to Calculate Starter SLOs We can round down these SLIs to manageable numbers (e.g., two significant figures of availability, or up to 50 ms5 of latency) to obtain our starting SLOs."
upvoted 1 times

hanweiCN 3 years, 8 months ago
from SRE book: start from a loosen SLO that u can tighten better than to choose an overly strict SLO that has to be relax after u discover it is unattainable.https://sre.google/sre-book/service-level-objectives/We can round down these SLIs to manageable numbers (e.g., two significant figures of availability, or up to 50 ms5 of latency) to obtain our starting SLOs. https://sre.google/workbook/implementing-slos/#continuous-improvement-of-slo-targets
upvoted 1 times

hanweiCN 3 years, 8 months ago
Selected Answer: C
from SRE book: start from a loosen SLO that u can tighten better than to choose an overly strict SLO that has to be relax after u discover it is unattainable.https://sre.google/sre-book/service-level-objectives/

upvoted 1 times

hanweiCN 3 years, 8 months ago

We can round down these SLIs to manageable numbers (e.g., two significant figures of availability, or up to 50 ms5 of latency) to obtain our starting SLOs. https://sre.google/workbook/implementing-slos/#continuous-improvement-of-slo-targets

upvoted 1 times

WhyIronMan 3 years, 8 months ago

going with B)

upvoted 1 times

DoodleDo 3 years, 8 months ago

Historical trend dictates customers are happy at the current performance level.

upvoted 1 times

enter_co 3 years, 9 months ago

According to this material, if the users are happy with the current not that stellar performance, the SLOs should be achievable, based on the historical values: https://cloud.google.com/blog/products/management-tools/practical-guide-to-setting-slos "It's important to set a target that is achievable so that alerts are meaningful. Normally, when choosing an SLO, it's best to start from historical trends and assume if enough people are happy with the service now, you're probably doing OK. Eventually, it's ideal to converge those numbers with aspirational targets that your business may want you to meet." Even though the SRE book says to avoid picking target based on the current performance https://sre.google/sre-book/service-level-objectives/ , this is in a particular context: extremely stable and high performing apps with high cost.

upvoted 2 times

joec 4 years ago

Agree with C

upvoted 2 times

buldas 4 years, 6 months ago

SLI != SLO

upvoted 2 times

Sekierer 4 years, 7 months ago

C is correct

upvoted 2 times

Manh 4 years, 9 months ago

Agree with answer C

upvoted 1 times

giammydell 4 years, 9 months ago

SLO is selected well if User are happy. B is correct

upvoted 2 times

Question #66 *Topic 1*

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

A. Develop an appropriate error budget policy in cooperation with all service stakeholders.

B. Negotiate with the product team to always prioritize service reliability over releasing new features.

C. Negotiate with the development team to reduce the release frequency to no more than once a week.

D. Add a plugin to your Jenkins pipeline that prevents new releases whenever your service is out of SLO.

Suggested Answer: *B*
Official Answer: B

Community vote distribution
A (100%)

KyubiBlaze Highly Voted 4 years, 1 month ago
Reason : Incident has not occurred yet, even when development team is already pushing new features multiple times a week. The option A says, to define an error budget "policy", not to define error budget(It is already present). Just simple means to bring in all stakeholders, and decide how to consume the error budget effectively that could bring balance between feature deployment and reliability.Good explanation
upvoted 12 times

KyubiBlaze Highly Voted 4 years, 1 month ago
Selected Answer: A
Reason : Incident has not occurred yet, even when development team is already pushing new features multiple times a week. The option A says, to define an error budget "policy", not to define error budget(It is already present). Just simple means to bring in all stakeholders, and decide how to consume the error budget effectively that could bring balance between feature deployment and reliability.
upvoted 12 times

AzureDP900 3 years, 3 months ago
Good explanation
upvoted 1 times

jomonkp 2 years, 2 months ago
Option A
upvoted 1 times

WhyIronMan 3 years, 2 months ago
going with A)
upvoted 1 times

mudot 3 years, 3 months ago
SLO Miss Policy If the service is performing at or above its SLO, then releases (including data changes) will proceed according to the release policy. If the service has exceeded its error budget for the preceding four-week window, we will halt all changes and releases other than P01 issues or security fixes until the service is back within its SLO. Depending upon the cause of the SLO miss, the team may devote additional resources to working on reliability instead of feature work.
upvoted 1 times

eliC 3 years, 8 months ago
I vote for A
upvoted 2 times

Sekierer 4 years, 1 month ago
I vote for A
upvoted 3 times

smoaz 4 years, 1 month ago
Correct answer: A

upvoted 1 times

TNT87 4 years, 1 month ago

the reason why A is the answer we want something that is in relation to A....Service reliability, and who needs that service to be reliable its the customer, remember as Devops engneer you must put yourself inthe shoes of a customer.

upvoted 2 times

TNT87 4 years, 1 month ago

thats the enswer

upvoted 2 times

not_thanos 4 years, 2 months ago

I don't think the answer should be A. Error Budget = 100% - SLO. Since the SLO is "well defined", ideally the Error Budget is going to already be well defined as well.

upvoted 2 times

Biden 4 years, 2 months ago

Correct Answer - A Error Budget Definition: A "Quantitative" measurement shared between SRE & product teams to balance innovation and stability. Hence first step is defining this "quantitative" error budget(Answer A). Only when there is risk of exceeding error budget then negotiate with product team (stakeholder) to prioritize reliability over features (option B)

upvoted 2 times

ESP_SAP 4 years, 2 months ago

Correct Answer is (B): When developing for your system, think of reliability as a feature, and in fact your most important feature, as customers simply expect that things 'just work'. Once you have your objectives, reliability engineering provides practices to help you reach them, including: • Redundancy systems: Such as contingencies for using backup servers • Fault tolerance: Such as error correction algorithms for incoming network data • Preventative maintenance: Such as cycling through hardware resources before failure through overuse • Human error prevention: Such as cleaning and validating human input into the system • Reliability optimization: Such as writing code optimized for quick and consistent loading Keeping these reliability practices in mind as you develop will make your code acceptably reliable. At the same time, you're able to confidently accelerate development by evaluating it against SLOs.

upvoted 1 times

TNT87 4 years, 2 months ago

Ans B https://www.fullstory.com/blog/bootstrapping-an-availability-program-using-sre-principles/ Lets not forget the principles of availbility & reliabilty

upvoted 1 times

TNT87 4 years, 2 months ago

Ans B https://www.fullstory.com/blog/bootstrapping-an-availability-program-using-sre-principles/ Lets not forget the principles of availbility & reliabiltyAns B. To define how is an error budget policy or how the error budget is being consumed is not going to avoid a major incident. The question here is: What should you do before a major incident occurs?. Nevertheless, to relase new features during a week is a bit heavy, the service realibitiy is being forgotten by that team.

upvoted 1 times

Feliphus Most Recent 2 years, 1 month ago

Ans B. To define how is an error budget policy or how the error budget is being consumed is not going to avoid a major incident. The question here is: What should you do before a major incident occurs?. Nevertheless, to relase new features during a week is a bit heavy, the service realibitiy is being forgotten by that team.

upvoted 1 times

Manh 4 years, 3 months ago

Should be A

upvoted 3 times

MF2C 4 years, 3 months ago

the key is "shift its focus from working on features to improving service reliability" which means to reduce risky deployment and the answer should be B

upvoted 2 times

sindra 4 years, 3 months ago

I think it should be A

upvoted 3 times

giammydell 4 years, 3 months ago

Could be A but also D: the request is preventing deployment of new release when SLO is missingNot D for sure.. D is preventing deployment if the service is running out of SLO already.. but the question here is to prevent before SLO miss..

upvoted 2 times

giammydell 4 years, 3 months ago

Could be A but also D: the request is preventing deployment of new release when SLO is missing

upvoted 2 times

maddy94 3 years, 11 months ago

Not D for sure.. D is preventing deployment if the service is running out of SLO already.. but the question here is to prevent before SLO miss..

upvoted 1 times

Question #67 *Topic 1*

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.

Suggested Answer: *D*
Reference: https://kubernetes.io/docs/reference/access-authn-authz/rbac/

Community vote distribution
D (100%)

TNT87 Highly Voted 3 years, 9 months ago
https://cloud.google.com/architecture/prep-kubernetes-engine-for-prod#roles_and_groups Ans D
upvoted 16 times

TNT87 Highly Voted 3 years, 9 months ago
https://cloud.google.com/architecture/prep-kubernetes-engine-for-prod#roles_and_groups Ans DI think the en version is missing for some reason, translate page: https://cloud.google.com/architecture/prep-kubernetes-engine-for-prod?hl=fr#roles_and_groups
upvoted 16 times

mburuian6 2 years, 7 months ago
I think the en version is missing for some reason, translate page: https://cloud.google.com/architecture/prep-kubernetes-engine-for-prod?hl=fr#roles_and_groups
upvoted 1 times

jomonkp Most Recent 1 year, 8 months ago
Option D
upvoted 2 times

CloudKida 1 year, 11 months ago
GKE has two access control systems: Identity and Access Management (IAM) and role-based access control (RBAC). IAM is Google Cloud's access control system for managing authentication and authorization for Google Cloud resources. You use IAM to grant users access to GKE and Kubernetes resources. RBAC is built into Kubernetes and grants granular permissions for specific resources and operations within your clusters.
upvoted 3 times

jacknguyen 1 year, 11 months ago
D is correct
upvoted 1 times

mudot 2 years, 9 months ago
Appeard in 7/11/2022 exam keys : least no of clusters + separation
upvoted 1 times

AzureDP900 2 years, 9 months ago
D is correct
upvoted 1 times

ssmb 2 years, 10 months ago
I vote for D as the correct answer, considering cost reduction is stated in the question.
upvoted 1 times

PhilipKoku 3 years, 6 months ago

D - Different project for Prod and UAT. RBAC to access each app team GKE area.

upvoted 3 times

Sekierer 3 years, 7 months ago

D is correct

upvoted 2 times

guid1984 3 years, 8 months ago

The answer is D (enterprise multi-tenancy using GKE, reduce costs and control access by RBAC)

upvoted 2 times

muk5658 3 years, 8 months ago

D -100%

upvoted 2 times

Manh 3 years, 9 months ago

Sound like B is good ans

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.

1 project and 1 cluster per team with 2 namespace

upvoted 2 times

Manh 3 years, 9 months ago

Sound like B is good ans 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. 1 project and 1 cluster per team with 2 namespaceRe-read the question again....and we cant mix development and production project NO!Mate you have one project per app per environment. You explaination clearly violates this basic principle

upvoted 2 times

TNT87 3 years, 8 months ago

Re-read the question again....and we cant mix development and production project NO!

upvoted 1 times

PJ_2021 3 years, 9 months ago

Mate you have one project per app per environment. You explaination clearly violates this basic principle

upvoted 2 times

NXD 3 years, 9 months ago

I go with A.

upvoted 2 times

NXD 3 years, 9 months ago

I go with A.but you have to minimize the costwrong.I will go with A since PROD and DEV/TEST environments need to kept separate always. Having them in the same cluster, may impact the resources for PROD when there is a need to do a load testing spinning off multiple PODs and nodeshttps://www.google.com/search?q=For+almost+all+cases%2C+Kubernetes+RBAC+can+be+used+instead+of+IAM.+GKE+users+require+at+minimum%2C+the+container.clusters.get+IAM+permission+in+the+project+...&oq=For+almost+all+cases%2C+Kubernetes+RBAC+can+be+used+instead+of+IAM.+GKE+users+require+at+minimum%2C+the+container.clusters.get+IAM+permission+in+the+project+...&aqs=chrome..69i57.479j0j7&sourceid=chrome&ie=UTF-8

upvoted 2 times

giammydell 3 years, 9 months ago

but you have to minimize the cost

upvoted 1 times

gcpz 3 years, 7 months ago

wrong.

upvoted 1 times

Biden 3 years, 9 months ago

I will go with A since PROD and DEV/TEST environments need to kept separate always. Having them in the same cluster, may impact the resources for PROD when there is a need to do a load testing spinning off multiple PODs and nodeshttps://www.google.com/search?q=For+almost+all+cases%2C+Kubernetes+RBAC+can+be+used+instead+of+IAM.+GKE+users+require+at+minimum%2C+the+container.clusters.get+IAM+permission+in+the+project+...&oq=For+almost+all+cases%2C+Kubernetes+RBAC+can+be+used+instead+of+IAM.+GKE+users+require+at+minimum%2C+the+container.clusters.get+IAM+permission+in+the+project+...&aqs=chrome..69i57.479j0j7&sourceid=chrome&ie=UTF-8

upvoted 2 times

Biden 3 years, 9 months ago

I will go with A since PROD and DEV/TEST environments need to kept separate always. Having them in the same cluster, may impact the resources for PROD when there is a need to do a load testing spinning off multiple PODs and nodes

upvoted 2 times

TNT87 3 years, 7 months ago

https://www.google.com/search?q=For+almost+all+cases%2C+Kubernetes+RBAC+can+be+used+instead+of+IAM.+GKE+users+require+at+minimum%2C+the+container.clusters.get+IAM+permission+in+the+project+...&oq=For+almost+all+cases%2C+Kubernetes+RBAC+can+be+used+instead+of+IAM.+GKE+users+require+at+minimum%2C+the+container.clusters.get+IAM+permission+in+the+project+...&aqs=chrome..69i57.479j0j7&sourceid=chrome&ie=UTF-8

upvoted 1 times

Alaaelanwr 3 years, 9 months ago

Ans: D

upvoted 3 times

Nik22 3 years, 9 months ago

Will go with D

upvoted 3 times

Question #68 *Topic 1*

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.

Suggested Answer: *B*

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

Community vote distribution

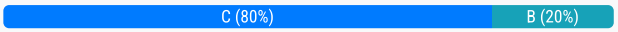

Nik22 Highly Voted 4 years, 9 months ago

To maximize the bandwidth for transferring the images to the cluster, one needs the the registry to be closer to the production or system where it needs to be deployed. I would go with C, since the production system is in Europe.

upvoted 18 times

Meyucho 2 years, 8 months ago

It says PUSH!!! The building system is in US so the best option to push is US.Nope, you forgot about 'maximize the bandwidth' for transferring the images to the cluster, and the best option is allocate imagens in the same cluster region ;)

upvoted 1 times

Meyucho 2 years, 8 months ago

Selected Answer: B

It says PUSH!!! The building system is in US so the best option to push is US.

upvoted 1 times

lucs1080p Most Recent 2 years ago

Nope, you forgot about 'maximize the bandwidth' for transferring the images to the cluster, and the best option is allocate imagens in the same cluster region ;)

upvoted 1 times

jomonkp 2 years, 8 months ago

Option C

upvoted 1 times

jacknguyen 2 years, 11 months ago

C is the answer.

upvoted 1 times

mudot 3 years, 9 months ago

Appeard in 7/11/2022 exam the answers were different in the exam. they are like a) us-docker.pkg.dev , b)europe-docker.pkg.dev c)europe-west1-docker.dev.pkg 4)private registry. answer : europe-docker.pkg.devc)europe-west1-docker.dev.pkg looks correctI think you should answer C), because it exists the specific $LOCATION look at -> https://cloud.google.com/artifact-registry/docs/repositories/repo-locations As well, you can find some examples how to use it -> https://cloud.google.com/artifact-registry/docs/docker/copy-images

upvoted 4 times

mudot 3 years, 9 months ago

Appeard in 7/11/2022 exam

the answers were different in the exam. they are like a) us-docker.pkg.dev , b)europe-docker.pkg.dev c)europe-west1-docker.dev.pkg 4)private registry. answer : europe-docker.pkg.dev

upvoted 4 times

adasdsa 3 years, 8 months ago

c)europe-west1-docker.dev.pkg looks correct

upvoted 2 times

Feliphus 2 years, 8 months ago

I think you should answer C), because it exists the specific $LOCATION look at -> https://cloud.google.com/artifact-registry/docs/repositories/repo-locations As well, you can find some examples how to use it -> https://cloud.google.com/artifact-registry/docs/docker/copy-images

upvoted 1 times

F1fans 3 years, 10 months ago

when i see this question,the option aren't those.The four option are eu.xxx,one of those, is eu-west-1.xxx,please notice.

upvoted 1 times

sid0127 3 years, 10 months ago

c is the ans

upvoted 1 times

Rupo7 4 years ago

It seems possible to choose the GCR region (I didn't know) https://cloud.google.com/container-registry/docs/pushing-and-pulling, therefore to maximize bandwidth to the the cluster (in europe) we must choose C (eu.gcr.io).

upvoted 1 times

buldas 4 years, 6 months ago

Same region.

upvoted 2 times

Sekierer 4 years, 7 months ago

C is correct

upvoted 2 times

Manh 4 years, 9 months ago

C for sure

upvoted 2 times

MBA_1 4 years, 9 months ago

It's C - because of GDPR compliance. 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

upvoted 4 times

giammydell 4 years, 9 months ago

C could be correct, but also D is nearer because in the same zone. A and B are are both in US

upvoted 3 times

Question #69 *Topic 1*

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

A. In the Google Cloud Platform Console, use the Cost Breakdown section to visualize the costs per system.

B. Assign all instances a label specific to the system they run. Configure BigQuery billing export and query costs per label.

C. Enrich all instances with metadata specific to the system they run. Configure Stackdriver Logging to export to BigQuery, and query costs based on the metadata.

D. Name each virtual machine (VM) after the system it runs. Set up a usage report export to a Cloud Storage bucket. Configure the bucket as a source in BigQuery to query costs based on VM name.

Suggested Answer: *D*

Reference: https://cloud.google.com/compute/docs/logging/usage-export

Community vote distribution

B (100%)

TNT87 Highly Voted 3 years, 9 months ago

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

upvoted 14 times

jomonkp Most Recent 1 year, 8 months ago

option B

upvoted 2 times

maxdanny 1 year, 9 months ago

B is correct, it is the fastest solution and allows you to group by label

upvoted 2 times

jacknguyen 1 year, 11 months ago

B is correct

upvoted 1 times

ramzez4815 2 years, 11 months ago

B is the correct answer, labes are Google's best practices to break down costs per units

upvoted 2 times

Sekierer 3 years, 7 months ago

B is correct

upvoted 2 times

not_thanos 3 years, 8 months ago

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

upvoted 1 times

Manh 3 years, 9 months ago

B. U NEED TO ADD LABELS TO INSTANCE FOR COST visibility

upvoted 4 times

job_search83 3 years, 9 months ago

B, labels

upvoted 3 times